UNIVERSITÀ DEGLI STUDI
DI TRENTO

DEPARTMENT OF INFORMATION ENGINEERING AND COMPUTER SCIENCE
ICT International Doctoral School

ENERGY NEUTRAL DESIGN OF EMBEDDED SYSTEMS FOR RESOURCE CONSTRAINED MONITORING APPLICATIONS

Maurizio Rossi

Advisor

Prof. Davide Brunelli

Università degli Studi di Trento

March 2016

The greatest threat to our planet is the belief
that someone else will save it.

Robert Swan

Abstract

Automatic monitoring of environments, resouces and human processes are crucial and foundamental tasks to improve people's quality of life and to safeguard the natural environment. Today, new technologies give us the possibility to shape a greener and safer future. The more specialized is the kind of monitoring we want to achieve, more tight are the constraints in terms of reliability, low energy and maintenance-free autonomy. The challenge in case of tight energy constraints is to find new techniques to save as much power as possible or to retrieve it from the very same environment where the system operates, towards the realization of energy neutral embedded monitoring systems.

Energy efficiency and battery autonomy of such devices are still the major problem impacting reliability and penetration of such systems in risk-related activities of our daily life. Energy management must not be optimized to the detriment of the quality of monitoring and sensors can not be operated without supply. In this thesis, I present different embedded system designs to bridge this gap, both from the hardware and software sides, considering specific resource constrained scenarios as case studies that have been used to develop solutions with much broader validity. Results achieved demonstrate that energy neutrality in monitoring under resource constrained conditions can be obtained without compromising efficiency and reliability of the outcomes.

Keywords

Embedded electronic systems, wireless sensor networks, environmental monitoring, energy harvesting, hardware software codesign.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Nowadays, technology advances are shaping and improving people's quality of life to the detriment of natural environment which is exploited in an uncontrolled way to provide energy and resources to sustain the consumer society. In this sense, ozone layer depletion, negative effect of pollution on human health and renewable energy sources integration in the electricity Grid, — to reduce the dependancy on fossil fuels — are nowadays issues of great relevance. Public awareness toward these topics has been fostered by scientists and no-profit organizations and it is now gaining a *real* attention from governments and regulators as it is demonstrated by newly issued laws and regulations or the organization of meetings like the COP21 [56] world summit.

New technologies give us today the capability to tackle the situation and move towards a greener and safer future. Consequently, the deployment of pervasive and unobtrusive monitoring networks in *any* field of human activity should be fostered. Knowledge comes from information and, in the digital era, we can rely on autonomous systems to achieve a capillary coverage, mitigating at the same time, costs and efforts to achieve the goal.

Environmental monitoring is the process of collecting etherogeneous information to evaluate, assess and control the harmful effects of processes, systems and devices to people and the natural environment. Under this broad umbrella falls a number of activities that span from the collection of meteorological data like temperature, humidity and light intensity, to the acquisition and characterization of chemical emissions like, for example, the by-products of industrial processes and the exhaust gases of transportation means.

The more specialized is the kind of monitoring we want to achieve, the higher the costs and the efforts to be done, in particular because of the need of costly and complex instruments that can be used only by trained people. Too often accidents and failures results in injuries and death to people because security systems are not working properly or not even present in homes and working environments; for example, in year 2014, only in Italy 126 accidents were registered related with Natural Gas distribution failures [19]. Other examples can be easily found in the number of fires or power outages, but the point is that, clearly, there is a huge gap between the need for monitoring instruments and their affordability and spread.

The research in embedded systems targeted to environmental monitoring tasks aims at developing cheap devices, which are affordable for everyone, to replace complex and expensive laboratory instrumentation. To this end, the use of modern low-cost sensors and devices available in the market should be preferred. Wireless nodes are widely used to collect environmental parameters using a number of low-cost and low-power sensors that can last for years embedded in battery-powered devices.

"Wireless Sensor Networks (WSN) are distributed systems typically composed of embedded devices, each equipped with a processing unit, a wireless communication interface, as well as sensors and/or actuators" [54]. They usually have small dimensions, and are portable and easily reconfigurable. Designed to build unobtrusive monitoring systems in any kind of condition — indoor, outdoor and harsh environments — their key features are the ease, the speed and the low-cost of development (with respect to wired devices). Portable electronic boards have very-low power consumption, since they are active for few hundreds of milliseconds to perform the measure, while the remaining time is spent in low power sleep state where all the peripherals can be turned-off. Standard peripherals included in all designs are temperature, humidity and light intensity sensors because of the cheap price and the very-low power required to work. On the contrary, more specialized sensors like gas sensors, accelerometers and gyroscopes, are less commonly included in wireless node with respect the previous ones, due to the higher power consumption, the demanding signal processing techniques required to filter the noise and the design efforts for the integration in embedded systems.

Today, WSNs are considered a "mature" technology, tens of examples and case studies can be found in the literature; on the contrary we are at the dawning of the Internet of Things (IoT) era, and the potentiality offered by this technology are countless. However, we are still facing the same issues that were limiting the widespread of distributed embedded systems in the past: energy efficiency and battery autonomy are still the major concerns impacting reliability and penetration in risk-related activities of our daily life.

Research and development in the embedded system world is currently driven by automotive, healtcare and the growing IoT sectors, for a total market share that is envisioned to be worth more than 214 Billion $ by year 2020 [26]. At the same time also gas sensor and batteries markets are expected to grow exponentially, the first one up to a production of 350 million units in year 2021 from the actual 1.2 million, thanks to their introduction in future generations of smartphones [66]; the latter one, because of the proliferation of new technologies that will be necessary to address the existing and future challenges that our society will require [35]. This is a huge opportunity for the electronic research community to develop new devices which foster a greener energy management, a sustainable utilization of resources and safer living conditions for people.

1.1 The Need

Whether it is a home weather station or a radiation detector, digital measurement instrumentation require energy to work. Portable embedded systems usually rely on disposable batteries as primary supply source, whose lifetime strictly depends on the set of sensors and monitoring tasks to perform. High end instrumentations instead rely on recharge-

able batteries, if they are occasionally used, or directly connected to the electricity Grid if they have to provide security services (like smoke detectors or surveillance cameras). As demonstrated by the market analysis on batteries research and development [35], many companies are working (and billion $ will be invested) to bridge the gap with what is requested by modern IoT ready devices (smartphones, computers and, in the near future, electric cars). In the same study, it is envisioned that the early results of such investments will be ready for the market in more that ten years, thus the question arise, what can we do in the meanwhile?

In the development of this work I tackled the problem of designing energy management techniques and devices to enable *today* monitoring applications in resource constrained environments. In such situations, requirements are case specific and no affordable solution exists yet to meet the reliablity and security of Grid powered respective conunterparts.

Specific constraints exist for many scenarios that must be taken into account to deploy a reliable system. For example, ordinary people are not willing to replace disposable batteries very often (less than two years) for domestic air quality monitoring equipments, while few months is the expected lifetime of current deployments. Conversely, in the case of controlled environments (laboratories, machine rooms or data centers), the presence of people should be avoided or reduced at most for security and safety reasons, thus it would be preferable to have energy neutral devices that do not use batteries at all.

Particularly, in the case of air quality monitoring systems, the main problem is represented by the on-board presence of volatile chemicals sensing devices (gas sensors). As it will be shown, these devices require a large amount of power to work thus their impact on monitoring networks' lifetime is quite remarkable and gas sensors are commonly used only in Grid powered smoke detectors or dedicated measurement instruments.

In the case of controlled environments like server rooms of data centers, the use of specific sensors allows to efficiently monitor many interesting information, such as: local room temperature, humidity, air flow, air quality, moisture leak and intrusion detection. Given the numbers involved in the data centers scenario, even a small reduction of electricity Grid dependancy lead to substantial environmental and economic benefits.

Energy management in this context is a multivariate optimization problem in which any μJ of energy counts. The challenge for the designer in case of tight energy constraints is to find innovative ways to save as much power as possible or to retrieve it from the surrounding environment, towards the realization of energy neutral embedded systems.

1.2 The Solution

By definition, energy neutral operations characterize those embedded systems in which all the supply energy is retrieved from the very same environment where the system operates. Thus, energy scavenging devices, conditioning circuits and proper storage systems must be included in the design. To achieve energy neutral operations, an embedded system must implement both hardware and software solutions that realize the best trade-off between minimization of energy used and reliability of the monitoring.

As described above, monitoring in resource constrained environments introduces a series of tight requirements for which strict energy neutrality features can not be com-

pletely fullfilled in all applications, but significant improvements to the state-of-the-art can be achieved. This thesis presents the results achieved focusing on real case studies, highlighting intrinsic limits and challenges of current technologies, to demonstrate the relevance of proposed approaches. As it will be shown, these results can have a much broader impact than the test cases analyzed.

In particular three case studies were identified to worth the effort for investigation:

- Volatile pollutant monitoring in indoor environments;
- Gas leackage detection in outdoor environments;
- Security and safety assessment for IT service rooms.

On the monitoring of volatile pollutant side, an innovative methodology for immediate environmental safety assessment has been developed and demonstrated. The technique exploits a characteristic features of low-cost commercially-available gas sensors to immediately react to volatile chemicals, allowing timely detection of harmful situations like gas leakages. It is based on the analysis of the frequency response of gas sensors and has been optimized for the implementation and usage with mobile embedded systems, that allows to achieve a reduction of 10 times the energy required by these sensors with respect to the standard utilization strategy. For this reason it has been defined Ulta Low-Power (ULP) technique. Both gas detection and energy reduction performance have been demonstrated using a Wireless Gas Sensor Network (WGSN) deployed in a real indoor environment. Moreover, data collected with the experimental setup demonstrate that the expected lifetime of such distributed monitoring system can be extended up to 3 years, which is significantly longer than what has been reported in the literature for sensor networks with similar features.

The main advantage of this technique resides in the evaluation of the harmonic content of the Discrete Cosine Transform (DCT) applied to a very short interval of gas sensor samples (corresponding to ≈ 500 ms), thus it can be potentially implemented and rapidly computed by almost any microcontroller based embedded system, enabling air quality monitoring capability in a broad range of applications.

The experience in volatile chemicals detection gained in the development of the ULP technique has been also applied to the outdoor monitoring scenario by means of mobile robots, in particular Unmanned Aerial Vehicles (UAVs). The outdoor scenario presents completely different requirements and constraints, in fact static and single point of measure installations are not suitable for monitoring a rapidly varying dynamic environment, and opeator-based campaigns are costly and time-consuming. Mobile robots instead can alleviate the burden for operators by periodically conduct automatic measurement campaigns. In this sense, it has been developed a battery powered embedded system with gas sensor and GPS receiver in a compact board (5×5 cm) which can be installed almost on any kind of robot to work as a monitoring peripheral (thanks to usb and local logging) or as a real-time gas leak localization device (thanks to the ZigBee and GSM/GPRS auxiliary wireless interfaces). By means of field experiments using a professional UAV (80 cm diameter), it has been demonstrated that the system can help the precise localization of a gas leakage. Moreover, the careful design of hardware and firmware allowed to achieve an autonomy on battery ($\geq$ 40 min) longer then that of the host ($\leq$ 30 min) system. Additionally an energy optimization control scheme has been developed for

autonomous gas tracking applications, that allows to minimize the time and energy required by mobile systems to locate the gas leak. Finally, the impact of solar energy harvesters (photovoltaic panels) on top of UAVs has been deeply investigated to provide auxiliary supply and further extend the limited autonomy of such systems.

The main contributions of this work have been the demonstation of general purpose gas sensors' performance in outdoor environment with a minimal and lightweight box, and the development of the energy optimization control scheme for gas leak localization that is suitable for any kind of mobile carrier.

On the security assessment of controlled environments, data centers' server rooms have been focused since recently [25], data centers managers have expressed a growing interest in improving actual Data Center Infrastructure Management (DCIM) systems, collecting data to improve security and optimization of resources, alarmed by the risk of failures or security breaches that can cause millions of dollars losses [76]. Usually, in modern data center, wasted heat coming out from server rooms is used to generate hot water for surrounding services, however, it can also be recovered to generate electric energy and mitigate wastage [33]. In order to meet both goals a *smart heat-sink* monitoring device has been developed. This system, placed inside server units on top of CPUs, consists of a battery-less embedded devices designed to realize a maintenance-free WSN for data center monitoring. The energy necessary to operate is extracted by the thermal dissipation of the CPU, the most abundant kind of wasted energy in a data center. By exploiting it, the smart heat-sink system can provide *continuous* sampling of several environmental parameters — like temperature, humidity and light — and *enhanced security* — thank to the use of a gas sensor and a positional switch — thus demonstrating energy neutrality monitoring in such constrained environment.

The most important achievements in this sense have been the successful realization of an energy harvester able to completely sustain a portable monitoring node with a gas sensor embedded, and the successive miniaturization of the complete system (monitoring node plus harvester and conditioning circuit) in the form factor of a commercial heat-sink.

1.3 Structure of the Thesis

The Thesis is organized as follows, Chapter 2 presents a broad overview of the state-of-the-art in the fileds covered by this work.

Chapter 3 introduces volatile chemicals detection principles and focuses, in particular, on the relevant technologies for the development of embedded monitoring systems.

Chapter 4 presents the details of the ULP gas sensors sampling and monitoring methodology developed and illustrates the results obtained with experiments in controlled lab facilities and real WSN deployments.

Chapter 5 describes how the experience gained with gas sensors has been used to design an embedded system for environmental monitoring using mobile vehicles, in particular UAVs. Characterization of the board and results of the field experiments are deepely analyzed and discussed to provide an extensive performance evaluation of the design.

Chapter 6 introduces the energy harvesting principles and focuses on the reuse of heat generated by computing units, demonstrating that, although really limited, the recoverd energy can successfully sustain a gas monitoring system in controlled environments like a data center.

Chapter 7 discusses the *smart heat-sink*, an emebedded system designed and developed with energy neutral characteristics, to monitor data centers' server rooms. The comparison of simulations and results obtained with the prototype device are deeply analyzed to demonstrate the efficiency of the battery-less system.

Finally, Chapter 8 draws the conclusions and resumes the main achievement presented.

Chapter 2

State of the Art

Given the multidisciplinarity of the research topic, ranging from chemical sensing to energy management and embedded systems design, in the following is presented a review of the state of the art subdivided in the three sections corresponding to the main research fields that has been investigated in the development of this thesis:

- Gas Sensor Networks
- Environmental Monitoring with Unmanned Aerial Vehicles
- Thermoelectric Energy Harvesting

2.1 Volatile Chemicals Detection with Embedded Systems

Many research groups are working on the development of air monitoring systems based on distributed embedded systems; each solution differs in the type of application scenario, sensors used and the energy-optimization achieved.

On the volatile chemical sensing side, also known as eNose applications [55], three main technologies exist that are suitable for distributed monitoring using autonomous equipment, that require compactness, low-weight and low-power characteristics:

- Chemo-optical [80];
- Chemoresistive [75];
- Electrochemical [90].

Chemo-optical devices are the most expensive and bulky while electrochemical sensors need periodic refill and maintenance of reagent substances. On the other hand, chemoresistive technology — also known as Metal OXide semiconductor sensors (MOX) — offers fast response and long term stability without maintenance even if their power consumption is not negligible [93].

On the Wireless Gas Sensor Network (WGSN) side, Wang et al. in [89] present a sensor network for outdoor installations with a solar panel mounted on top of the sensor node box to provide constant supply. In this WSN they employed electro-chemical

Table 2.1: Comparison of power consumption of devices commonly used in sensor node design.

Device	**Manufacturer**	**Consumption**
Gas Sensors		
MICS-5121	e2v	32mA
TGS2610	Figaro	56mA
AS-MLK	Applied Sensors	15mA
Radio Transceivers		
JN5148-001	NXP	TX: 15mA RX: 17.5mA
CC2500	Texas Instruments	TX: 21.2mA RX:13.3mA
Microcontrollers		
JN5148-001	NXP	Active mode: 8mA@3V/32MHz Deep Sleep mode: 1.3μA
MSP430F247	Texas Instruments	Active mode: 321μA @3V/1Mhz Low power mode: 1μA@3V/32.768KHz

sensors to reduce the consumption, so that, the solar panel provides enough energy to sustain the monitoring requirements and the recharge of the backup batteries. They also compensate the gas measurement considering the environmental conditions of humidity and temperature which severely affect the sensors response. However, data reported in this article are not sufficient to evaluate the lifetime of the network (only 10 days provided).

In [79] a miniaturized system has been realized with the shape of a wrist watch to achieve a large scale monitoring network by letting the people wear it. Here the main efforts are posed on the custom design of the bracelets; power consumption are reduced by employing an electro-chemical sensor but periodic maintenance is required and no information on the sensor lifetime is provided.

Hardware optimization to reduce the amount of wasted power has been implemented successfully by some groups. In this way significant reduction has been achieved when the electronic boards are in idle or deep sleep state. No improvements are proposed to address the consumption of chemo-resistive sensors [40, 10].

Another approach to reduce energy requirements of WGSNs is to develop new sensors for ad-hoc applications. This choice, although interesting, is not cost effective in terms of both time required, to design the new device, and industrial production. In [82] the authors developed a low-power SAW sensor and its frequency readout circuitry. SAW technology is not usually adopted in chemical sensors because it is very expensive in comparison to others. The electronic circuitry required to interface the SAW device with the commercial sensor board makes the use of the SAW sensor complex, because of the physical dimensions and the need of powerful microcontrollers to perform the required computations.

Somov et al. [81] have presented a new device that necessitates very little energy:

those have developed a new chemo-resistive sensor and the adequate sampling and processing algorithm. Authors have focused their efforts on the development of the device to reach, theoretically, two years of lifetime. On the other hand the characterization of the sensor response under different environmental conditions of temperature and humidity is missed.

Few researches exist where a new strategy and a new algorithm, to periodically switch-on the device and evaluate the air-quality, is applied to commercial sensors. Usually new algorithms are employed in the optimization of tasks to perform and their scheduling. The goal is to modify the network behavior considering environmental factors (presence/absence of humans, day/night scenario, temperature and humidity) and available resources (battery level, presence of harvesters).

In [40, 39, 87, 17] the solution adopted to reduce the power consumption, consists on duty-cycle and results in extending the life of a node (as far as these are left asleep), still using sensors as indicated by the manufacturer. There are also some reports where the efforts are focused on the reliability of the measuring system, here, complex processing based on neural networks and fuzzy algorithms are implemented in sensing nodes to reveal if a dangerous situation occurs, hence the data are transmitted only when it is necessary thus reducing the radio consumption [69, 42]. The drawback of these kind of complex and demanding algorithms, is the time required to accomplish the computation on a wireless nodes.

Batteries are usually the only and unavoidable power-source that can be used with Wireless Sensor Networks, even though, energy-harvesting systems are recently introduced to increase the autonomy of the nodes [53, 23, 43]. The choice of introducing more hardware to harvest energy from the environment, however, is not so straightforward considering that the dimension are comparable, if not bigger, than the board itself; as a consequence, the decreasing of portability alone does not justify the effort.

Finally, Kuncoro et al. in [43], present a review of the state of the art in the field of WGSN. They describe the current results obtained in the integration between traditional WSN and e-Nose devices. The term e-Nose commonly is referred to arrays of electronic devices used to discriminate and measure volatile chemicals. Kuncoro reviewed some works presented in the last ten years to focus the key features of wireless sensor networks for air-quality monitoring. These articles do not focus on the energy optimization and solutions to extend the network lifetime. In the conclusion the reviewers underline the previous sentence: despite the increasing interest in this field and the proposed solutions, the lifetime of WGSNs is still bounded and finite.

The future of the research in this field is in the development of techniques to reduce energy consumption. Innovative hardware solutions must be designed to reduce the energy wasted and batteries alone are not able to guarantee long lasting performance. Harvesting techniques, also, should be introduced to compensate energy lack as well as ad-hoc network protocols should be investigated. The main issue in this background is how to reach two years (at least) of autonomy with battery-powered wireless sensor networks, with embedded MOX sensors.

2.2 Environmental Monitoring with Unmanned Aerial Vehicles

Unmanned Aerial Vehicles (UAVs) or drones, have overcome their first military uses, becoming today one of the most important technologies in data collection and remote sensing. This is highlighted by several researches, which list among its advantages a high resolution and positional accuracy. Their relevance for Earth system understanding and environmental science research has been pointed out for example in [32]. Among the others, localization and mapping of geographical areas by means of advanced 3D imaging techniques have been demonstrated for example in [77, 58].

The usage of mobile flying robots for civil purposes include various applications of gas detection, such as obtaining gas distribution mapping, monitoring emissions and gas source localization in geographical areas where environmental concern is a hot topic Caltabiano et al. in [13] designed and tested a plane-like UAV for automatic sampling and analysis of chemicals inside volcanic plumes; the specific scenario required a mobile system capable to reach 4000 m above sea level and to carry payloads as heavy as 5 kg, thus the final prototype had 3 m wingspan and 2 kW electric motor. Obviously, in such case is not possible to use a multi-rotor copter-like drone, but smaller flight control and gas sensing systems could guarantee longer flight aunomy (currently 30 min max).

Fleets of aerial vehicles have been proposed to extend the total area covered by the monitoring system, creating a kind of wireless sensor network where static nodes are replaced by small plane-like UAVs [4]. Authors of this work focused mostly the radio communication aspects and the self-management of the newtwork of unmanned vehicles, rather than the development and testing of the sensing infrastructure. However, this is one of the first works that demonstrates the feasibility of sensors flocks. A similar approach has been followed also by authors of [20] that present design and system specifications to employ radio communication channels already in place (GPRS and WCDMA).

Most of the works on gas mapping and leak localization focus on plane-like unmanned aerial vehicles given the longer flight autonomy and the larger payload capacity of such carriers. However multi-rotor copters offer a much more precise positionig capability and even a very small form factor, and for these motivations they were also proposed for emergency handling and fire localization in indoor monitoring applications [71].

Another advantage provided by multi-rotor system is presented in [57]. Authors of this work propose to use the information on actual trajectory followed by the drone and provided by the GPS receiver against the motion originally planned by the flight controller to estimate the wind direction and speed-up the gas mapping in a predefined geographic area. Altough significant, the advantages in terms of energy saving provided by this approach are still limited by the 20 min maximum flight autonomy of multi-rotor copter-like UAVs.

The main problems of all the current implementation are:

- the limited flight autonomy;
- the gas sampling rate which affect autonomy;
- the size-to-payload ratio.

Few efforts have been dedicated so far to the joint optimization of the flight path and energy consumption of chemical sensors.

Currently, some research teams have demonstrated solar power UAVs for remote sensing and wide-area monitoring applications [46] to solve the problem of the limited flight-autonomy; as far as we know those systems are based on airplane-like drones that offers big surfaces (wings and fuselage) to host photovoltaic (PV) panels but precise positioning and hoovering can not be achieved with these devices.

Following this research trend we propose an analysis on the expected performance boost of employing PV panels on multi-rotor copters against the power-to-payload ratio, considering numbers from commercial drones' manufacturers and panels respectively.

2.3 Thermoelectric Energy Harvesting

The key element for harvesting thermal energy is represented by Thermoelectric Generators (TEG) that exploit a phenomenon called the Seebeck effect to convert heat (or temperature gradient) into electrical energy [88]. They have been widely used for a variety of applications, including energy harvesting from heat dissipated in combustion engine vehicles. TEGs generally exhibit low efficiency, even in the high temperature context of automotive applications [96].

According to analysts forecast, the thermoelectric-generation market will increase from $40 millions in 2014 to $950 millions in 2024 [36]. This trend is mainly driven by the automotive sector [31], where the thermal energy generated by combustion engines is so high that by using integrated TEG it is possible to scavenge hundreds of Watts, and to reduce fuel consumption and carbon emissions at the same time.

However, thermoelectric scavengers can be useful in any scenario where thermal energy is lost in large amounts, such as power electronics and high performance computing. Here, the energy harvested from otherwise wasted heat can be reused for boosting standard cooling systems and/or supply related tasks. In the IT sector, for instance, part of the energy spent in computing is wasted through heat, and major efforts are made to keep the equipment at a safe temperature.

Wireless sensors powered by the dissipated heat can lower the system cost by boosting the heat extraction and providing a monitoring infrastructure to enforce the existing DCIM systems. By eliminating power wires, the installation costs decrease, moreover, not having to periodically replace their batteries allows the system to be installed in yet unexplored environments.

The integration of TEC and TEG on the CPU's cooling devices to convert wasted heat into electricity and boost the heat exchange at the same time is an active research field [45, 38] but current solutions are usually large and bulky.

The amount of scavenged energy is proportional to the heat generated by the CPU that is directly dependent on the kind of task the processor is running, its duration, and the environmental temperature.The recovered energy is small compared to the energy consumed by the CPU itself, therefore it is not enough to supply components of the device it scavenges from, nor to recharge the system's batteries [95, 28]. However, it can be used to provide active cooling to the processor, so as to increase the PUE of the system

reducing the energy used for cooling (or, conversely, allows the system to run at a higher clock frequency without causing overheating) [41].

Moreover, the rapid growth of networked smart sensors and actuators is about to hit the data center domain, promising the control of environmental parameters that are crucial for the system monitoring. Providing monitoring with free energy means having a security infrastructure which works in harmony with the data centers without increasing the energy spent, requiring no maintenance, and no additional wiring. Given the needs of the data center operators to collect environmental data to monitor safety in their infrastructure [25], and considering the fact that the WSN are very well suited for this purpose [74] — since they consume a very limited amount of power — we focused on the development of a wireless monitoring device that can be powered with the energy that the CPUs waste through the heat.

The very-low efficiency of thermo-electric generators make them suitable only as power supply of monitoring systems with relaxed constraints and small duty-cycles. In this sense we exploited this limit to investigate and develop very-low powered wireless sensor nodes that can be deployed in data center facilities, and run thanks solely to the energy converted from the excess of heat, which is constant and abundant.

Chapter 3

Volatile Chemicals Detection with Emebdded Systems

Pollution and urban air quality are majors environmental risks to public health. Gas emissions are responsible for a variety of deseases — respiratory illnesses, asthma and cancer — climate changes and environmental problems. Most of volatile compounds are colourless, tasteless and odourless hence, human beings are not able to recognize these leakages early enough to activate proper counter measures without auxiliary tools. The increasingly strict regulation in this regard requires technological improvements to prevent the risk of high concentration of pollutants in air, to keep a clean living environment, to avoid any risk for people and goods and ensure safety in industrial and public sites.

This chapter presents an overview on volatile chemicals detection principles and focuses, in particular, on the relevant technologies for the development of embedded monitoring systems. The most suitable sensing technology for portable system is selected and deeply analyzed through experimental evaluation, to highlight features and limitations that will be the target of further investigation provided in Chapter 4.

3.1 Volatile Chemicals Detection

The problem of measuring volatile pollutant presence in air and safety condition assessment in indoor environments, requires to embed chemical sensors in monitoring instruments. Classification of gas sensors is accomplished following the principles of signal transduction [34], available techonologies are:

Optical sensors following absorbance, reflectance, luminescence, fluorescence, refractive index, optothermal effect and light scattering;

Electrochemical sensors among them voltammetric and potentiometric devices, chemically sensitized field effect transistor (CHEMFET) and potentiometric solid electrolyte gas sensors;

Electrical sensors including those with metal oxide and organic semiconductors as well as electrolytic conductivity sensors;

Mass sensitive sensors i.e. piezoelectric devices and those based on surface acoustic waves;

Magnetic sensors based on paramagnetic gas properties (mainly for oxygen);

Thermometric sensors based on the measurement of the heat effect of a specific chemical reaction or adsorption which involves the analyte;

Other sensors mainly based on emission or absorption of radiation.

Among these, in the embedded system world only a subset of technologies are suitable for miniaturization, mass market production and interfacing with low-power microcontroller units. Currently, commercial sensors that can be embedded in wireless nodes belongs to three families:

- Chemo-optical sensors
- Electrochemical sensors
- Electrical sensors, in particular chemo-resistive (or Semiconductor)

The first two exploit a non-reversible chemical reaction that occurs between the target gas and the detector, thus a periodic maintenance is required to refill reacting elements. Semiconductor sensors instead, exploit a reversible reaction, triggered by heat, which results in higher energy absorption with respect to the former ones. Semiconductor sensors based on chemo-resistive principle (also called Metal OXide semicondutor MOX sensors) have intrinsic advantages since they are very low-cost (thanks to the fabrication in CMOS process), commercially off-the-shelf available and long lasting with no maintenance required. The drawback of this technology is the great amount of energy drained to activate the reaction.

3.1.1 MOX advantages for Embedded Applications

MOXs drain more power with respect to any other peripheral in embedded electronic world, radio transceivers included. In Table 2.1 we present a series of devices commonly used in the development of wireless nodes for comparison. The advantage of MOXs resides in their life-span, as stated above. Huge difference compared with other technologies like chemo-optical and electro-chemical ones, for which periodic refill and calibration are required, and lifetime depends on the maximum amount of gas they can absorb (in harsh environment can be few months). All the MOX sensors drains more current than the microcontroller devices used to implement the monitoring network. Moreover, considering the costs, both elctro-chemical and semiconductor technologies are inexpensive with respect to the optical family.

A standard characteristic of chemo-resistive gas sensors is the very long time required to perform the measure, in comparison with sensors targeted to the other environmental parameters. For example temperature sensors with digital interface require few hundreds of milliseconds to provide a measure, while a MOX one takes more than 5 seconds to heat-up itself and reach a steady state response.

Figure 3.1: AS-MLK Natural Gas sensor [8] (left), MICS-5121 CO/VOC sensor [2] (center) and MICS-5525 CO/VOC sensor [3] (right).

Currently, available WGSNs in the literature can last from three weeks up to few months [43], while, to the best of our knowledge, no commercial solution exists yet. In order for these devices to reach the target two years of autonomy, the energy presently drained must be *reduced 35 times*. For this reasone, embedded and battery-powered environmental monitoring devices, available in the market nowadays, do not include chemical sensors, because customers are not interested in security systems that require battery replacement every three weeks. The only commercial monitoring devices that uses MOXs are smoke detectors installed in buildings since they require the sensors to be always switched-on and consequently they are Grid connected.

3.2 Standard MOX response

3.2.1 MOX Working Principle

Chemoresistive sensors are essentially MOSFET transistors with a catalyst coated gate that interacts with gas vapors resulting in a modification of the gate charge and of the channel's conductive characteristics [55]. They are composed of several stacked layers made of different materials, depending on the target chemical compound. They are built on a silicon substrate in CMOS technology which results in a very low cost for industrial production. As previously introduced, they exploit a reversible chemical reaction that occurs in the sensitive layer, triggered by the heat. Moreover they require stable supply to precisely control the working temperature, as demonstrated for example in [22]. The heat is provided by an inner metal layer which requires high power to heat up in few milliseconds and maintain the whole structure around 400^oC. The sensitive layer is made of semiconductive material which changes its conductivity depending on the amount of particles that are deposited on it. The volatile analyte in the surrounding environment reaches the sensitive layer passing through a filter which prevents extraneous substances (e.g. dust) to enter the device [30].

The main issues of MOXs are the filter that let almost any kind of volatile to deposit on the sensitive layer; and the low selectivity of the alloy used as sensitive layer. MOXs are sensitive to family of chemical compounds and their response depends on environ-

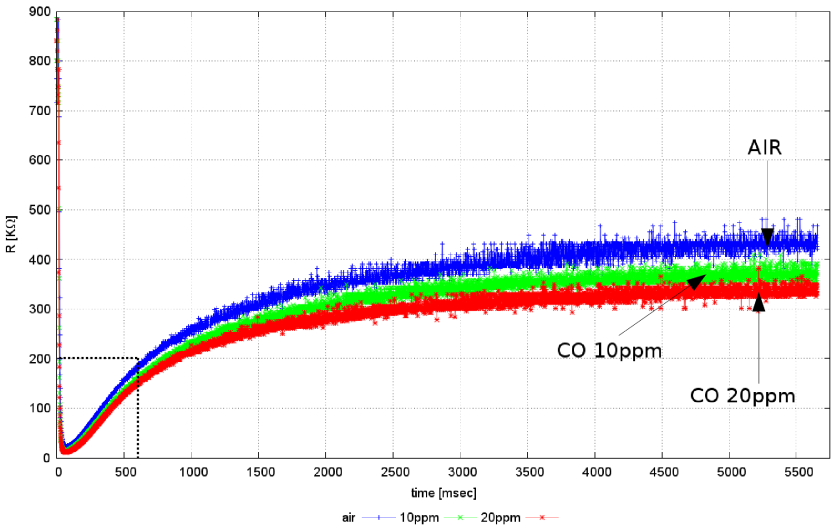

Figure 3.2: Standard MICS-5121 sensor response 20° C and 50% RH. Resistance versus time for different gas concentrations.

mental conditions too. Temperature and humidity strongly affect sensor's response since particle mobility (temperature) and the presence of water (humidity) interact with target chemicals during the reversible reaction. The output signal of MOXs depends on all the above considerations and those must be taken into account in the definition of the characteristic response.

3.2.2 Sensor Response

The characteristic response is generally obtained empirically by means of experiments, this can be a mathematical expression (obtained by numerical regression) or a table (often also a graph) that correlates the input concentration with the sensor's output signal. Thus, to define the application scenario of interest in order to build a corresponding experimental setup is of primary importance. Generally, two different monitoring scenario can be addressed:

Continuous, the sensor is kept always switched on, hence the temperature of the device is almost stable, the power consumption is very high and the effect of environmental conditions (i.e. humidity, particles deposition) very low. This is the operating condition of smoke detectors where a fast and reliable response is needed.

Duty-Cycled, typical of WSNs, the sensor is switched-on for a given time interval (τ) shorter than a time period (t) which defines the repetition interval of the measurement. The duty-cycle (D) is then computed as $D = \tau/t$. Moreover, the role of the environmental parameters and the deposition of any kind of material strongly slows down the sensor response because particles on the sensitive layer must be *burned*.

In continuous monitoring applications MOX sensors reacts almost immediately with

chemical compounds, exhibiting a fast drop in the response and providing an immediate feedback.

In duty-cycled applications, the MOXs considered require at least 5 seconds of continuous supply to reach the steady state response, for repetition intervals of few minutes (as low as 5 minutes, a medium-high interval in WSNs for environmental monitoring applications). An example of standard response of the MICS-5121 CO/VOC sensor is shown in Fig. 3.2. In this case three signals are overlapped, one for each gas concentration in the fluxed mixture, the 5000 samples are taken every $\approx$1 ms for a total of about 5.8 s. In this test can be noticed that the steady state response (90% of the final value) is reached after approximately 3.5 s, although the noise makes it not clear to understand in autonomous manner. A resource consuming low-pass filtering stage is necessary to correclty evaluate this signal.

