ScuDo
Scuola di Dottorato ~ Doctoral School
WHAT YOU ARE, TAKES YOU FAR

Doctoral Dissertation
Doctoral Program in Electronics Engineering (29^{th}cycle)

Reconfigurable Antenna Systems:
Platform implementation and low-power matters

By

Simone Ciccia

******

Supervisor(s):
Prof. Giuseppe Vecchi

Doctoral Examination Committee:
Prof. Andrea Massa, Università di Trento, Trento, Italy
Prof. Paolo Rocca, Università di Trento, Trento, Italy
Prof. Stefano Maci, Università di Siena, Siena, Italy
Dr. Riccardo Maggiora, Politecnico di Torino, Torino, Italy
Dr. Vittorio Rampa, Consorzio Nazionale delle Ricerche, Milano, Italy

Politecnico di Torino
2017

Declaration

I hereby declare that, the contents and organization of this dissertation constitute my own original work and does not compromise in any way the rights of third parties, including those relating to the security of personal data.

Simone Ciccia
2017

* This dissertation is presented in partial fulfillment of the requirements for **Ph.D. degree** in the Graduate School of Politecnico di Torino (ScuDo).

I would like to dedicate this thesis to my loving parents and my little Bibi, Giorgia

Acknowledgements

Thanksgiving and gratitude to Prof. Giuseppe Vecchi who has believed in my potential offering me the possibility of conducting the PhD under his supervision. He not only guided me in research activity, but was also able to offer his theoretical and methodological background that made me grow, personally and professionally.

Furthermore, I really appreciate the Istituto Superiore Mario Boella that guested me. In particular, I would like to thanks Dr. Giorgio Giordanengo with whom I worked closely and all those who work in the Advanced Computing and Electromagnetics area with which I collaborated during the Ph.D. period. This experience has given optimum scientific results, publications and a patent, as well as a job contract to continue the research with them.

I also acknowledge the positive comments and suggestions of two reviewers Prof. Andrea Massa and Prof. Paolo Rocca of Università di Trento that have helped improve the thesis.

Abstract

Antennas are a necessary and often critical component of all wireless systems, of which they share the ever-increasing complexity and the challenges of present and emerging trends. 5G, massive low-orbit satellite architectures (e.g. OneWeb), industry 4.0, Internet of Things (IoT), satcom on-the-move, Advanced Driver Assistance Systems (ADAS) and Autonomous Vehicles, all call for highly flexible systems, and antenna reconfigurability is an enabling part of these advances. The terminal segment is particularly crucial in this sense, encompassing both very compact antennas or low-profile antennas, all with various adaptability/reconfigurability requirements.

This thesis work has dealt with hardware implementation issues of Radio Frequency (RF) antenna reconfigurability, and in particular with low-power General Purpose Platforms (GPP); the work has encompassed Software Defined Radio (SDR) implementation, as well as embedded low-power platforms (in particular on STM32 Nucleo family of micro-controller). The hardware-software platform work has been complemented with design and fabrication of reconfigurable antennas in standard technology, and the resulting systems tested. The selected antenna technology was antenna array with continuously steerable beam, controlled by voltage-driven phase shifting circuits. Applications included notably Wireless Sensor Network (WSN) deployed in the Italian scientific mission in Antarctica, in a traffic-monitoring case study (EU H2020 project), and into an innovative Global Navigation Satellite Systems (GNSS) antenna concept (patent application submitted).

The SDR implementation focused on a low-cost and low-power Software-defined radio open-source platform with IEEE 802.11 a/g/p wireless communication capability. In a second embodiment, the flexibility of the SDR paradigm has been traded off to avoid the power consumption associated to the relevant operating system.

Application field of reconfigurable antenna is, however, not limited to a better management of the energy consumption. The analysis has also been extended to satel-

lites positioning application. A novel beamforming method has presented demonstrating improvements in the quality of signals received from satellites. Regarding those who deal with positioning algorithms, this advancement help improving precision on the estimated position.

Contents

List of Figures **ix**

List of Acronyms **xi**

1 Reconfigurable Antennas: Software-defined Radio Implementation and Integration on System-on-chip Based Platforms **1**

1.1 Introduction to Reconfigurable Antennas 1

1.1.1 Applications of Reconfigurable Antennas 2

1.1.2 Adaptive Reconfigurable Antennas 3

1.2 Software-defined Radio; A low-power approach 4

1.3 Software-defined Reconfigurable Antenna 5

1.4 Integration of Reconfigurable Antennas in Ultra Low-power System-on-chip Based Platforms 6

2 Involved Projects **8**

2.1 The National Program of Research in Antarctica 8

2.2 LOw Power Heterogeneous Architecture for NExt Generation of SmaRt Infrastructure and Platforms in Industrial and Societal Applications 10

3 Introduction to GNSS Reconfigurable Antenna System **12**

3.1 State-of-the-art 13

3.2 Fundamentals 13
3.3 General Description of the Invention (Patent Pending) 16
3.4 A First Actuation Form 18
3.4.1 The Designed antenna v1.0 18
3.4.2 Control Criteria and Simulation Results 23
3.4.3 Case of study 1 - Single satellite tracking 24
3.4.4 Case of study 2 - Multiple satellite tracking 25
3.5 An Advanced Actuation Form 26
3.5.1 The Designed Antenna v2.0 26
3.5.2 Advanced Control Criteria 27
3.5.3 Simulation Results 28
3.6 Conclusions 32

4 Conclusions **33**

References **34**

Appendix A (Publications) **39**

List of Figures

1.1 RHCP and/or LHCP patch antenna 2

1.2 Frequency agile antenna 3

3.1 Schema of the proposed GNSS system 16

3.2 A detailed schema of the GNSS array antenna 17

3.3 Detailed schema of the control unit 17

3.4 TOKODAK antenna, a commercially available circular polarized microstrip patch 19

3.5 Simulation of the GNSS array antenna 20

3.6 An actuation form of the GNSS beamformer 20

3.7 A possible realization of a phase shifter 21

3.8 A possible realization of the GNSS beamformer 22

3.9 A possible realization of the phase shifter for the proposed beamformer 23

3.10 Estimated satellites orbits 24

3.11 Single satellite tracking 25

3.12 Multiple satellite tracking - Mean criteria 26

3.13 General purpose beamforming network 27

3.14 Satellite tracking 8:15 a.m. 29

3.15 Satellite tracking 8:45 a.m. 29

3.16 Satellite tracking 9:15 a.m. 30

3.17 Satellite tracking 9:45 a.m. 30
3.18 Satellite tracking 10:15 a.m. 31
3.19 Satellite tracking 10:45 a.m. 31

List of Acronyms

Acronyms / Abbreviations

ACE Advanced Computing and Electromagnetics

ADAS Advanced Driver Assistance Systems

ADC Analog to Digital Converters

ARM Advanced Reduced instruction set computing Machine

ASIC Application Specific Integrated Circuits

BeiDou Chinese satellite navigation system

BFN Beam Forming Network

C/N0 Carrier to Noise Ratio

CPU Central Processing Unit

DAC Digital to Analog Converter

DSP Digital Signal Processors

FE Front End

FPGA Field Programmable Gate Array

Galileo European global satellite-based navigation system

GLONASS GLObal NAvigation Satellite System

GNSS Global Navigation Satellite Systems

GPP General Purpose Platforms

GPS Global Positioning System

HPBW Half Power Beam Width

I/Q In-phase and Quadrature components

IEEE Institute of Electrical and Electronic Engineers

IoT Internet of Things

IRNSS Indian Regional Navigational Satellite System

ISM Industrial, Scientific and Medical

ISMB Istituto Superiore Mario Boella

LHCP Left Hand Circular Polarization

LOS Line Of Sight

LUT Look Up Tables

OPERA LOw Power Heterogeneous Architecture for NExt Generation of SmaRt Infrastructure and Platforms in Industrial and Societal Applications

OS Operative System

PC Personal Computer

PoliTO Politecnico di Torino

PSO Particle Swarm Optimizer

PWM Pulse Width Modulator

RF Radio Frequency

RHCP Right Hand Circular Polarization

RSSI Receive Signal Strength Indication

SDR Software Defined Radio

SIMD Single Instruction Multiple Data

SoC System on Chip

SSID Service Set IDentifier

UHF Ultra-High Frequency

ULP Ultra Low-power

URSS Union of Soviet Socialist Republics

WSN Wireless Sensor Network

Chapter 1

Reconfigurable Antennas: Software-defined Radio Implementation and Integration on System-on-chip Based Platforms

1.1 Introduction to Reconfigurable Antennas

According to the Institute of Electrical and Electronic Engineers (IEEE) definition, antennas are the transmitter or receiver part of a radio-apparatus able to radiate or capture electromagnetic waves [1]. State-of-the-art antenna design presents fixed characteristics like bandwidth, radiation pattern, polarization. The ability to reconfigure an antenna attributes dynamicity to these properties. Reconfiguration is achieved by perturbing mechanically or electrically the antenna structure, for this reason the antenna system is termed ”reconfigurable”.

An example of reconfigurable antenna is depicted in Fig. 1.1. In this context the desired polarization of the field between Right Hand Circular Polarization (RHCP) and Left Hand Circular Polarization (LHCP) can be obtained by switching the excitation signal into the intended port.

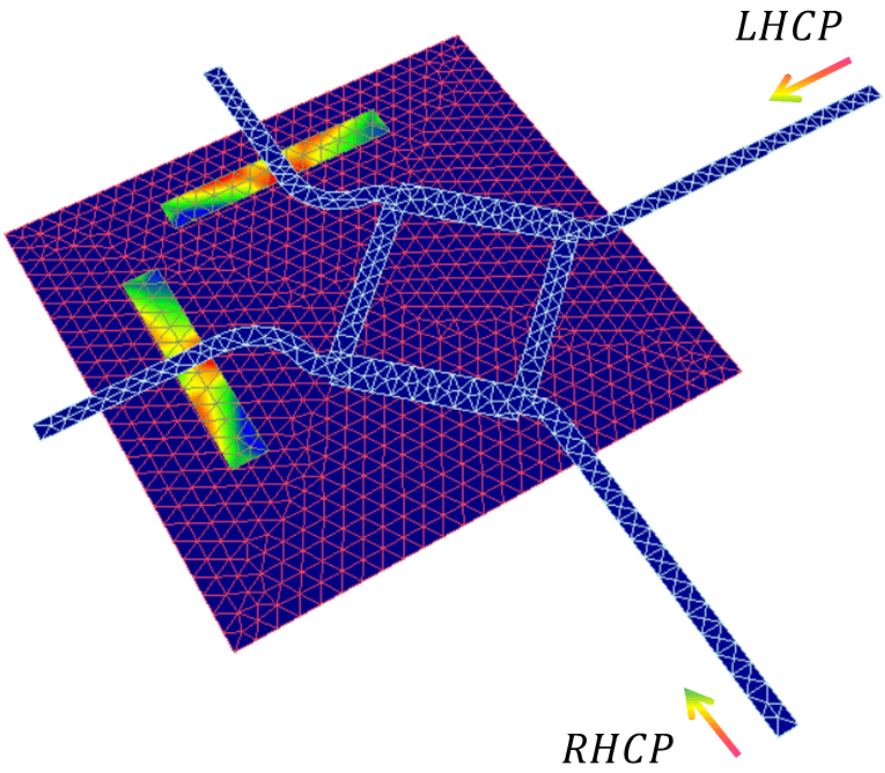

Fig. 1.1 A designed and tested circular polarized aperture coupled microstrip antenna

1.1.1 Applications of Reconfigurable Antennas

Nowadays, communication systems demand extreme requirements in terms of high throughput, energy efficiency, reliability, low latency, and low-cost devices. Reconfigurable antennas play a crucial role in this context since they are the last frontiers of wireless technologies that can meet these stringent requirements [2]. They are widely used in cognitive radio systems which demand for a single antenna to be configured on different operative bandwidths (one at a time) in order to efficiently and opportunistically exploit the available spectrum (dynamic spectrum management) [3–6]. A good example of frequency reconfigurable antenna can be found in Fig. 1.2 where the operative bandwidth is changed by acting on PIN switches [7].