Energetic requirements of MOX sensors are closely related with the time they need to reach stability, once switched-on. For chemo-resistive sensors, used in duty-cylced fashion, this time interval is between 5 and 6 seconds, depending on the environmental conditions. To save power, this time has to be shorten with a specific strategy to avoid the risck of affecting the device's reliability.

In the next Chapter these considerations will be used for the analysis of MOXs sensor response both in time and frequency domain, particularly focusing on the transient phase of the response, highlighted in the lower left corner of Fig. 3.2.

Chapter 4

Ultra Low-Power Volatile Chemicals Detection with Portable Embedded Systems

As described in Chapter 2, the research target of indoor air-quality assessment using portable and low-cost embedded electronic devices, is to reach two years of autonomy using both commercial MOX sensors and battery powered devices [43]. The use of wireless nodes is often the preferred choice to deploy such systems in indoor environments, with respect to wired ones, because of the ease, speed and reduced costs of installation.

Chapter 3 introduces the most suitable and widely used sensing technology to realize a portable wireless systems with embedded gas detection feature. However these devices have a remarkable power consumption that impacts their widespread in embedded systems. The challenge must be faced in the design phase, it is required to find the best trade-off between the remarkable power consumption and the maintenance required to work.

The ultra-low power methodology presented in this chapter, has been developed thanks to the collaboration with the Fondazione Bruno Kessler which provided laboratory facilities and the sealed chamber where the tests were conducted. Results preseted in the following have led to publications [J3,J4].

4.1 Research Setup

We started from the study of MOX sensors to analyze their features and characteristic responses to find a methodology suitable to reduce at most the energy consumption. In this research, the target methodology is a combination of power supply strategies and signal processing algorithms to be implemented on any kind of cheap and low-power embedded system.

The low-cost requirement of the target sensor network is met choosing commercial devices, both for sensors and for the development board (wireless node). Three MOX

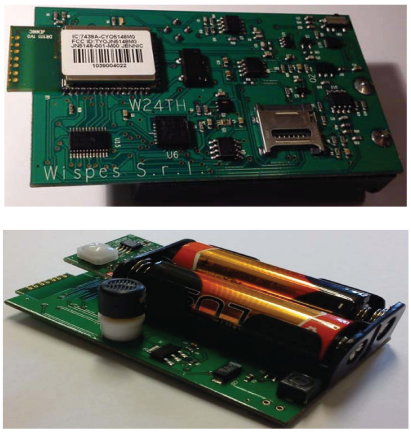

Figure 4.1: W24TH, IEEE 802.15.4 compliant wireless MCU [91].

devices were selected:

AS-MLK: a Natural Gas sensor from the Applied Sensor [8] (Fig. 3.1-left);

MICS-5121: a CO/VOC sensor from the from SGX Sensortech Ltd [2] (Fig. 3.1-center).

MICS-5525: a CO sensor from the from SGX Sensortech Ltd [3] (Fig. 3.1-right).

The MiCS-5121 is a general purpose CO/VOC sensor. This sensor and the mode of operation are designed to take measures of reducing gases such as carbon monoxide (CO), hydrocarbons (HC), and volatile organic compounds (VOC). The MiCS-5525 contains exactly the same sensitive element as the MiCS-5121 sensor, the only difference is the presence of a charcoal filter, which has been added on top of the sensitive element. Although we noticed that this charcoal filter increases the sensor reaction time and reduces the sensor sensitivity, conversely it enhances the selectivity. The AS-MLK sensor is intended for contiuous monitoring of Natural Gas transmission systems and leak detection applications. These devices are intended for mass market application whose key requirements are long lifetime, low cross sensitivity and long term stability. They are targeted at real-time monitoring applications, in the sense that *they must be always switched on* to have a prompt response.

The wireless node used in the whole process of analysis and characterization, up to the deployment of the real network for the tests on the field is the W24TH mote from Wispes [91], a state-of-the-art electronic board, designed to reduce at most the energy wasted during idle state (down to 8 μW, Fig. 4.1). The CPU of the sensor node is the Jennic JN5148 module [1] which is an ultra low-power, high performance wireless microcontroller (MCU). The most remarkable features are the computational architecture, a 32 bit MCU at 32 MHz useful to perform on line processing; integrated RF module compliant to the IEEE 802.15.4 standard; sensors for temperature, relative humidity, light and dock for the catalytic sensors. It holds two AA batteries as power supply and the current consumption is 15 mA for TX and 18 mA for RX, compared to other similar and commercial sensor nodes, it gives about 35% power savings. Current consumption of the selected hardware is also reported in Table 2.1.

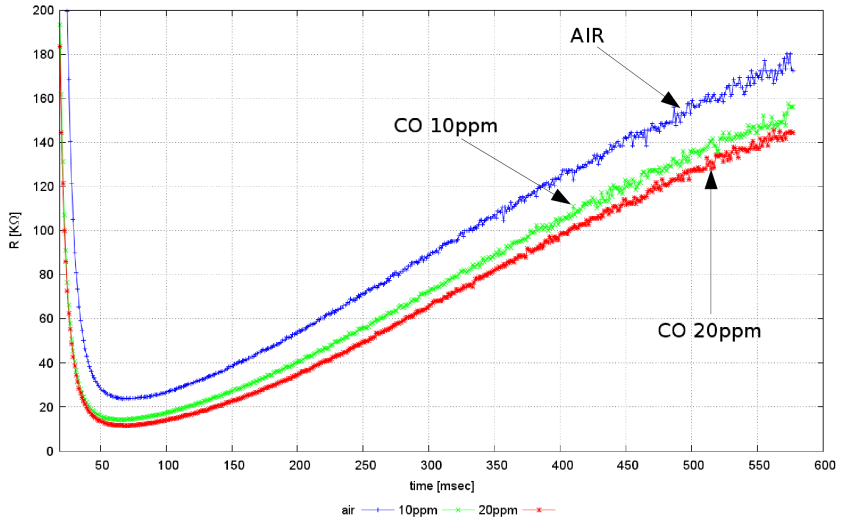

Figure 4.2: Initial transient response of MICS-5121 sensor, first 550 samples of Fig. 3.2.

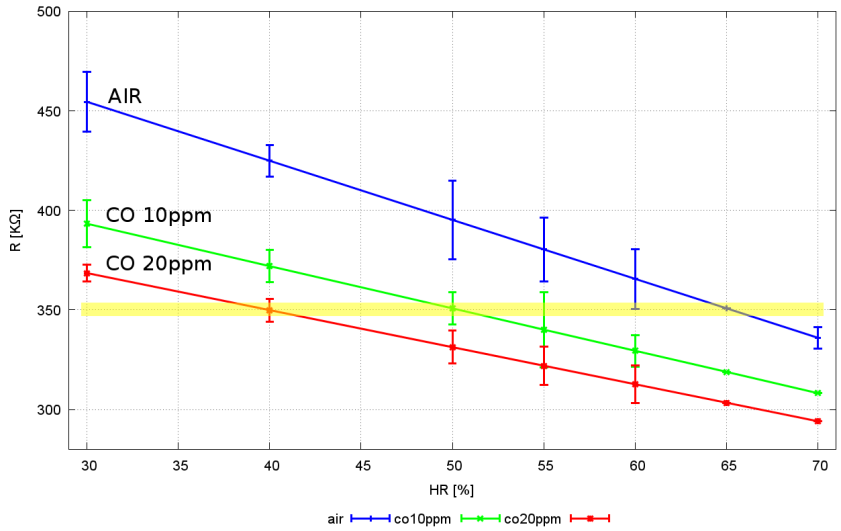

Figure 4.3: Input-Output characteristic of CO/VOC sensor. Standard measurement technique. Resistance versus humidity for different concentrations.

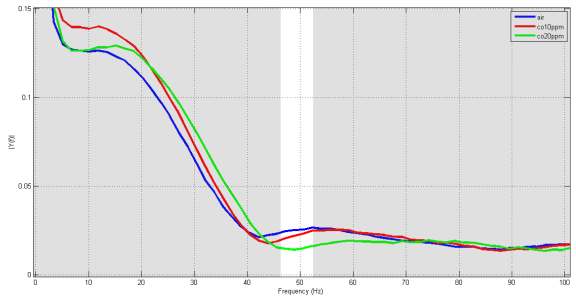

Figure 4.4: Amplitude spectrum of the interpolated signal depicted in Fig. 3.2.

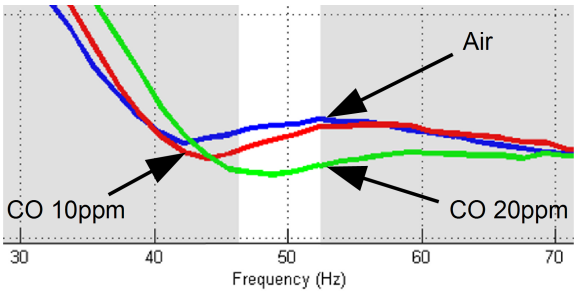

Figure 4.5: Detail of Fig. 4.4 around the bandwidth of interest.

4.2 Ultra Low-Power CO/VOC Measurement

The standard characteristic extracted for the CO/VOC sensor is shown in Fig. 4.3. Here, the horizontal line underlines that without the knowledge about relative humidity (RH) in the environment, the resistance itself does not provide any useful information.

It is possible to observe the two phases of evolution in the resistance sampled at the output port of the device: the *transient phase*, characterized by a fast drop of the signal up to $100\,ms$, and a *slow increase* toward the stability region. In Fig. 4.2 it is shown the detail of the transient response already shown in Fig. 3.2. Here, in the very first samples of the output resistance, it is clearly possible to discriminate the gas concentration in the fluxed mixture, even if some noise is present. In case of very-low duty-cycles (e.g. a measure every 15 minutes) or in harsh environments, the sensor's transient response can easily decrease to zero despite the analyte concentration, hence it is not possible to use the minimum as target feature for the ultra low-power methodology. Anyway we focus this phase to define a strategy useful to extract features from few initial samples.

It has been empirically observed that the transient behavior strictly depends on the time between successive measurements. Longer the time between two switching-on, slower the increase of the resistance. This is due to the amount of particles, present in air, that are deposited on the sensitive layer and must be *burned*, delaying the useful chemical reaction. Despite this, the evolution pattern represented as a steep decrease followed by a slow increase, can be always observed in all the experiments. We try to exploit these observations applying strategies like duty-cycle and input power modulation,

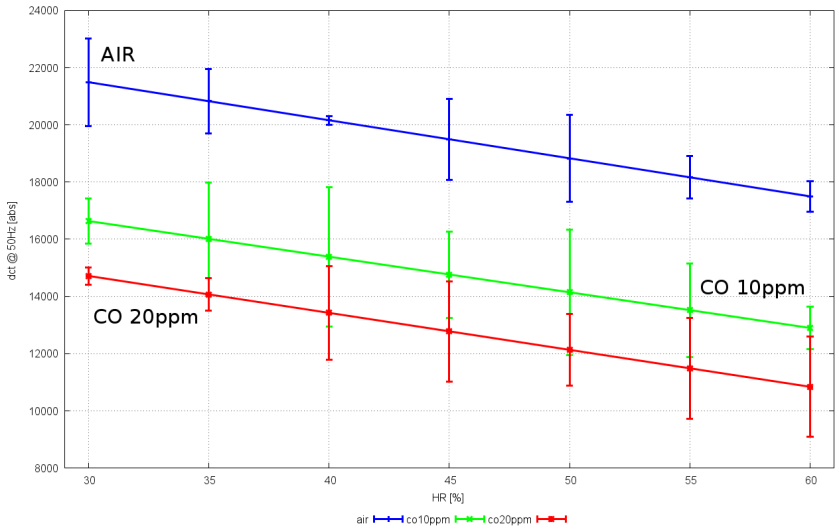

Figure 4.6: Input-Output characteristic of CO/VOC sensor. Ultra Low-power measurement methodology. Resistance versus humidity for different concentrations.

as it has been reported in previous works, however with useless results.

The contribution of my work overcomes the limitation of the state of art strategies applying a completely different approach, shifting from the time based analysis to the frequency domain of observation. The resulting methodology employs a complex processing based on the Discrete Cosine Transform (DCT) to extract the features from the sensor's transient response. Since the output resistance of the catalytic sensor is an aperiodic signal it can be interpolated to extract a continuous spectrum through the Fast Fourier Transform algorithm. Fig. 4.4 shows a portion of the spectrum of the signal in Fig. 3.2. From the analysis of amplitude spectrum has been observed that in the neighborhood of $50\,Hz$ (Fig. 4.5) the components show a behavior proportional to the gas concentration more or less evident, despite temperature, humidity and duty-cycle. This feature was used to perform the characterization of the device.

We analyzed the harmonic components of the first 512 samples instead of waiting for the steady state response. Therefore we reduce the operating time of the device of 10x with respect to the recommended time (≈0.5 s against ≈5 s). The transient signal is processed using the DCT to extract the amplitude of a single spectrum's component using the following equation:

$$\hat{X}_k = \sum_{n=0}^{N-1} x_n \cdot e^{-i2\pi \frac{k}{N} n}$$

The main advantage of the DCT is the lack of complex computation, thus reducing the arithmetical operations, the amount of necessary memory and resulting in a fast execution of the task. This solution introduces a computational burden which results in a severe drain of current but still lower with respect to keep the sensor switched-on for equivalent time intervals. Using the wireless nodes presented, with two standard $2500\ mAh$ battery, it is possible to extend the life-time of the network from 3 to 24

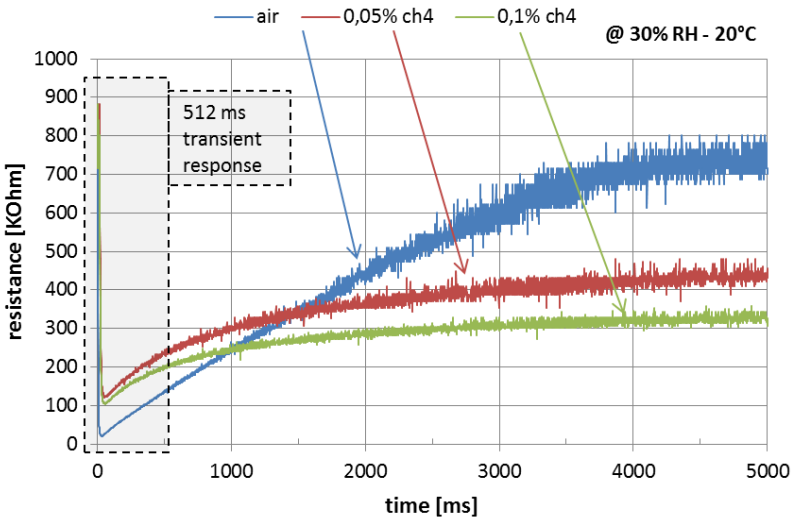

Figure 4.7: AS-MLK standard response to different gas concentrations @ 20° C and 30% RH, 5 s switch-on time required to reach a steady state condition for normal reading. The box highlights the transient response used by our algorithm.

weeks, which is far from the target goal but still an achievement. We chose 15 minutes as measurement's frequency to reduce at most the consumption without impairing the safety. In indoor scenario gas takes time to propagate between rooms and dangerous situations are unlikely to happen.

Implementing the algorithm, we have characterized the sensors at different environmental conditions (i.e. concentration and humidity, temperature cannot be controlled in the available test bench), with 0,057% duty-cycle (≈0.5 s every 15 min) as shown in Fig. 4.6. In the picture we also reported the standard deviation for the points of the linear regression, computed by averaging the results of multiple experiments. The fact that there is an overlap in the standard deviations (as happens between 10 ppm and 20 ppm curves) determines the qualitative characteristic of this methodology, since it is not possible to discriminate univocally different concentrations. Finally we obtained a clear separation between clean air and contaminated environment. Indeed, the approach is striking in fast determination of threshold levels of gas presence in the environment, strongly reducing the overall power consumption.

4.3 Ultra Low-Power CH_4 Measurement

Similarly, we evaluated the ultra low-power methodology with the natural gas sensor, starting from the analysis of the standard measurement techhnique to verify the conclusion drawn in the CO/VOC case. Fig. 4.7 and 4.8, show the AS-MLK response when used in duty-cycled fashion with "single-shot" 5 s continuous and constant supply, as indicated by the manufacturer. Moreover it is highlighted the comparison between the 5 s

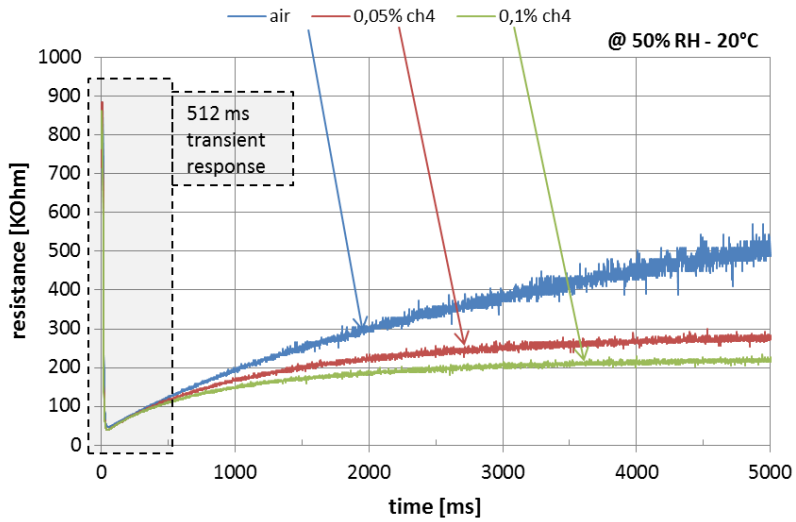

Figure 4.8: AS-MLK standard response @ 20° C and 50% RH, with 5 s switch-on time. The box highlights the transient response used by our algorithm.

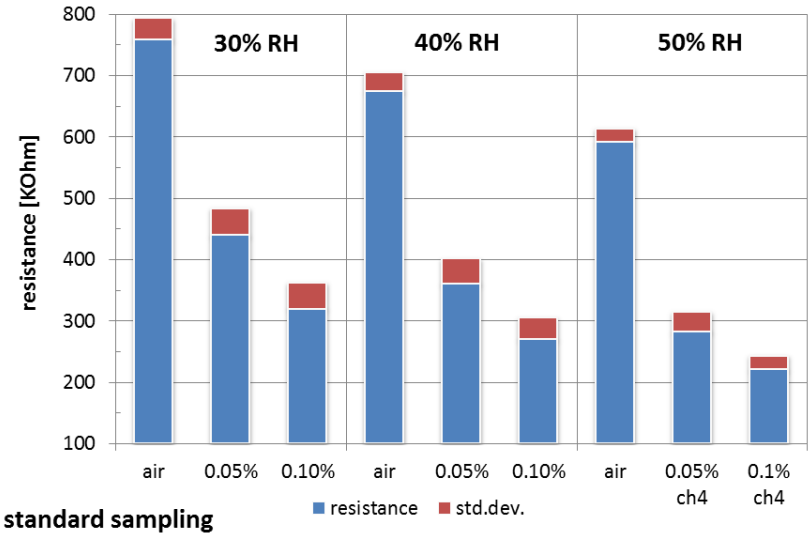

Figure 4.9: AS-MLK characterization extracted for sensor operated in duty-cycled fashion using standard measurement technique.

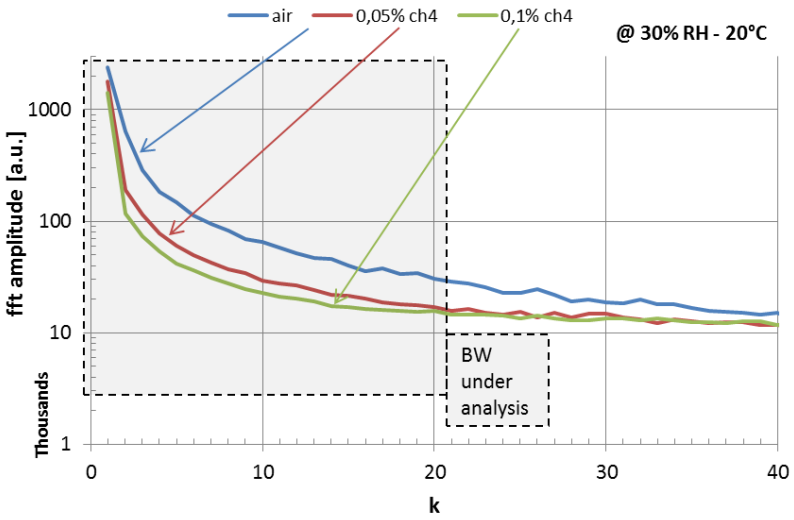

Figure 4.10: FFT analysis of the standard response @ 20° C and 30% RH. The box highlights the bandwidth of interest.

duty-cycled response's length and the 512 ms ultra low-power duty-cycle proposed. In addition to the different length in time, these pictures show the influence of the others environmental parameters to assess the sensor's response. Notice that a difference of almost 250 kΩ in the response to clean air when the relative humidity (RH) concentration in air changes from 30% to 50%. Moreover 0,05% natural gas concentration at 30% RH is similar to sensor response to clean air at 50% RH. Fig. 4.9 present the device's characteristic response in kΩ with standard 5 s sampling strategy, obtained by averaging several experiments - maximum value for the standard deviation is ≈40. These results once again underline the relevance of the environmental condition to correctly interpret the sensor response.

Fig. 4.10 and Fig. 4.11 depict the FFT transform of the two cases presented above. The boxes highlight the first 20 components of the amplitude response, where the difference is clearly evident. We noticed that, in the frequency domain of observation, the difference between characteristic sensor's response to natural gas concentrations are amplified in low frequencies with respect the CO/VOC case. To this reason, the 512 ms output signal has been processed using the DCT analysis to extract the amplitude of a single spectrum's component, normalized with respect to the DC, using the following equation, where N is 512, k is the component of interest: 1 to compute the `DCT[1]/DC`, 20 for the `DCT[20]/DC` and n is the index of the input time series:

$$\hat{X}_k = \frac{\sum_{n=0}^{N-1} x_n \cdot e^{-i2\pi \frac{k}{N} n}}{\sum_{n=0}^{N-1} x_n}$$

This further increase the number of computations to be performed (2 components required) but we achieved more interesting results, namely lower standard deviation thanks to the normalization.

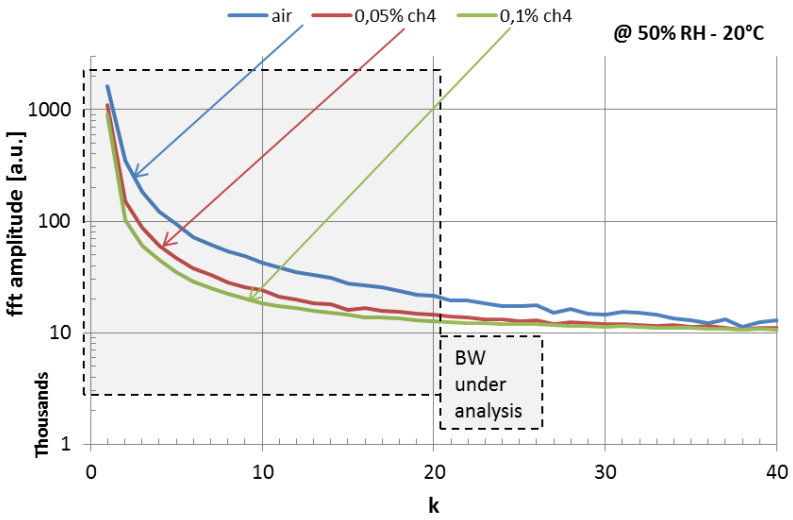

Figure 4.11: FFT analysis of the standard response @ 20° C and 50% RH. The box highlights the bandwidth of interest.

In the natural gas case we evaluated two different mode of operation:

- Medium-low duty-cycle (0.43%), where a measurement is taken every 2 minutes;
- Very-low duty-cycle (0.057%) where the interval is set to 15 minutes.

Fig. 4.12 and Fig. 4.13 both for 0.43% duty cycle, depict the characteristic response for the first normalized component (`DCT[1]/DC`) and the twentieth (`DCT[20]/DC`). Similarly Fig. 4.14 and Fig. 4.15 for 0.057% duty-cycle. These 4 pictures underline that it is possible to extract a useful characterization considering also multiple components, opening the possibility to a more complex kind of post-processing which is out of the scope of the present work. The reduced standard deviation on top of each bar of the graphs — below 0.75% in 0.43% duty-cycle and below 3% in 0.057% case — confirms the validity of the strategy.

The characterization has been performed considering a range of humidity conditions of the environment and several repetition intervals. The presented graphs have been selected to remark the fundamental role of the calibration in the proposed approach. Indeed, notice that in case of 0.43% duty-cycle the response to clean air is the lowest (A) with respect to the presence of gas, while the opposite in case of lower duty-cycle (B). This inverse behavior depends on the time the sensor spends in off state - the duty-cycle - and can be related with the previous considerations on the deposition of particle on the sensitive layer. The following discussion holds for the analysis of the transient behavior of the sensor response that is limited in time to 500 ms and the steady state is not reached. (A) A flat response in time (high gas concentration) has a higher harmonic content with respect to a fast increasing time series (pure technical air). Thus in case of medium speed duty-cycle (0.43% case) the sensor is almost clean and the resistance decreases with gas concentration. (B) In case of slow speed duty-cycle (0.057% case)

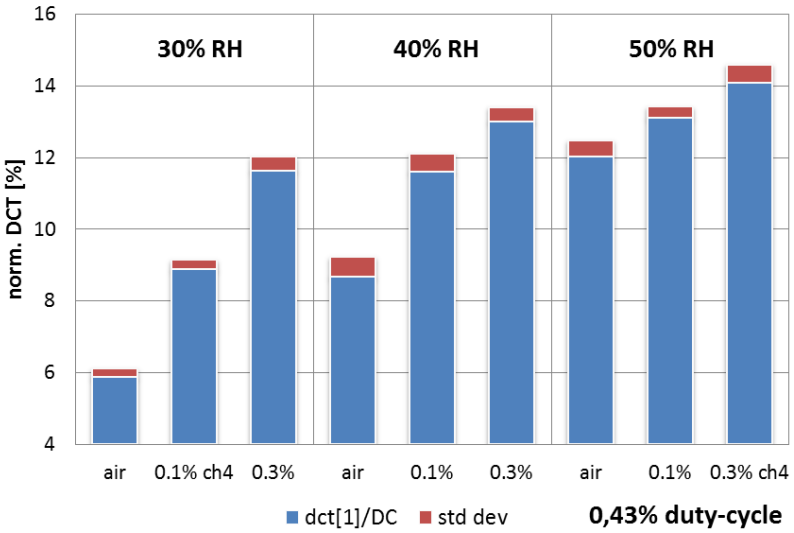

Figure 4.12: Characteristic response to natural gas concentration with 0.43% duty-cycle @ 20° C and different humidity conditions. 1st component of the normalized DCT extracted with transient response analysis.

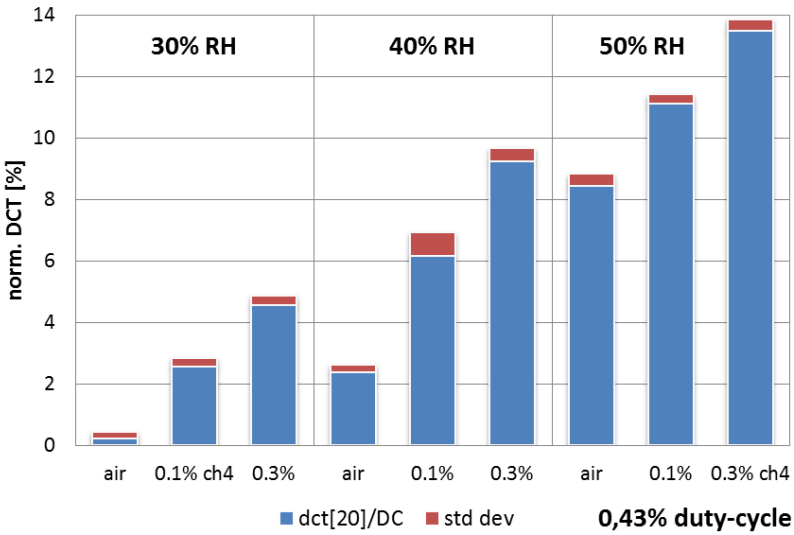

Figure 4.13: Characteristic response to natural gas concentration with 0.43% duty-cycle @ 20° C and different humidity conditions. 20th component of the normalized DCT extracted with transient response analysis.

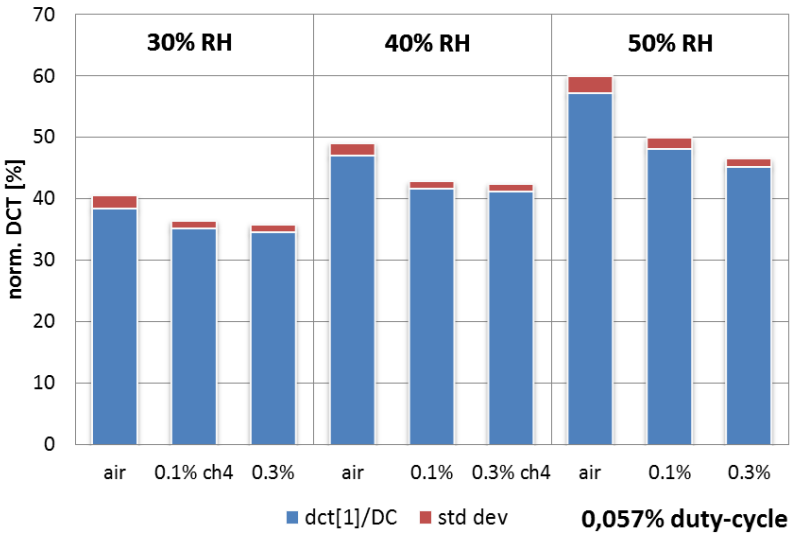

Figure 4.14: Characteristic response to natural gas concentration with 0.057% duty-cycle @ 20° C and different humidity conditions. 1st component of the normalized DCT extracted with transient response analysis.

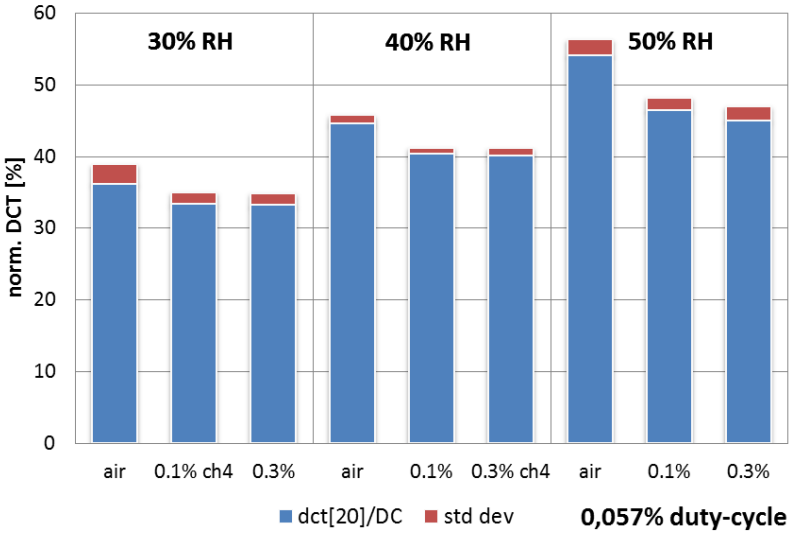

Figure 4.15: Characteristic response to natural gas concentration with 0.057% duty-cycle @ 20° C and different humidity conditions. 20th component of the normalized DCT extracted with transient response analysis.

it happens that the presence of target gas in the heterogeneous mixture of deposited particles, "speeds up" the chemical reaction in time that exhibits lower harmonic content: the higher the concentration the faster the reaction, the faster the resistance evolution thus the lower the harmonic content at higher frequencies.

The sensor response in real environment is quite different with respect to what is obtained in controlled lab facilities. The presence of humidity, pollutant, dust and unpredictable volatile substances, affect the MOX response in terms of amplitude and speed to reach stability. These in turn affect the outcomes of the proposed approach. For example in `DCT[20]/DC` case (Fig. 4.15) for 0.057% duty-cycle there is an overlap for gas concentration above 0.1%, while using the first component all the considered gas concentrations are remarkably separated. In the next section will be shown the impact of real environmental conditions on the capability of detecting harmful situations that is the goal of this work. Indeed the `DCT[20]/DC` resulted unuseful to this end while `DCT[1]/DC` provides a prompt response despite the unavoidable attenuation.

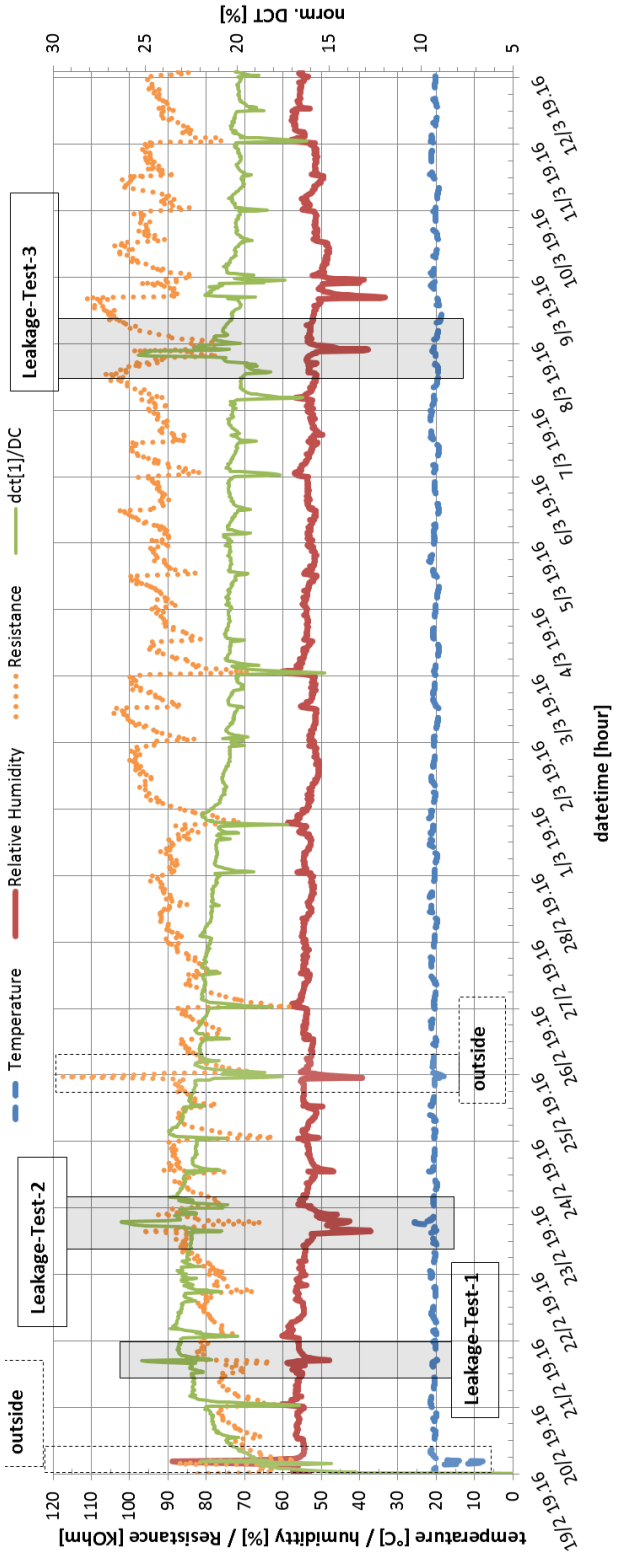

Figure 4.16: Results of the deployment with 0.057% duty-cycle and the proposed sampling methodology.