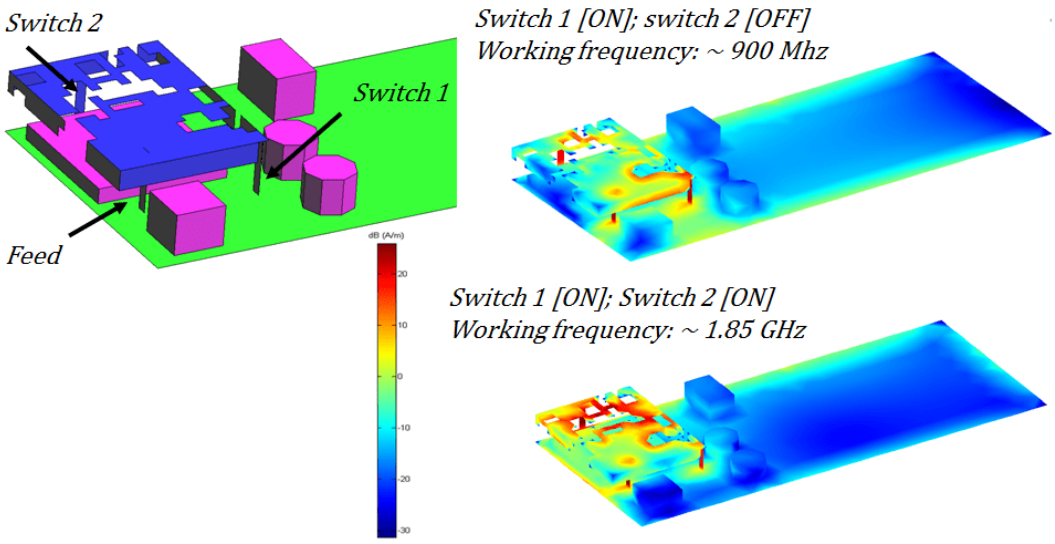

Fig. 1.2 Frequency reconfigurable antenna by means of PIN diode (a Courtesy of Prof. G. Vecchi)

Another important application in which reconfigurable antennas find large employment is in the coverage of cellular radio system. Basically, they are used in each base station to better sectorize the cell radio in order to provide high capacity and high interferer rejection [8–10].

Among all applications of reconfigurable antennas the thesis extensively cover WSN case and focuses the attention on radiation pattern reconfigurability most concerning the minimization of radiation emitted in unwanted directions in order to provide energy efficiency at the node side that usually has critical-power constraints.

1.1.2 Adaptive Reconfigurable Antennas

The method by which a set of antennas (i.e. array) forms the radiation pattern is named beamforming. This technique has appeared in the technological world since 1905, thanks to the own pioneer Karl F. Braun who elaborated an array antenna able to spread radiation in one direction [11]. The idea of having an antenna capable of adapting its radiation diagram based on an optimization of the received signal is a form of adaptive antenna beamforming. With the advent of processors, beamforming has become smart thanks to the ability to associate it with a control/optimization algorithm [12]. SDR plays a crucial role in this scenario as it offer a high degree of flexibility on the control of the radio parameters, such as the power of the received

signal. The first goal was achieving the communication standard and antenna control algorithm on the same GPP (i.e. fully implemented in software).

1.2 Software-defined Radio; A low-power approach

The GPPs, generally defined as embedded personal computers, are an excellent candidate for SDR technology because they provide a great programming flexibility to dynamically deploy computational complexity to the resources available and most conform to signal processing. Furthermore, they allow to fully control the radio parameters in each section of the data processing of a digital transmitter and/or receiver. These factors are of particular relevance when experimenting algorithms to best configure an antenna. However, these benefits come at the expense of energy consumption, that is, unfortunately, higher than System on Chip (SoC)-based radios. For this reason, it is necessary to minimize the computational complexity of the SDR requirement in order to reduce the energy consumption. To achieve this goal, we started from a open-source code of IEEE 802.11 a/g/p wireless standard for GNURadio, a free and open-source framework for signal processing. The nature of this code, which is for scientific-proof, is unsuitable to operate on low-power GPPs. In fact, the code fails to encode and/or decode signals on low-power boards, expecially in the case of Advanced Reduced instruction set computing Machine (ARM) processor. Practically speaking, due to the high workload demand in a limited computational power board, processors saturate and frequent loss of data happens. According to this critical issue, the code has been restructured as described in the Journal publication entitled "WLAN 802.11a/g/p open-source software defined radio on low-power and low-cost platforms", that can be found in appendix, in order to provide a lighter and more efficient version suitable for low-power boards [13]. It also illustrates a review of digital signal processing techniques to significantly improve the performance of the software and minimize the power to run. The restructured software, which we leaved as an open-source, is optimized for low-cost and low-power ARM GPPs as well as for the Intel counterpart.

1.3 Software-defined Reconfigurable Antenna

The antenna is unable to reconfigure itself, but the processor provides the instructions for the optimal spatial selectivity of the radiation pattern. This leads to significant positive effects on communication systems. Some good illustrations of this are:

- increased covered range
- decreased power required to transmit the information
- interferers attenuation (at the receiver) and therefore the link can exploit maximum throughput

In order to test the potential effects of reconfigurable antenna radio system, the software radio (i.e. coding/decoding functions) are integrated along with the software algorithm that controls the antenna on a GPP architecture.

A practical approach is demonstrated in the conference paper entitled "Reconfigurable Antenna System for Wireless Applications" [14]. This work has developed a fully software defined reconfigurable antenna radio system able to control the radiation pattern of the antenna based on the received signal power. It is important to mention that the novel aspect is related to the software that control the beamformer. The system shows low-power properties as the RF beamformer needs a single radio Front End (FE) to operate. Moreover, this solution is low-cost as it is characterized by only software components that can be used on any GPP. The realized system has been applied in an autonomous WSN in order to validate the advantages of using reconfigurable antennas.

Generally, a WSN is a set of sensors distributed in the environment to be monitored. The node are termed "autonomous" in the sense that, once placed in the field they send the information of interest directly to the base station (e.g. a data elaboration center) without requiring the user intervention for many years, until the battery goes down. This discussion highlights the importance of battery life, which should be as long as possible. This leads to a desirable decrease in the energy consumption dedicated to a wireless node. The publication entitled "Software-defined Reconfigurable Antenna for Energy Efficient Wireless Links" shows how the ability to reconfigure the antenna radiation can reduce the energy consumption in a point-to-point communication, between sensor node and base station [15]. The manuscript

reports a comparative analysis of the daily consumption between a wireless node configured with standard antenna and another configured with reconfigurable antenna. With the same amount of information transmitted daily it turns out that the reconfigurable antenna system, which is able of to direct the radiated energy in the direction of the base station and vice versa, allows to consume less energy. The battery life of the wireless devices is prolonged thanks to the following advantages:

- At the same distance, reconfigurable antenna system achieves higher rates, then data transfer is faster (the device is in transmission state for less time)
- Power consumption of the wireless device is proportional to the transmitted power, so the higher the antenna directivity the lesser the required transmit power. As a consequence, the device consumes less power when it is in transmit mode

Another publication entitled "Low Power Advanced Wireless Communication Exploiting Reconfigurable Antennas" introduces a practical point of view [16]. It provides the implementation of a programmable power supply controller able of turning on and off the radio components in order to be more energy efficient. When the radio is turned on, the software configures the antenna main beam towards the desired direction. The system is also able of turning on the radio when it is necessary to transmit the information and then turn it off during inactivity periods.

In conclusion, a significant reduction of consumption can be obtained through the possibility of reconfiguring the antenna.

1.4 Integration of Reconfigurable Antennas in Ultra Low-power System-on-chip Based Platforms

This section deals with energy consumption reduction in WSN nodes by means of antenna pattern reconfiguration at RF and its consistent implementation on low-power hardware platforms.

In the Journal publication entitled "Integration of Reconfigurable Antennas in Ultra Low-power Radio Platforms Based on System-on-chip for Wireless Sensor Network", we employed an antenna array with steerable beam, controlled by an

algorithm embedded in the firmware of the selected Ultra Low-power (ULP) hardware platform (i.e. a STM32 Nucleo micro-controller), eliminating the consumption overhead required to run an Operative System (OS) [17]. The developed software maximizes the received power by adaptively directing the beam in the direction of the line-of-sight link. The work reports both estimates of energy savings and experimental results to prove the effectiveness of the proposed approach. In particular, we explicitly considered the consumption overhead involved for the steering process in the global energy efficiency of the proposed system in order to validate a true energy gain, especially in a dynamic scenario where the re-location of nodes happens frequently. Such analysis still lacks in literature and is addressed in this paper for the first time. Furthermore, the paper shows the feasibility and the possibility of saving energy with a commercial ULP wireless platform, showing a more application-oriented view.

Chapter 2

Involved Projects

This chapter discusses the research activity over two project which I involved, the "National Program of Research in Antarctica - PNRA", whose purpose is to improve the observation system in the Italian base, and the project "LOw Power Heterogeneous Architecture for NExt Generation Of Smart Infrastructure and Platforms in Industrial and Societal Applications - OPERA", a European project with a strong industrial impact on the new generation of ULP computing systems.

2.1 The National Program of Research in Antarctica

Installing data acquisition system in remote areas that are distant from the data elaboration center provides the use case for testing low-power communication and computing systems. Antarctica region falls into one of these area. In this strategic region a vast community of researches (which amounts in 28 different countries) have installed their own permanent base station to analyze various natural phenomena. In this critical context the installation of cables to establish communication between a sensor and base station is difficult to achieve and can not be maintained. Temperatures in this area are unstable and in some months they reach temperature as low as $-70°$. During these periods human intervention on data acquisition systems is impossible and two serious problem arises:

- The stiffness of cable installations are under severe risk of break-down due to wind perturbations

- Battery life is of vital importance in order to ensure the operativity of the observatory system during these critical periods

It become clear that an energy efficient wireless link needs to be achieved in order to overcomes these limitations. For this reason, the Advanced Computing and Electromagnetics team of Istituto Superiore Mario Boella (ISMB) has developed a fully autonomous sensor powered by renewable resources. The prolonged battery life, also supported from the energy saving provided by means of the reconfigurable antenna radio system, guarantees the operativity of the link during the most critical environmental conditions. The prototype of the data acquisition system is discussed in the work entitled "Low Power Computing and Communication System for Critical Environments" [18]. The paper includes a detailed description of the following designed components:

- the battery pack
- renewable energy components (solar panel and wind turbines)
- energy supply controller, which manages the activities of each components (duty cycle)
- the cluster computing employed for local analysis of data (raw data pre-processing)
- the smart wireless communication realized with reconfigurable antenna technology in order to minimize the energy consumption

and describes each part of the acquisition system providing also an accurate analysis of the Antarctica climate all over the year, a valuable information which explains the implementation choices especially to size the various modules of the apparatus. It also shows the tests carried out in the climatic chamber to emulate, as close as possible, the critical environment conditions of the Antarctica. This system guarantees the operativity of the observatory system also in periods when the solar radiation is practically null.

The publication entitled "GreenLab: autonomous low power system extending multi-constellation GNSS acquisition in Antarctica" describes in details the phenomenon of Ionospheric scintillation, the reason for which the measuring system has

been designed [19]. This natural phenomena is a noise source that afflicts satellite communications all over the world. In fact, in satellite positioning systems, it may lead to an incorrect estimation of the position. This publication is intended as an introduction to the measurement system that observes, analyzes, locally processes and sent to the base station the data related to ionospheric scintillation. Dr. G. Giordanengo, who had the honor to place his feet on the Antarctica territory, in particular in the Italian permanent observatory entitled Mario Zucchelli station, it had the opportunity to install and test the realized data acquisition system. The paper entitled "Energy Efficient System for Critical Environment Observation" shows the effectiveness and robustness of the proposed solution in providing a real-time acquisition and processing of data along the entire Antarctica winter, period where the state-of-the-art system failed for the critical condition of energy starving and freezing electronics. Of particular relevance, this publication shows the real experimental results obtained by testing the wireless communication, which extends, by means of reconfigurable antenna, the IEEE 802.11 b/g/p SDR in a 500m link with excellent results (i.e. negligible communication losses) [20]. The wireless communication is also optimized in term of consumption thanks to the power management module which activates the radio with a well-defined scheduled profile.

2.2 LOw Power Heterogeneous Architecture for NExt Generation of SmaRt Infrastructure and Platforms in Industrial and Societal Applications

LOw Power Heterogeneous Architecture for NExt Generation of SmaRt Infrastructure and Platforms in Industrial and Societal Applications (OPERA) is a Europen funded project, whose objective is to exploit heterogeneous computing technologies over the whole computing continuum (i.e., spanning from data center servers to the edge of the Cloud), aiming at improving the energy efficiency and performance. The solution provided by OPERA is intended to be representative of the architecture of the next generation data centers. Large effort has been devoted to the seamless integration of ULP devices within the Cloud infrastructure: through the creation of dedicated services, which can fully exploit the benefits of a large and heterogeneous computing infrastructure, it becomes possible to remotely accelerate computing

intensive tasks. Moreover, big attention has been reserved to improve the RF communication sub-system adopted by such ULP devices, with the aim of globally improve the energy efficiency. The book chapter entitled "The Green Computing Continuum : the OPERA Perspective", provides an insight of the set of low power solutions devised in the context of the OPERA project [21]. In this context we addressed the traffic monitoring use case. An ULP computing system equipped with an embedded surveillance camera acquires and locally elaborates images to detect traffic events (e.g. traffic congestion and wrong way vehicles detection to name a few) and send alarms to the elaboration center via wireless communication. For this purpose we proposed the system described in (Sec. 1.4) since it is an optimum candidate with excellent wireless communication capabilities at a very low-power consumption.

Chapter 3

Introduction to GNSS Reconfigurable Antenna System

In 1945 the Wireless World Magazine published Arthur C. Clarke's letter which theorized the worldwide radio coverage with rocket stations[22]. In the late 1957 the first artificial Earth satellite, called SPUTNIK, was launched by the ex - Union of Soviet Socialist Republics (URSS) for space exploration [23, 24]. Only in 1960 National Aeronautics and Space Administration launched ECHO1, the first artificial satellite in orbit dedicated to communications, proving that worldwide radio coverage was possible[25]. This experiment has captured the attention of industries and governments all over the world which stimulated a large variety of studies and researches in military and civil application. In 1973 the Department of Defense of United State established the NAVSTAR-GPS program that lead to today's Global Positioning System (GPS). Concurrently the ex Soviet Union approved GLObal NAvigation Satellite System (GLONASS) [26, 27]. In recent days every international organization has approved their own satellite navigation program, European global satellite-based navigation system (Galileo) in European Union, Chinese satellite navigation system (BeiDou) in China, Indian Regional Navigational Satellite System (IRNSS) in India and others. All these names constitute the GNSS.

3.1 State-of-the-art

In GNSS array antennas can improve Carrier to Noise Ratio (C/N0) by directing their beam toward satellites and reduce interference (e.g. placing nulls in the pattern). In order to be effective they require the flexibility of both phase and amplitude excitation which are processed digitally by means of dedicated powerful Digital Signal Processors (DSP) components. This technique is referred as Digital Beamforming [28–37]. Although very effective, such solution implies energy inefficiency and high fabrication cost which therefore, from the outset, it is unfeasible for the most general GNSS applications.

The present work introduces a newly approach to control the radiation pattern of an antenna array. The proposed system shows energy efficiency since able to increase the C/N0 by means of single FE for the received signal and a minimal software-defined control metric that can operate in a control unit, like a low-power GPP. This novel idea proposes chance to overcome the limitation poses by state-of-the-art system (like costs, energy consumption etc.) which uses multiple RF front-ends and dedicated computational hardware.

3.2 Fundamentals

It is well known that the availability of low-cost satellite receivers (i.e. positioning devices) have guaranteed the development of commercial services and products, giving rise to localization of persons and vehicles all over the world. However, these low-cost devices can not be employed in applications requiring a precision lower than 1m on the estimated position since they receive a weak C/N0 level which is not suitable to estimate with accuracy the travel time a signal takes from the satellite to the receiver. In fact, the positioning device is configured to determine its position (latitude, longitude, and height) by means of solving the following system

of equations:

$$\begin{cases} (X_1 - X_u)^2 + (Y_1 - Y_u)^2 + (Z_1 - Z_u)^2 = (R_1 - C_u)^2 \\ (X_2 - X_u)^2 + (Y_2 - Y_u)^2 + (Z_2 - Z_u)^2 = (R_2 - C_u)^2 \\ (X_3 - X_u)^2 + (Y_3 - Y_u)^2 + (Z_3 - Z_u)^2 = (R_3 - C_u)^2 \\ (X_4 - X_u)^2 + (Y_4 - Y_u)^2 + (Z_4 - Z_u)^2 = (R_4 - C_u)^2 \end{cases} \tag{3.1}$$

where Xu, Yu, Zu e Cu are the unknown terms that should be determined, latitude, longitude, height and clock deviation respectively. Clock deviation is defined as the mismatch between the receiver's clock and the satellites' clocks. These latter are all synchronous. The terms $X1$,$Y1$,$Z1$;...;$X4$,$Y4$,$Z4$ are the positions of the four satellites involved to determine the position of the GPS device. These variables are derived from the information contained in each signals (i.e. the ephemeris each satellites transmits at a precise time). The terms $R1$,...,$R4$ are determined in the following way:

$$R_i = c\, T_i\, [m] \tag{3.2}$$

where c is the speed of light and Ti corresponds to the travel time of the signal emitted from the satellite i. As explained above, the travel time Ti is the arithmetic difference between the instant a positioning device received the signal and the instant the satellite transmitted that signal. This information is encapsulated in the signal. It is deduced form this brief discussion how the C/N0 affects the precision with which the positioning device is able to determine the instant of time the satellite signal is received. In fact, the greater the C/N0 the greater the precision the device detected such instant of time. A possible technique to improve the C/N0 consists in the use of two GPS devices placed at a known position. In addition, a computational device (like a PC or smartphone) collects the information from both devices in order to estimates the position and, in particular, it is able to minimize the noise by taking account of the formed triangulation between GPS devices and satellites. However, this solution is not feasible where encumbrance and consumption are relevant factors. In fact, this way doubles energy consumption and encumbrance since the system is redundant. Noise sources can be categorized in:

- Multipath effect (reflections)
- Interference sources (jamming)

- Spoofing, that is a fake signal originated to deceive the estimated position

As described by the American patent demand [38], noise sources can be mitigated by means of multiple receivers spatially distributed and controlled through a beamforming network that is in charge of reducing the noise sources in a particular direction. The beamformer carry out this task by means of obscuring the set of receivers introducing the maximum contribution of noise in the positioning device. However, this solution fails to increment the C/N0 since, in a particular situation, it prevents the reception of signals to the positioning device. Furthermore, the issued American patent [39] describes a solution to reduce the effect of interference sources (i.e. jamming) by means of a beam forming network which operates at the downstream of the frequency down-converters, each one placed behind an antenna. This solution needs a modification of the receiving system, thus can not be employed in association with existing receivers. Moreover, from the point of view of fidelity the publication presents some critical points. Indeed, the presence of multiple down-converters decrease the mean time to failure. Furthermore, the mentioned solution estimates a position that is always affected by a minimum error that can not be reduced. This is due to the nature of down-converter circuitry which introduce noise floor and an unpredictable (i.e. random) phase delay. As a matter of fact, the presence of oscillatory circuits are inevitably affected by thermal phenomena. The proposed system overcomes these and other issues providing an apparatus along with a method for the reception of satellite positioning signals. The idea of the present invention suggests the use of multiple antennas for the reception of satellite signals, where each antenna is directly connected to phase shifters and these latter are joined together with an output collector which generates:

- Constructive interference among signals incoming from satellites
- Destructive interference among the reflections of the above mentioned signals and other source like jamming and spoofing

In this way, the C/N0 is effectively incremented and can be acquired by the positioning device that is directly connected to the output collector. In fact, this device acquires a signal with minimized multipath effect since it creates a constructive interference only among signals received from satellites and not on their reflections that interfere in a destructive manner.