4.4 Tests on Field with CH_4-targeted system

4.4.1 Experimental Setup

We developed a test bench to evaluate the sensing performance and the energy saving of the proposed methodology, in a real home using a real WSN based on the W24TH wireless sensor nodes (see Fig. 4.1). The wireless network protocol used in these simulations is based on the IEEE 802.15.4 standard, we used two identical wireless boards equipped with the same set of sensors to evaluate the comparative performance in case of natural gas leakages. One node (a) implements the proposed ultra low-power (ULP) approach to evaluate natural gas concentration while the second (b) implements the standard sampling technique of 5 s sampling interval and it has been used as ground-truth. A third node has been used as network coordinator (or gateway) to provide synchronization and to collect data. The two where placed next to each other, 5 m far from the gateway. We performed several tests of controlled leakages using a custom sealed chamber — containing nodes, placed on top of a kitchen's gas nozzle — while in between standard kitchen activities happened (e.g. cooking).

This test presents the results for measurement taken every 15 m (DC = 0.057%), it lasted for three weeks, from February 19th to March 12th. Fig. 4.16 depicts the sequences of environmental temperature, relative humidity, natural gas concentration, evaluated as `DCT[1]/DC` and measured by node (a), and the resistance measured by node (b) which represents the ground truth since the test was performed in a real environment without control of actual gas flux. In this case we had also the possibility to evaluated the stability in time of the proposed methodology. Standard home activities were conducted during the test and the results are equivalently very good.

4.4.2 Leak Detection Perfomance

We simulated three leakages during this campaing but we focus only on the leakage test number two, highlighted in Fig. 4.17 for node (a) and Fig. 4.18 for node (b), for performance assessment. This fragment of the whole sequence represents an interval of standard activity, where all the measured parameters change synchronously, and the controlled leakage test. We reported the details of this case only due to space reasons and because the same considerations holds for the other two tests.

In both the responses we can notice that in case of standard activity (cooking in this case) all the reported parameters change simultaneously: a change in temperature and RH result in a decrease of the `DCT[1]/DC` response (a) and in an increase of the resistance (b), in accordance with the previous characterization.

In case of natural gas leakage instead we can notice that the `DCT[1]/DC` (Fig. 4.17) and the resistance (Fig. 4.18) change earlier with respect to temperature and RH. In particular we notice that the at the 12.31 of the Feb. 23rd, when the leakage starts, the `DCT[1]/DC` is slightly faster than the resistance to react, while the environmental change happens half an hour later (2 samples) the detection of the leakage. Finally we can notice that the `DCT[1]/DC` response is more stable (less variations due to other activities) with respect to the resistive response.

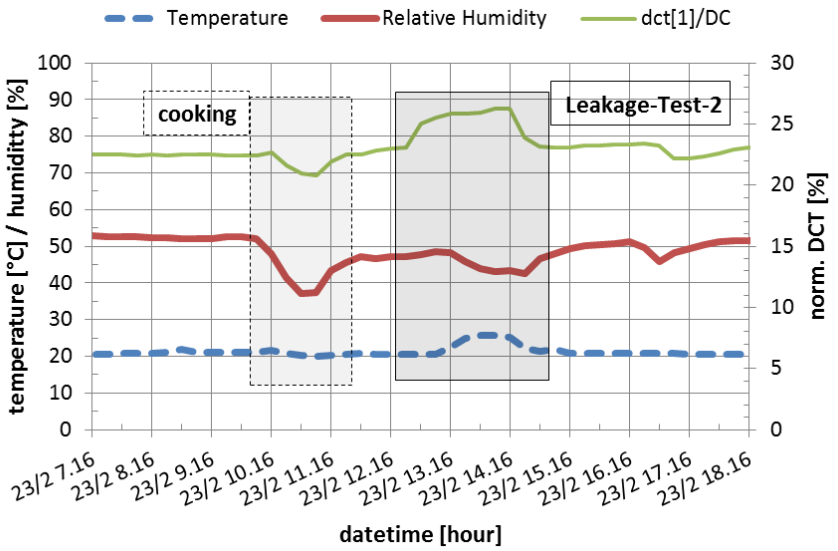

Figure 4.17: Detail of the second leakage with 0.057% duty-cycle and the proposed sampling methodology.

Fig. 4.19 present the response of the node (a) in terms of `DCT[20]/DC` component. In this picture the arrow highlights the instant where the leakage starts but evidently this is not caught. This test underlines that the cross sensitivity to other volatile substances is high and must be taken into consideration,in fact outcomes are very different than those obtained in controlled lab environment.

Similar considerations can be drawn by other experimental tests we have done, even if the results are not shown here. In particular we conducted tests using 2 m sampling interval (DC = 0.43%), from February 16th to 18th 2014, and again observing that `DCT[1]/DC` resulted faster than nominal reading method and `DCT[20]/DC` confirmed not to be effective in some conditions.

4.4.3 Results

The measurement campaigns clearly validate the ultra low-power approach to natural gas monitoring presented in this work. The DCT analysis allows to obtain an immediate feedback on the air quality and it resulted even faster than the standard approach in reacting to leakage events (particularly in very low duty-cycle scenario). The difference in the output response between results of the characterization performed in lab facilities, with respect to the real deployment ones, is to be addressed to the presence of many volatile compounds in indoor environments. As previously stated MOXs have high cross sensitivity to other volatiles. Even if we could not address empirically the impact of cross sensitivity using lab facilities, the results shown demonstrate its impact, but they also show that, despite it, the proposed approach exhibits the same performance of the standard measurement strategy.

In particular considering the case of leakage test 2 with 0.057% duty-cycle, we have

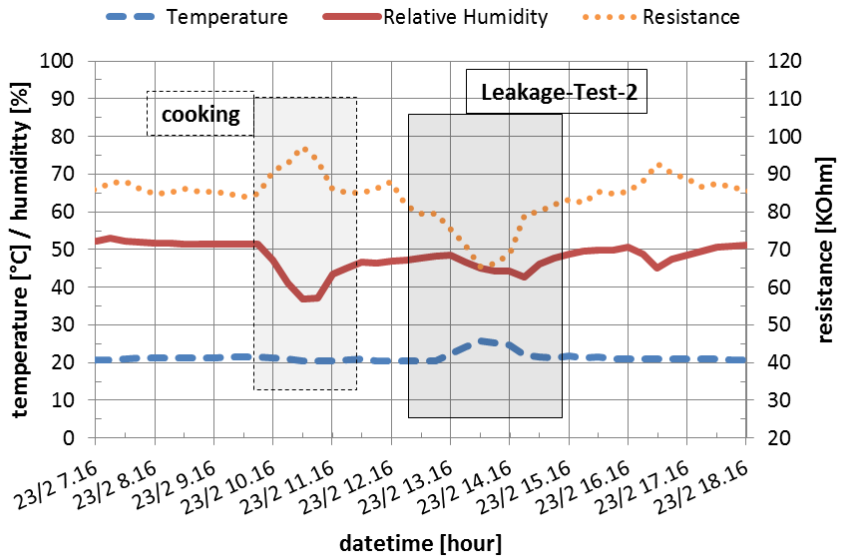

Figure 4.18: Detail of the second leakage with 0.057% duty-cycle and the standard sampling technique.

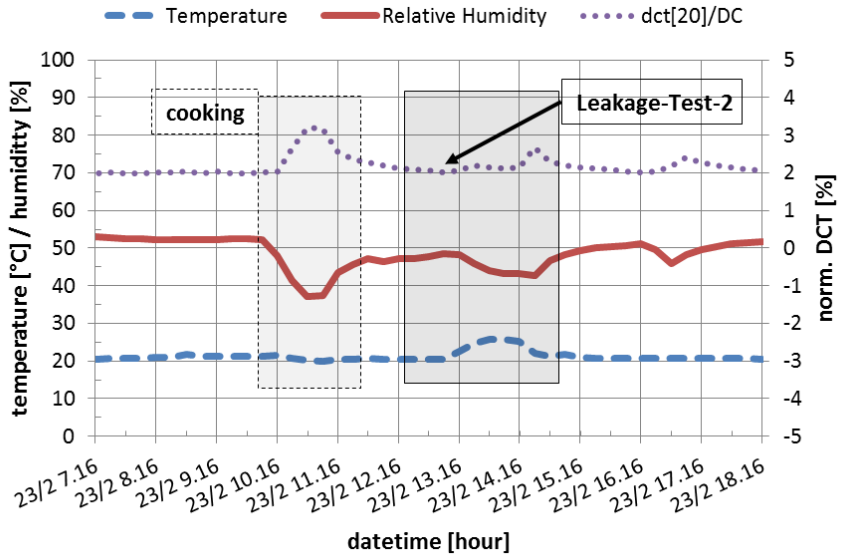

Figure 4.19: Detail of the second leakage with 0.057% duty-cycle and the proposed sampling methodology. Here the dct[20]/DC is shown to be unuseful.

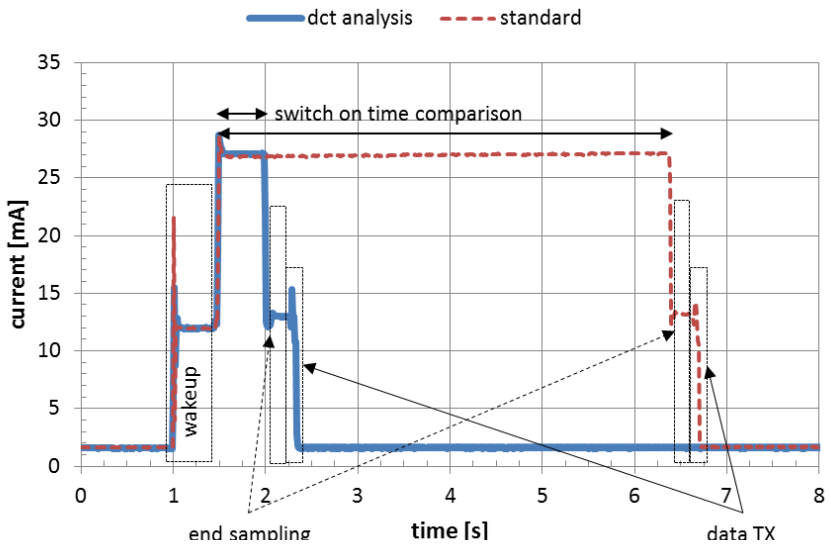

Figure 4.20: Energy consumption comparison of the proposed strategy with respect to the standard one. The boxes highlight the instruction executed by the W24TH. Measured using 1 Ω shunt resistor.

around 85 $K\Omega$ resistance @ 20°C and 50% RH (see results for Feb. 23rd around 09:00 in Fig. 4.18), while we expect to be in the neighbor of 600 KΩ considering the characteristic response reported in Fig. 4.9. A very similar consideration can be done for the normalized DCT approach. In this case for the same leakage test we can see a `DCT[1]/DC` value around 23% (see results for Feb. 23rd around 09:00 in Fig. 4.17) that should be around 57% from the characteristic reported in Fig. 4.14. These outcomes underline once again the need of a careful MOX characterization not only in terms of absolute response but particularly in terms of cross sensitivity to achieve quantitative assessment of gas concentration.

4.5 Evaluation of Energy Consumption Performance

We evaluated the performance of the ULP approach to volatile chemicals measurement considering four scenario:

1. Real time monitoring without network;
2. Duty-cycled monitoring without network;
3. Duty-cycled monitoring with Jennic's proprietary network stack;
4. IEEE 802.15.4 based network protocols.

In cases (2) and (3) the whole set of sensors (temperature, humidity, light and MOX) were switched on contemporary with the wireless node, while in case (4) we improved the implementation by enabling one sensor at the time to further reduce the energy drained. The network protocol of case (4) is based on a star topology with no joining

Table 4.1: Estimated lifetime of monitoring systems equipped with gas sensors implemententing the proposed ultra low-power strategy.

Configuration		Current [mA]	Lifetime [weeks]	Description
Real Time Monitoring	standard	**92.32**	0.16	Continuous sampling of Temp, RH, Batt and Natural Gas, w/o network.
	ULP	**11.15**	1.33	
Duty-cycled Monitoring	ULP - 0.43%	0.58	**25.44**	Duty-cycle based monitoring of Temp, RH, Batt and Natural Gas, w/o network.
	ULP - 0.057%	0.19	**75.16**	
	ULP - 0.43%	1.04	14.38	Duty-cycle based monitoring of Temp, RH, Batt and Natural Gas, w/ Jennic proprietary protocol.
	ULP - 0.057%	0.39	37.20	
	ULP - 0.43%	0.23	**65.57**	Duty-cycle based monitoring of Temp, RH, Batt and Natural Gas, w/ 802.15.4 based protocol.
	ULP - 0.057%	0.081	**185.14**	

nor packet routing management while the Jennic's stack, still based on the IEEE 802.15.4 standard allows more complex topology and provides network management facilities. Summary of the results are listed in Table 4.1. The first scenario permits to evaluate the energy savings in situations where MOX sensors are used for real-time monitoring (like in smoke alarms). The remaining are used estimate the lifetime achievable in classical wireless sensor networks applications including the impact of the network protocol on the power consumption.

Network overhead due to complex protocols is one of the main source of energy inefficiency in WSNs. Fig. 4.20 represents the current drained by the sensor node during the experiments, measured using 1 Ω shunt resistor in case (4). This picture also compares the impact of the proposed approach with the standard 5 s MOX sampling in terms of energy saving.

The values of the current drained by the nodes have been averaged over several days of simulations. Concerning the lifetime, roughly half of the time spent is used for networking tasks as it is evident from the results of Tab. 4.1—from 75 in case (2) to 37 weeks in case (3). By carefully designing the node's tasks and using a low power network protocol the consumption are further reduced—up to 185 weeks in case (4). In the duty-cycled monitoring case, the lifetime expectancy it is roughly 3 years and half in the best case. The most remarkable results are the 87% energy saving in case of real time monitoring (sensor always on) and the lifetime expectancy of 185 weeks in case of duty-cycled monitoring. This is a remarkable improvement both for standard applications like real-time smoke detectors and for wireless gas sensor networks.

Finally we can notice that the normalized DCT computation introduces an overhead which is comparable to the computation of the average resistance in case of standard 5s sampling strategy. These tasks are highlighted using boxes and named "end sampling" in Fig. 4.20.

Chapter 5

Outdoor Gas Leak Localization with Mobile Robots

The gas distribution mapping and the localization of a static gas source in outdoor environments are complex tasks, due to the rapidly varying dynamic conditions and the turbulent nature of gas transportation under natural phenomena.

Currently, many monitoring systems consist of a static network of sensors which are distributed at key locations but, in order to obtain a truthful representation of the gas distribution and be able to locate gas sources, it is essential to collect spatially distributed concentration measurements. However, the response of low-cost, chemoresistive gas sensors is caused by direct interaction of the chemical compounds with the sensor surface (1 cm^2), and thus represents a volume (few dm^3) around it. Hence, for economical and deployment-related reasons, a stationary sensor network is not a viable solution in many cases.

Vehicles can make a significant contribution in this area providing versatile systems for autonomous monitoring of diverse environments; a mobile robot equipped with dedicated embedded systems can collect environmental samples with a much denser spatio-temporal resolution than a human operator, resulting also in an safer working condition. However, mobile robots — UAVs in particular — represent a big challenge from the energy autonomy point of view. The limited resources available on-board are shared between several safety-critical and concurrent tasks — motion management, communication, sensing — thus the optimization of the gas sensing task is of primary importance.

This work has been carried out by my research group with the collaboration of Fondazione Bruno Kessler and a visiting student that I supervised during the development of her master thesis. My main contribution was the design and characterization of the first prototype of embedded gas sensing system and the definition of the optimal sensing strategy to use with gas sensors. Resutls presented in this chapter have been published in [J1,C5].

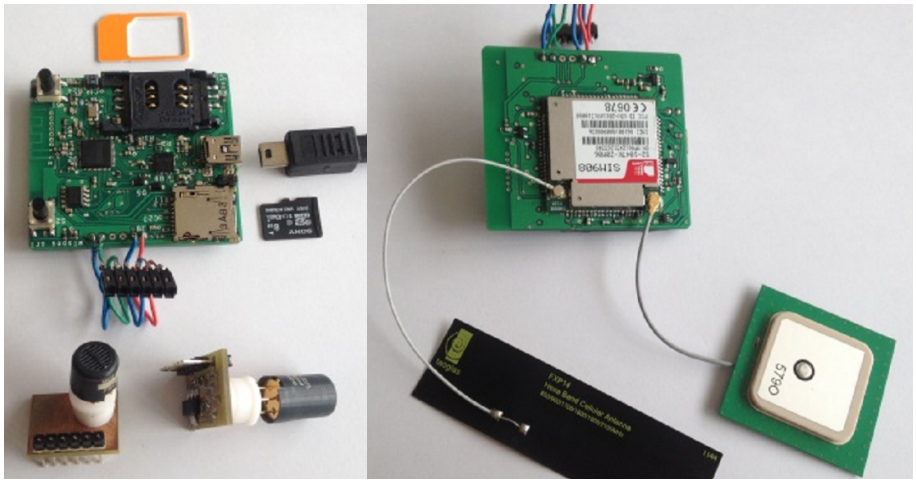

Figure 5.1: The prototype 2-side readout board, on the left is shown the MCU and gas sensor side while on the right the combined GSM/GPS and antennas.

5.1 Challenges of Outdoor Environment

Typically, turbulent transport is considerably faster compared to molecular diffusion. Apart from very small distances, where turbulence is not effective, molecular diffusion can be neglected concerning the spread of gas. A second important transport mechanism for gases, is advection due to prevailing fluid flow. Relatively constant air currents are typically found, even in an indoor environment without ventilation, as a result of pressure (draught) and temperature non-homogeneities (convection flow). Complexity increases considering outdoor environments where wind can play a major role in gas transportation in the atmosphere.

To this end, we developed an embedded electronic platform able to instrument any kind of drone with a chemical sensing system for environmental monitoring, gas leakage detection and pollution mapping applications. We developed a lightweight electronic system, based on the 32MHz/32bit JN5148 MCU that manages a GSM/GPS modem, a digitally controlled analog interface to drive MOX gas sensors and a micro SD-card for local logging. The two wireless interfaces, a short range 2.4GHz radio based on the IEEE 802.15.4 standard and the GSM/GPRS one, allow the development of custom and smart applications for real-time data sharing and/or alarm notification. The complete 16 cm^2 board (4x4 cm) is battery powered (Li-ion 1800 mAh) and the total weight is less than 30 g. The remarkably low weight allows the choice of any kind of carrier, not only UAVs, but also wheeled robots, public transportation means and so on, taking advantage of their tour to create real-time mapping of the urban air quality.

The most important design parameter in this applications is the energy autonomy of the drone for gas leakage localization and mapping which deals with weight of the system (the payload) and electrical power consumption. No commercial drones exists with autonomy longer than one hour and MOX gas sensors, like GSM modems, consume a remarkable amount of power. Differently from other works in the field we focused on

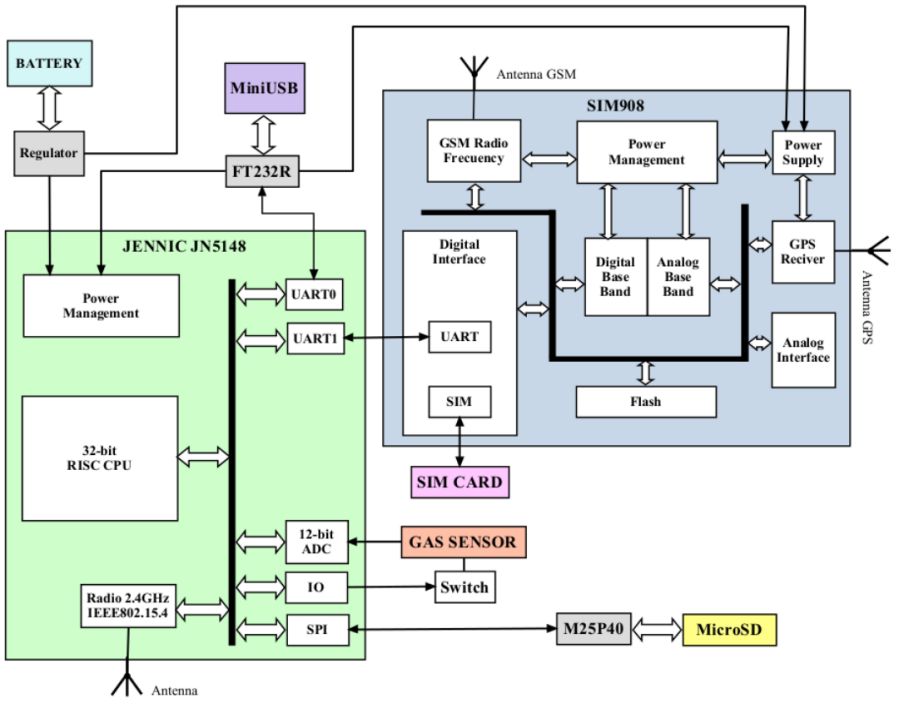

Figure 5.2: Block scheme of the readout board with MCU, modem and external components.

increasing the energy autonomy of the whole platform, carrier and sensing system. For this reason we designed the embedded sensing system to be modular and, on top of it, we developed and simulated a novel algorithm that exploits adaptive sampling to locate the gas leakage, minimizing at the same time the energy required by the sensing element. We conducted a set of simulations to evaluate the optimal trade-off between speed of the carrier, energy consumption and gas localization capability. Finally, we conducted a design space exploration to evaluate the impact of introducing solar panels on top of commercially available drones. As a result we were able to determine the trade-offs among UAV size, payload capacity, energy consumption, solar panel size and expected autonomy gain.

5.2 Embedded Platform

The architecture of the proposed mobile sensing system consists of two hardware units and antennas, sd-card and sim-card for a total occupancy of 4x4 cm and a weight less than 30 g (Fig. 5.1), optimized to be used with mobile carriers. The primary hardware unit hosts the MCU and the GSM/GPS modem and a bunch of auxiliary peripherals. The secondary unit is a gas sensor board designed to easily interface with the MCU.

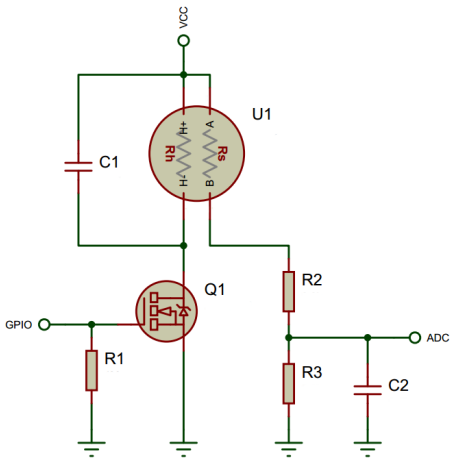

Figure 5.3: Electrical scheme of the digitally controlled MOX sensors readout module.

5.2.1 Primary Unit

The MCU used is the same JN5148[1] chip described in previous chapter, which offers high performance and low power consumption. The integrated IEEE 802.15.4 compliant radio transceiver exhibits high power consumption and a short range of communication, thus we used it in test phase for real-time synchronization of the data with the base station, but we avoided its use in the final algorithm to minimize energy consumption.

Implementation of tracking system requires receiving the GPS location, and the integration with GSM/GPRS modem, to be able to send or receive data over long distances. The current position of the object is provided by the SIM-908[2] integrated GSM/GPS transceiver, and then the data is sent via GPRS connection towards a dedicate web server with static IP, using TCP/IP protocol. The GPS engine is controlled by the GSM engine, so when it is necessary to run GPS the GSM engine must be powered on and not in SLEEP mode.

The power unit consists of a rechargeable, 3.6 V, 1800 mAh, Lithium-ion battery, from a Nokia N73 smartphone. Fig. 7.2 depicts the mobile gas sensing system's block scheme and the connections between modules.

5.2.2 Secondary Unit

The small secondary unit has been designed to easily interface the gas sensor with the MCU using four lines, two for the supply, one digital enable signal and the sensor's analog output. We realized a dedicated printed circuit board (PCB - shown in Fig. 5.1-left), to host the necessary signal conditioning circuitry. Fig. 5.3 shows its design where `GPIO` stands for the digital enabling line, `ADC` is the analog output signal and finally `VCC` is a 3 V line form the regulated supply of the MCU. Alternatively, it could be possible to use one Digital to Analog (DAC) converter module of the MCU to provide modulated supply

[1] `http://www.jennic.com/products/wireless_microcontrollers/`
[2] `http://www.simcom.us/product_detail.php?cid=1&pid=38`

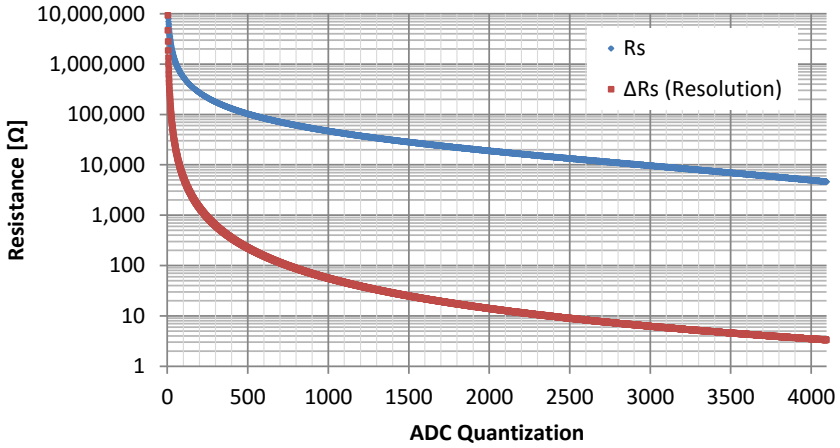

Figure 5.4: Numerical simulation of the non-linear response and resolution of the readout circuit over ADC quantization levels.

to the sensing element, if required; which is actually useless with the commercial devices we considered in this work.

The sensing system developed for measuring gases includes a single off-the-shelf available MOX sensor. This unit has been designed with standard socket for the sensing element to be interchangeable, to measure different gases and to fit the specific application. A stage of conditioning (`R2`, `R3` and `C2`) is required to adapt the analog signal of the MOX gas sensor to levels compatible with the ADC. Finally, the installation of a n-type MOSFET (`Q1`), used to switch the power signal on/off, keeps the power drained by the sensor under control. In this configuration, the MOX turns off at the same time as the device, saving power.

We tuned the circuit considering the MICS sensors already presented, that are nominally driven at about 76 mW. The temperature rise of the heating resistor (R_h) is proportional to the applied power and the nominal working point is at about 340° C. The gas measurement is taken by sampling the sensing resistance (R_s) which variation is inversely proportional to the volatile concentration, the higher the concentration, the higher the number of charge carriers in the sensitive layer, the lower the resistance. The operating range of chemoresistive sensors is quite high (from few MΩ to tens of kΩ), thus a high resolution ADC module is required to precisely measure both supply voltage and sensor response. We used the 12-bit SAR ADC integrated in the MCU.

Thanks to the inverse relationship between V and R and the voltage divider readout circuit, we obtained a non-linear resolution in the two decade range under evaluation; while the output parallel capacitor helps to stabilize the signal. By carefully dimensioning the components (`R2`=2.2kΩ, `R3`=6.8kΩ, `Vcc`=3.2V) and the ADC reference (`Vref_adc`=1.6V) we achieved an adequate resolution for the application scenario. Numerical estimation of the readout circuit behavior is presented in Fig. 5.4, output resistance R_S and resolution ΔR_S have been computed as a function of the ADC quantization levels. For example, we get ≈20kΩ resolution when the sensor's response is in the 1MΩ range, 300Ω @ 100kΩ down to 6Ω @ 10kΩ. In the whole working range we

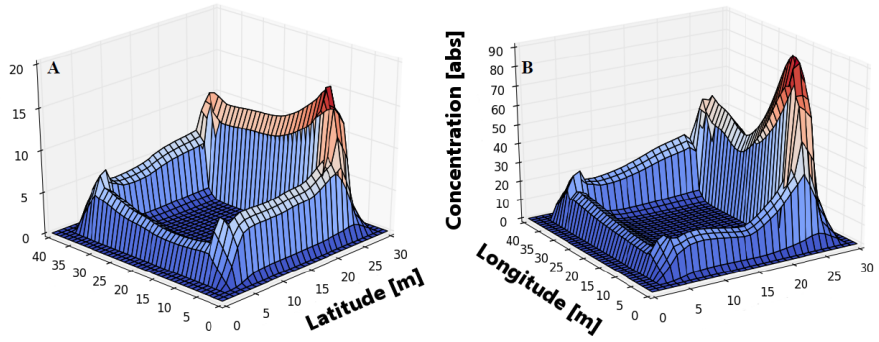

Figure 5.5: Preliminary sensitivity test outdoor n.1, the MOX sensor was kept vertically above the gas sources 20 cm (case-A) and 10 cm (case-B) respectively.

get a resolution of at least 2 orders of magnitude lower than the signal of interest. The power drained by these devices is higher than any other peripheral in the module, as it is shown in Table 5.1. This motivates the need of a careful optimization of the monitoring algorithm to maximize the autonomy of the complete module.

5.3 Preliminary Embedded Platform Characterization

To evaluate the capability of MOX sensors to timely react with leakages in outdoor environments we conducted two tests on the field. These experiments were carried using the prototype acquisition board, equipped with the MiCS-5121 VOCs targeted sensor. We carried the system along a predefined rectangular path (20 m x 30 m long) where four gas sources were located at the corners: two bottles filled with different amount of acetone (coordinates 5:5 - latitude:longitude - for small amount of reagent and 25:5 for large quantity in the Fig. 5.5 and 5.6) and two with different amount of isopropyl alcohol (coordinates 5:35 for the bottle with small amount and 25:35 for the large quantity in the pictures). The path was covered starting from position 0:0 toward location 5:5 and then following a counter-clockwise direction in all the experiments, with constant speed.

The first test aimed at verifying the sensor response in outdoor environment. Results shown in Fig. 5.5 illustrate the sensor response of two consecutive laps along the path with the sensing element vertically above the source and close approximately 20 cm during the first round (A) and 10 cm during the second (B), respectively. The reference concentration (flat area) represent the stable response of the sensor after the warm-up phase following the system switch-on. Obviously the shorter the distance the higher the response (notice the higher absolute value in picture 5.5-B), also with respect to the amount of reagent in the bottle. Moreover the long reaction and recovery time of sensor generates the wrong localization of the source that is particularly evident in case-B where the bottle of acetone located in position 25:5 results in a response spike in position 25:20.

In the second test, summarized in Fig. 5.6, we evaluated the sensitivity increasing

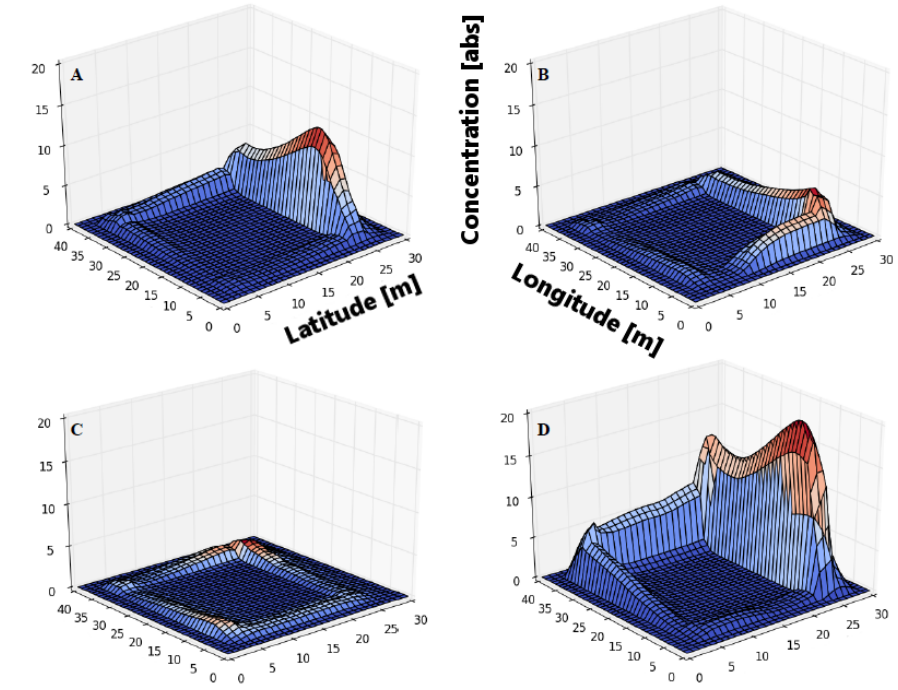

Figure 5.6: Preliminary sensitivity test outdoor n.2, the MOX sensor was kept at different distances from the gas sources, 50 cm vertically (case-A), ≈ 70 cm diagonally (case-B), ≈ 110 cm diagonally (case-C) and 30 cm vertically (case-D).

Table 5.1: Average current consumption required by the operating modes, computed using 1 Ω shunt resistor.