3.3 General Description of the Invention (Patent Pending)

This section provides a generic description of the invention which will cover all parts of the realized system. Next, some actuation forms will be also provided. It is important to note that the presented form of actuation are for an illustrative purpose and they do not pose any limitation to any part of the system.

Fig.3.1 illustrates the block diagram of a system for the reception of satellite positioning signals according to the invention.

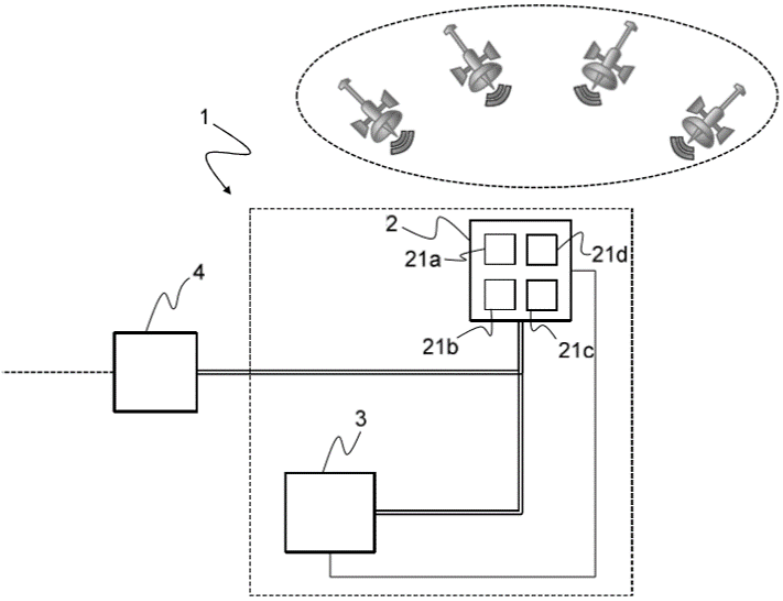

Fig. 3.1 Block diagram of the invention

Fig.3.2 shows the main elements of the array antenna included in the system of Fig.3.1.

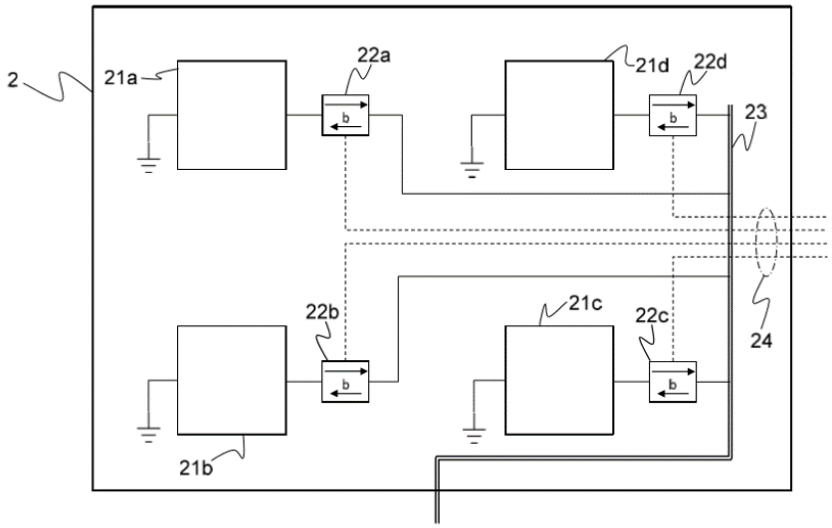

Fig. 3.2 Main Elements of the array antenna included in the GNSS system of Fig.3.1

Fig.3.3 depicts the main elements of a control unit that could be added to the system of Fig.3.1;

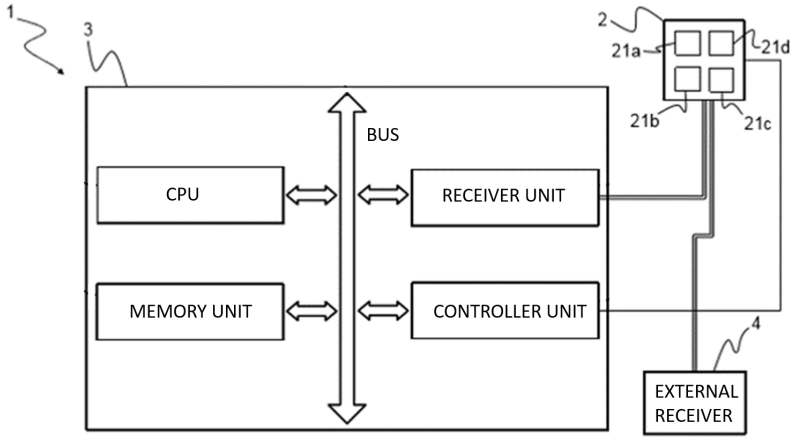

Fig. 3.3 Details of a control unit that could be included in the system of Fig.3.1

The following tips describe the positioning system (1) according to the invention:

- An array antenna (2) composed by a minimum of two radiating elements 21a, 21b, 21c, 21d designed for receiving positioning signals of at least one artificial constellation of satellites (e.g. GPS, GLONASS, Galileo);

- At least two phase shifters 22a, 22b, 22c, 22d, each connected one-by-one directly to an antenna 21a - 21d;
- An output collector (23) preferably a RF waveguide of known impedance that can be easily connected with an external positioning device (4) and where the phase shifters are joined together, best in parallel and/or in series, by means of the output collector (23). That way a constructive interference can be generated among the received satellite positioning signals, while a disruptive interference is generated among the reflections of these signals;
- A control unit (3), preferably a computational unit like a Central Processing Unit (CPU), which is able to execute instructions in order to control the phase shifters 22a - 22d by means of an optimization algorithm finalized to increase the C/N0 of at least one of the received signals through the antennas 21a - 21d. The controller unit is usually a Digital to Analog Converter (DAC) which provides an electrical signal to phase shifters;

3.4 A First Actuation Form

In a first actuation form, the preferred realization of the array antenna (2) includes four radiating elements 21a, 21b, 21c, 21d and an output collector that can be connected with transmission line, like a coaxial cable, to an external GNSS receiver (4) that is know to the state-of-the-art (e.g. UBX-G7020, UBX-G8020 to name a few commercially available GNSS receivers). The aim of this test is the tracking simulation of at least one satellite in order to validate the improvements in C/N0.

3.4.1 The Designed antenna v1.0

For this first validation a one-dimension phased array of four radiating elements is used which can steer the main beam in only one axis. Thus, the best case will be when the satellites fly-over the steering axis, while in the other cases the beam assumes the direction computed as the orthogonal projection of the satellites (i.e. followed from the main beam side lateral), as will be explained later in (Sec. 3.4.2).

Ground multipath is the major source of errors and its rejection have to be accounted in the antenna design. A comparative studies on GPS antenna suitable for

multipath minimization has been conducted in [40]. Since patch antenna presents a radiation pattern with a broad lobe over the hemisphere and a sharp slope for low elevation angle, it has been the right choice for this project[41]. The type used in this study is depicted in Fig.3.4, commercially available in ceramic substrate, resonating at GPS-L1 band (i.e $1575.42 \pm 10MHz$), 50Ω feed, $7dBi$ maximum gain and $70mmx70mm$ encumbrance. The measured pattern and all other details are reported in [42].

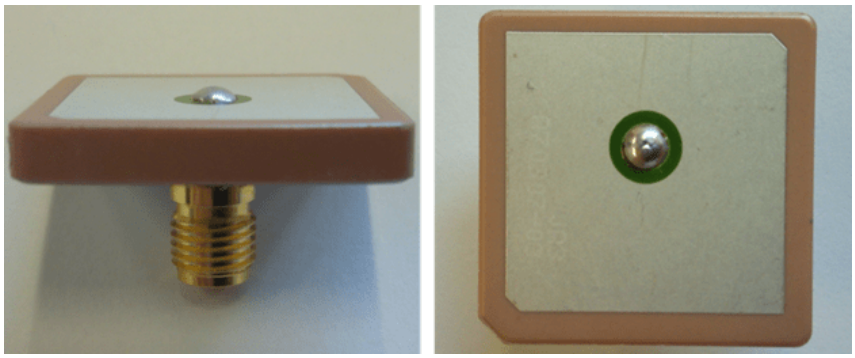

Fig. 3.4 ,
A circular polarized microstrip patch antenna that could be used in array configuration

To ensure that the visible range of the radiation pattern ends at least zero before the adjacent main lobe, the following grating lobe limit formulation has used:

$$d \leqslant \frac{1-1/N}{1+\mid cos(\alpha_{max}) \mid}\lambda_0 \; [m] \tag{3.3}$$

where d is the inter-element distance, N is the number of elements constituting the one-dimensional array, α_{max} is the desired maximum scan angle, and λ_0 is the free space wavelength. For the purpose of our test we choose a steering capability in the range $\pm\angle 45$ with 4 radiating elements, thus (3.3) provides $d = 85mm$. The simulated pattern is reported in Fig.3.5.

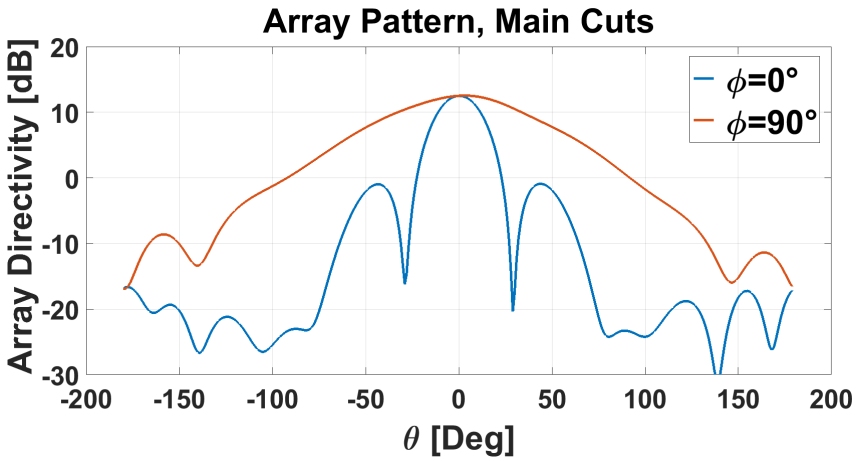

Fig. 3.5 Simulation of the GNSS array antenna - expected pattern in main planes

The Beam Forming Network (BFN), depicted in Fig.3.6, has simulated in IS400 substrate with the following characteristics:

- $\varepsilon_r = 3.9$ dielectric constant.
- $h = 1.55mm$ substrate thickness.
- $Tang\delta = 0.022$ dielectric perturbation/losses .