Status			**Consumption**
MCU	**GPS/GSM**	**MOX**	**[mA]**
sleep	off	off	<1
on	off	off	12.50
on	off	on	51.25
on	on	off	33.75
on	on	on	75.00
on (LogSD)	off	off	35.60
on (LogSD)	on	off	56.25

Figure 5.7: Pictures of the UAV carrier (hexacopter with 80 cm diameter) and gas source taken during tests on the field.

the relative distance from sensor and sources. We started with the sensor 50 cm vertically above the bottle (case-A), then 50 cm both vertically and horizontally ($\approx$ 70 cm of distance, case-B), 50 cm vertically and 100 cm horizontally apart from the bottles during the third round (case-C, $\approx$ 110 cm of distance) and the last lap was evaluated with 30 cm vertical distance from the bottles (case-D). Despite the low concentration of reagents (with respect to what we can expect from a pipe's leakage) the results show that the system is able to detect target volatile substances even without a forced convection device (pump or tube for active and passive convection, respectively).

Both tests were conducted outdoors, under moderate wind, variable environmental conditions (temperature, sun irradiance, humidity, etc.) and a non-constant aerial transportation and speed of propagation of volatiles. That results in fluctuations of the response, notice for example the difference (2 to 3 points in absolute value) between Fig. 5.5-B and 5.6-D. Both experiments demonstrate the capability of the system to detect and locate the region of the gas-leakages (characterized by a maximum) in a real environment, which is the major goal, while a quantitative measurement is marginal because it varies in the medium-long term.

5.4 Field Tests with UAV

The purpose of the real field experiments were to evaluate the sensitivity of MOX sensors in presence of turbulent air flow generated by the drone's propellers and the placement of the embedded system that was fastened below the main body of the UAV, to have the sensing unit perpendicularly oriented with respect to the ground.

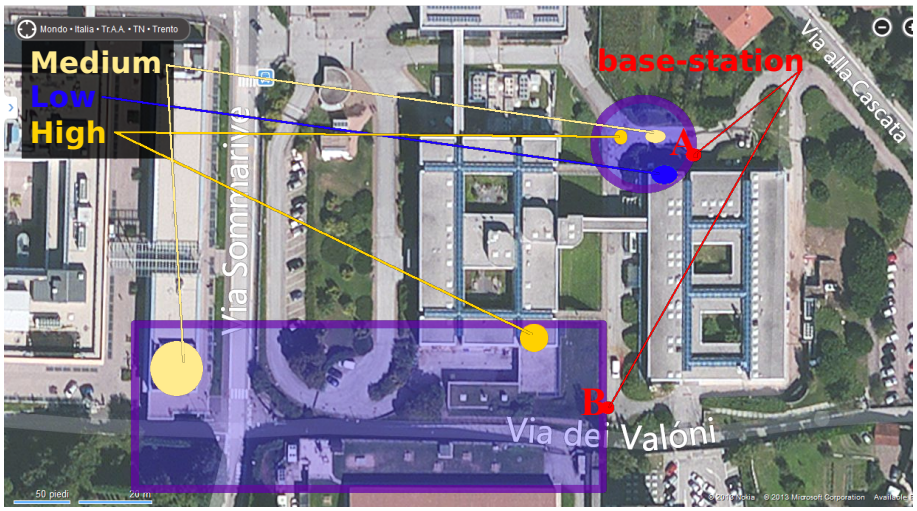

Figure 5.8: Summary of the field tests with hexacopter. Superposition of the measured gas distribution map and aerial view of the test location.

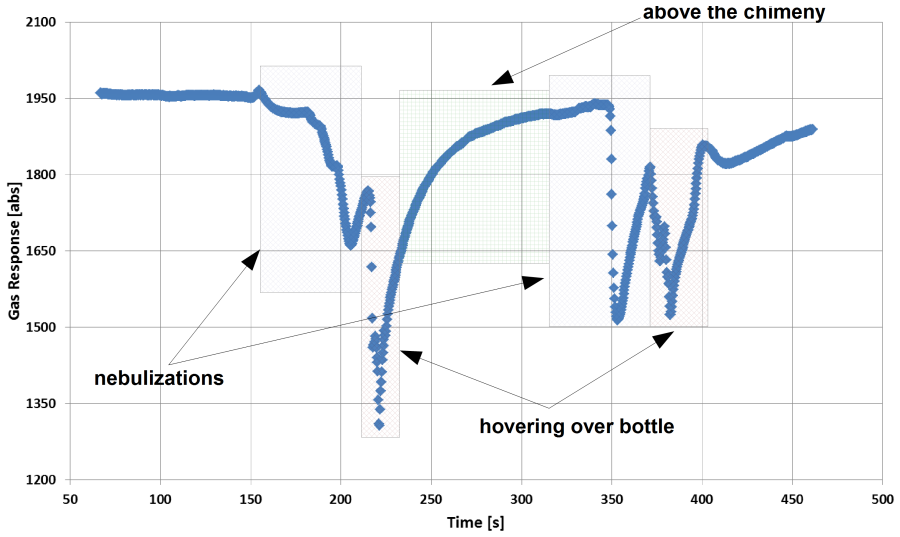

Figure 5.9: Detailed results from the first flight test with hexacopter carrying the gas mapping instrumentation.

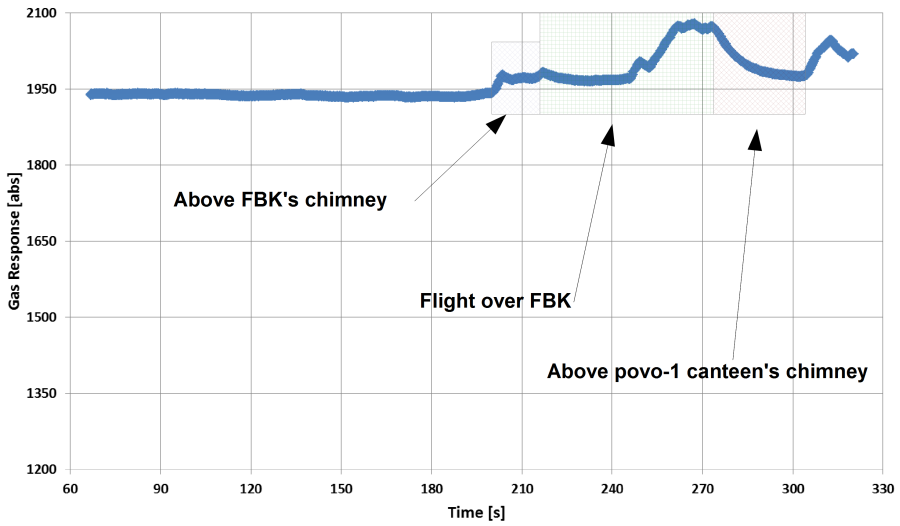

Figure 5.10: Detailed results from the second flight test with hexacopter carrying the gas mapping instrumentation.

We run several experiments using a DJI hexacopter[3]; this drone is suitable for limited region gas distribution mapping thanks to its 80 cm diameter and 10 Ah battery capacity which assures 15 min autonomy with up to 2.5 kg maximum payload. The experiments took place in the surroundings of the University and Fondazione Bruno Kessler (FBK) buildings (the technological center in the city of Trento) using the MICS-5121 VOCs targeted sensor. Results of two experiments are summarized using the map shown in Fig. 5.8. In the first one we tested the system's sensitivity during the flight using a bottle of isopropyl alcohol in a parking slot as leakage source (Fig. 5.7). In the second experiment the UAV surveyed the kitchens of the FBK center (smaller spot in the rectangle of Fig. 5.8 marked with high concentration) and the University canteen (bigger spot in the same rectangle).

Detailed responses are reported in Fig. 5.9 and 5.10. In both pictures, the decreasing spikes highlight the sensitivity of the system to react to volatile substances in air (marked with arrows and boxes). The second experiment, that covered a wider spatial range, demonstrates the capability of the system to react to changes in the environment, the increase in the response can be noticed when the drone is in between the two buildings (less VOCs concentrations can be assumed) while a decrease occurs just on top of the University canteen's chimneys that demonstrates the gas sources localization during movements.

These results confirms that environmental conditions and cross correlation of volatile chemicals strongly affect MOX's response, moreover the characterization of these responses in open environments is a challenging task, to this reason we plan to move to multiparametric custom MOX sensors to eventually achieve volatile discrimination. We observed that the air flow generated by propellers does not blow away the volatile com-

[3] `http://www.dji.com/product/spreading-wings-s800/feature`

Table 5.2: Measured time required by the single sampling stages

Task	**Time [ms]**
LogSD	150
ADC Measure	1026
GPS request	5125
Mean Measure	5
Total	**6805**

pounds diffused in air thus, from a sensitivity perspective, the measurement unit's placement and the passive transportation, resulted in a good choice and was decided not to use any auxiliary convection equipment (active or semi-active) also in this configuration.

In these experiments, the prototype unit demonstrated ≈30 min autonomy of continuous gas sampling, data logging and wireless data streaming, which is longer than the autonomy of the employed UAV. Hence, we decided to further increase the energy autonomy by introducing improved sampling strategies based on duty-cycled sensor usage replacing real-time streaming, but aggressive data compression using compressive sensing techniques [12] can be also useful.

5.5 Energy Optimization

The definition of the monitoring strategy to optimize the lifetime of the mobile sensing system has been carried out using a simulating framework developed ad hoc, in Matlab environment, considering an Unmanned Aerial Vehicle (UAV) as carrier. This framework has been designed using generic models (except for the gas monitoring prototype) to be modular, customizable and ready for improvements. We firstly defined the monitoring strategy for the sensing module, then we characterized behavior and performance of each single stage of the system during monitoring to finally build a detailed model of execution time and energy requirements. We simulated the UAV considering a set of constant-speed experiments along straight paths. This simplistic approach perfectly fits leakage detection tasks for example in outdoor pipelines monitoring. In the following, firstly the details of the monitoring strategy and model of the mobile sensing system are provided; secondly the energy optimization algorithm for chemical concentration measurement is discussed.

5.5.1 Model Development

We developed the model of the mobile sensing system by characterizing the execution of the embedded application. This is based on a finite state machine where firstly the MCU and the GPS/GSM module are initialized, along with other peripherals like the ADC and the SD-card controller for logging. Then starts the main loop where the gas measurement are collected and time-stamped with the GPS provided information (UTC

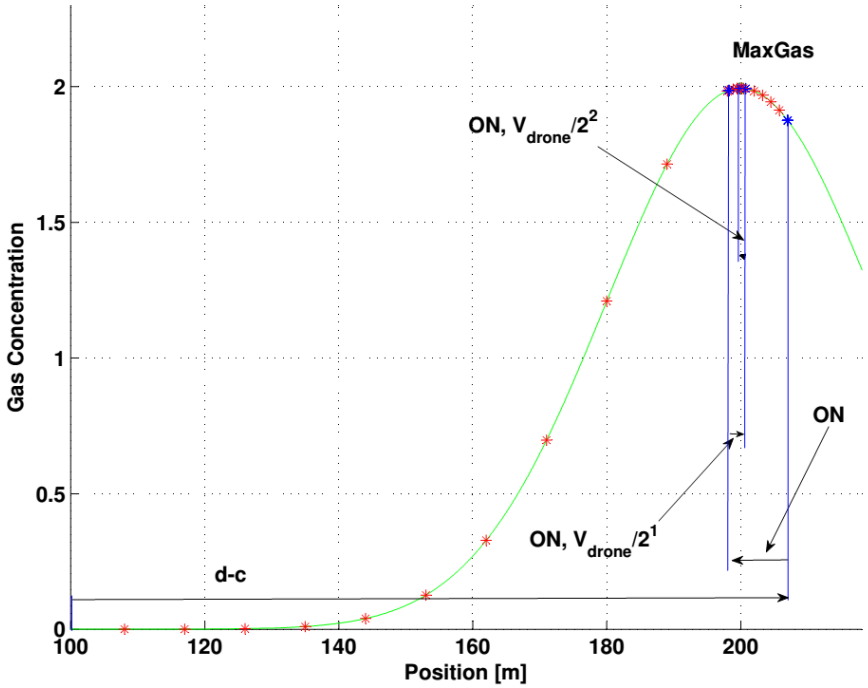

Figure 5.11: Graphical representation of the system evolution (monitoring instrument and carrier) during the optimization routine.

time and coordinates). Gas measurement occurs for 1 s with a fixed rate of 500 μs per sample, for a total of 2000 $samples$ per measurement.

The heating time of 1 s is considered sufficient for MOX devices when the duty-cycle is higher than $\approx$ 1% (e.g. 1 s against 1 min sleep-time), to ensure the quality of service and fast response in outdoor environments where we can expect higher concentrations of pollutants (exahust gases and dust) with respect to indoor environments. We have chosen this standard routine based on the results of several experiments not reported for the sake of summary and based on the observations made in the development of the indoor monitoring strategies. In the end the mobile system is set to sleep-mode for a number of seconds which depends on the duty-cycle defined by the optimization algorithm. We also included an Alert System to prevent risks and damages to people or properties, a routine that sends text messages to a predefined mobile phone; although, we did not modeled this sub-routine for the evaluation of the optimization algorithm. It consists on a plain text message containing relevant information about the system, such as the vehicle location and gas concentration. In addition, the system is also able to notify the user of other potential risks, like low battery.

The model used in the simulating framework is based on the sum of the energy required by each state, computed as the product of time and average power (considering the reference voltage). Notice that the current required by the sleep-mode is less than 1 mA (in the order of few hundreds of μA) so it is possible to neglect this contribution in

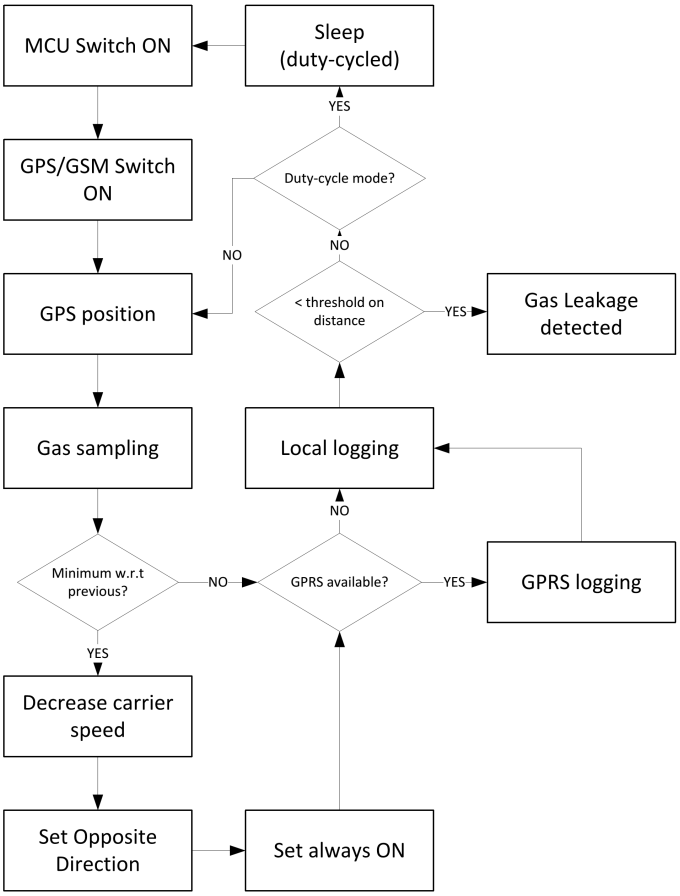

Figure 5.12: Block scheme of the tracking algorithm based on the hill-climbing approach.

the sum without loss of generality. Tables 5.1 and 5.2 present the mean values used to compute the model, obtained by averaging a dozen of experiments. In this characterization we used a custom rechargeable battery pack made of 3, AA-type batteries connected in series (1.5 V, 2500mAh each). The current has been measured using 1 Ω shunt resistor in series with the negative supply pole, while the time has been measured using an Agilent DSO7032A Oscilloscope[4]. If the system works continuously in ON state, with a mean consumption of 75 mA, it lasts for 1 day and 9 hours approximately. Which is longer than the expected UAV autonomy (for commercial devices). The main tuning parameter, that is the optimization target, is the duty-cycle (dc subscript in the following pictures) defined as the overall ON-time, regardless of the state, divided by the total time (ON-time and sleep-interval).

[4] http://www.keysight.com/en/pd-1293547-pn-DSO7032A/oscilloscope-350-mhz-2-channels

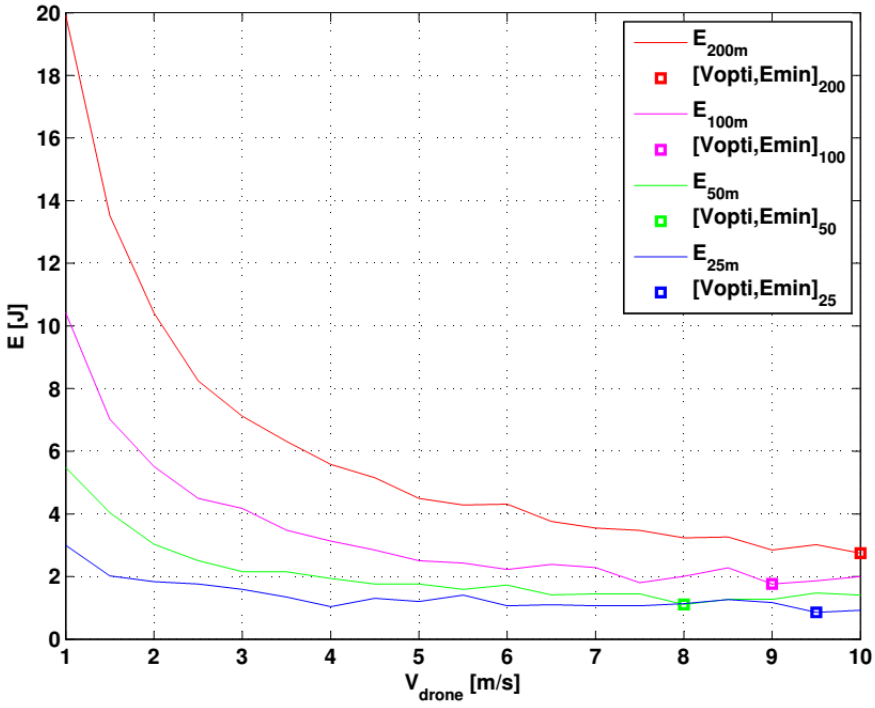

Figure 5.13: Results of the simulations in terms of optimal carrier speed against mobile system energy consumption and gas leakage locations.

5.5.2 Gas Leak Localization Algorithm

The gas source location strategy optimizes the speed of the robot while minimizing the energy consumption of the whole monitoring system. We implemented the mobile sensing system and UAV models in a discrete-time Matlab framework, considering a 2D representation of the environment where the UAV carrier moves. The first model has been presented above, while for the second one we used a linear motion equation. The shape of the gas distribution must be roughly circular with a strong central peak, being possible to use the maximum response peak of the gas sensor as an approximation of the gas source location. Previous studies conclude that the various theories of contaminants' distribution function tend to the Gaussian function. Thus we modeled the profile of gas concentration measures, using a 2D Gaussian distribution model. Fig. 5.11 shows the qualitative behavior of the algorithm, here the green line illustrates the 2D projection of the concentration along the path.

The searching is accomplished by applying a exploration strategy based on the "hill climbing" algorithm, summarized in the block scheme of Fig. 5.12. This strategy always converges to the optimal solution in the case of a single gas source. The optimal speed of the carrier is found after an exhaustive search of the state space determined by the set of initial carrier speeds (from 1 m/s up to 10 m/s) and duty-cycles. As described above the ON-time is fixed to ensure a reliable gas measurement, while the sleep-time was set as a

parameter to compare the performance in different configurations (employed values are 11%, 20%, 33%, 50%, 55% and 71% duty-cycles). The gas leakage is placed at fixed distances from the carrier's starting position (reference 0 in the following) in the range from 25 m to 200 m.

The exploration algorithm operates by simulating the movements of the UAV carrier following the linear motion model, every second the position is updated while the gas reading is updated according to the duty-cycle. Once a new gas sample is obtained it is compared with the previous one. If this is greater or equal than the previous one, then make it the current reference state called maximum (red points in Fig. 5.11). The algorithm continues until a decrease in the response is found. At this point the current state is marked as minimum (blue points in Fig. 5.11), the speed of the carrier is halved and the movement is oriented in the opposite direction. This process iterates until the distance between two successive minima is less than a fixed threshold (1 m). The loop is performed for all possible starting velocity of the carrier which decrease according to an exponential behavior described by the following equation:

$$V_{drone}^{next} = V_{drone}^{actual}/2^n$$

where n is the iteration number, counting from 0 to N. For example, with a starting speed $V_{drone}^{start} = 10\ [m/s]$, if we assume to locate the gas leakage after 5 iterations we get a final speed of:

$$V_{drone}^{end} = 10/2^4 = 0.625\ [m/s] \tag{5.1}$$

Finally the algorithm calculates the energy consumption to locate the gas leakage, and compares the energy spent for all possible initial speeds. Finally it returns the optimal speed that minimizes the energy consumption.

5.5.3 Simulated Results

A selection of results is presented from Fig. 5.13 to 5.15. We can observe (Fig. 5.13) that the optimal speed is usually fast, which is intuitive, less time implies less consumed energy for the mobile sensing system. Moreover the optimal one depends on the distance from the gas source. Clearly when the distance to a gas source increases, the energy consumption increases as well. It is important to notice that the difference of consumed energy between high and low speeds decrease when the distance to a gas source decreases. Fig. 5.14 compares the duty-cycled strategy with a continuous monitoring approach, to underline its effectiveness especially at low speeds. Finally, Fig. 5.15 is useful to underline the remarkable difference between duty-cycles effects on current consumption which obviously require a careful evaluation. From these results we can conclude that the sampling resolution must be as quickly as possible to rapidly locate the area of interest and eventually to cope with the rapid dynamics intrinsic to gas propagation in real environments. However, to find the maximum with high accuracy, this speed must be reduced promptly, that is the reason for the exponential decrease adopted in the proposed approach.

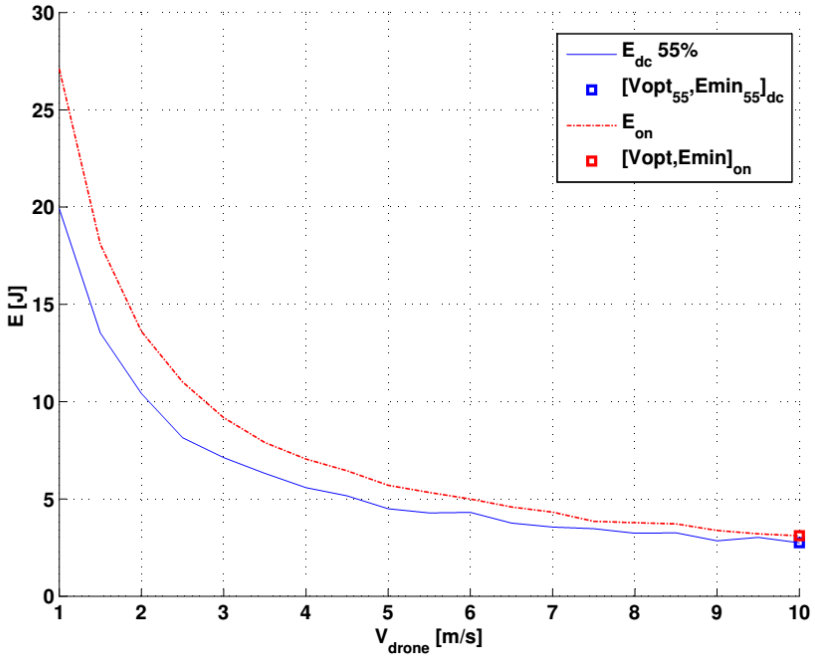

Figure 5.14: Results of the simulations in terms of energy consumption of duty-cycled and continuous monitoring strategies as a function of the carrier speed.

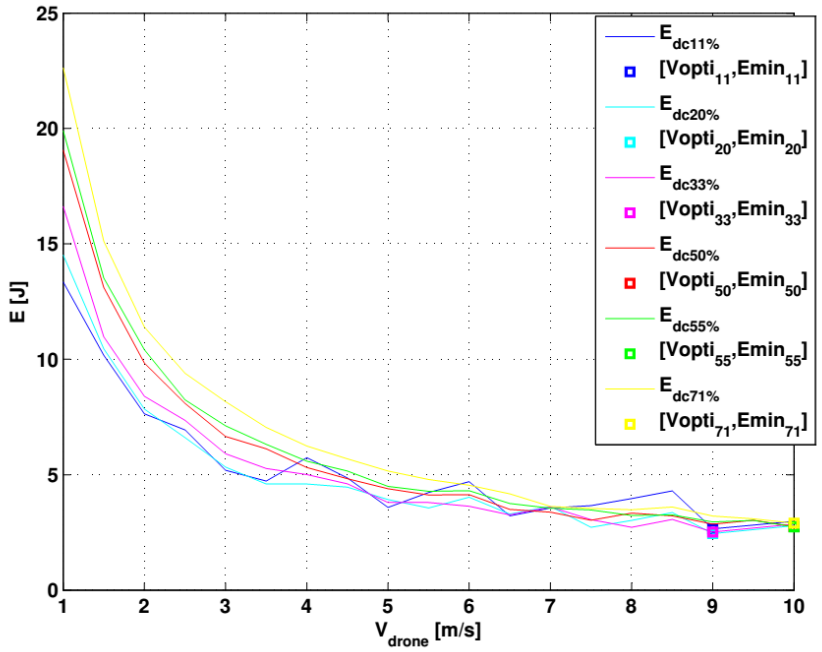

Figure 5.15: Results of the simulations in terms of energy consumption of different duty-cycled monitoring approaches as a function of the carrier speed.

Table 5.3: Comparison of main features of commercial copter-like drones

Copter Name	Weight [g]	Payload Max [g]	Voltage [V_{BAT}]	Capacity [mAh]	Autonomy [min]	Size [mm]	**Power [W]**
Hubsan nano q4	20	0	3.7	100	5	50x50	**4.44**
Crazyflie 2.0	27	15	3.7	240	7	92x92x29	**7.61**
DJI Phantom-2 v+	1242	800	11.1	5200	25	350 (d)	**138.53**
Hubsan x4 pro	1060	360	11.1	7000	30	370 (d)	**155.40**
Parrot ar.2.0 classic	400	<100	11.1	1000	12	570 (d)	**55.50**
3dr iris+	1282	400	11.1	5100	16	550 (d)	**212.29**
3dr x8+	2560	800	14.8	10000	15	350x510x200	**592.00**

5.6 Renewable Energy Sources Integration

In this last section we propose the design space exploration for the study of solar powered copter-like drones. This analysis is meant to evaluate the possible performance boost, in terms of extension of the flight autonomy, that could be achieved introducing solar energy harvesting capabilities in the drone carrier.

Copter-like UAVs, available in the market, have limited flight autonomy which severely limits their employment in wide-area unmanned monitoring applications. The main characteristics of a selection of commercial UAVs, provided by manufacturers, is presented in Table 5.3 for comparison, ranging from the toy-like Hubsan Nano Q4 up to the professional 3dr x8+ used by film-makers for stable aerial shooting. Maximum payload is the first parameter to consider in the choice of a carrier for custom monitoring applications, since it poses an upper bound to the size of the equipment that can be used without jeopardizing take-off capability and without reducing the flight autonomy below 50% of the nominal value.

By analyzing the information provided by manufacturers in data-sheets we computed the nominal power consumption of the copters (highlighted in the table) according to the following equation:

$$P_{[W]} = \frac{Capacity_{[C]}}{Autonomy_{[s]}} \cdot V_{BAT}$$

Those values are representative of nominal flying conditions (no wind, medium and constant speed) and account only for the drone's motion manually driven using a remote controller without payload, no auto-pilot (using GPS) and no video streaming/recording (where available). Results demonstrate that power consumption is remarkably high even for the very small copters.

Considering the negligible payload introduced by the proposed gas sensing equipment with respect to the nominal weight of the drones (excluding the first two in the list), we propose to extend the flight autonomy with the method already demonstrated with airplane-like drones, namely including photovoltaic panels in the structure of the

Table 5.4: Expected PV panels performance gain

Copter Name	Area [cm^2]	PV Power [W]	PV/P [%]	Gain [s]
DJI Phantom-2 v+	612.50	4.56	3.29	49
Hubsan x4 pro	684.50	5.09	3.28	59
Parrot ar.2.0 classic	1624.50	12.09	**21.78**	**157**
3dr iris+	1512.50	11.25	5.30	51
3dr x8+	1785.00	13.28	2.24	20

copter to harvest solar energy during the flight, as shown in the concept depicted in Fig. 5.16, where a polycristalline pv panel has been stuck between body and propellers of a Crazyflie 2.0 micro-drone (withoout gas sensing board).

Drones can fly only in clear sky conditions with low wind intensity, due to strict safety regulations promoted in the last years. Optimal conditions also for PV panels to achieve the maximum power production. Several questions arise in this context, for example, how can we modify the shape of the copter to host one or more panels? Which is the best orientation to maximize the conversion efficiency in any displacement? And what is the impact on the payload of the energy conditioning circuitry? Should we add another accumulator or not? According to our point of view, the first thing to evaluate is the impact on the autonomy that can be achieved. In other words, how long the flight-time can be extended. To address this issue, we analyzed several PV panels available on the market with the aim to select suitable components for the evaluation. Table 5.5 lists few commercial devices that are arranged in two groups, traditional and thin-film (bendable) photovoltaic modules. Numbers show that thin-film panels are extremely lightweight and could be a good candidate for the UAV's upgrade, however their efficiency (expressed here as power to size ratio) it is slightly more than half of the traditional ones.

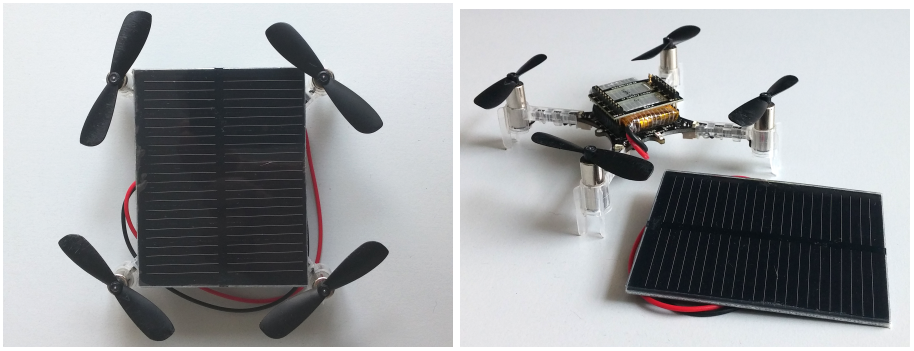

Figure 5.16: Experiment of a solar powered drone using a CrazyFlie 2.0 and a 70x55 mm solar panel.

Table 5.5: Comparison of commercial photovoltaic panels features

Panel Model	Power [W]	V_{OC} [V]	I_{SC} [mA]	V_{MPP} [V]	I_{MPP} [mA]	Size [mm] LxWxH	Weight [g]	Power/area [W/cm^2]
Traditional								
Multicomp MC-SP0.8-NF-GCS	0.8	4.8	2.30	3.85	210	140x100x4.9	134	0.0057
BP Solar MSX-005F	0.446	4.6	160	3.3	150	114.3x66.8x3	160	0.0006
BP Solar MSX-01F	1.2	10.3	160	7.5	150	127x127x3	340	0.0074
BP Solar SX-5M	4.5	20.5	300	16.5	270	245x269x22.6	800	0.0068
Solar Technology STP005BP	5	21	390	16.8	300	306x218x25	1000	0.0075
Solar Technology STP010BP	10	21	660	16.8	590	397x280x25	1500	0.0090
ThinFilm								
PowerFilm MP7.2-150	1.44	10.5	150	7.2	220	253x146x0.61	25.9	0.0039
PowerFilm MPT15-150	1.54	19	120	15.4	110	253x146x0.61	26	0.0042

Table 5.4 presents the ration between the runtime power consumption and the expected power contribution from the PV harvester. It indicates the percentage of power which can be extracted during the flight from the environment, using PV panels with 0.0074 W/cm^2 efficiency (as for MSX-01F panels) on the available surface of the drone. This method contributes to the overall power budget and allows to extend the flight autonomy correspondingly. Simulations show that, using drones that could carry the weight of such structure, the maximum achievable gain is less than 6% which means longer flight autonomy as reported in the last column of the table. In this comparison, only the lightweight Parrot AR 2.0 Drone could significantly boost its autonomy up to 20%.

Chapter 6

Monitoring for Free Exploiting Energy from Heat Dissipation

Controlled facilities like data centers' server rooms can be considered resource constrained environments since the presence of people must be avoided, thus maintenance and security must be carefully managed and controlled. In this chapter is presented the research effort of my group, in collaboration with master students, to enable energy neutral monitoring solutions in this specific case. The discussion of results continues also in Chapter 7 which presents the most recent achievements of this research.

6.1 Energy Harvesting in Constrained Environments

Technology advances enable the development of innovative architectures for embedded systems towards orthogonal directions. From one side, systems characterized by high computing performance exhibit lower power consumption and thermal dissipation compared to previous generation of personal computers. In fact, ARM-based devices are considered a promising solution for building computing server for cloud service providers and Web 2.0 applications [73, 94]. On the other side, technology advances offer also low-power microcontrollers that are able to manage multiple sensors and to communicate wirelessly to others with few mJ, and, thanks to their characteristics, it is possible to supply them with non conventional power sources [14].