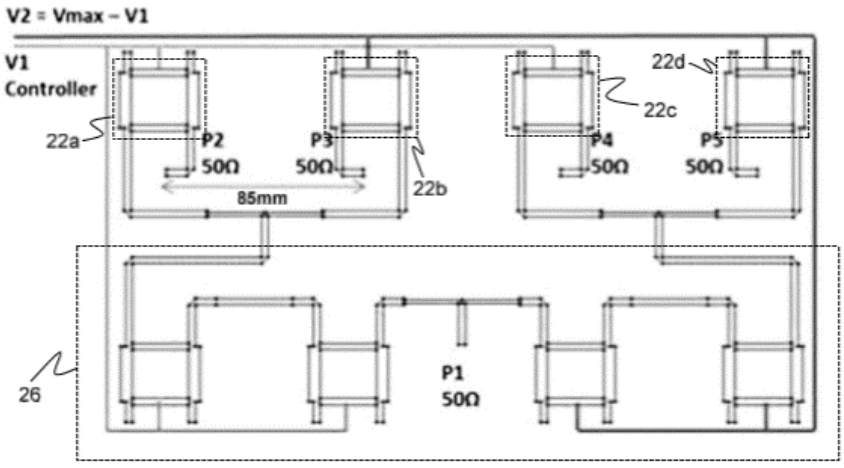

Fig. 3.6 Illustration of a beamformer optimized for steering which can be the part (2) of fig. 3.2

The latter includes phase shifters in order to steer the main beam in a single plane. The preferred form, depicted in Fig.3.7, is a RF 90° hybrid coupler loaded

with two (nearly) identical voltage controlled varactor diodes that allow (ideally) to delay the input signal within the range 0° - 150°. The input/output relationship can be described by the following equation:

$$Io = Ii\, e^{-j\Phi(Vdc)} \tag{3.4}$$

where Ii is the input signal while $\Phi(Vdc)$ maps a linear relationship between the controlled voltage and the provided phase shift [43].

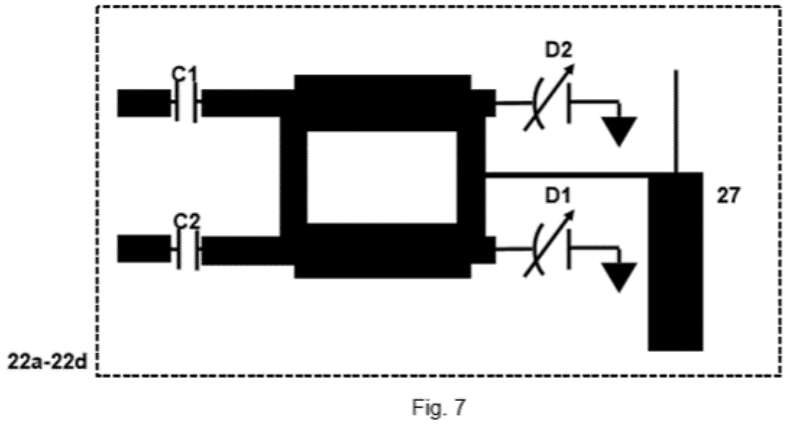

Fig. 7

Fig. 3.7 Detailed circuital schema of 3.6 illustrating a possible physical realization of a phase shifter

Fig.3.9 shows the physical details of a phase shifter. The flag connected to the center of this component represents a RF stop circuit. This filter ensure that the RF signal does not disturb the control signal. It is composed by two quarter wavelengths with a large difference in impedance. Varactor diodes are choosen as equal as possible in order to avoid losses and circuit mismatching. A preferred component for this realization is the Skyworks SMVA1248-079LF diode that works with a maximum control voltage of $5V$ [44].

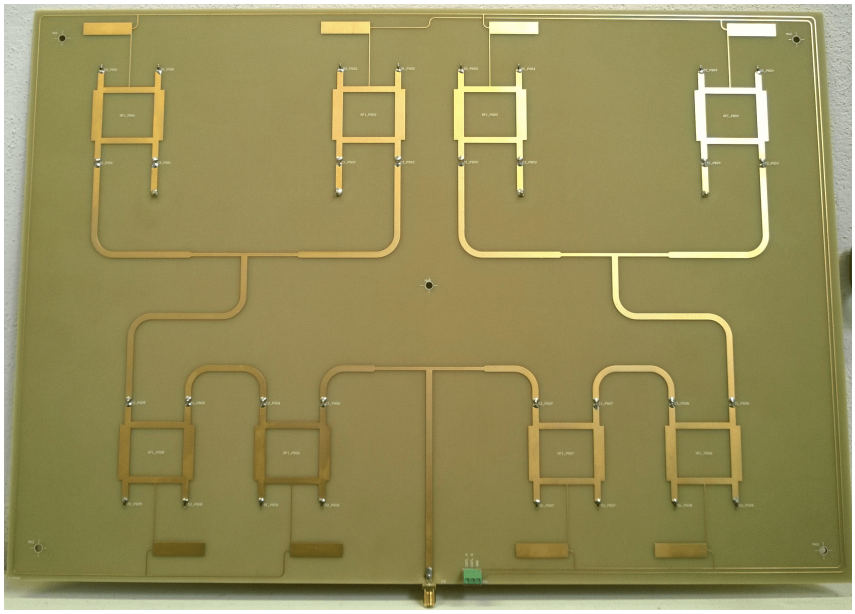

Fig. 3.8 A realized beamforming network in microstrip technology with steering capability

To perform the scan in one axis ($\pm 45°$) with these implementation choices, then 8 phase shifters and two complementary analog voltages are necessary, as depicted in Fig.3.6 which illustrates the physical realization of the proposed BFN. The first couple of phase shifters (labeled 26 in Fig.3.6) provide a controlled phase (ideally) from 0° to 300°. Since the couple are controlled with complementary voltages, then when a brunch provides 300° (controlled with $V1 = 5V$) the other provides 0° (controlled with $V2 = 0V$). Another phase shifter is added to each brunch to provide a linear relationship among radiating elements. When the main beam is in broadside $\theta_{steer} = 0°$, $V1 = V2 = 2.5V$ and all radiating elements assume (ideally) the set of phase (225°,225°,225°,225°); With the example of 3.6, when $V1 = 5V$ and $V2 = 0V$ control signal are provided, then the main beam is directed at $\theta_{steer} = 45°$ and radiating elements 21a - 21d assume the following set of phase (450°, 300°, 150°, 0°); and finally, when when $V1 = 0V$ and $V2 = 5V$ the beam is oriented at $\theta_{steer} = -45°$ and radiating elements assume the set (0°, 150°, 300°, 450°).

Signals received by the four radiating elements are coherent added at RF by means of microstrip lines that joints to a SMA connector which is linked to an external receiver. In addition to the above discussion, the system (1) could also include a programmable control unit (3) that controls the phase shifters 22a - 22d

with at least a control logic, better discussed in next section, able to increase the C/N0 of at least one signal received by means of radiating elements 21a - 21d.

The realized system is actually under test. Following the discussion simulation results will be provided.

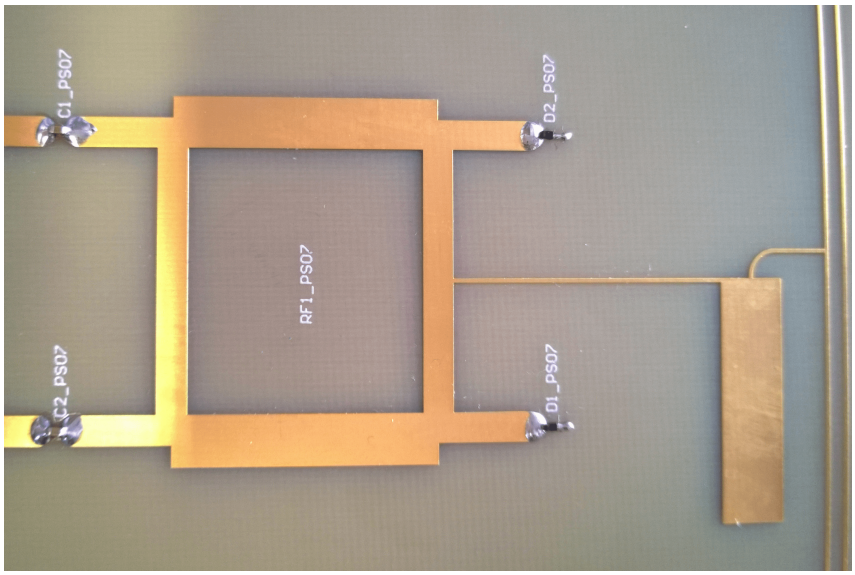

Fig. 3.9 Phase shifter and RF Stop circuit in microstrip technology

3.4.2 Control Criteria and Simulation Results

The following section analyses some optimization algorithms that could be adopted in the control unit (3) of this first actuation form. The performed analysis simulates the realized array antenna placed on the roof of Politecnico di Torino. The broadside direction is toward the Zenith, while beam steering is applied in one axis (i.e. North($\phi = 0°$)/South($\phi = 180°$) directions). The obtained results (with the array) are compared to a single radiating element [42].

Orbitron, a free satellite orbit simulator software, has used in order to predict satellites trajectories in November 2015, from 8:00 a.m. to 12:00 a.m. [45]. Based on this data we performed the satellite tracking simulation by means of configuring the antenna radiation pattern and analyzing the resulted C/N0.

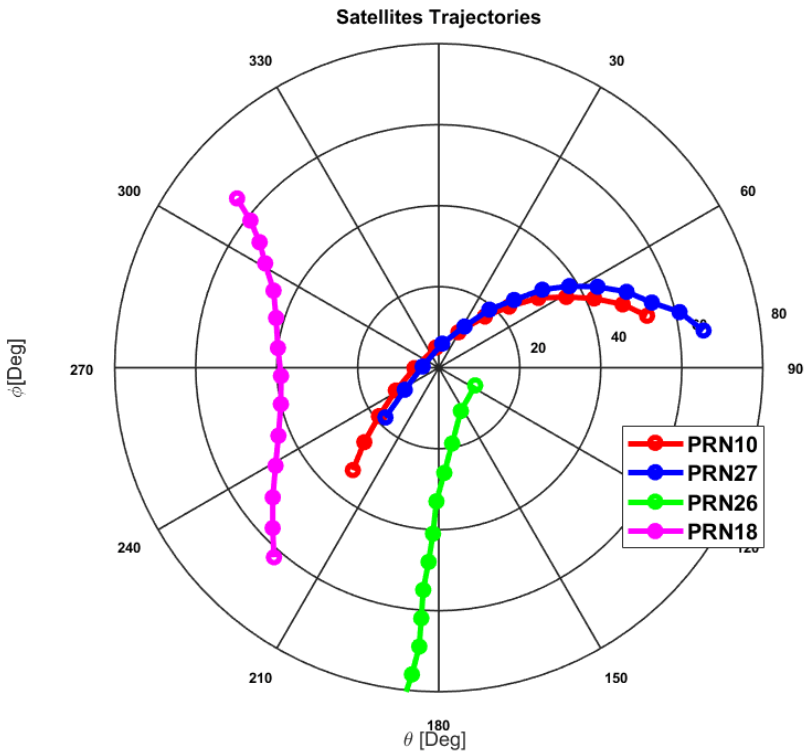

Fig. 3.10 Estimated satellites orbits in November 2015, from 8:00 a.m. to 12:00 a.m.

3.4.3 Case of study 1 - Single satellite tracking

Satellite tracking is provided by means of C/N0 at the receiver, the objective function to maximize. The main beam of the array under test can be steered in only one direction (i.e. North and South in this test) ad establishes the direction of maximum radiation through a polar coordinates projection over the steering axis, so that:

$$\begin{aligned} \theta_{steer}(t) &= \theta_{sat}(t)cos(\phi_{sat}(t))\ [rad] \\ \theta_{sat} &= \frac{pi}{2} - elv_{sat}\ [rad] \\ \phi_{sat} &= azm_{sat}\ [rad] \end{aligned} \tag{3.5}$$

where $\theta_{steer}(t)$ is the angle in which the main beam has to be steered, while elv_{sat} and azm_{sat} are satellite elevation and azimuth (ephemeris) at the time t, respectively.

A Look Up Tables (LUT) provides the conversion of the steering angle with the corresponding controlled voltage to produce the appropriate phase shift.

As illustrated in the graph matrix of Fig.3.11, the proposed tracking mechanism is able to raise the C/N0 level of the intended satellite.

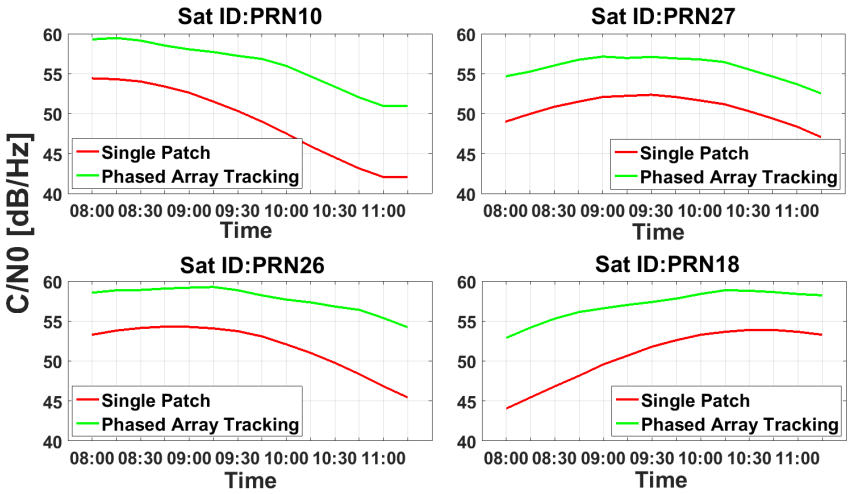

Fig. 3.11 Single satellite tracking with steerable array - simulations

3.4.4 Case of study 2 - Multiple satellite tracking

In this second case, the adopted metric try to optimize globally the mean C/N0 level of all visible satellites. The concept is summarized by the following formula:

$$Opt\theta_{steer}(t) = Argmax_{\theta_{steer}}\{Mean_i\{CN0(i,\theta_{steer},t)\}\}\ [rad]. \tag{3.6}$$

where one have to point the main beam in each possible direction θ_{steer} and evaluate the overall C/N0 of all visible satellites. Then the best direction θ_{steer} that provides the maximum averaged C/N0 is chosen. The analysis reported in Fig.3.12 proposes the simultaneous tracking of four satellites (i.e. PRN10, PRN27, PRN18 and PRN21) and demonstrates improvements in the average C/N0. I found that this method degrades with the number of considered satellites (improvements up to 4 satellites), due to the steering capability that is bounded on one-axis. Furthermore, this methods improves the average C/N0 with respect the use of a single radiating element.

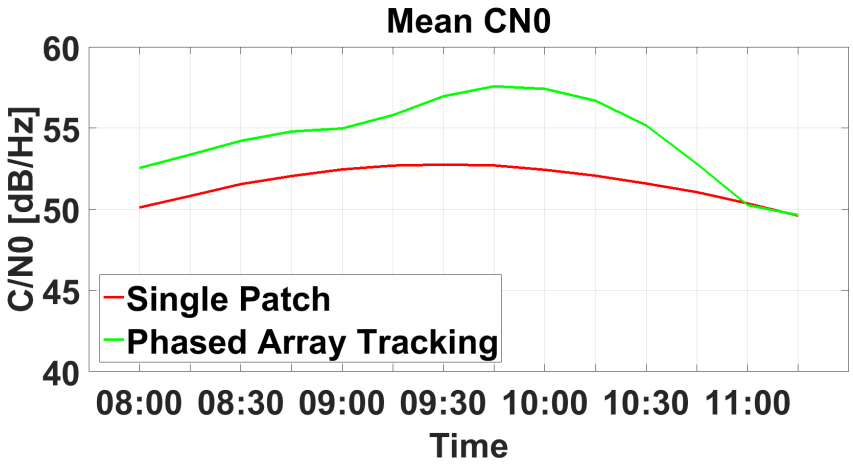

Fig. 3.12 Multiple satellite tracking - Mean criteria

3.5 An Advanced Actuation Form

This advanced actuation form still employ a one-dimension array with radiating elements of the type depicted in Fig.3.4, but with a novel beamforming approach. In fact, this new realization presents independent control channel that allow to assign a phase to each radiating element of the array independently. In this way, a custom multi-beam can be generated to track more than one satellites independently.

3.5.1 The Designed Antenna v2.0

With respect the state-of-the-art introduced in (Sec. 3.1), this work is intended to overcome the limitations posed by actual digital beamforming systems in term of complexity, cost and consumption. In fact, with the presented BFN I would like to demonstrate the possibility of tracking more then one satellites at a time by generating a least two custom beam with minimum complexity. This latter statement is intended as the use of a single front-end to acquire the RF signal and a low-power processing unit is in charge to change the phase of the BFN by means of a LUT. This way the antenna can be associated to any commercial receiver and will provide better performance. The present work does not cover interferer suppression, although with the only phase excitation the maximization of the pattern in the intended directions can, in some conditions, reduce it.

As depicted in Fig.3.13 which shows an example of simulated BFN, each channel is made up by a series of three phase shifters directly connected (i.e. 1:1) to an antenna of the array. The series of phase shifters provides a phase control that spans, ideally, from 0° to 360° for each channel.

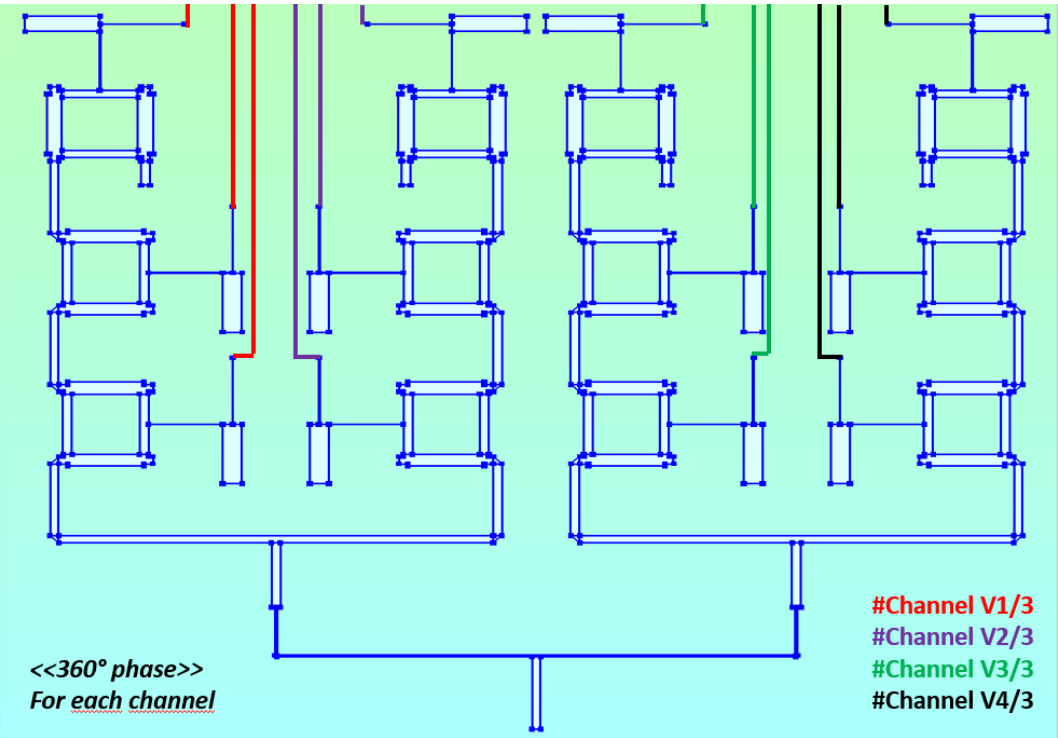

Fig. 3.13 Separate control channels beamforming network, Simulated Geometry

3.5.2 Advanced Control Criteria

In the GNSS array, all elements are fed with a signal of the same amplitude (no amplitude excitation), but different phase. To determine the phases, which, in turn, determine the directions of maximum directivity, we use a simple generic optimization technique known as Particle Swarm [46]. The basic idea of the technique is to numerically simulate the behavior of moving particles to explore the solution space. The position of each particle encodes the values of the optimization variables. The velocity of each particle, hence its motion, has both a random component (to explore the space) and a component linked to the current best of the objective function (to move towards the optimum). In this implementation, the position of a particle are the possible phases of the feeds. The objective function to be maximized is the amplitude

of the field radiated in the desired directions. The radiated field is evaluated applying the array factor formulation to the known field radiated by each element in isolation. To favour speed, no normalization on the directivity is performed, so the evaluation of the objective function is very fast. The problem is anyway bounded due to the constant feeds amplitudes. High side lobes can be present, but in the examined cases the algorithm was able to find solutions within the requested constraints. No additional bounds are placed on the values of the phases, but the values found at the end of the algorithm are then normalized in the interval [0°, 360°). The Particle Swarm Optimizer works with discrete phase values.