There exists several souces of renewable (or free) energy that can be exploited to supply low-power embedded systems for monitoring applications. The most widely used one in outdoor applications are photovoltaic cells, thanks to their wide availability with different shapes and sizes in the maket. Despite the very low efficiency, $\leq$ 20 %, a single panel with the size of an A4 sheet can provide up to few hundreds of mW in clear sky and

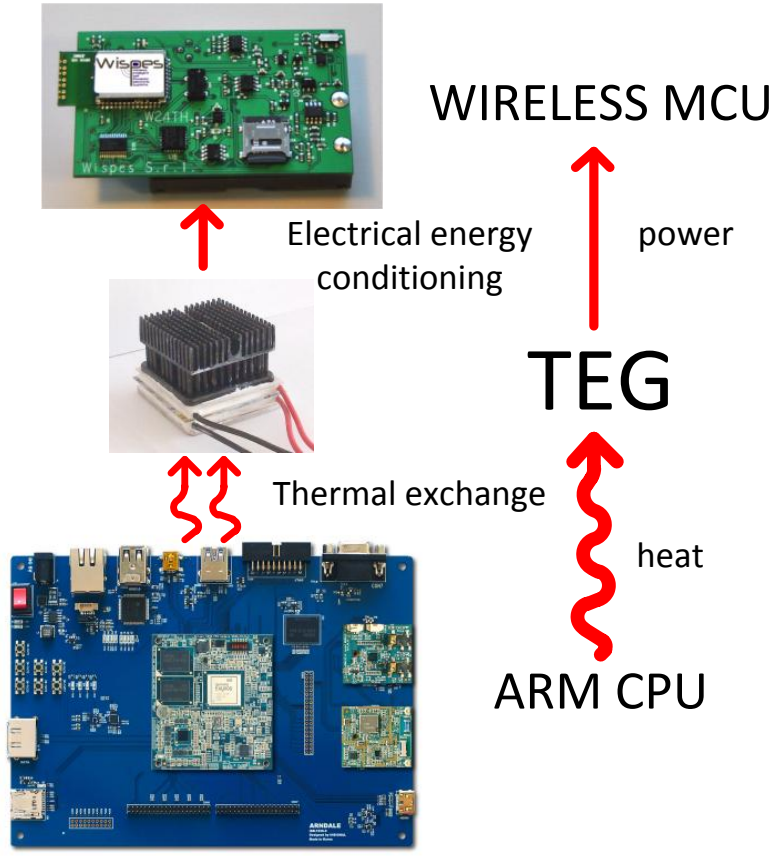

Figure 6.1: Schematic view of the target application scenario. The energy wasted as heat can be collected and transformed into electrical energy to supply wireless embedded systems.

direct irrandiance conditions, making this solution suitable for most of monitoring applications. Other harvesting technologies that has been proved effective in the literature, exploit for example piezoelectric properties of materials subject to mechanical stress to generate electricity; or electromagnetic radiation from antennas and magnetic field of currents flowing in a transmission line, to extract a fraction of electric energy. Finally, heat — or better thermal gradients — are interesting and promising sources of μJ energy to supply dedicated tasks.

Among the solutions presented we focused on heat recovery since it is the most promising solution to enable enegy neutral monitoring in contolled laboratory facilities and, in particular, data centers' server rooms. Wherever a computing unit is running, heat is generated as by-products and currently considered a wastage.

We demonstrated an energy neutral monitoring device supplied by electrical power harvested from heat dissipated by surrounding systems. In particular, the proposed design harvests heat wasted by high performance devices (such as server boards), while it

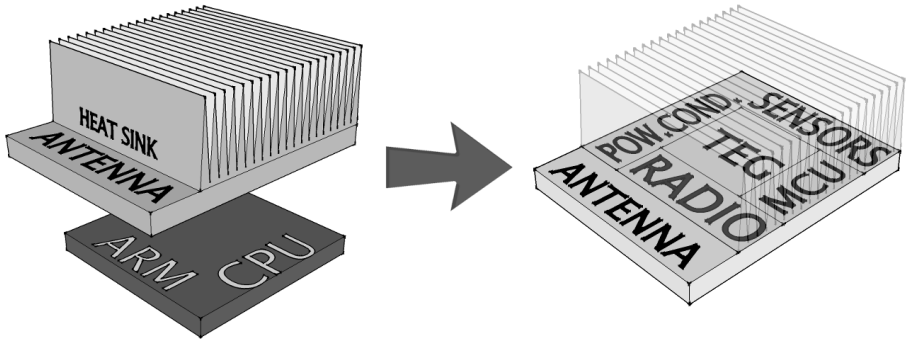

Figure 6.2: Exploded view of the envisioned SoC. Harvesting, conditioning and monitoring systems are embedded in the heat sink.

is able to monitor environmental parameters and valuable data to ensure safety inside server rooms. The system works in symbiosis with the server, since it relies on the server to guarantee its power supply, and acts as a data collector to monitor interesting features, including temperature and other quantities related to thermal and resources wastage.

In a generic computing device there are some hot points that represent a thermal source for harvesting, for example CPUs or GPUs. Power and thermal management constitutes an important trade-off between the performance and lifetime characteristic of a computing device. In fact, CPUs are designed to avoid high temperature by adopting dynamic voltage and frequency scaling, or other techniques to evade the necessity of using bulky cooling systems [11, 92]. Moreover, to guarantee durability, monitoring system temperature, and eventually the use of strategies for passive or active heat dissipation (heat sinks, fans, liquid cooling, Peltier cells, etc.) are necessary.

In addition, for safety and reliability reasons, there is the need to monitor restricted environments like server rooms, and Wireless Sensor Networks are a promising candidate to undertake this task [61]. Through the use of specific WSNs, it is possible to efficiently monitor many interesting information, such as: local room temperature, humidity, air flow, air quality, moisture leak and intrusion detection.

The objective of this study is to take advantage of the heat dissipated by the servers to power the monitoring sub-system, thus coupling technologies that are conceptually separated. Fig. 6.1 illustrates the basic components and the energy chain in our vision.

The design described provides monitoring service for the server and thanks to its low-power characteristics it is able to do this task for free. Contrary to battery powered systems, it does not require the intervention of an operator for battery replacement, and being independent from the supply Grid, it does not contribute to the energy bill, and can be deployed or moved easily.

The critical challenge in the design of such a system is to demonstrate that the presence of extra hardware does not affect the performance of the server. The proposed system composed of a thermoelectric harvesting module, an energy buffer and a wireless microcontroller (MCU) has been realized and demonstrated to be effective from both the monitoring and the heat dissipation points of view. To develop and validate the design, we focused on a single ARM based board — ARM CPU — that represents a single

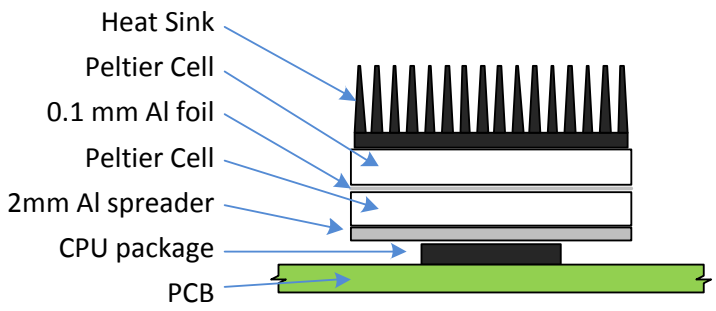

Figure 6.3: Thermoharvester schematic made of 2 Peltier cells, spreaders and sink in stacked fashion. We put thermal grease in between each layer of different materials to guarantee maximum thermal contact.

CPU inside a server room. On top of it we placed the thermo-harvesting generator (TEG a device that converts thermal gradients into electricity) to scavenge energy and supply a stand alone embedded wireless node - WSN NODE - through a dedicated conditioning circuit.

Analysis and design demonstrates the validity of the idea that will be expanded in the next Chapter where the whole embedded system into a *System on Chip* (SoC) device with the size and shape of a commercial heat sink, where harvesting, sensing and radio are embedded in the base of the sink itself (cfr. Fig. 6.2). In the future such systems will take advantage also from μ-scale TEGs [88] and μW radio transceivers [52] fields of research.

6.2 Thermoelectric Generators

Thermoelectric generators (TEG) are devices that convert thermal gradients into electrical energy, through a phenomenon called the Seebeck effect. Basically, the temperature difference across a junction between two different conductive materials (connected in a closed loop), induces a current flowing across the circuit. Such a principle is the opposite of the Peltier effect, for which a current injected across the metal junction generate a difference in temperature between the two sides of the circuit.

The TEG represents the basic component for the conversion of wasted heat into electrical power. It consists of a number of semiconductor elements (p-n junctions) that are thermally connected in parallel (but electrically in series), and enclosed in a package designed to spread the heat on the whole surface (can be metallic or ceramic). Each of these elements exhibits an electromotive force when exposed to a thermal difference (or thermal gradient) due to the Seebeck effect. Since these elements are electrically connected in series, each contribution sums up to build a voltage at the output metal connectors of the device. The generated voltage is proportional to the thermal difference at which the device is exposed to.

When operated as a generator, one side (or hot-side) of the device is heated to a temperature greater than the other side (or cold-side), and as a result a difference in voltage is observed at the output. However, a well-designed Peltier cooler is a mediocre

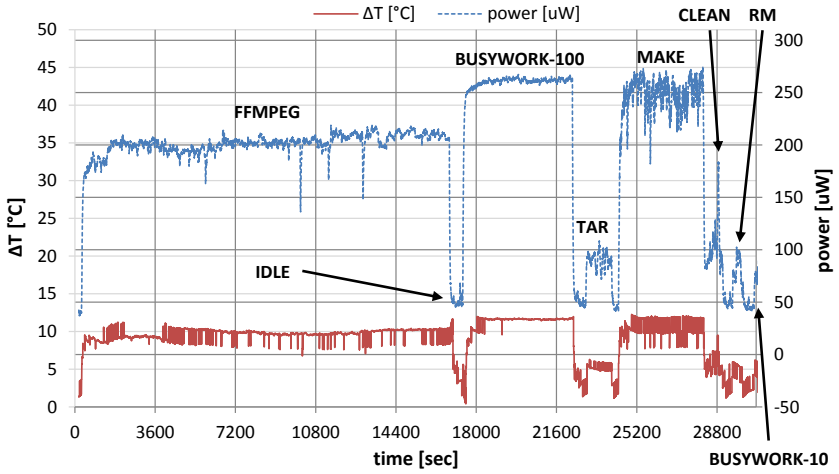

Figure 6.4: Output power of the Nextreme TEG placed on top of the Pandaboard.

thermoelectric generator, and vice versa due to different design and packaging requirements.

To estimate the order of magnitude of the power generated by thermoelectric generators in the context of computing systems, we have compared three different devices provided by different manufacturers. These devices differ in terms of aspect ratio, number of cells and their arrangement.

Two of these are specifically designed as thermoelectric generators, in fact they are made by microfabricated thermocouples optimized for the Seeback effect. The third one is a generic Peltier cell — generally used for cooling applications — that we used in two different configuration: a single cell, and two cells connected in series. The latter configuration is shown in Fig. 6.3, in the other case the Peltier/thin-spreader/Peltier stack is substituted by a single element. The devices we chose as TEGs are:

- *Nextreme* eTEG HV56 Thermoelectric Power Generator with 32 $\times$ 32 mm squared footprint, and a 31 mm $\times$ 33 mm power generator without its output power regulator [59];

- *Micropelt* TE-CORE7 TGP-751 ThermoHarvesting
  Power Module with a 43 $\times$ 33 mm heat sink, and a circular footprint with 5 mm radius [49].

- *Peltier-Cell* PE1-12706AC 40×40 mm squared cells.

We built the stacked structure presented here to optimize the thermal exchange between ARM CPUs and Peltier cells. The lower spreader (thickness 2 mm) allows the heat to distribute on the whole surface, while the one in between (thickness 0.1 mm) speeds up the exchange between cells, and spreads the heat once again. The spreaders are made of aluminum because of its thermal efficiency, low cost and market availability. The last layer of the stack is a commercial heat spreader that optimizes the heat exchange with the environment, to guarantee a lower temperature on the cold-side.

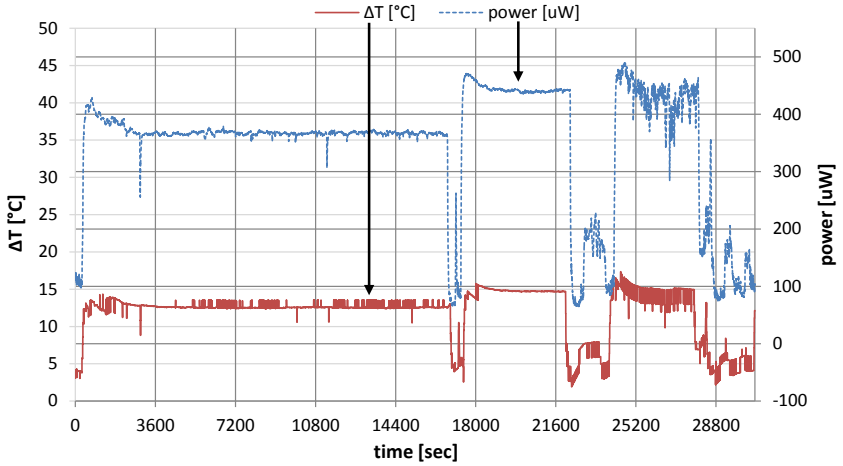

Figure 6.5: Output power of the Micropelt on Pandaboard.

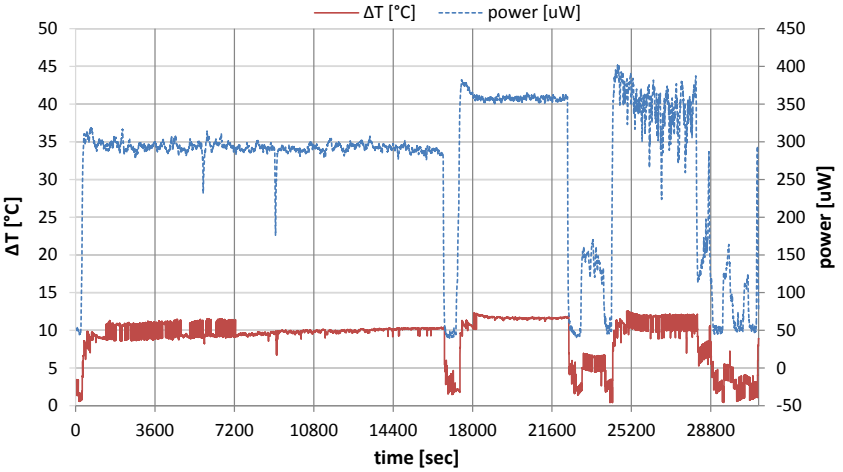

Figure 6.6: Output power of a single Peltier cell on Pandaboard.

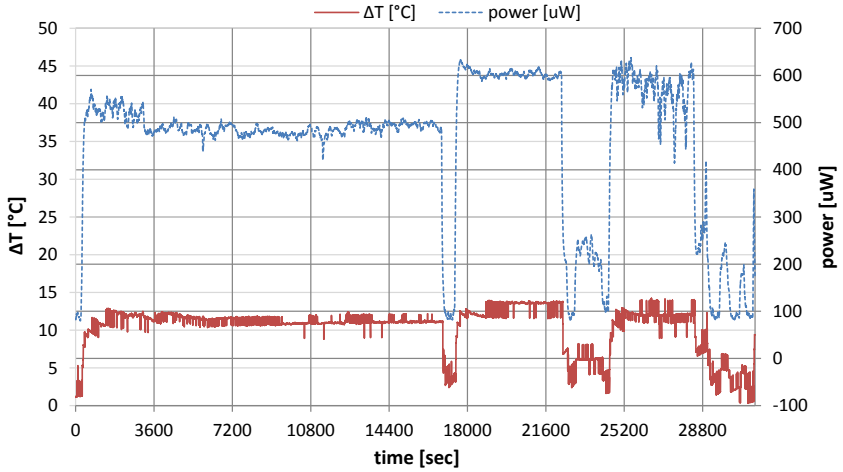

Figure 6.7: Output power of a series of two peltier cells on Pandaboard.

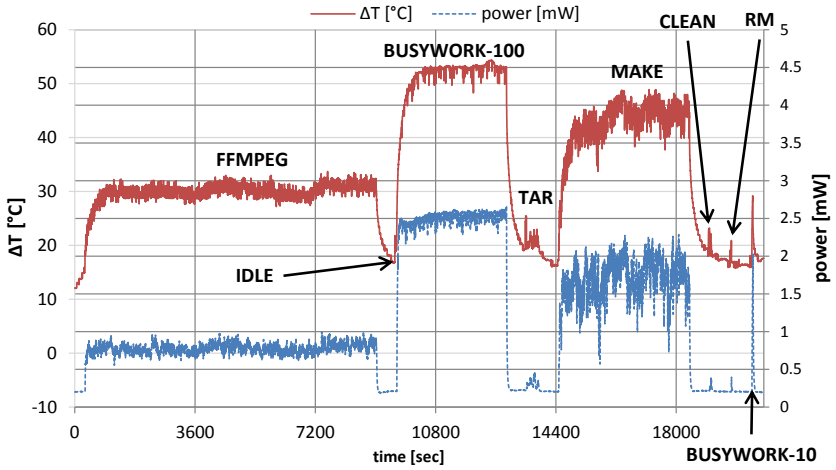

Figure 6.8: Output power of the Nextreme on Arndale.

6.2.1 Setup for TEGs Characterization

The characterization of harvesters performance are obtained using open circuit voltage and impedance matched power. The former one is used to evaluate the Seebeck coefficient of a given cell with N number of p-n junctions, useful for device characterization and performance evaluation, using the equation:

$$\alpha = \frac{V_{oc}}{N\Delta T}$$

The latter states the maximum power that the cell can provide when is connected to a matched load — a load with exactly the same impedance of the generator — which is of interest to correctly size the design of the monitoring embedded system.

Given the target of implementing an energy neutral monitoring system based on wireless embedded devices, we were mostly interested about performance in terms of maximum power and in different working conditions of the selected devices. To this end, we used two embedded system boards based on ARM processors, as "heater":

- a Pandaboard with a 1 GHz ARM Cortex A9 CPU and 1 GB of low power DDR2 RAM [65];
- a Samsung Arndale equipped with an Exynos 1.7 GHz dual-core ARM Cortex A15 processor with 2 GB DDR3 RAM 800 MHz [9].

Both systems run Linux kernel version 3.10.

We conducted the characterization of the TEGs listed above by executing different tasks in sequence in a benchmark fashion. In particular the tasks were selected to achieve a range of time length and CPU loads (with fixed clock frequency set as the maximum of each ARM CPU board):

- **video encoding** using `ffmpeg`[1] with four threads, which converts a two hour long movie;

[1] ffmpeg.org

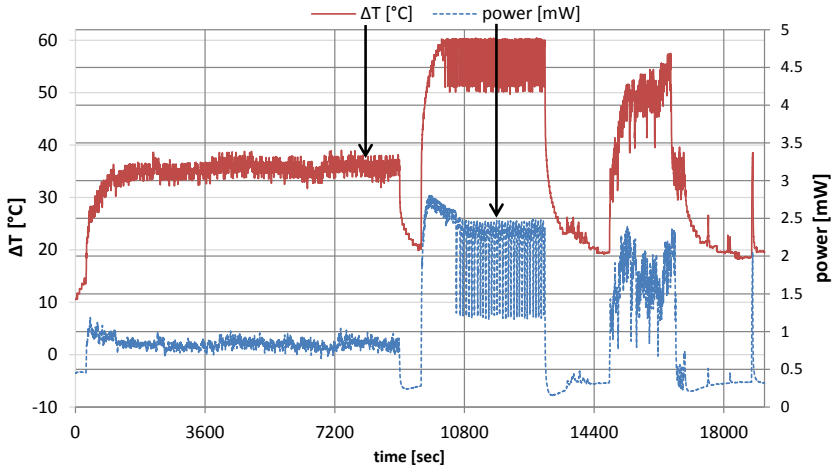

Figure 6.9: Output power of the Micropelt on Arndale.

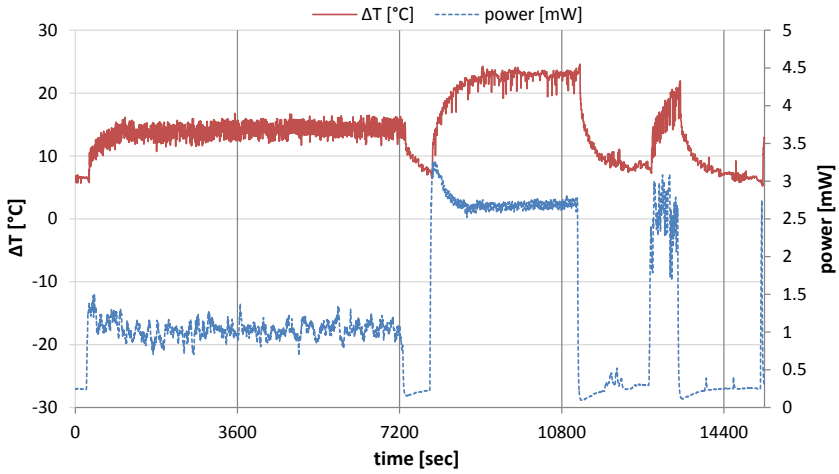

Figure 6.10: Output power of a single Peltier cell on Arndale.

- **multithread application** that performs millions of algebraic and trigonometric computations of floating point numbers using a user defined number of threads (called `busywork` in pictures);
- **kernel operations** in this task a Linux kernel is uncompressed, compiled and then cleaned, then the folder is removed.

Output voltage and current presented in the output responses depicted have been measured over a matched load (evaluated with further measurement campaigns) with a 1 s period using multimeters[2] controlled via software.

The benchmark applications allowed us to obtain an output power profile for several configurations (ARM CPUs plus TEGs). Detailed results of the experiments are reported in Fig. 6.4 to Fig. 6.11. These pictures present the output impedance matched power of each TEG under the ΔT produced by the heat dissipated by the CPUs, computed as the product of the measured voltage and current. ΔT in turn is strictly related with the CPU

[2]Agilent 34401A and 34411A Digital Multimeters.

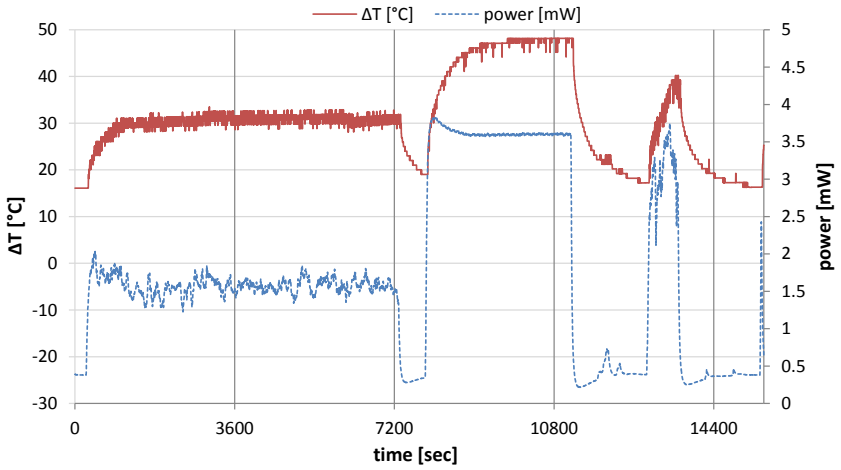

Figure 6.11: Output power of a series of two Peltier cells on Arndale.

load. The air conditioning system of the lab kept the room temperature almost constant in the range 22° C to 25° C, allowing to emulate the server room working condition.

6.2.2 Characterization Results

Generally speaking, the longer the task the higher the CPU temperature as well as the number of memory accesses and switches (algebraic computations). Pandaboard's CPU builds up lower temperature with respect to the Arndale's (37° C vs. 80° C), resulting in a 10x difference in the TEG measured output power in our experiments (μW range for Pandaboard while mW for Arndale's). All but Nextreme exhibit a good reactivity to fast temperature spikes, this one produces smooth power profiles while others result in sharp peaks at the beginning of each task. Micropelt has the worst thermal efficiency with respect to the others: we can notice the effect of the voltage scaling in the middle of the benchmark — around 10800 seconds — clearly depicted in Fig. 6.9 and less evidently in Fig. 6.8, where the output power "flickering" reflects the temperature behavior, which is regulated automatically by the Linux kernel governor. The two Peltier cells almost double the output power of the single cell, the spikes move from 350 μW to 600 μW in Pandaboard's case and from 2.6 mW to 3.6 mW with Arndale.

The harvesting system composed by two Peltier cells in series, arranged in stacked fashion, exhibits higher performance with respect to other configurations. The comparative results, pictured in Fig. 6.12, demonstrate that the performance of this TEG setup are better for tasks that last more than 30 seconds: `idle`, `ffmpeg`, `busywork-100`, `tar`, `make` and `clean`. Here we can notice the maximum mean impedance matched power of almost 3.6 mW on Arndale and 600 μW in case of `busywork-100`. Only in case of very short tasks, `rm` and `busywork-10` last few seconds each, the Micropelt TEG performs better and only on the Arndale board.

The stacked TEG outperforms the other solutions also in terms of total energy. Considering the results of both Fig. 6.14 and 6.15 together, the conclusion for Arndale board is that, in lower time, the stacked architecture is able to harvest the highest amount of energy (26 J in 4^h 30^m with double Peltier cell vs. 24 J in 5^h 45^m with Nextreme and

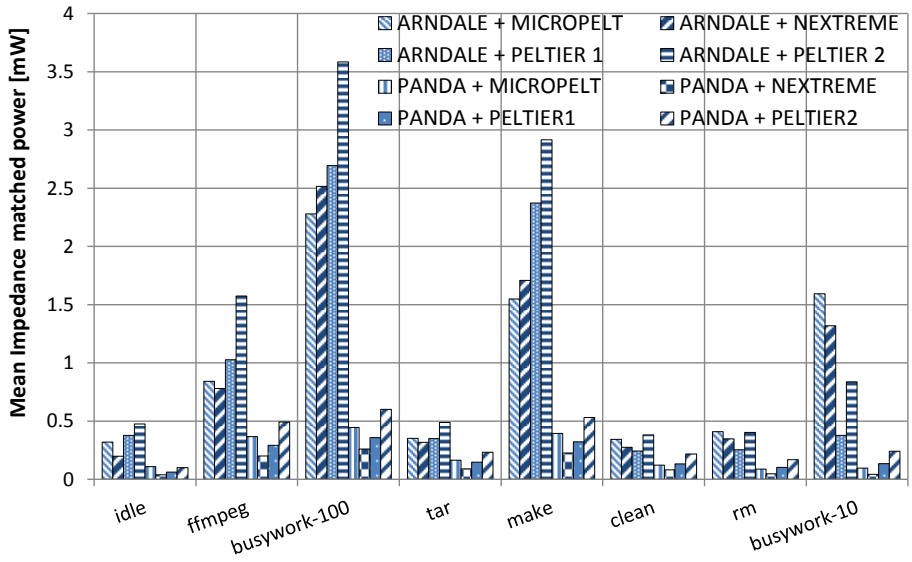

Figure 6.12: Average output power per benchmark's stage.

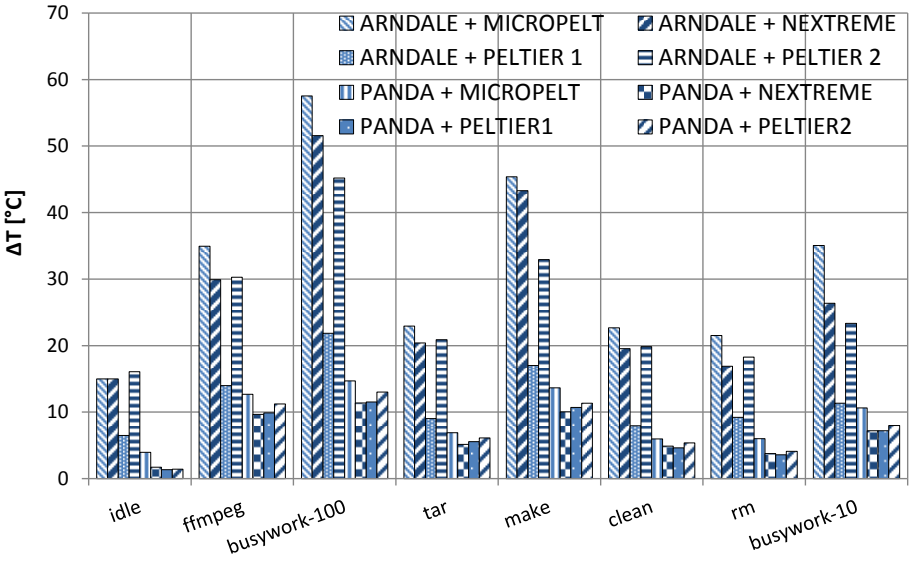

Figure 6.13: Average ΔT per benchmark's stage.

20 J in 5^h 20^m); while for Pandaboard the execution time is almost constant (26 J with double Peltier, 6 J with Nextreme and 11 J with Micropelt in 8^h 20^m). The difference in the Arndale's results is due to the effect of the voltage scaling performed by the kernel governor, which reduces the execution speed to reduce the CPU temperature. As stated above, Nextreme and Micropelt TEGs are not able to dissipate the heat generated by the Arndale's CPU and also to efficiently exploit the high variation in the resulting ΔT.

6.2.3 Stability and Security Analysis

Having a device on top of the CPU package limits the natural heat exchange of the microprocessor. If the system overheats, high temperature may cause structural damages to the ARM CPU board, and may reduce the device lifetime. Also the TEG may suffer structural damages if exposed to excessive thermal gradients [44]. We evaluated the impact of the harvesting system on the CPU temperature. The results are presented in Fig. 6.16 where the CPU temperature is depicted against time for three different scenario:

(a) natural dissipation (CPU used as it is);

(b) heat sink on top of the CPU package;

(c) the proposed harvester.

Without any thermal dissipation (case-a) the CPU immediately reaches a dangerous temperature and the voltage scaling decreases its performance for safety (dash-dot line, the temperature reaches immediately the critical threshold of 82° C and then lays bounded around 80° C). In the other two cases (b and c) the rate of increase is smoother and the presence of the harvester slows down the thermal exchange; at steady state both cases exhibit comparable performance (stability around 70° C). This once again underlines the good performance of the double Peltier stacked structure specifically built for this work, which demonstrates the best thermal dissipation among all, as evident from the comparison in Fig. 6.13.

Even if the result that TEGs are equivalent to heat sinks in dissipating heat, could look like counter-intuitive, one must take into consideation that those devices are ment to extract heat and convert it into electricity, thus heat is allowed to flow through the device and evetually reach the environment through the proper sink.

We can state that the stacked series of two Peltier is the most suited solution to build an energy neutral embedded system to monitor the state of the server and its environment. For this reason we continued the development of the design considering only this TEG.

6.3 μW Sensor Network

After the feature extration and the selection of the best TEG, we investigate a distributed system which is completely self sustained by μW harvested power. Each node in the monitoring network embeds sensors to collect data that are in turn routed to a central gateway (aggregator or coordinator) responsible for data processing, storage, and for the actuation of predefined policies. Nodes can also be designed as active devices to implement the control instruction sent by the policy manager. Key feature and enabling

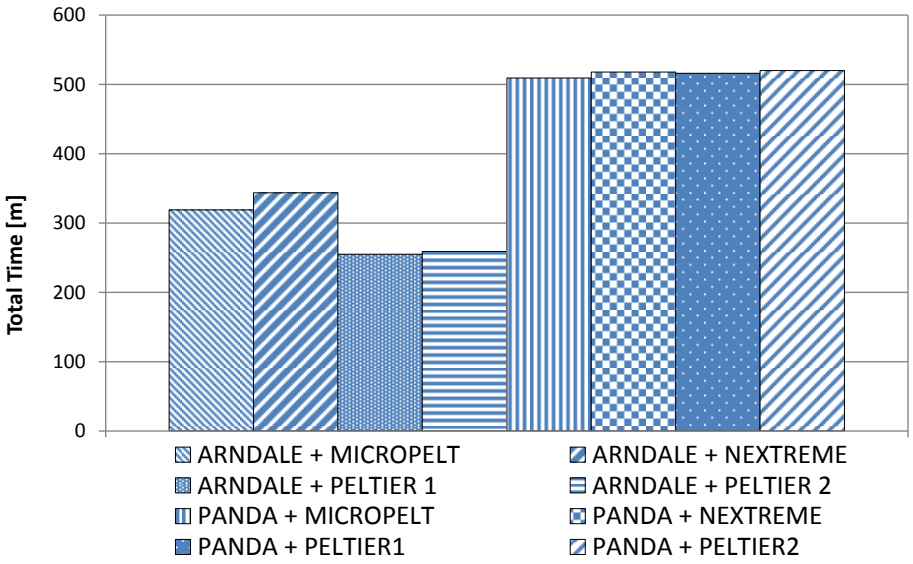

Figure 6.14: Total benchmark execution time per configuration.

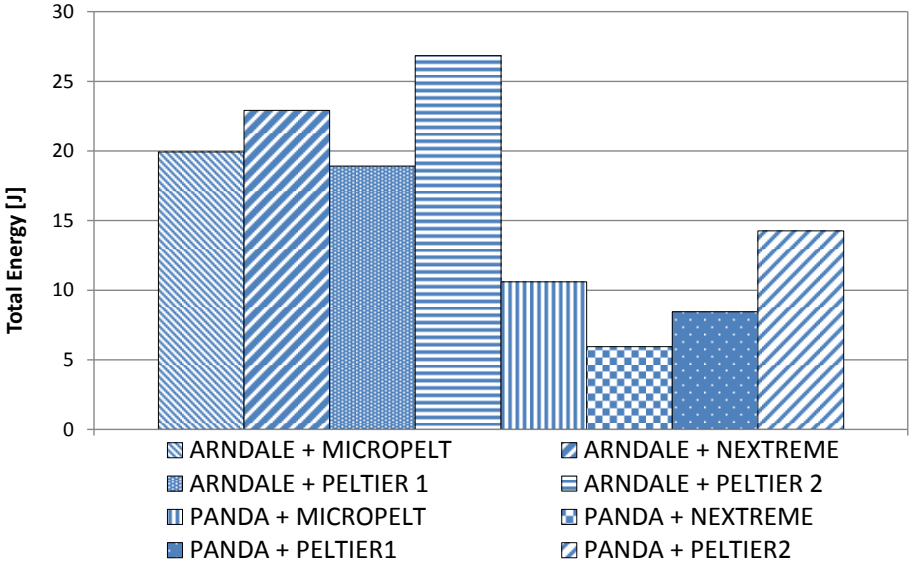

Figure 6.15: Total impedance matched energy per configuration.

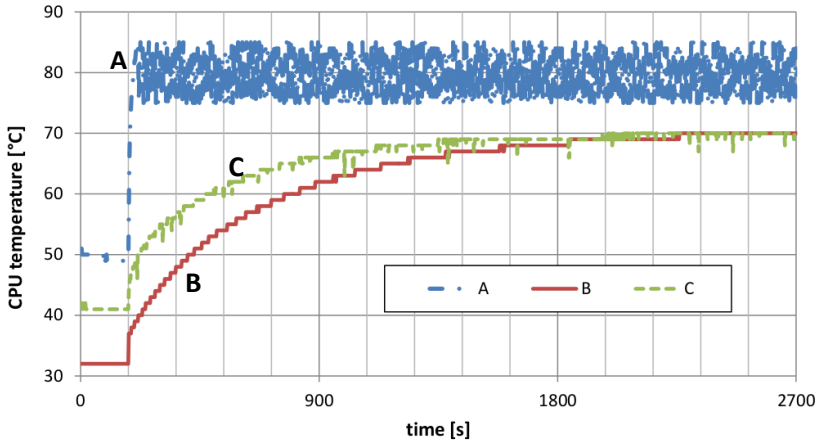

Figure 6.16: Comparison of CPU temperature rise with 100 % CPU load @ 1.7 GHz: (A) natural heat dissipation, (B) heat sink, (C) double Peltier stacked structure above CPU package.

characteristic of the nodes is the very low power required both in active (sensor sampling) and sleep modes of operation. The main limits in this scenario are the power required by the communication task and the set of sensors used.

Reducing the power needed by a transceiver to μW range definitely results in a lower range of communication. The introduction of energy harvesting techniques from TEG provides state-of-the-art powering solutions replacing disposable batteries with dedicated electronics; thus lowering costs and maintenance efforts, to the detriment of slightly increasing the design complexity. We used a commercial wireless node to implement a self-sustained environmental monitoring system and evaluate its feasibility, that is the very same W24TH wireless boad already presented and discussed in previous chapters. On top of this, we developed an ad-hoc monitoring application using all the available sensors (temperature, humidity, ligth intensity and MOX for volatile chemicals), since all those parameters can be useful in the detection of damages and failures in the server room and trigger an alarm.

The faster the execution the lower the energy required, so we put most of the effort in the design of this application to reduce the execution time, while assuring reliability of the monitoring. We configured the device to use a very simple network protocol, built on top of the IEEE 802.15.4 MAC layer. Simply, the device sends data to the coordinator as soon as enough energy is available to complete the task. Sensors require some specific time to reach a steady state response. We turn them on at once and perform the measurement as soon as they are stable. This approach reduces the time the node is active, and hence reduces the total power consumed by the node. Fig. 6.17 shows the current drained by the embedded system while executing the application.

The required energy has been measured using a 1 Ω shunt resistor using an oscilloscope. The whole execution requires almost 800 ms and the maximum required power is around 160 mW (@3 V), 70 mW on average.

The target issue of the proposed work is to match an unstable and unpredictable μW

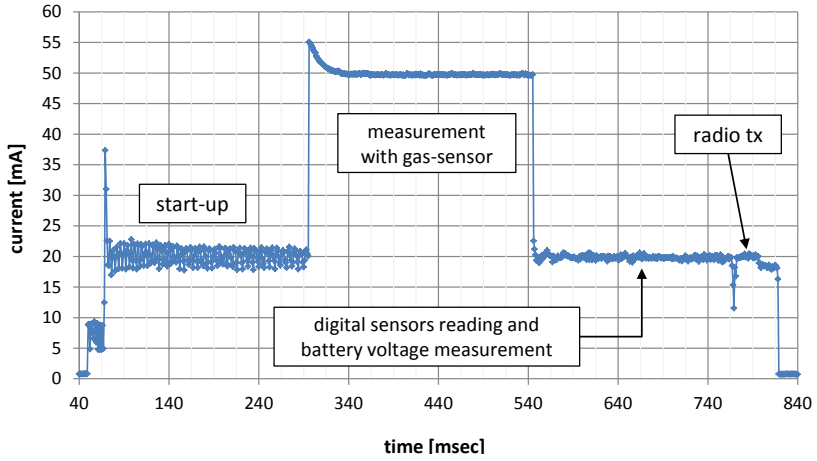

Figure 6.17: Wireless sensor node, current consumption during monitoring activity.

range power source with a device that requires mW, as shown above. To this reason a conditioning conditioning circuit is required to fulfill this gap and to adaptively choose a suitable duty-cycle, based on the charge collected using TEGs.

Particular care must be given to the presence of a chemoresistive gas sensor (MOX) which is responsible for the current increment in the middle of the firmware execution (from 300 to 550 msec in Fig. 6.17). This increment is higher than the radio transmission that occurs at the end of the monitoring task (from 760 to 800 msec in Fig. 6.17) and requires almost 51 mW (17 mA @3 V). The profile depicted is due to the MICS-5121, which requires 76 mW of power to work. However, the ULP strategy presented in Chapter 4 allows to significantly reduce the energy required by the MOX sensors.

We evaluated the impact of this kind of monitoring device in terms of energy taken from the grid. We compared two situations:

(a) Continuous, when energy is provided by means of batteries or dedicated wires, also during sleep;

(b) Discontinuous supply, the proposed and implemented choice where the node is switched on only when the reserve can sustain it.

In case of 0.25% duty-cycle (one measure every 5 minutes) the W24TH requires 2.7 mW on average in case (a) against the 225 μW on average in case (b). When duty cycling is 0.02% (one measure per hour) we get 2.4 mW for (a) and 18.6 μW for (b). Translating these into energy per year (Wh) we can estimate a saving of 92% (1.4 KWh against 117 Wh) with 0.25% duty-cycle, and 99.25% with 0.02% duty-cycle (1.3 KWh against 10 Wh). These are tiny numbers compared to the energy required by server farms but they must be scaled with the numbers of nodes in the monitoring network. The resulting energy saving combined with the lack of need for maintenance makes the designed system inexpensive.

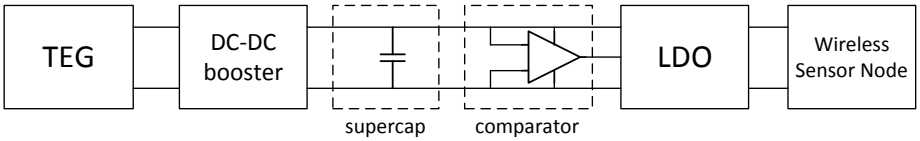

Figure 6.18: Block diagram of the conditioning circuit.

6.3.1 Simulation

The preliminary results and information presented in previous sections were used to evaluate the feasibility of an energy neutral monitoring device with wireless connectivity, and to design it. The system we designed and developed was made of three parts:

(i) the TEG;

(ii) the conditioning circuit;

(iii) the WSN node.

We kept only the Arndale Board as testing environment for the neutral monitoring node evaluation since:

(a) it has higher computational performance;

(b) it generates more heat;

(c) its CPU architecture is more recent with respect to the Pandaboard, hence more realistic to simulate an ARM based data center;

(d) to speed up the development due to the length in time of our tests.

Simulations were used to predict the amount of energy harvested from the heat dissipated through the CPU package of the device under different loads, and various clock frequencies. Due to the high variability of the output quantities, and the uncertainty of the cooling capabilities, it is likely that the simulation would overestimate the real output power profile (cfr. [16]), thus we have taken it into account, especially at low frequencies.

The output power provided by TEGs depends on the ΔT as previously shown, so we decided to operate in terms of server CPU load and clock frequencies (f_{CLK}) that are directly related with the dissipated heat. The available Arndale Board's f_{CLK} ranges from 200 MHz up to 1.7 GHz, with 100 MHz steps. In our simulations we selected the range $1.2 - 1.7$ GHz because not enough heat is dissipated below this interval. In addition, we chose a sequence of CPU loads of 0, 30, 50, 70 and 100%, for a total of 30 different working conditions.

The electric power generated by the TEG fluctuates, therefore it needs to be conditioned so that it can be stored, or used. We build a conditioning and storage circuit based on commercial components, arranged as shown in the block diagram of Fig. 6.18 and deeply discussed in the following chapter. A supercapacitor (3 F, 2.7 V) was chosen as storage unit because of the energy required by the monitoring wireless node. Fig. 6.19 depicts the lifetime of the WSN node supplied by direct connection to this supercap. The reason for using a supercapacitor stems from its low internal resistance, which results in

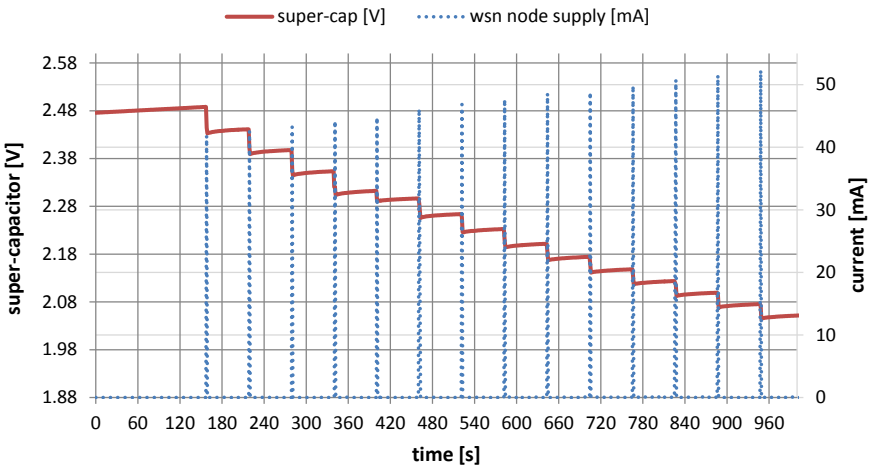

Figure 6.19: WSN node lifetime when supplied with direct connection to the supercap charged at V_{max} (bypassing comparator and LDO).

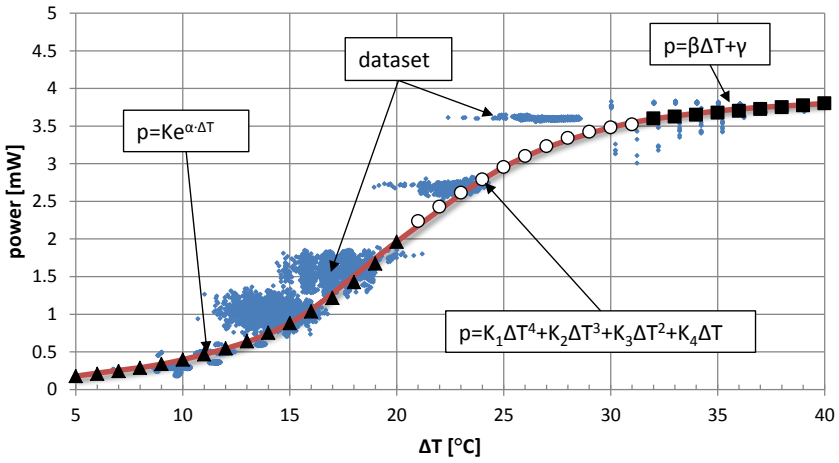

Figure 6.20: Input output relation of the proposed harvester consisting of booster and supercap.

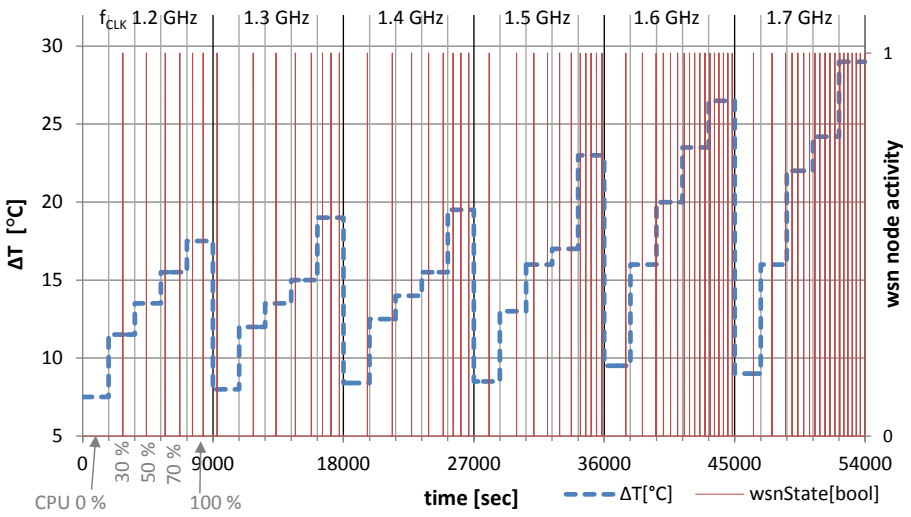

Figure 6.21: Energy budget simulation, wsn node completed tasks as function of the TEG's ΔT, tuned plying with CPU loads and f_{CLK}.

a high surge current, as well as its low self-discharge characteristic that make it desirable in energy harvesting application.

A comparator is used to enable the current flow only when the charge stored in the supercap is enough to power the node. The comparator has been realized with an integrated circuit and some components to tune the range. The high threshold V_{max} is set to 2.48 V, that is also the W24TH node supply voltage. The lower threshold $V_{min} \simeq 2.41$ V guarantees to completely disable output power, so that the node is isolated and unable to drain current. This choice does not affect the node network functioning when the node is turned back on. The above two thresholds were chosen for two reasons:

- With 2.48 V the super capacitor has enough energy to supply more than a single node activity (that has been demonstrated sufficient in Fig. 6.19), and it is charged up to a value lower than the nominal, providing a security margin.
- The tiny gap (≈70 mV) allows the comparator to switch off the LDO (linear voltage egulato) immediately after the node completes its duty.

The influence of the CPU load in the generation of thermal gradient must be understood to predict the availability of electrical power to supply the small computing device. The profile, based on the knowledge of the behavior of the TEG exposed to different thermal gradients driven by the CPU activity, is determined by the measurements described earlier in this chapter.

We processed the data of Fig. 6.7 and 6.11 with different regression scheme (provided by commercial software) to obtain the input-ΔT vs. output power. Our approach differs from the ones proposed in the literature [16, 60, 64, 6] since we consider both TEG and dc-dc boost circuit efficiency as a black-box. Fig. 6.20 depicts the empirical input-output relation of our harvesting device. The resulting model is a composite piecewise continuous function of the temperature, and switches between an initial exponential segment between 5° and 20° and 20° ΔT, followed by a polynomial segment be-

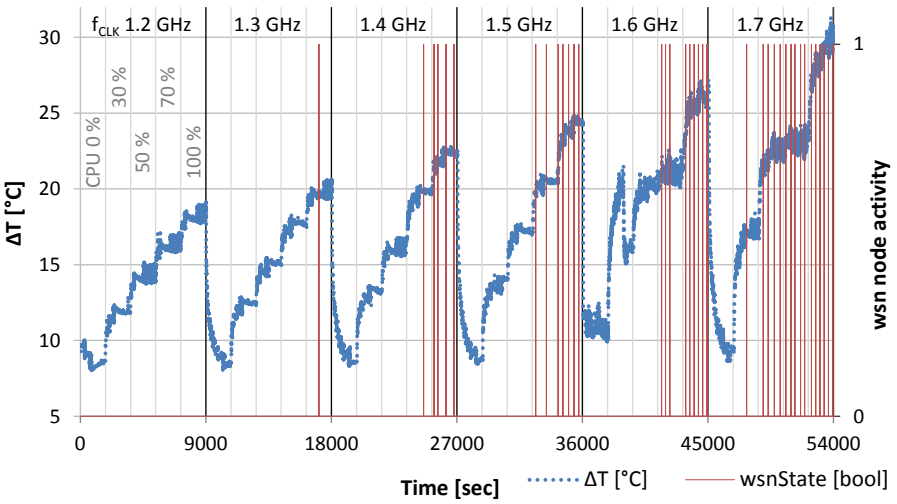

Figure 6.22: Real test of energy neutral monitoring using WSN nodes. Same CPU load and f_{CLK} used in simulation.

tween 21° and 31°, and finally by a linear segment from gradients ranging from 32° up to 40°.

This arrangement was required due to the highly non-linear regulation between the input temperature and the efficiency of both the TEG and the booster. The estimated parameters of the model are $K = 0.08$ and $\alpha = 0.16$ for the exponential scheme in the 5° to 20° C of ΔT range; $K_1 = -0.000005$, $K_2 = 0.00001$, $K_3 = 0.0104$, $K_4 = -0.7$ for the polynomial component in the 21° to 31° C of ΔT range; and finally $\beta = 0.025$ and $\gamma = 2.8$ from 32° up to 40° C of ΔT. We fed this model into the simulation software along with empirically estimated thermal gradients for a range of CPU loads and frequencies, obtained with further experiments.

The simulation is based on a simple energy budget evaluation for each fixed time-slot (1 second). The input power is provided to the harvesting circuit by the TEG, and the conditioning circuitry manages the wireless node. We fixed the efficiency of the complete harvesting system equal to 5%. This value takes into account the mismatch between the TEG and the DC-DC booster, (based on an inductive transformer [68]), the efficiency of the supercap during charge and the overall performances of the integrated circuits and components (Comparator, LDO, ...).

The result of the simulation is reported in Fig. 6.21. In this case we used a fixed time intervals of 1800 s (30 min) for each CPU load, hence 2.5 hours for each f_{CLK}. The graph shows the occurrence of monitoring tasks that the wireless node performs as the time elapses, in binary fashion, on top of the ΔT sequence, function of the CPU activity. The simulation clearly shows the strong correlation between temperature and frequency of monitoring, that is summarized in Fig. 6.23. These results justify the effort of building a real energy neutral system since it would be possible to achieve 5 min monitoring frequency at most, when the CPU is under a heavy load, and the monitoring can be crucial to prevent failures.

6.4 Prototype Performance Evaluation

We realized the harvesting circuit described above to validate the theoretical results of the simulation. We used the same sequence of CPU loads and f_{CLK} presented in the simulation's description (Sec. 6.3.1). We achieved the target load activities by interleaving mathematical computations with sleep intervals to generate variable load profiles. The results of the 15-hour running of the experiment are presented in Fig. 6.22. In this picture the average current drained by the node is depicted on top of the TEG's ΔT. We can immediately notice the reduced number of monitoring tasks performed by the node with respect to simulated estimate, as we were expecting, however, during high demanding tasks, we observe a tight correspondence between simulated and real occurrences.

Comparative results are summarized in Fig. 6.23. In this picture we compared the resulting monitoring frequencies, in terms of time between two sampling, for f_{CLK} in the $1.5 - 1.7$ GHz range. The extracted TEG and circuit model does not fit well the real efficiency for ΔT lower than $\approx 20^{\circ}$ C, but it perfectly fits in the higher range, where the worst case mismatch is within 15 s. The difference is due to two reasons:

(i) the fixed thermal gradients used in the model, which constitutes a very strong simplification of the CPU behavior;

(ii) the fixed efficiency used to describe the control circuit (comparator and LDO), which is overestimated, since also here efficiency depends on the working point.

CPU frequencies lower than $100\%@1.5$ GHz results in monitoring frequencies that are very difficult to evaluate due to extremely large time requirements. The efficiency of the control circuit is definitely lower than the expected with very low input power and result in very slow supercap recharging below $\approx 20^{\circ}$ C. This arises from system non-idealities and, as a consequence, the efficiency of the real control circuit is definetly lower than what has been predicted by the model, for temperatures below 20°.

When demanding CPU tasks are running we can reach a very good 5-minute monitoring interval (≈ 300 seconds in case of $100\%@1.6$ and 1.7 GHz). In some conditions (for example $70\%@1.6$ and 1.7 GHz) we have better performance in real tests than with simulations. This phenomenon can be related to the fact that very fast thermal variations result in higher output power with respect to a constant thermal profile, as can be noticed considering the ΔT behavior in Fig. 6.22. Lastly, we can state that the limit of the prototype implementation is 5 minutes.

The presence of the TEG in the proposed configuration does not affect the CPU performances at all. Considering the initial benchmarks, in Fig. 6.8 and 6.9 we can notice the very high ΔT and the effect of the frequency scaling on the power harvested in case of demanding tasks. We can notice a higher efficiency of the proposed thermo harvester for lower thermal gradients compared with the result shown in Fig 6.11. In this last case there is no need for frequency regulation since the thermal dissipation provided by our solution overcomes also the standard package, as shown in Fig. 6.16.

For the sake of completeness, we evaluated also the "*boundary conditions*" of the proposed system, namely the cold-start time and the autonomy on supercap after CPU switch-off.

With cold-start time we refer to the amount of activity time of the host system (the "heater") required to charge the supercap from 0 V up to V_{max} that corresponds to

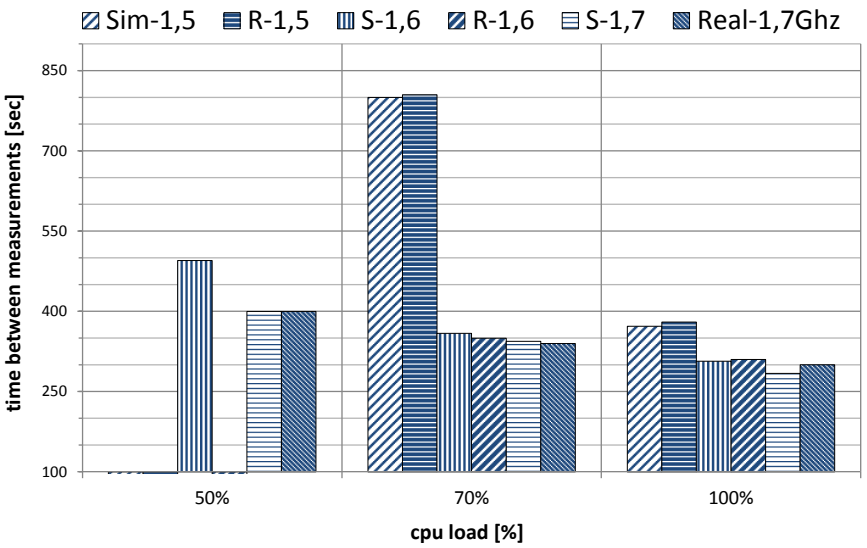

Figure 6.23: Monitoring intervals, real (R) vs. simulated (S) results, as function of load and f_{CLK}.

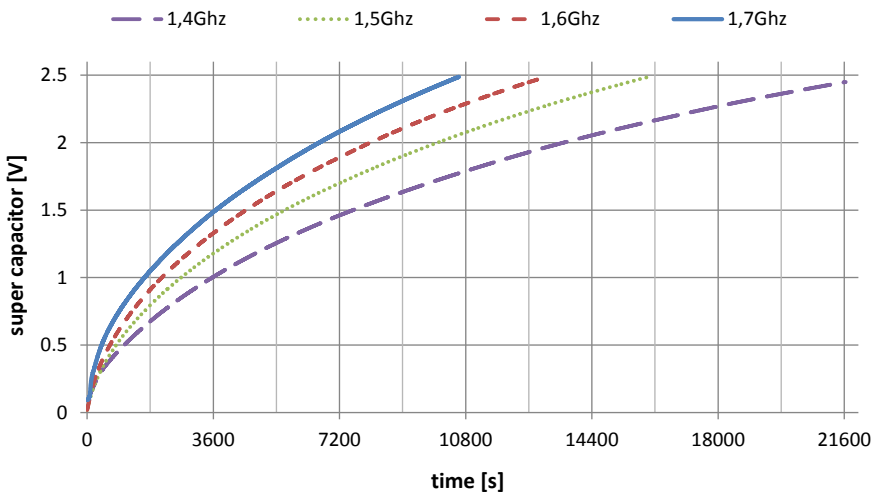

Figure 6.24: Charge time with the proposed architecture consisiting of TEG, booster and 3 F supercap.

the first activity of the monitoing node. Fig. 6.24 presents the cold start time for four different CPU conditions, ranging from 1.4 GHz up to the maximum clock frequency of 1.7 GHz, always with 100% of CPU load. These curves show how much time the ARM CPU board is required to run in order to have the monitoring subsystem working.

We can notice that the time required at 1.4 GHz is roughly double the time required at a clock frequency of 1.7 GHz. To improve the cold-start time, different solutions can be adopted: on the storage side, the usage of bootstrap capacitors [18, 15], or reconfigurable capacitor bank block (CBB or flying capacitors) [62, 78], allow the rate of the transferred energy to be increased, using multiple capacitors charged in parallel, that are then switched into series configuration to achieve the required voltage on the load. On the boosting interface side, both the capacitance of the charge pump can be decreased to speed up the switching frequency of the oscillating circuit [47], and a maximum power point tracking stage can be introduced to maximize the power transferred [48].

In case of switch-off or, equivalently, very low CPU load, the supercap can sustain the monitoring node for only one extra measurement with the V_{min} threshold set as above. To be able to completely drain the energy from the supercap (as in the case depicted in Fig. 6.19 where it is demonstrated that it is possible to sustain the node for more than 10 minutes with monitoring interval of 1 minute), extra hardware would be required to bypass the comparator that is actually off after the last node activity. Otherwise, rising the supercap to the maximum rated voltage, or replacing it with another rated to higher voltage can withstand this limitation.

Chapter 7

Design of an Energy Neutral *Smart Heat-Sink* Powered by Thermoelectric Energy

As introduced and discussed in previous chapter, monitoring for free exploiting renewable energy sources available in the environment requires highly customized solutions to achieve reliability and unobtrusiveness in the final design. For this reason any solution in this field is stricly related with the target application.

From this consideration, in this chapter, we present the desing of an energy neutral monitoring system that exploit considerations and results discussed throughout this thesis. Once again we started from the very specific Data Center case study, since design requirements and specifications come from real targets and have to be suited to them.

Results discussed here and in previous chapter have been published in [J2,J5].

7.1 Data Center Case Study

Nowadays, IT systems are more distributed than ever and Data Centers are spread all over the world providing services for infrastructure virtualization, cloud and IoT applications. Moreover, in the near future, all the small to medium internally company-owned IT infrastructure will be replaced by, so-called, "mega-data centers", since outsourcing these facilities is much more cost-effective rather than maintaining an own infrastructure and hiring IT managers [70]. In this scenario, the need for security, reliability and the maximization of service up-time can be achieved only relying on automatic incident remediation solutions, both on the cyber-security and on the physical side.

Another important challenge for service providers is to adapt their infrastructure to a more environmentally friendly and fossil fuel free economy [85], following and complying with regulations that are currently being issued by governments all over the world. Generally, data centers facilities are always turned on, and are sources of high thermal waste and energy consumption. In fact, only 0.5% of total fossil fuel power is spent in

useful computation, while more than 65% is used for cooling or lost as heat [29]. Server rooms have plenty of wasted heat, which can be recovered to generate electric energy, and mitigate wastage [33]. Given the numbers involved in the data centers scenario, even a small improvement that diminishes the power acquisition from the grid can introduce environmental and economic benefits. In this regard, all the biggest companies providing cloud services in the IT market [67] (Amazon [5], Google [27], Apple [7], Microsoft [51], Rackspace [72], and IBM [37] for instance) have already introduced renewable energy sources in their supply chain.

Clearly, future "green" data centers will require a completely different design also on the Data Center Infrastructure Management (DCIM) side, to achieve capillary monitoring of the thermal environment and the energy resources, as well as to guarantee security. These features are best achieved using approaches which are as neutral as possible relative to the energy impact, i.e., which work transparently by taking advantage of freely available power from the environment, without using additional energy from batteries or from the wired infrastructure. In fact, one of the most used metrics for evaluating the energy efficiency of a data center, the Power Usage Effectiveness (PUE), is defined as the ratio between the total energy used to supply the data center, and the energy used for computation alone. The lower the PUE, the higher the efficiency. The additional energy used from the monitoring infrastructure increases the amount of total energy spent, thus the PUE increases, resulting in a worse utilization of the provided electrical power. The aim of this research is to provide a solution that is beyond the industrial standard and to push forward the state-of-art in terms of energy neutral networked portable electronic devices for monitoring tasks.

Monitoring system already available in the market usually are powered from the grid, burdening the data center PUE. In fact, the energy spent by a monitoring system increases the power consumption of the infrastructure, and consequently worsens the datacenter energy efficiency. Moreover, their sampling rate is lower compared to the proposed system. The Data Center Genome Project, a monitoring system developed by Microsoft Research, samples the environment every 30 seconds [50]. The sampling period of the solution proposed by Dell, called OpenManage Power Center Temperature Monitoring, can be set to 1, 3, or 6 minutes [21]. Given the number of nodes that must be installed to instantiate the monitoring network, it would be impracticable to use wires for power supply and for communication. A connector to interface the proposed monitoring device with the data center motherboard does not exist, therefore, new dedicated wires must be placed to deploy a wired sensor network. Using batteries poses the issue of replacing them when exhausted. For these reasons, we explored the utilization of renewable energy sources, and we adopted a wireless solution for our monitoring infrastructure.

7.2 The Smart Heat-Sink

We present a batteryless embedded system that has been designed to realize a maintenance-free wireless sensor monitoring infrastructure for data centers packaged in the shape of a commercial heat-sink, which retrace the SoC envisioned and presented in previous chapter, namely the smart heat-sink.

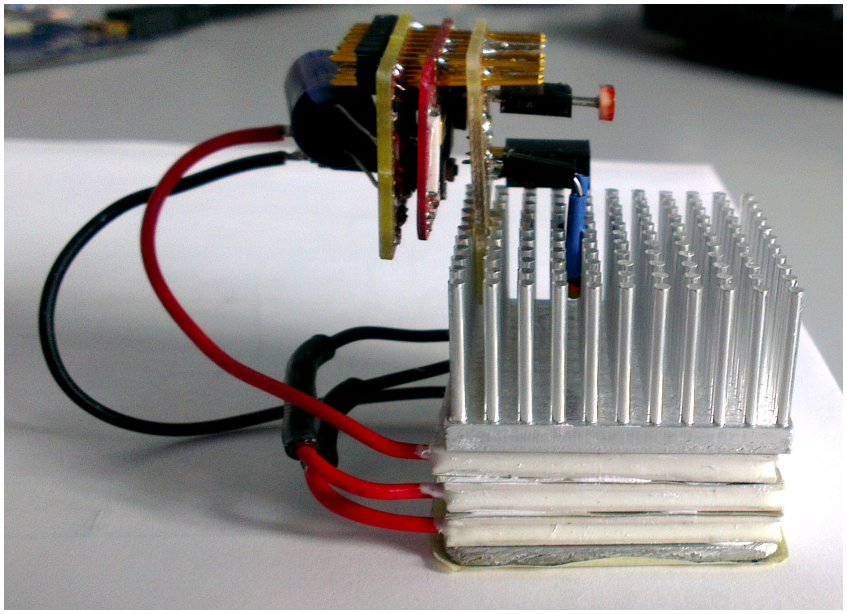

(a) Side view of the TEG with heat sink, sensors, and transceiver board

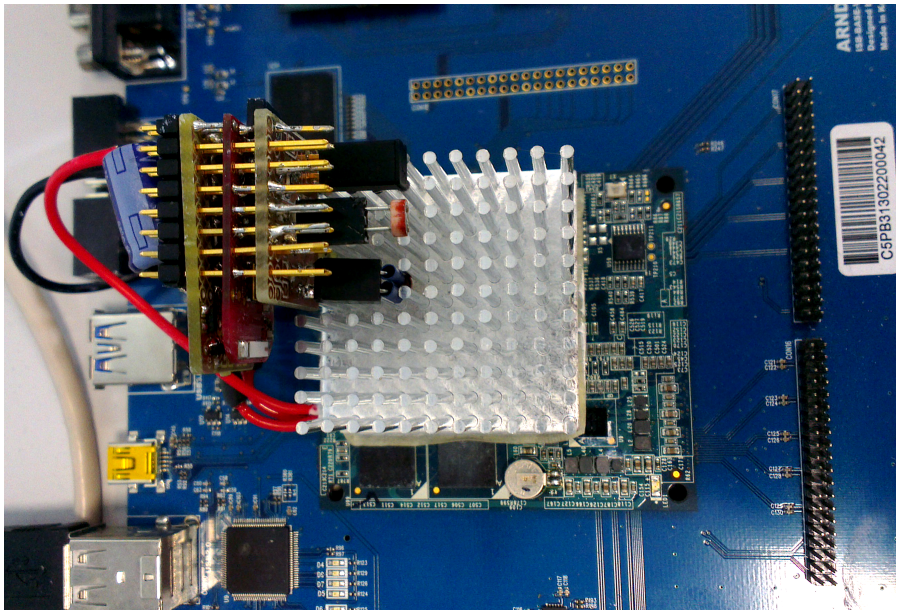

(b) Top view of the system prototype mounted on top of the target CPU board.

Figure 7.1: Pictures of the system prototype.

By exploiting wasted heat energy inside data center facilities, the most abundant kind of wasted energy, our monitoring system can provide *continuous* sampling of several environmental parameters—like temperature, humidity and light—and *enhanced security*, thanks to the use of various sensors.

The system we propose introduces significant improvements, because it is totally powered by free energy it scavenges from wasted energy, and it monitors the environmental parameters with a shorter sampling period. Moreover, thanks to the smart sinks reduced size, a large number of them can be placed inside a server room, therefore it is possible to perform distributed monitoring of the environment, and to provide a wider cumulative data rate. In case of overheating, the proposed batteryless monitoring device scavenges enough power to operate with less than one second data rate.

We designed, built, and tested a TEG based energy transformer with dedicated conditioning circuitry able to guarantee never-ending supply to the wireless sensing node with as little as 10° C temperature difference across the harvesting stage. Each CPU in the data-center can be equipped with such a device, setting up a complete wireless sensor network able to track the environmental parameters.

7.3 Customized Thermoelectric Harvesting Unit

The system is used to monitor the health of the host machine. It exploits a thermoelectric generator to convert the thermal gradient of the heat, dissipated from a high performance processor (*host*), into electrical energy to power the smart sensors.

The smart heat-sink, consisting of three stacked TEG, an electronic circuit for the conditioning and storage of the harvested energy and the monitoring node based on the MSP430 MCU is depicted in Fig. 7.1(a). The harvester module is placed on top of the host CPU, and is completed by an aluminum heat sink, as in Fig. 7.4(b). The heat sink is needed to keep the heat flowing thought the TEG and ease the thermal exchange (of the excess heat not converted into electricity) with the environment. Removing the heat from the cold side of the TEG is essential to ensure a continuous electric current from the generator. If not removed, the heat would be retained in the TEG until a thermal equilibrium which stops the generation of electrical power.

Each wireless smart heat-sink consists of the thermoelectric harvester, a circuit board for the conditioning circuit and the storage capacitance, a circuit board that houses the sensors, and another board for the microcontroller and the radio chip. Distributed monitoring is therefore achieved by instrumenting the datacenter with hundreds of nodes, one for each CPU or heat generating device. The overall dimension of the node is $40 \times 40 \times 50$ mm.