3.5.3 Simulation Results

Simulation has been performed via software, analyzing the trajectories of the satellites PRN16 and PRN18 from the 08:00 to 12:00 a.m. The following set of slide shows a time sequence optimization of the radiation pattern according to the satellites movements. The iterative procedure (simulation) is summarized here:

- azimuth and elevation of PRN16 and PRN18 are analyzed at the time t
- by means of (3.5), the direction of the beam $\theta_{steer}(t)$ is computed for both PRN16 and PRN18. These two direction are the inputs for the Particle Swarm Optimizer (PSO).
- The PSO generates the optimum set of phases which maximizes the radiation towards PRN16 and PRN18 simultaneously

For this simulation we employed an array of 8 elements with inter-elements distance $d = 0.5\lambda = 86mm$ for a better spatial selectivity and understanding of the proposed method.

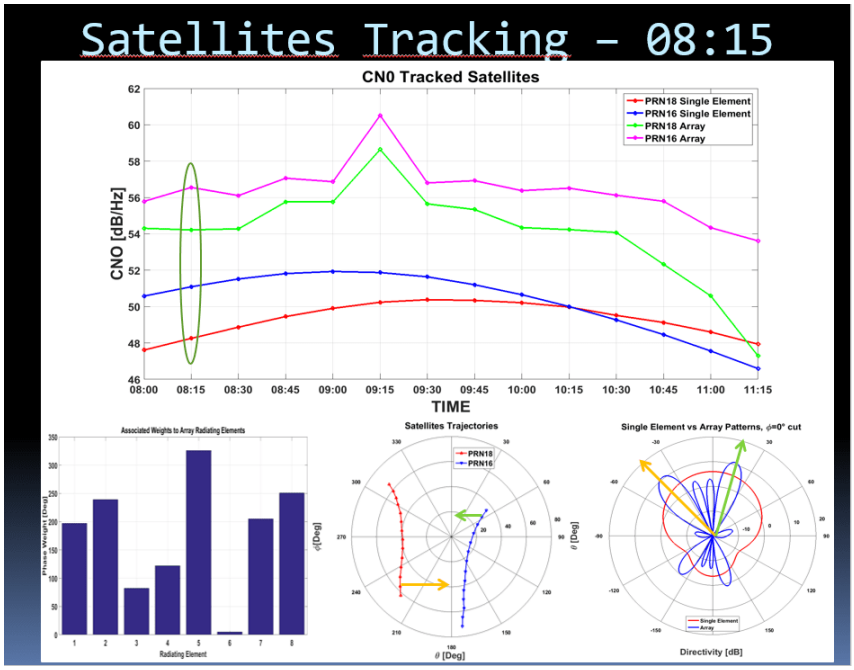

Fig. 3.14 Satellite tracking simulation 8:15 a.m.

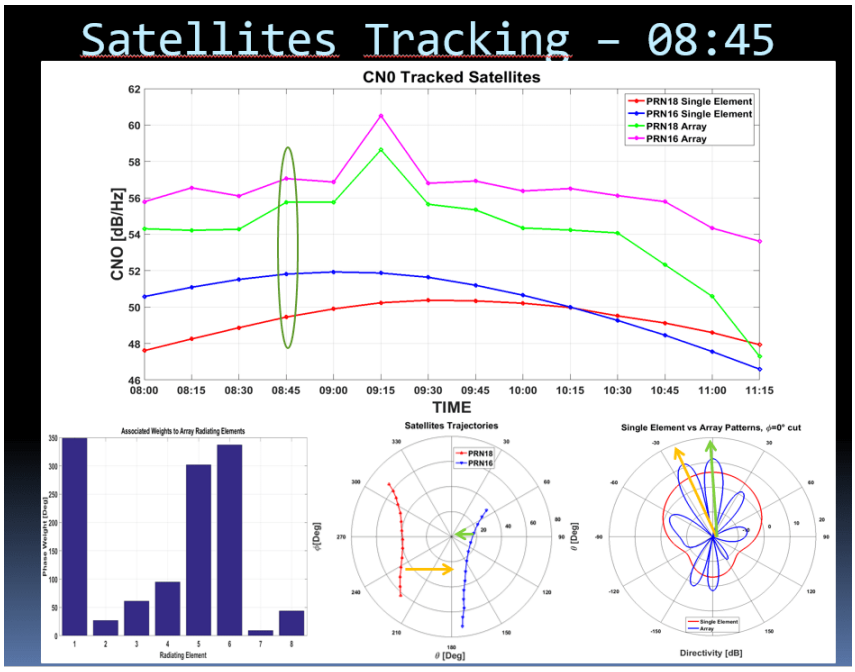

Fig. 3.15 Satellite tracking simulation 8:45 a.m.

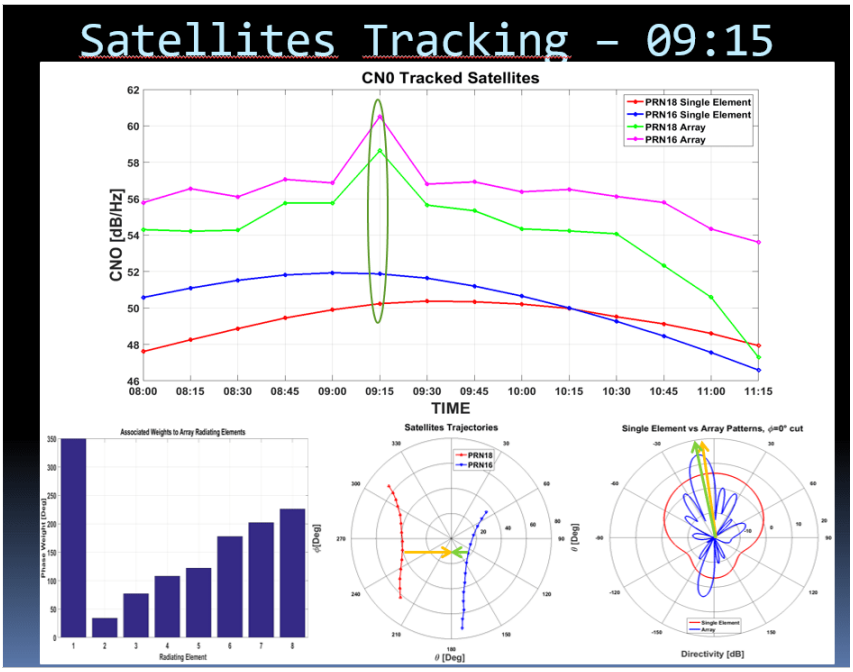

Fig. 3.16 Satellite tracking simulation 9:15 a.m.

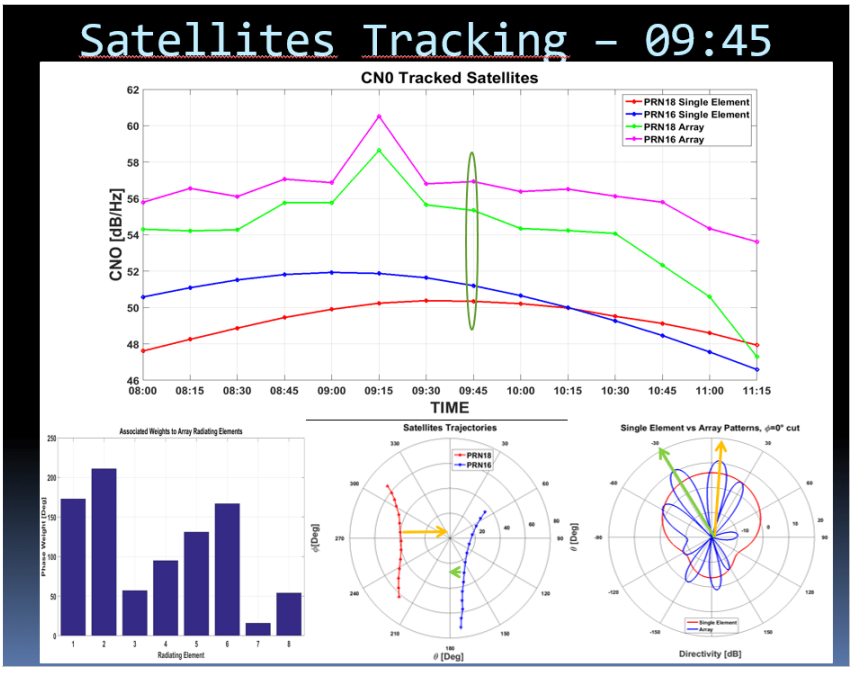

Fig. 3.17 Satellite tracking simulation 9:45 a.m.

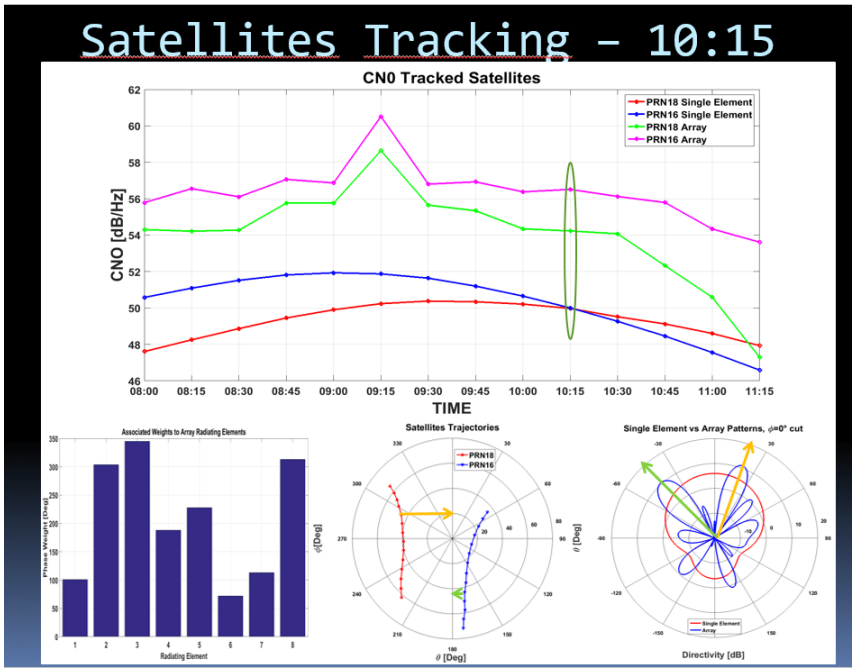

Fig. 3.18 ,
Satellite tracking simulation 10:15 a.m.