7.3.1 Conditioning Hardware

The electrical power generated by TEG is generally lower than the supply required by the wireless sensor nodes. Typically, it stands in the order of hundreds of microwatts with voltages in the hundred of millivolts. Therefore, an efficient conditioning circuit is required to raise the generated voltage up to an acceptable level to feed the system, and to accumulate energy. The sensing node operates sporadically in short bursts with

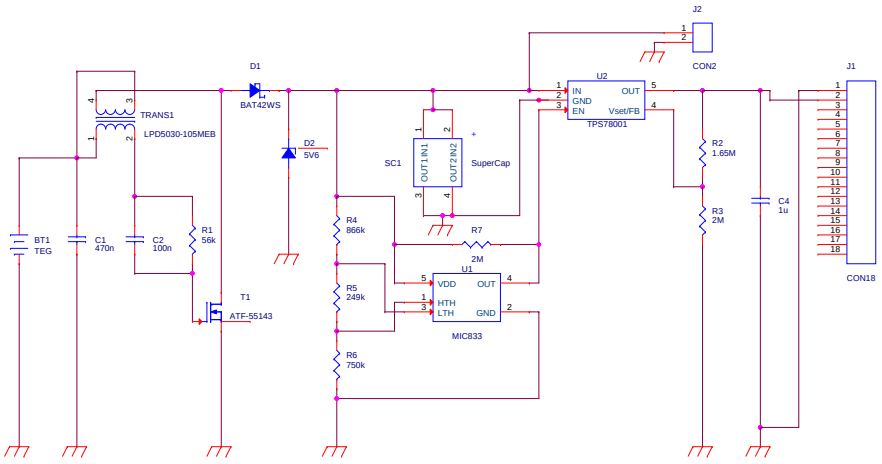

Figure 7.2: Electronic diagram of the conditioning hardware circuit that converts and stores the energy generated by the TEG, and supplies the WSN.

a given duty cycle, therefore, the conditioning circuit must contain a storage element to accumulate the energy, and to provide energy on demand. The sizing of the storage elements depends on the kind of application implemented on the wireless node, and on the expected system lifetime.

Operating under a low duty cycle, the monitoring system consumes, on average, a low amount of power, and the application can be supported using a small energy buffer (a capacitance of few millifarads). On the other side, a larger storage unit guarantees the node operations for longer time, even in case there is no power at its input. In any case, the size of the storage elements is driven by the application requirements, because it depends on the amount of time necessary to accumulate enough energy to operate.

The complete conditioning circuit, shown in Fig. 7.2, is composed of a cascade of several components:

- step-up (boost) converter;
- rectifier;
- storage unit (capacitor);
- voltage comparator;
- LDO (linear voltage regulator).

The initial stage consists of a resonant step-up oscillator that converts a very low-power input voltage (hundreds of millivolts) into a higher voltage output. The resonator circuit multiplicative factor (σ) depends on the voltage at its input. The resonator requires at least 220 mV input voltage to let the power through the output with the amount of energy required to charge the storage unit. With an input larger than 290 mV the step-up converter provides a 5 V output voltage, with $\sigma \geq 15$.

Then, the signal is filtered, decoupled, and rectified to be used by the subsequent stage: the storage unit *SC1* which consists of a capacitive element. The storage unit may be made with supercapacitors or solid state batteries depending on the robustness

required by the application. For example, with less-frequent, but energy-demanding tasks (e.g., wireless communication) a supercapacitor of a few Farads represents a good solution, since it takes a few hours to charge, with high energy density to supply powerful embedded platforms, as demonstrated in previous chapter. In this case, we selected a single capacitor aiming to boost recharge time rather than energy density. In the following we describes in details the design choices made with respect to the harvester performance, and the application requirements.

The LDO (*U2*) maintains a steady output voltage which is lower than the input voltage source. It is needed to supply the node with a constant, proper voltage even when the voltage stored in the supercapacitor is higher, so to limit the current drop. The LDO is controlled by a comparator (*U1*) that drives its enable signal. When the charge accumulated in the storage unit is above a fixed threshold (*HTH*, expressed in Volts), the enable signal rises, and the LDO starts feeding the load. The enable signal is kept high until the charge drops below the lower fixed threshold (*LTH*). This configuration allows the output conditioning unit to be decoupled during the recharge phase. The two thresholds are selected by means of a voltage divider corresponding to three resistor elements in the schematics: *R4*, *R5*, and *R6*.

7.3.2 Harvester Characterization

A common approach to increase the power harnessed using thermo-electric generator devices is to take several TEG, and place them thermally and electrically in series, where each TEG represents a *stage*. In this way, each of the stages is exposed to a fraction of the total thermal gradient, therefore it generates less electrical power. However, the combination of the stages results in a greater total output power.

Starting from the results discussed in previous chapter, we changed TEG devices to achieve a more compact design and we compared several configurations ranging from one up to five stages to find the best trade-off between maximum output power and overall size of the smart heat-sink. The time required to recharge an empty storage capacitor up to *HTH* with the two-stage configuration is $2.6\times$ the time required by the three-stage solution. On the other hand, using the five-stage harvester, the recharge time is $1.5\times$ faster with respect to the three-stage. We found that the *three-stage* setup realizes the best compromise for the specific target application. In the rest of this work, we present results obtained with this configuration.

The characterization has been performed directly on top of our target Arndaleboard's CPU (cf. Chapter 6). We measured the power generated for the whole range of working temperatures of the processor by varying the *host* CPU workload. Fig. 7.3 shows the measured values of the generated output power in relation to the thermal gradient between the CPU package and the heat sink of the harvester. Following the TEG's model extraction approach based on the regression from experimental results,also in this case we extracted, by cubic interpolation, the characteristic curve (in red) of the new harvester (represented on top of the new data set – blue dots).

The temperature of the ARM CPU has been measured by the *host* itself using its integrated temperature sensor, while the temperature on the heat sink is monitored by the wireless sensor node reading a thermistor placed in between heat sink fins. We

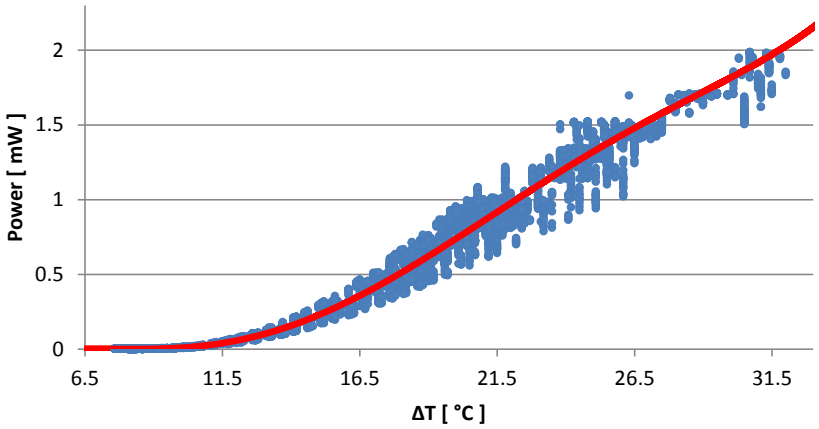

Figure 7.3: Generated output power (in Watts) *versus* Thermal gradient (in Celsius degrees).

verified the validity of those readings using a calibrated thermal imaging camera. An example of the temperature picture is reported in Fig. 7.4. These experiments allowed us to confirm the thermal stability of the system. The results show that the harvesting module does not affect the thermal stability of the ARM CPU; conversely, the module allows the host to keep the maximum working temperature below the warning limit of 82° C (the threshold used by the DVFS module of the Linux Kernel to scale down the working frequency) even with 100% load at its maximum clock speed (1.7 GHz). More details are given in Section 7.5.

To control the host CPU workload, we let the host processor run benchmark applications with different combinations of the clock frequency f_{CLK}, the percentage of activity dedicated to the task CPU_{load}, and the duration I_{task}, where the parameter CPU_{load} considers the user CPU usage without the operating system overhead. Differently from previous charaterization, in which we modeled some reasonable application scenarios for a general purpose computing unit, in this case we preferred a to use a generic algebraic task, and to play with the frquency governor of Unix's kernel to achieve a finer control of systems' performance.

From the measured data, we extracted the mathematical formula (see Fig. 7.3 above) that models the input-output relation of the harvester. Moreover, we obtained the charge-discharge model of the storage unit, and the circuit efficiency (η), and we studied the relation between tasks (expressed as a combination of f_{CLK}, CPU_{load}, I_{task}) and harvested power.

To characterize the harvester, we focus our considerations on $CPU_{load} = \{0,\ 25,\ 50,\ 75,\ 100\}\%$, and $f_{CLK} = \{1.2, 1.3, 1.4, 1.5, 1.6\}$ GHz. Whatever the clock frequency, with no task in execution (meaning the CPU is Idle, $CPU_{load} = 0\%$), the thermal gradient stabilizes around 8° C. When CPU_{load} is 25% or larger, the thermal gradient ΔT grows with the clock frequency. Also the task duration influences the amount of energy the system can generate. Generally, the transients of the CPU temperature last less than 120 s, but they are faster with higher clock frequencies. In fact, short tasks duration

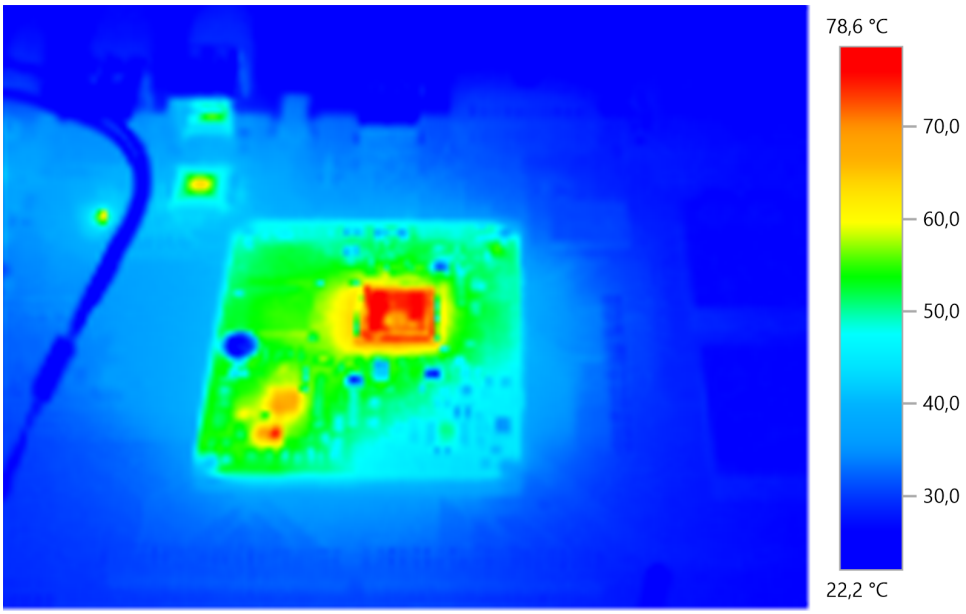

(**a**)Top view.

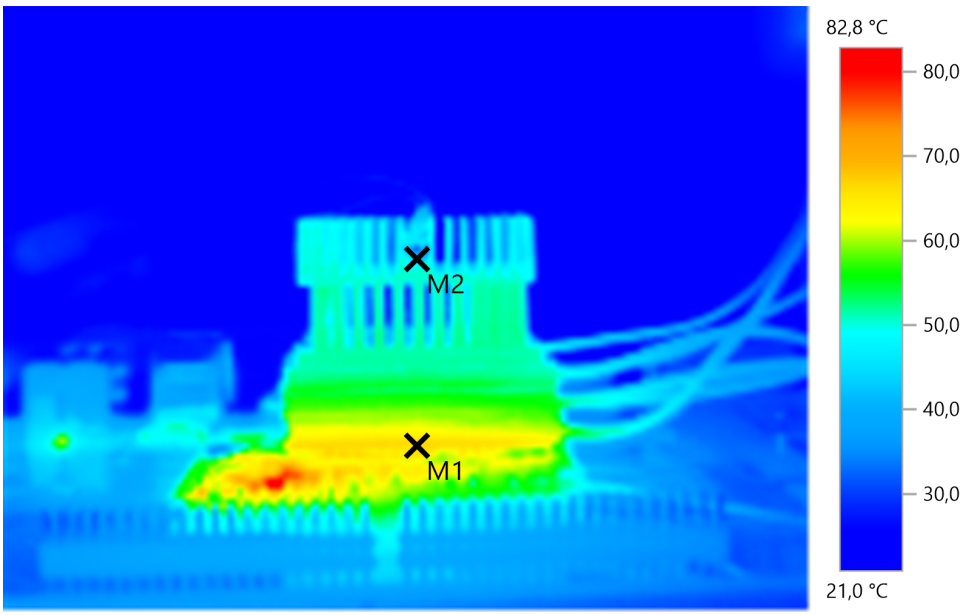

(**b**)Side view.

Figure 7.4: Thermal picture of the host equipped with the harvesting system during the experimental characterization.

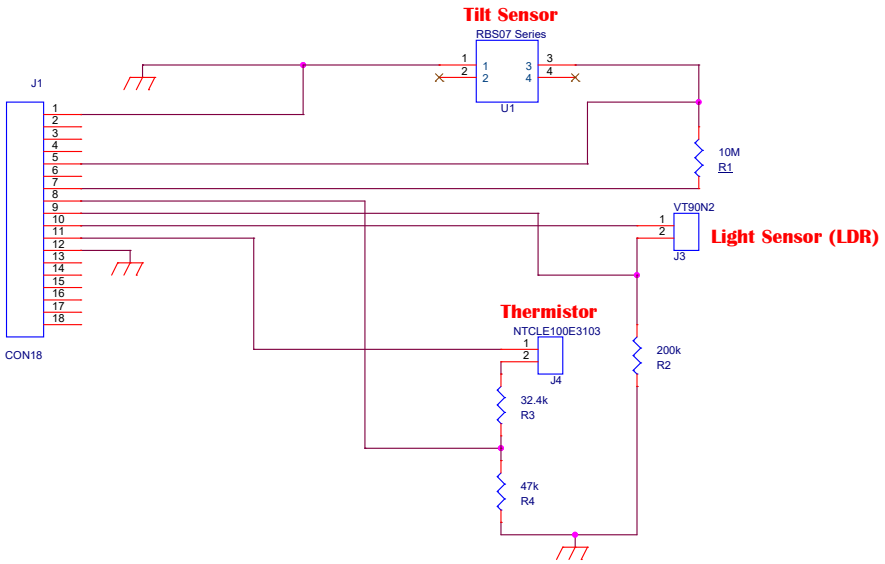

Figure 7.5: Schematic representation of the interconnection between the wireless node and the sensors.

cannot increase the temperature as much as needed in order to observe a noticeable amount of power at the harvester output. For a CPU load $CPU_{load} = [25\%, 50\%, 75\%]$, whatever the CPU speed, a task that runs for more than two minutes generates around 3° C greater thermal gradients with respect to short computational bursts. Under the same conditions, for a task with $CPU_{load} = 100\%$, the thermal gradient may increase as much as 5° C. When the CPU clock is set at its maximum speed, we observe the maximum value for ΔT reaching 30.31° C, that matches with the maximum generated electrical power of about 2 mW (as shown in Fig. 7.3). These values represent the boundary conditions that we used to evaluate the performance of the platform.

7.4 The Smart Heat-Sink Design

A monitoring system for data centers is used to measure several environmental parameters whose value, when out of the ordinary, could denote that a hazard may be incurred. Different types of sensors, such as light or proximity detectors, can reveal the presence of unauthorized personnel. Moreover, by controlling the temperature value it is possible to infer whether there is an overheating problem, due for example to a fault in the cooling system. In our specific case, we aimed at the control of the ambient temperature, the heat sink temperature, the ambient light, the node orientation, and the state of charge of the supercapacitor that fed power to the monitoring device.

7.4.1 Control Unit

Among the microcontroller based devices that lend themselves to this type of applications, that are fitted with radio, and show an energy budget in line with the availability provided from the TEG, we decided to switch to the *TI eZ430-RF2500*, based on a MSP430 MCU and the cc2500 radio chip, mainly because of reduced sizes and low-power features.

Several sensors have been connected to the microcontroller/ transceiver board according to the diagram reported in Fig. 7.5. A light dependant resistor (LDR), the *Excelitas Tech VT90N2* [24], has been used to measure light variations, and to spot the presence of personnel inside the server room. The LDR is read out through a voltage divider connected to internal ADC of the MSP430. To minimize the power consumption, the LDR circuit is powered only during measurement. The sensor stabilizes its value after only 76 ms.

A tilt-switch sensor, the *DIP - RBS07 Series* [63], has been attached to a General Purpose Input-Output (GPIO) port (configured as input) of the MCU, and can be used, for example, to detect if the server rack door has been opened. When the sensor stands on its correct position, the connection is closed. If it has been tilted, for example because the door has been opened, the sensor switches and closes the electrical connection. It consumes negligible power in both stages, when the contact is open and thanks to the 10 MΩ series resistor also when is closed.

A thermoresistor ([86]) is placed within the heat sink central fins to measure the temperature. The wireless node collects the environmental parameters, and sends the aggregated data to a central node. Moreover, the node can exploit the temperature information to infer the thermal gradient, and estimate the expected generated power of the TEG. Exploiting this prediction, the sensor node can determine autonomously at which rate to operate in order to stay alive. To reduce the power consumption, the monitoring node powers the sensors only when it is sampling their value.

7.4.2 Prototype Development

The *TI eZ430* board implements the proprietary *Simplicity* communication protocol stack, which provides advanced networking capabilities. We implemented a communication scheme that does not require setting up the radio link, the packet acknowledge, and token forwarding, since for the monitoring purposes a star-topology is most suitable, and allows the node to consume less energy when transmitting data, assuming for example to deploy a dozen of nodes and a single data collectro in each rack.

The monitoring node is fed as soon as the voltage stored in the supercap reaches $HTH = 3.08$ V. The node boots with its registers settled for network configuration, timer usage and ADCs, and set to Low-Power Mode 3 (LPM3) with interrupts enabled. After a given amount of time, I_s, the microcontroller wakes up, and broadcasts a packet to a remote node (or to a PC) with the measured value for the heat sink temperature, light, orientation, and its voltage supply. Then in returns to *Idle* mode, and reiterates after I_s seconds, assuming that the amount of energy available is enough to power it. When the voltage of the capacitor decreases under the threshold $LTH = 2.4$ V, the comparator opens the path between the storage unit and the microcontroller, which consequently

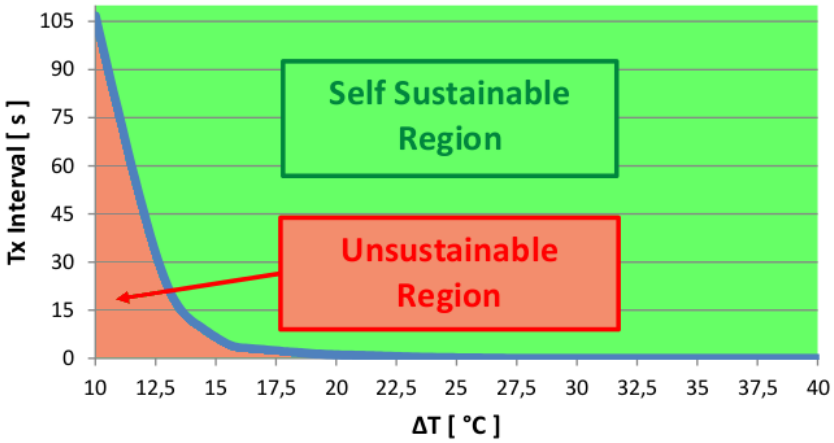

Figure 7.6: The Self-Sustainability curve represents the union of the points that separates the region of the plot where the nodes can operate autonomously from the region in which it fails to operate.

switches off.

We implemented two different policies to choose the value of I_s:

Fixed time: in this implementation I_s is fixed, and set equal to one second in our experiments. Therefore, if the supercapacitor contains enough power to supply the node, the node continuously transmits with a constant data rate;

Temperature dependent: in the second implementation, the time between two successive transmissions is chosen by the sensor node according to the temperature measured on the heat sink of the harvester, which provides an indication of the amount of scavenged energy, used to estimate an optimal communication interval as described below.

In the first implementation, the node goes to deep low power mode in the time interval between two successive transmitting events, limiting the total power consumption. However, this implementation makes it possible for the node to drain enough power to bring the voltage on the storage unit below the LTH threshold, in case the harnessed energy is not enough. If this occurs, the node cannot operate until the available power returns larger than HTH.

In the second implementation, the value of I_s is a function of the heat sink temperature. Since the temperature on the heat sink is a function of the TEG system temperature, we can infer from it the amount of power the harvester is generating, and then predict the expected amount of charge that can supply the node. The node estimates the amount of power available through a look-up-table containing the data that express the relationship between temperature on the heat sink, and generated voltage, which was derived from characterization data extracted.

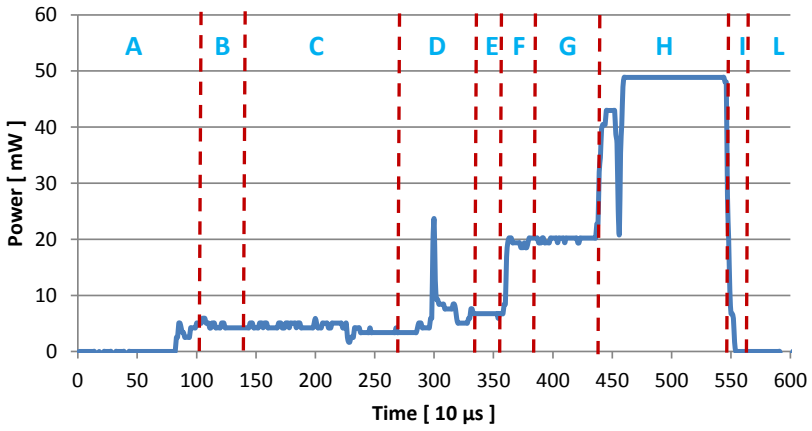

Figure 7.7: Measured power consumption of the prototype node during a sense and send event.

7.4.3 Prototype Energy Consumption

The energy spent by the node has been characterized by measuring the voltage across a resistor R_{meas} placed in series between the sensing node and the power supply circuitry. Through this approach, and by knowing the resistance $R_{meas} = 10.2\ \Omega$, we computed the effective power consumption. By analyzing how the values of the power consumption evolve with time, we also associated the specific task the node is performing to the corresponding amount of energy spent. Table 7.1 summarizes the task specific power consumption value as extracted from Fig. 7.7. From Table 7.1 it is possible to compute the amount of energy required by the system to perform a single sense-and-send operation, that is $E_{TX} = 84.69\ \mu$J. Moreover, knowing that for each send operation the radio is used for $I_{TX} = 4.86$ ms, the total current consumption of the *MSP430* board in *Idle* mode together with its radio module *CC2500* is [84, 83]: $MSP430_{IDLE} + CC2500_{IDLE} = 1.3\ \mu$A. The power consumption is: $P_{IDLE} = 2.2$ V $\times 1.3\ \mu$A $= 2.86\ \mu$W. To supply the system in *Sleep* mode for one second $2.84\ \mu$J are needed. Summing all together, to run the application on the board for one second, $E_{tot} = 87.53\ \mu$J is the required total energy.

7.4.4 Self Sustainability Analisys

To evaluate the range of values within which the node can be self-powered, we compute the self-sustainability curve. The curve in Fig. 7.6 represents the maximum transmission interval time that is feasible for thermal gradients in the range 10-40° C.

The region above the curve represents the self-sustainable region (S-S.R.), and corresponds to the condition in which the node can operate continuously, while the region below the curve corresponds to the unsustainable region (U.R.) and represents the transmission rates the node cannot operate at for longer periods of time. Knowing this data, it is possible to identify the optimal communication interval depending on the heat generated by the thermal source. For example, with 10° C thermal gradient the shorter transmission interval is approximately 110 s, while at 15° C it is possible to set a trans-

Table 7.1: Power consumption of the *eZ430-RF2500* board during a sense-and-send event.

Symbol	**Event**	$\overline{I}$ **[mA]**	**Time [ms]**	**P [mW]**	**E [μJ]**
A-L	Idle	1.3×10^{-3}	-	-	-
B	Temp & Vcc	1.94	0.38	4.26	1.61
C	NTC Temp	1.94	1.37	4.26	5.83
D	Calc & XOSC Startup	2.56	0.70	5.63	3.94
E	Ripple Counter	3.02	0.21	6.64	1.39
F-G	Msg & PLL calib.	8.29	0.87	18.23	15.86
H	Tx Mode	20.95	1.19	46.09	54.84
I	Switch LPM	3.95	0.14	8.69	1.21

mission rate lower than 10 s.

Tasks running on the host processor have been described according to three parameters: the clock frequency f_{CLK}, the percentage CPU_{load} of CPU time dedicated to the task, and the duration I_{task} of the process. This allows us to model every software as a combination of these three values, and to associate the corresponding CPU temperature. With $f_{CLK} = 1.2$ GHz the mean heat sink temperature stabilizes around 31.8° C, while the average CPU temperature is near 45° C. The node transmits every 30 s even if it is possible to send every 13 s with a $\Delta T = 14^\circ$ C. As a consequence the system consumes less power than available, and the storage voltage increases. With the adjustable transmission rate implementation, after a certain period of time, the voltage in the storage capacitors tends to stabilize around 4.6 V, therefore the system is self-sustainable.

Every time the supply voltage crosses the high threshold, the node activates, and starts sampling and sending. In Fig. 7.8, and Fig. 7.9 we represented the amount of time the node is active, and the number of transmissions it is able to perform with two different values of f_{CLK}. When the host processor runs with $f_{CLK} = 1.2$ GHz and $CPU_{load} = 100\%$, the recharge time required to reach the voltage threshold is about 25 s, and after being activated, the monitoring unit is able to send data for 7 or 8 times before discharging the supercapacitor, and becoming inactive. This condition is represented in Fig. 7.8, while Fig. 7.9 shows measured data when the host processor clocks at its maximum speed: $f_{CLK} = 1.7$ GHz and $CPU_{load} = 100\%$. In the latter case, it is possible to perform continuous transmissions (every second) without discharging the supercapacitor once the temperature reaches its steady state. Moreover, just 6 seconds are required to recharge the selected storage unit.

In Fig. 7.10 and Fig. 7.11 we show the activity of the node when the $CPU_{load} = 50\%$ for $f_{CLK} = 1.2$ GHz, and for $f_{CLK} = 1.7$ GHz respectively, together with the voltage across the storage supercapacitor. Even with half the CPU load, the energy harvesting system is able to recharge the storage supercapacitor. For $f_{CLK} = 1.2$ GHz, the monitoring system can sample the sensors and send the data six times before going under threshold. After that, 33 seconds are required to restore the charge. For $f_{CLK} = 1.7$ GHz, the monitoring

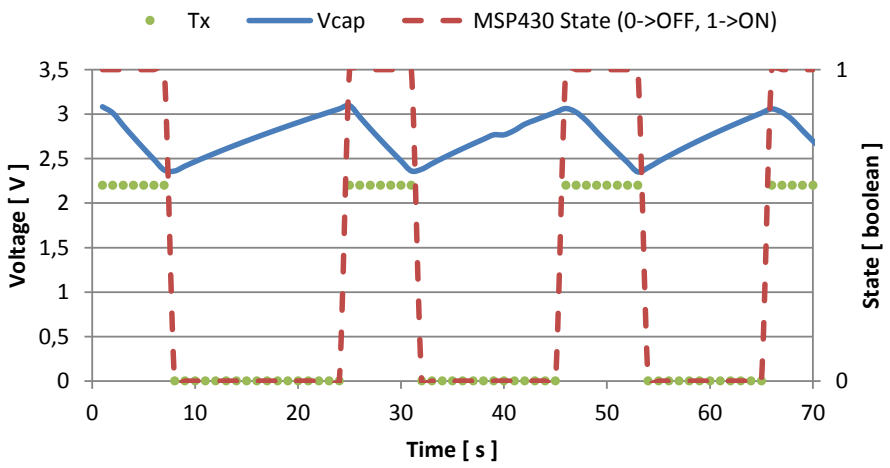

Figure 7.8: Voltage in the storage capacitance and transmission events with $f_{CLK} = 1.2$ GHz, and $CPU_{load} = 100\%$.

unit senses and sends eleven times and switches off for 8 or 9 seconds before resuming.

7.5 Smart Heat-Sink Performance Evaluation

To ensure that the proposed systems does not adversely affect the reliability of the server, we evaluated how the harvester and the sensor node influence the thermal dissipation characteristic of the target host CPU. In our experiment, we launched a process with (f_{CLK}, CPU_{load}, I_{task}) = (1.7 GHz, 100%, 120 s) and measured the CPU temperature in three scenarios: the CPU without passive dissipation (as sold by the vendor), with the heat sink on top (the same heat sink mounted on top of the harvesting module), and with the complete proposed prototype. The experiment is used to understand whether the system causes the CPU to overheat. One the contrary, it turns out that our prototype allows the CPU to work at lower, safer temperatures. Results are depicted in Fig. 7.12. As the reader can easily notice, the temperature of the CPU without heat sink rapidly reaches a dangerous value (about 90° C). The presence of the heat sink allows the temperature to not exceed 62° C. Even better, when the complete harvester is mounted the temperature reaches only 55° C. Therefore, the proposed system is an excellent passive heat sink which ensures a reduction of temperature of about 5° C in Idle mode, and up to 7° C when the target host runs at its maximum performance. Compared to the host deprived of the heat sink, the harvester guarantees to cool the CPU by about 35 degrees. This is contrary to other similar applications where the presence of a TEG in between a desktop CPU package and the heat sink raises the CPU temperature by about 10–20° C [38, 45].

The application with fixed sampling period $I_s = 1$ s is self-sustainable when the thermal gradient is larger than $\Delta T = 12^\circ$ C, corresponding to a CPU temperature of about 35° C. The dynamic (or adaptive) I_s implementation preserves energy when the temperatures are low, and allows the system to increase the sampling rate when temperatures

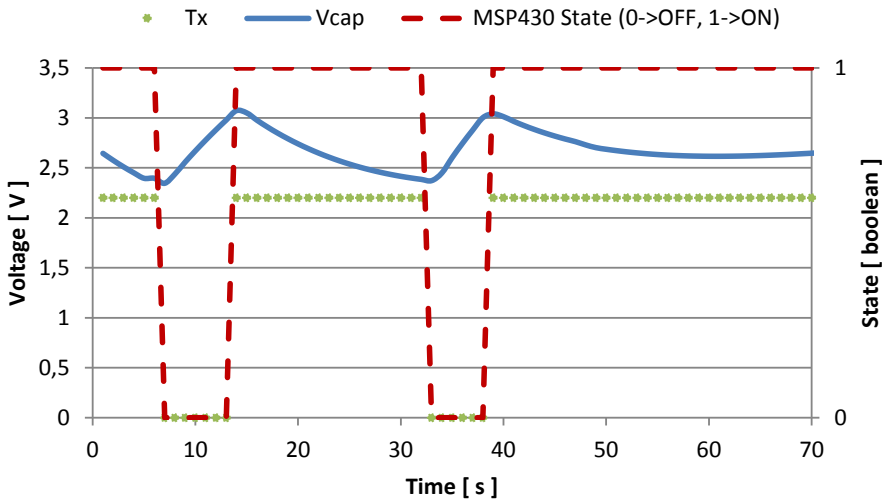

Figure 7.9: Voltage in the storage capacitance and transmission events with $f_{CLK} = 1.7$ GHz, and $CPU_{load} = 100\%$.

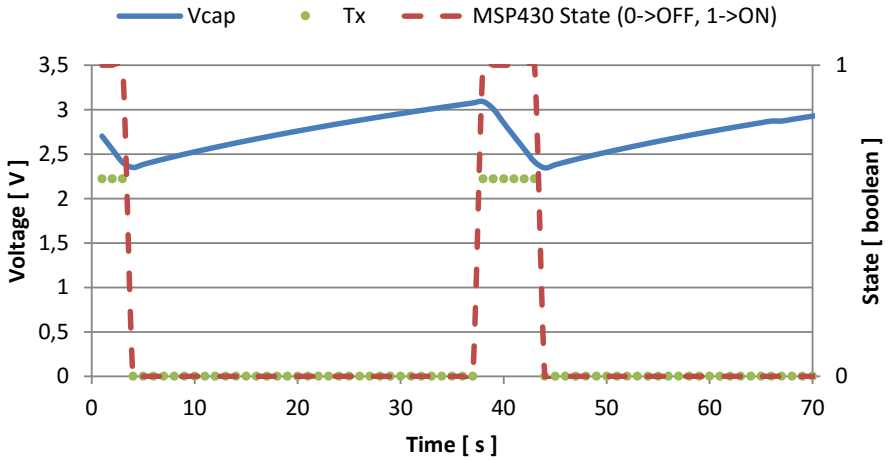

Figure 7.10: Voltage in the storage capacitance and transmission events with $f_{CLK} = 1.2$ GHz, and $CPU_{load} = 50\%$.

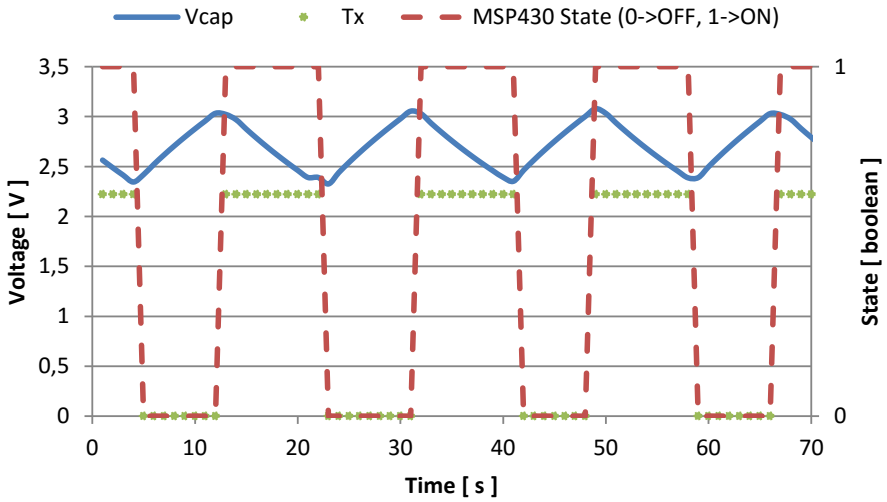

Figure 7.11: Voltage in the storage capacitance and transmission events with $f_{CLK} = 1.7$ GHz, and $CPU_{load} = 50\%$.

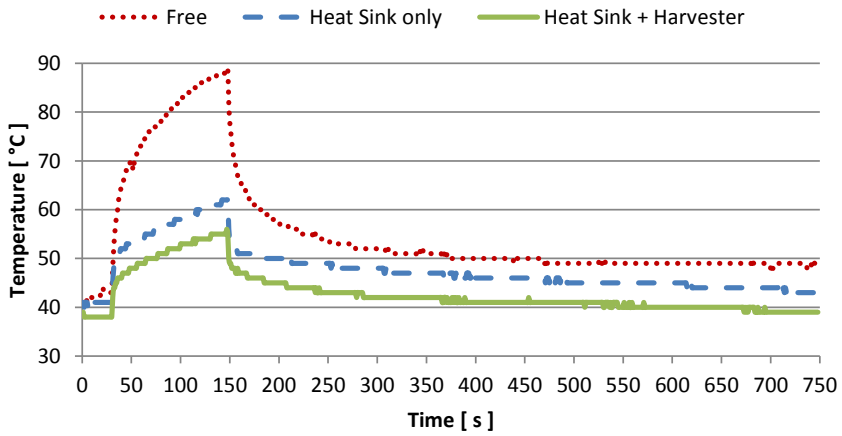

Figure 7.12: CPU temperature trend in three configuration: with and without the heat sink, and with the proposed prototype on top.

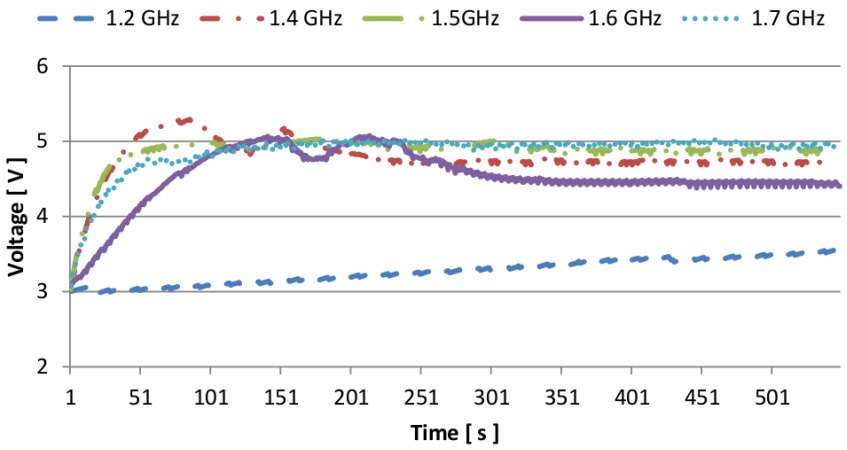

Figure 7.13: Measured supercapacitor voltage using adaptive sampling rate for different clock frequencies of the host processor.

increase. For thermal gradients smaller than 12.5° C, the system works with sampling period $I_s = 109$ s, while with 15° C the period decreases to around 6 seconds, which is reasonable for the target application of the WSN. When $\Delta T \geq 22.5^{\circ}$ C, the time interval is set to $I_s = 1.27$ s, therefore the data rate is comparable with the fixed sampling period implementation. For thermal gradients above 25° C the system is self-sustainable with a transmission interval as small as $I_s = 0.35$ s.

The dynamic (or adaptive) application has been tested to evaluate its effectiveness under different working conditions. Results are summarized in Fig. 7.13, in which we focused on the trend of the capacitor charge upon reaching the threshold of operation of the node (i.e., *HTH*). The plot shows that the implementation requires some initial time to reach a steady state. During this time, of a few minutes, the heat has not propagated uniformly, therefore the temperature readings are not accurate and the voltage inside the storage unit tends to fluctuate, but never decreases below the *HTH* threshold. After the temperature stabilizes, the system is always switched on and the voltage across the supercapacitor increases continuously toward the maximum value of 5 V. Interestingly, for $f_{CLK} = 1.6$ GHz the curve does not follow the increasing trend, and instead stabilizes around 4.5 V. This happens because the resulting thermal gradient ΔT is mid way between two values in the look-up-table the microprocessor uses to choose the sampling rate I_s, which are selected alternatively. As a consequence, the voltage tends to fluctuate and never reaches 5 V, since in this case the node consumes more power than what the harvester generates.

The storage capacitance must be sized according to the application the node has been programmed to perform, the efficiency of the harvester, the thermal characteristics of the *host* processor, and the policy of data collection rate. In fact, in the prototype device described in previous sections, by using a storage supercapacitor $SC1 = 1$ F, 67 minutes are required to fully charge the empty storage unit with the *host* configured with $f_{CLK} = 1.7$ GHz and $CPU_{load} = 100\%$. When the capacitor is full the node is able to

perform 12 thousand transmissions (with one second interval) after the *host* is switched off. With the same conditions, the prototype equipped with $SC1 = 2.2$ mF recharges in 53 seconds, but is able to transmit only five times after the *host* has been switched off.

Chapter 8

Conclusion

Automatic monitoring of environment, resouces and human processes are crucial and foundamental tasks to improve people's quality of life and to safeguard the natural environment. Achieving energy neutral features with embedded electronic designs in these applications is complex, since energy management and reliable data acquisition require opposed optimization strategies, but mandatory to ensure a capillary spread of such systems. The task becomes harder when one moves to the field of resource constrained applications, since requirements are more tight and available renewable energy sources are limited. Energy management must not be optimized to the detriment of the quality of monitoring and sensors can not be operated without supply.

In this thesis, I presented some examples of embedded system designs to mitigate this gap, both from the hardware and software sides. Three specific monitoring scenarios have been taken into consideration to develop solutions and demonstrate that energy neutrality in monitoring under resource constrained conditions, can be achieved without compromising efficiency and reliability of the collected information.

The main results achieved in the development of this research have been:

- The definition and succesful implementation of a novel ultra low-power methodology to use with gas sensor in portable battery-powered wireless nodes.
- The demonstration of such methodology with two different target gases, such as Volatile Organic Compounds (CO/VOCs) and Natural Gas (CH_4).
- The performance demonstration of such methodology with indoor monitoring WSN for early gas leakage detection, for which three years of autonomy can be expected.
- The design and validation of an embedded module for gas leakage localization in outdoor enviroments suitable for installation as payload of autonomous vehicles.
- The development of a gas leak localization algorithm to minimize the energy required by mobile robots in outdoor monitoring missions.
- The development of an energy harvesting system to be used as primary supply for an embedded wireless node, exploiting thermoelectric energy harvesting (useful where any other kind of renewable energy source is not available and batteries

can not be used) which enables perpetual gas monitoring operation in controlled environments.

- The demonstration of a *smart heat-sink*, namely an energy neutral wireless node with embedded sensors powered by a thermoelectric energy generator, in the shape of a commercial aluminum heat-sink.

Techinques, results and designs described in this thesis, are a starting point to tackle the challenges of our moden society. Many research paths can be identified from the achievement discussed up to this point.

On the volatile pollutant side, there is the need to investigate and understand the impact of using array of sensors with different target chemicals. Diversity provides significant advantages in many fields of engineering and it has been indicated as the path to follow also in the particular case of eNose applications. This means that both new hardware and new software requirements arise to enable the deployment of portable embedded systems capable of running complex pattern recognition algorithms, able to analyze the huge amount of data povided by array of gas sensors, driven by highly optimized energy management techniques.

Obviously these advencements will affect both the indoor and outdoor application scenario, but, in particular for the mobile robot use case, there is also the need to tighten the interaction among sensing unit and motion control system. A feedback control loop from the sensing unit to the embedded system running the auto-pilot routine must be investigated to optimize the scarse resources of the platform. Moreover, a deep integration will pave the way to the realization of swarm of mobile monitoring vehicles, maybe heterogeneous, with wheeled robots and UAVs working together to maximize area coverage, reliability of measure, and in a single word security.

On the neutral monitoring of data centers the most interesting research challenge is to model and design a deeply integrated DCIM system. Basically a control scheme that integrates all the information available from traditional systems and from the smart heat-sinks to achieve a new level of resources optimization (computational time, energy, etc.). The natural evolution of this will be the integration of networks of data centers to jointly manage and optimize resources across different geographical location and a tight integration also with the Smart Grid.

List of Publications

Journal Publications

J1 M. Rossi and D. Brunelli, *Autonomous Gas Detection and Mapping With Unmanned Aerial Vehicles*, IEEE Transactions on Instrumentation and Measurement, 2016, pre-print.

J2 D. Brunelli, R. Passerone, L. Rizzon, M. Rossi and D. Sartori *Self-Powered WSN for Distributed Data Center Monitoring*, MDPI Sensors 2016, vol. 16, no. 57, pag. 1 – 15.

J3 M. Rossi and D. Brunelli, *Ultra Low Power MOX Sensor Reading for Natural Gas Wireless Monitoring*, IEEE Sensors Journal, 2014, vol. 14, no. 10, pag. 3433 – 3441.

J4 D. Brunelli and M. Rossi, *Enhancing lifetime of WSN for natural gas leakages detection*, Microelectronics Journal, Elsevier, 2014, vol. 45, no. 12, pag. 1665 – 1670.

J5 M. Rossi, L. Rizzon, M. Fait, R. Passerone and D. Brunelli, *Energy Neutral Wireless Sensing for Server Farms Monitoring*, IEEE Journal on Emerging and Selected Topics in Circuits and Systems, 2014, vol. 4, no. 3, pag. 324 – 334.

Conferences, Book Chapters and Workshop Publications

C1 D. Brunelli and M. Rossi, *Smart grid configuration tool for HEES systems in smart city districts*, Accepted for publication in proc. of the IEEE Speedam Symposium 2016, 22-24 June, Capri, Italy.

C2 M. Rossi, A. Khouia, L. Lorenzelli and D. Brunelli, *Energy neutral 32-channels embedded readout system for IoT-ready fitness equipments*, Accepted for publication in proc. of the IEEE Sensors Application Symposium 2016, 20-22 April, Catania, Italy.

C3 V. Gallego, M. Rossi and D. Brunelli, *Unmanned aerial gas leakage localization and mapping using microdrones*, Sensor Application Symposium, 2015 IEEE Conference, Zadar, Croatia.

C4 L. Rizzon, M. Rossi, R. Passerone and D. Brunelli, *Energy neutral hybrid cooling system for high performance processors*, Green Computing Conference (IGCC), 2014 IEEE International, Dallas, TX, USA.

C5 M. Rossi, D. Brunelli, A. Adami, L. Lorenzelli, F. Menna and F. Remondino, *Gas-Drone: Portable gas sensing system on UAVs for gas leakage localization*, Sensors, 2014 IEEE Conference, Valencia, Spain.

C6 L. Rizzon, M. Rossi, R. Passerone and D. Brunelli, *Self-powered heat-sink SoC as temperature sensors with wireless interface: Design and validation*, Sensors, 2014 IEEE Conference, Valencia, Spain.

C7 M. Rossi, A. Toppano and D. Brunelli, *Real-time optimization of the battery banks lifetime in Hybrid Residential Electrical Systems*, Design, Automation and Test in Europe Conference and Exhibition (DATE), 2014, Dresded, Germany.

C8 M. Rossi and D. Brunelli, *Ultra low power CH4 monitoring with wireless sensors*, Sensors, 2013 IEEE Conference, Baltimore, MD, USA.

B1 A.Pahlevan, M. Rossi, P. G. Del Valle, D. Brunelli and D. Atienza, *Joint Computing and Electric Systems Optimization for Green Datacenters*, Handbook of Hardware/Software Codesign, Springer, pre-print.

B2 M. Rossi, L. Rizzon, M. Fait, R. Passerone and D. Brunelli, *Self-powered Active Cooling System for High Performance Processors*, in A. De Gloria (a cura di), Applications in Electronics Pervading Industry, Environment and Society, Springer International Publishing, 2016, p. 25-33.

B3 L. Tamburini, M. Rossi and D. Brunelli, *Electronic and ICT Solutions for Smart Buildings and Urban Areas*, in Vesco Andrea, Ferrero Francesco (a cura di), Handbook of Research on Social, Economic, and Environmental Sustainability in the Development of Smart Cities, IGI Global, 2015.

B4 D. Brunelli and M. Rossi, *CH4 Monitoring with Ultra-Low Power Wireless Sensor Network*, in A. De Gloria (a cura di), Lecture Notes in Electrical Engineering: Applications in Electronics Pervading Industry, Environment and Society, Springer International Publishing, 2014, p. 13-25.

W1 M. Rossi and D. Brunelli, *Forecasting data centers power consumption with the Holt-Winters method*, Environmental Energy and Structural Monitoring Systems (EESMS), 2015 IEEE Workshop on, Trento, Italy.

W2 D. Brunelli, I. Minakov, R. Passerone and M. Rossi, *Smart monitoring for sustainable and energy-efficient buildings: A case study*, Environmental Energy and Structural Monitoring Systems (EESMS), 2015 IEEE Workshop on, Trento, Italy.

W3 D. Brunelli, D. Porcarelli, D. Balsamo and M. Rossi, *Self-powered wireless energy meter*, Proc. of the 1st International Workshop on Energy Neutral Sensing Systems, ACM ENSSys 2013, Rome,Italy.

W4 D. Brunelli, I. Minakov, R. Passerone and M. Rossi, *POVOMON: An Ad-hoc Wireless Sensor Network for indoor environmental monitoring*, Environmental Energy and Structural Monitoring Systems (EESMS), 2014 IEEE Workshop on, Naples, Italy.

W5 L. Rizzon, M. Rossi, R. Passerone and D. Brunelli, *Wireless Sensor Networks for Environmental Monitoring Powered by Microprocessors Heat Dissipation*, Proc. of the 1st International Workshop on Energy Neutral Sensing Systems, ACM ENSSys 2013, Rome, Italy.

W6 M. Rossi and D. Brunelli, *Electricity demand forecasting of single residential units*, Environmental Energy and Structural Monitoring Systems (EESMS), 2013 IEEE Workshop on, Trento, Italy.

W7 M. Rossi and D. Brunelli, *Analyzing the transient response of MOX gas sensors to improve the lifetime of distributed sensing systems*, Advances in Sensors and Interfaces (IWASI), 2013 5th IEEE International Workshop on, Bari, Italy.

W8 M. Rossi and D. Brunelli, *Ultra low power Wireless Gas Sensor Network for environmental monitoring applications*, Environmental Energy and Structural Monitoring Systems (EESMS), 2012 IEEE Workshop on, Perugia, Italy.

Bibliography

[1] *Data Sheet: JN5148-001, IEEE 802.15.4 Wireless Microcontroller, Jennic, U.K.* http://www.jennic.com/files/support_files/JN-DS-JN5148-1v8.pdf (Version of 20-07-2011).

[2] MICS-5121 data sheet.

[3] MICS-5525 data sheet.

[4] Jude Allred, Ahmad Bilal Hasan, Saroch Panichsakul, William Pisano, Peter Gray, Jyh Huang, Richard Han, Dale Lawrence, and Kamran Mohseni. Sensorflock: an airborne wireless sensor network of micro-air vehicles. In *Proceedings of the 5th international conference on Embedded networked sensor systems*, pages 117–129. ACM, 2007.

[5] Amazon.com, Inc. (Amazon web services). Amazon initiatives for sustainable energy. [Online] `http://aws.amazon.com/it/about-aws/sustainable-energy/`.

[6] Anand, M.L. and Ramesh, C. R. Assessment and impact study of surface soil on thermoelectric generator energy efficiency. *International Journal of Scientific and Engineering Research*, 4, August 2008.

[7] Apple Inc. Environmental responsibility. [Online] `http://www.apple.com/environment/our-progress/`.

[8] AppliedSensor. *Datasheet AS-MLK*. http://www.appliedsensor.com.

[9] Arndale Board [Technical Reference]. http://www.arndaleboard.org/wiki/-index.php/WiKi, 2013.

[10] P. Bhattacharyya, D. Verma, and D. Banerjee. Microcontroller based power efficient signal conditioning unit for detection of a single gas using mems based sensor. *International Journal on Smart Sensing and Intelligent Systems*, 3(4), December 2010.

[11] C. Bolchini, A. Miele, A. Das, A. Kumar, and B. Veeravalli. Combined DVFS and mapping exploration for lifetime and soft-error susceptibility improvement in mpsocs. In *Design, Automation and Test in Europe Conference and Exhibition (DATE), 2014*, March 2014.

[12] C. Caione, D. Brunelli, and L. Benini. Compressive sensing optimization for signal ensembles in wsns. *Industrial Informatics, IEEE Transactions on*, 10(1):382–392, Feb 2014.

[13] Daniele Caltabiano, Giovanni Muscato, Angelo Orlando, Cinzia Federico, Gaetano Giudice, and Sergio Guerrieri. Architecture of a uav for volcanic gas sampling. In *Emerging Technologies and Factory Automation, 2005. ETFA 2005. 10th IEEE Conference on*, volume 1, pages 744–749. IEEE, 2005.

[14] David Carli, Davide Brunelli, Luca Benini, and Massimiliano Ruggeri. An effective multi-source energy harvester for low power applications. In *Design, Automation Test in Europe Conference Exhibition (DATE), 2011*, March 2011.

[15] Chien-Ying Chen and P.H. Chou. DuraCap: A supercapacitor-based, power-bootstrapping, maximum power point tracking energy-harvesting system. In *Low-Power Electronics and Design (ISLPED), 2010 ACM/IEEE International Symposium on*, pages 313–318, Aug 2010.

[16] Min Chen, Lasse A Rosendahl, Thomas J Condra, and John K Pedersen. Numerical modeling of thermoelectric generators with varing material properties in a circuit simulator. *Energy Conversion, IEEE Transactions on*, 24(1):112–124, 2009.

[17] Sukwon Choi, Nakyoung Kim, Hojung Cha, and Rhan Ha. Micro sensor node for air pollutant monitoring: Hardware and software issues. *Sensors*, 9(10):7970–7987, 2009.

[18] Pai H Chou and Sehwan Kim. Techniques for maximizing efficiency of solar energy harvesting systems. *Proc ICMU*, 10:26–28, 2010.

[19] Comitato Italiano Gas. Incidenti da gas combustibili 2014.

[20] Kai Daniel, Bjoern Dusza, Andreas Lewandowski, and Christian Wietfeld. Airshield: A system-of-systems muav remote sensing architecture for disaster response. In *Systems Conference, 2009 3rd Annual IEEE*, pages 196–200. IEEE, 2009.

[21] Dell Inc. Openmanage power center temperature monitoring. [Online] `http://www.dell.com/support/article/us/en/19/SLN283321/EN`.

[22] A. Depari, A. Flammini, D. Marioli, E. Sisinni, E. Comini, and A. Ponzoni. An electronic system to heat mox sensors with synchronized and programmable thermal profiles. *Instrumentation and Measurement, IEEE Transactions on*, 61(9):2374–2383, Sept 2012.

[23] D. Dondi, A. Bertacchini, L. Larcher, P. Pavan, D. Brunelli, and L. Benini. A solar energy harvesting circuit for low power applications. In *Sustainable Energy Technologies, 2008. ICSET 2008. IEEE International Conference on*, pages 945 –949, nov. 2008.

[24] Excelitas Tech. Photoconductive cell, vt900 series. [Online] `http://www.farnell.com/datasheets/612931.pdf`.

[25] Robert Gates. Data center physical security gets a tougher look. [Online] `http://searchdatacenter.techtarget.com/news/4500248374/Data-center-physical-security-gets-a-tougher-look`.

[26] Globe Newswire. Embedded system market size worth $214.39 billion by 2020: Radiant insights inc. [Accessed Feb 19th 2016] `https://globenewswire.com/news-release/2015/11/10/785476/0/en/Embedded-System-Market-Size-Worth-214-39-Billion-By-2020-Radiant-Insights-Inc.html`.

[27] Google Inc. Google green. [Online] http://www.google.com/green/.

[28] CA Gould, NYA Shammas, Steven Grainger, and Ian Taylor. Thermoelectric cooling of microelectronic circuits and waste heat electrical power generation in a desktop personal computer. *Materials Science and Engineering: B*, 176(4):316–325, 2011.

[29] Data Centre Specialist Group. Meeting the energy efficiency and financial challenges in it. Technical report, BCS, 2007. http://www.oxon.bcs.org/downloads/BCS-DCSG 20070927.pdf.

[30] Peter Gründler. *Chemical Sensors: An Introduction for Scientists and Engineers*. Springer, 2007.

[31] Harry Zervos. Energy harvesting for automotive applications. Cambridge, MA: IDTechEx, 2011. http://media.idtechex.com/pdfs/en/S2142B5538.pdf.

[32] Jane K Hart and Kirk Martinez. Environmental sensor networks: A revolution in the earth system science? *Earth-Science Reviews*, 78(3):177–191, 2006.

[33] Anna Haywood, Jon Sherbeck, Patrick Phelan, Georgios Varsamopoulos, and Sandeep K.S. Gupta. Thermodynamic feasibility of harvesting data center waste heat to drive an absorption chiller. *Energy Conversion and Management*, 58(0):26 – 34, 2012.

[34] A. Hulanicki, S. Glab, and F. Ingman. Chemical sensors: definitions and classification. *Pure and Applied Chemistry*, 63(9), 1991. http://www.iupac.org/-publications/pac/63/9/1247/.

[35] IDTechEx Ltd. Advanced and post lithium ion batteries 2016 - 2026 markets forrecasts. [Online] http://www.idtechex.com/research/reports/advanced-and-post-lithium-ion-batteries-2016-2026-technologies-markets-forecasts-000449.asp.

[36] IDTechEx Ltd. Thermoelectric energy harvesting 2014-2024: Devices, applications, opportunities. http://www.marketresearch.com/IDTechEx-Ltd-v3153/Thermoelectric-Energy-Harvesting-Devices-Applications-8506004/.

[37] International Business Machines Corporation (IBM). Energy, the environment and ibm. [Online] http://www.ibm.com/ibm/green/.

[38] Sriram Jayakumar and Sherief Reda. Making sense of thermoelectrics for processor thermal management and energy harvesting. In *International Symposium on Low Power Electronics and Design (ISLPED)*, 2015.

[39] V. Jelicic, M. Magno, D. Brunelli, G. Paci, and L. Benini. Context-adaptive multimodal wireless sensor network for energy-efficient gas monitoring. *Sensors Journal, IEEE*, 13(1):328–338, 2013.

[40] V. Jelicic, M. Magno, G. Paci, D. Brunelli, and L. Benini. Design, characterization and management of a wireless sensor network for smart gas monitoring. In *Advances in Sensors and Interfaces (IWASI), 2011 4th IEEE International Workshop on*, pages 115 –120, june 2011.

[41] Robel Kiflemariam and Cheng-Xian Lin. Numerical simulation and parametric study of heat-driven self-cooling of electronic devices. *Journal of Thermal Science and Engineering Applications*, 7(1):011008, 2015.

[42] Y.W. Kim, S.J. Lee, G.H. Kim, and G.J. Jeon. Wireless electronic nose network for real-time gas monitoring system. *IEEE International Workshop on Robotics and Sensors Environment*, 2009.

[43] C. Bambang Dwi Kuncoro, Nor Hayati Saad, Ahmed Jaffar, Cheng Yee Low, and Salmiah Kasolang. Wireless e-Nose Sensor Node: State of the Art. *Procedia Engineering*, 41(Iris):1405–1411, January 2012.

[44] Laird Technologies, Inc. Thermoelectric handbook. www.lairdtech.com/temhandbook/.

[45] Soochan Lee, Patrick E Phelan, and Carole-Jean Wu. Hot spot cooling and harvesting cpu waste heat using thermoelectric modules. In *ASME 2014 International Mechanical Engineering Congress and Exposition*. American Society of Mechanical Engineers, 2014.

[46] Alexander Malaver, Nunzio Motta, Peter Corke, and Felipe Gonzalez. Development and integration of a solar powered unmanned aerial vehicle and a wireless sensor network to monitor greenhouse gases. *Sensors*, 15(2):4072–4096, 2015.

[47] Loreto Mateu, Cosmin Codrea, Néstor Lucas, Markus Pollak, and Peter Spies. Energy harvesting for wireless communication systems using thermogenerators. In *Proceeding of the XXI Conference on Design of Circuits and Integrated Systems (DCIS), Barcelona, Spain*, 2006.

[48] Loreto Mateu, Markus Pollak, and Peter Spies. Analog maximum power point circuit applied to thermogenerators. *PowerMEMS 2008*, pages 9–12, 2008.

[49] Micropelt. TE-CORE6 — TE-CORE7 ThermoHarvesting Power Module, October 2012.

[50] Microsoft Corporation. Data center genome project. [Online] `http://research.microsoft.com/en-us/projects/dcgenome/`.

[51] Microsoft Corporation. Microsoft environment. [Online] `https://www.microsoft.com/environment/`.

[52] A. Moradi and M. Sawan. A new FSK-based transmitter dedicated for low-power wireless medical transceivers. In *Electrical and Computer Engineering (CCECE), 2011 24th Canadian Conference on*, pages 1238–1241, May 2011.

[53] C. Moser, L. Thiele, D. Brunelli, and L. Benini. Adaptive power management for environmentally powered systems. *Computers, IEEE Transactions on*, 59(4):478 –491, april 2010.

[54] Luca Mottola and Gian Pietro Picco. Programming wireless sensor networks: Fundamental concepts and state of the art. *ACM Comput. Surv.*, 43(3):19:1–19:51, April 2011.

[55] H.T. Nagle, R. Gutierrez-Osuna, and S.S. Schiffman. The how and why of electronic noses. *Spectrum, IEEE*, 35(9):22–31, Sep 1998.

[56] United nations conference on climate change. [Online] `http://www.cop21.gouv.fr/en`.

[57] Patrick P Neumann, Sahar Asadi, Achim J Lilienthal, Matthias Bartholmai, and Jochen H Schiller. Autonomous gas-sensitive microdrone: wind vector estimation and gas distribution mapping. *Robotics & Automation Magazine, IEEE*, 19(1):50–61, 2012.

[58] Francesco Nex and Fabio Remondino. Uav for 3d mapping applications: a review. *Applied Geomatics*, 6(1):1–15, 2014.

[59] Nextreme Thermal Solutions. Thermomobility wireless power generator (WPG-1).

[60] Xing Niu, Jianlin Yu, and Shuzhong Wang. Experimental study on low-temperature waste heat thermoelectric generator. *Journal of power sources*, 188(2):621–626, 2009.

[61] Luís ML Oliveira and Joel JPC Rodrigues. Wireless sensor networks: a survey on environmental monitoring. *J. of Communications*, 6(2), 2011.

[62] Jacopo Olivo, Davide Brunelli, and Luca Benini. A kinetic energy harvester with fast start-up for wearable body-monitoring sensors. In *Pervasive Computing Technologies for Healthcare (PervasiveHealth), 2010 4th International Conference on*, pages 1–7. IEEE, 2010.

[63] Oncque. Dip - rbs07 series. [Online] `http://www.oncque.com/05-tilt_switch.html?CID=1`.

[64] EI Ortiz-Rivera, A Salazar-Llinas, and J Gonzalez-Llorente. A mathematical model for online electrical characterization of thermoelectric generators using the PI curves at different temperatures. In *Applied Power Electronics Conference and Exposition (APEC), 2010 Twenty-Fifth Annual IEEE*, pages 2226–2230. IEEE, 2010.

[65] pandaboard.org. OMAP4 Pandaboard System Reference Manual, September 2010.

[66] Peter Clarke. Updated: smartphones to kickstart gas sensor market. In *EETimes Europe*. [Accessed Feb 19th 2016] `http://www.electronics-eetimes.com/en/updated-smartphones-to-kickstart-gas-sensor-market.html?cmp_id=7&news_id=222927543`.

[67] Piper Jaffray CIO Survey. Technology, media & telecom 2015. [Online] `https://piper2.bluematrix.com/sellside/EmailDocViewer?encrypt=7856c68e-3f1a-4ce9-a7e7-99fe25145cd9&mime=pdf&co=Piper&id=jyarow@businessinsider.com&source=mail`.

[68] Markus Pollak, Loreto Mateu, and Peter Spies. Step-up DC-DC converter with coupled inductors for low input voltages. *Fraunhofer IIS*, 2008.

[69] O.A. Postolache, J.M.D. Pereira, and P.M.B.S. Girao. Smart sensors network for air quality monitoring applications. *IEEE Transactions on Instrumentation and Measurement*, 58(9), 2009.

[70] Press Release. Idc finds growth, consolidation, and changing ownership patterns in worldwide datacenter forecast. In *IDC Corporate*, 10 Nov. 2014. [Online] `http://www.idc.com/getdoc.jsp?containerId=prUS25237514`.

[71] Aveek Purohit, Zheng Sun, Frank Mokaya, and Pei Zhang. Sensorfly: Controlled-mobile sensing platform for indoor emergency response applications. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, pages 223–234. IEEE, 2011.

[72] Rackspace Global. Understanding our environmental impact. [Online] `http://responsibility.rackspace.com/planet/energy/renewables`.

[73] Rick Merritt. Dell, IBM give thumbs up to ARM Servers. In *EE Times*, 2010.

[74] M.G. Rodriguez, L.E. Ortiz Uriarte, Yi Jia, K. Yoshii, R. Ross, and P.H. Beckman. Wireless sensor network for data-center environmental monitoring. In *Sensing Technology (ICST), 2011 Fifth International Conference on*, pages 533–537, Nov 2011.

[75] M. Rossi and D. Brunelli. Analyzing the transient response of mox gas sensors to improve the lifetime of distributed sensing systems. In *Advances in Sensors and Interfaces (IWASI), 2013 5th IEEE International Workshop on*, pages 211–216, June 2013.

[76] Gianni Rusconi. L'incendio ad aruba.it provoca il più grande black out internet in Italia. In tilt posta e siti di molte aziende. *Il Sole 24 Ore*, April 2011.

[77] Eduard Santamaria, Florian Segor, Igor Tchouchenkov, and Rainer Schoenbein. Rapid aerial mapping with multiple heterogeneous unmanned vehicles. *International Journal On Advances in Systems and Measurements*, 6(3 and 4):384–393, 2013.

[78] Delong Shang, Xuefu Zhang, Fei Xia, and Alex Yakovlev. Asynchronous design for new on-chip wide dynamic range power electronics. In *Proceedings of the conference on Design, Automation & Test in Europe*, page 138. European Design and Automation Association, 2014.

[79] L.V. Shum, P. Rajalakshmi, A. Afonja, G. McPhillips, R. Binions, L. Cheng, and S. Hailes. On the development of a sensor module for real-time pollution monitoring. *International Conference on Information Science and Applications*, 2011.

[80] Stephen So, Farinaz Koushanfar, Anatoliy Kosterev, and Frank Tittel. Laserspecks:: laser spectroscopic trace-gas sensor networks-sensor integration and applications. In *Proceedings of the 6th international conference on Information processing in sensor networks*, pages 226–235. ACM, 2007.

[81] Andrey Somov, Alexander Baranov, Alexey Savkin, Mikhail Ivanov, Lucia Calliari, Roberto Passerone, Evgeny Karpov, and Alexey Suchkov. Energy-aware gas sensing using wireless sensor networks. In *Wireless Sensor Networks, 9th European Conference, EWSN 2012*, volume 7158 of *Lecture Notes in Computer Science*, pages 245–260. Springer Berlin / Heidelberg, 2012.

[82] K. Tang, S. Chiu, H. Hao, S. Wei, T. Lin, C. Yang, D. Yao, and W. Yeh. An electronic-nose sensor node based on polymer-coated surface acoustic wave array for environmental monitoring. *International Conference on Green Circuits and Systems*, 2010.

[83] Texas Instruments. Cc2500 low-cost low-power 2.4 ghz rf transceiver. [Online] `http://www.ti.com.cn/cn/lit/ds/swrs040c/swrs040c.pdf`.

[84] Texas Instruments. Mixed signal microcontroller. [Online] `http://www.ti.com.cn/cn/lit/ds/slas504g/slas504g.pdf`.

[85] The Guardian, 8 June, 2015. G7 leaders agree to phase out fossil fuel use by end of century. [Online] http://www.theguardian.com/world/2015/jun/08/g7-leaders-agree-phase-out-fossil-fuel-use-end-of-century.

[86] Vishay. Ntcle100e3 datasheet. [Online] http://www.vishay.com/docs/29049/ntcle100.pdf.

[87] S. De Vito, P. Di Palma, C. Ambrosino, E. Massera, G. Burrasca, M.L. Miglietta, and G. Di Francia. Wireless sensor networks for distributed chemical sensing: Addressing power consumption limits with on-board intelligence. *IEEE Sensors Journal*, 11(4), 2011.

[88] R.J.M. Vullers, R. van Schaijk, I. Doms, C. Van Hoof, and R. Mertens. Micropower energy harvesting. *Solid-State Electronics*, 53(7):684 – 693, 2009.

[89] D. Wang, D. P. Agrawal, W. Toruksa, C. Chaiwatpongsakorn, M. Lu, and T. C. Keener. monitoring ambient air quality with carbon monoxide wireless network. *Communication of the ACM*, 53(5):138–141, 2010.

[90] Demin Wang, Dharma P Agrawal, Wassana Toruksa, Chaichana Chaiwatpongsakorn, Mingming Lu, and Tim C Keener. Monitoring ambient air quality with carbon monoxide sensor-based wireless network. *Communications of the ACM*, 53(5):138–141, 2010.

[91] Wispes. W24th. http://www.wispes.com/products-page/wireless-sensors/wispes-w24th/.

[92] Heng Yu, Rizwan Syed, and Yajun Ha. Thermal-aware frequency scaling for adaptive workloads on heterogeneous mpsocs. In *Design, Automation and Test in Europe Conference and Exhibition (DATE)*, March 2014.

[93] Jinbiao Yu, Jie Li, Qilin Dai, Dong Li, Xiaobei Ma, and Yunpeng Lv. Temperature compensation and data fusion based on a multifunctional gas detector. *Instrumentation and Measurement, IEEE Transactions on*, 64(1):204–211, Jan 2015.

[94] Zewde Yeraswork. AMD Launches First ARM-based Server CPU. In *EE Times*, 2014.

[95] Yu Zhou, S. Paul, and S. Bhunia. Harvesting wasted heat in a microprocessor using thermoelectric generators: Modeling, analysis and measurement. In *Design, Automation and Test in Europe*, pages 98–103. DATE, March 2009.

[96] KT Zorbas, E Hatzikraniotis, and KM Paraskevopoulos. Power and efficiency calculation in commercial TEG and application in wasted heat recovery in automobile. In *Proc. of 5th European Conference on Thermoelectrics*, page T8, 2007.