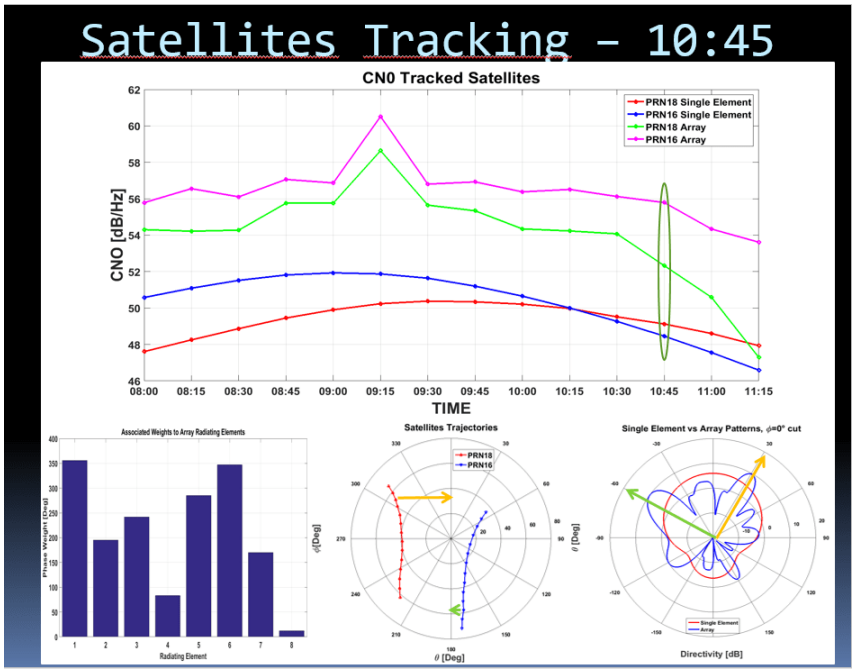

Fig. 3.19 Satellite tracking simulation 10:45 a.m.

3.6 Conclusions

This work proves that a simple phased array could improve C/N0 of more than one satellite, without requiring expensive dedicated hardware and complex iterative software methods. Currently, both antennas v1.0 and v2.0 are fabricated and under test. Then, real data acquisition will be performed to confirm the simulations results. Future works will extend the study to multidimensional array tracking.

Chapter 4

Conclusions

The thesis has addressed the implementation of reconfigurable antennas in all of its aspects. In particular the work has addressed in detail the implementation on low-power GPP as well as the integration on ULP SoC, showing at the same time their limitations.

The experimentation has demonstrated that the employment of reconfigurable antennas helps to reduce the energy consumption related to the radio as well as to provide a more efficient wireless communication, in particular the improvements are in terms of achievable coverage, data rate and interference reduction.

The research conducted on GPP has revealed that they provide large flexibility as full (re-)configurability and (re-)programmability. The disadvantage of these system lies in the energy consumption of the operative system.

On the other hand, SoC-based radios has proven to achieve very low-power consumption although they are less flexible.

In this thesis work, the reception of GPS signals with reconfigurable antennas has also been addressed, and the innovative approach has been the subject of a patent application.

References

[1] "Ieee standard for definitions of terms for antennas," *IEEE Std 145-2013 (Revision of IEEE Std 145-1993)*, pp. 1 – 50, Mar. 2014, 10.1109/IEEESTD.2014.6758443.

[2] D. Choudhury and T. Inoue, "Guest editorial special issue on 5g wireless communication systems and technologies," *Published in the IEEE Transactions on Microwave Theory and Techniques Journal*, vol. 64, no. 7, pp. 2205 – 2206, Feb. 2016, 10.1109/TMTT.2016.2575999.

[3] J. Aberle, S.-H. Oh, D. Auckland, and S. Rogers, "Reconfigurable antennas for wireless devices," *Published in the IEEE Antennas and Propagation Magazine*, vol. 45, pp. 148 – 154, Dec. 2003, doi:10.1109/MAP.2003.1282191.

[4] P. Gardner, P. S. Hall, M. Hamid, and F. Ghanem, "Reconfigurable antennas for cognitive radio," *Proceeding of the IEEE Topical Conference on Antennas and Propagation in Wireless Communications (APWC), Torino, Italy*, pp. 1225 – 1228, Dec. 2011, doi:10.1109/APWC.2011.6046836.

[5] N. Ryu, Y. Yun, S. Choi, R. C. Palat, and J. H. Reed, "Smart antenna base station open architecture for sdr networks," *Published in the IEEE Wireless Communications Journal*, vol. 13, no. 3, pp. 58 – 69, Sept. 2006, 10.1109/MWC.2006.1700072.

[6] D. Raychaudhuri and N. B. Mandayam, "Frontiers of wireless and mobile communications," *Published in the proceedings of the IEEE*, vol. 100, no. 4, pp. 824–840, Feb. 2012, 10.1109/JPROC.2011.2182095.

[7] J. L. A. Quijano and G. Vecchi, "Optimization of a compact frequency- and environment-reconfigurable antenna," *Published in the IEEE Transactions on Antennas and Propagation*, vol. 60, pp. 2682 – 2689, Apr. 2012, 10.1109/TAP.2012.2194634.

[8] N. Ryu, Y. Yun, S. Choi, R. C. Palat, and J. H. Reed, "Smart antenna base station open architecture for sdr networks," *Published in the IEEE Wireless Communications Journal*, vol. 13, no. 3, pp. 58 – 69, Sept. 2006, doi:10.1109/MWC.2006.1700072.

[9] G. V. Tsoulos, "Smart antennas for mobile communication systems: benefits and challenges," *Published in the IET Electronics and Communication Engineering Journal*, vol. 11, no. 2, pp. 84 – 94, Feb. 1999, doi:10.1049/ecej:19990204.

[10] A. P.-N. X. M. J. R. Fonollosa, "Smart antennas in software radio base stations," *Published in the IEEE Communications Magazine*, vol. 39, no. 2, pp. 166 – 173, Feb. 2001, doi:10.1109/35.900648.

[11] G. Marconi and F. Braun, "The nobel prize in physics," 1909. [Online]. Available: http://www.nobelprize.org/

[12] N. Anselmi, P. Rocca, A. Massa, and E. Giaccari, "Synthesis of robust beamforming weights in linear antenna arrays," *Proceedings of the IEEE Conference on Antenna Measurements and Applications (CAMA), Antibes Juan-les-Pins, France*, pp. 1 – 3, Nov. 2014, doi:10.1109/CAMA.2014.7003325.

[13] S. Ciccia, G. Giordanengo, and G. Vecchi, "Open-source implementation of an ad-hoc ieee802.11a/g/p software-defined radio on low-power and low-cost general purpose processors," *Published in the Radioengineering Journal*, June 2017, 10.13164/re.2017.0001.

[14] S. Ciccia, G. Giordanengo, S. Arianos, F. Renga, P. Ruiu, A. Scionti, L. Mossucca, O. Terzo, and G. Vecchi, "Reconfigurable antenna system for wireless applications," *Proceeding of the IEEE 1st International Forum on Research and Technologies for Society and Industry Leveraging a better tomorrow (RTSI), Torino, Italy*, pp. 111 – 116, Nov. 2015, doi:10.1109/RTSI.2015.7325080.

[15] S. Ciccia, G. Giordanengo, F. Renga, and G. Vecchi, "Software-defined reconfigurable antenna for energy efficient wireless links," *Proceeding of the IEEE International Symposium on Antennas and Propagation (APSURSI), Fajardo, Puerto Rico*, pp. 1241 – 1242, July 2016, doi:10.1109/APS.2016.7696328.

[16] ——, "Low power advanced wireless communication exploiting reconfigurable antennas," *Presented in the 10th European Conference on Antennas and Propagation (EuCAP), Davos, Switzerland*, 2016.

[17] S. Ciccia, G. Giordanengo, and G. Vecchi, "Integration of reconfigurable antennas in ultra low-power radio platforms based on system-on-chip for wireless sensor network," *Submitted to the IEEE Access Journal*, 2017.

[18] L. Pilosu, L. Mossucca, A. Scionti, G. Giordanengo, F. Renga, P. Ruiu, O. Terzo, S. Ciccia, and G. Vecchi, "Low power computing and communication system for critical environments," *Proceeding of the 11-th International Conference on P2P, Parallel, Grid, Cloud and Internet Computing (3PGCIC), Asan, Korea*, pp. 221 – 232, Oct. 2016, doi:10.1007/978-3-319-49109-7-21.

[19] L. Mossucca, L. Pilosu, P. Ruiu, G. G. adnd S. Ciccia, G. Vecchi, O. Terzo, V. Romano, L. Spogli, C. Cesaroni, I. Hunstad, and A. Serratore, "Greenlab: autonomous low power system extending multi-constellation gnss acquisition in antarctica," *Presented at the 32nd International Union of Radio Science General Assembly and Scientific Symposium (URSI GASS), Montréal, Québec, Canada*, Aug. 2017.

[20] G. Giordanengo, L. Pilosu, L. Mossucca, F. Renga, S. Ciccia, O. Terzo, G. Vecchi, V. Romano, and I. Hunstad, "Energy efficient system for environment observation," *Proceeding of the International Conference on Complex, Intelligent, and Software Intensive Systems (CISIS), Torino, Italy*, pp. 987 – 999, July 2017, doi:10.1007/978-3-319-61566-0-93.

[21] A. Scionti, O. Terzo, P. Ruiu, G. Giordanengo, S. Ciccia, G. Urlini, J. Nider, M. Rapoport, C. Petrie, R. Chamberlain, G. Renaud, D. Tsafrir, I. Yaniv, and D. Harryvan, *The Green Computing Continuum: the OPERA Perspective*, June 2017.

[22] A. C. Clarke, "Extra-terrestrial relays:" can rocket stations give world-wide radio coverage?" *Published in the Wireless World Megazine*, pp. 305 – 308, Oct. 1945.

[23] W. Davenport, "Sputnik, et cetera," *Published in the IEEE Transactions on Information Theory*, vol. 3, pp. 213 – 213, Dec. 1957.

[24] F. Krieger, "Announcement of the first satellite, from pravada, october 5, 1957," pp. 213 – 213, 1958.

[25] J. R. Hansen, *Spaceflight Revolution: NASA Langley Research Center from Sputnik to Apollo*, 1995.

[26] L. B. Slater, "Ffrom minitrack to navstar: The early development of the global positioning system, 1955–1975," *Proceeding of the IEEE International Microwave Symposium Digest (MTT-S), Baltimore, MD, USA*, pp. 1 – 4, June 2011, 10.1109/MWSYM.2011.5972582.

[27] P. Daly, "Navstar gps and glonass: Global satellite navigation systems," vol. 5, pp. 349 – 357, Dec. 1993, 10.1049/ecej:19930069.

[28] C. Junqueira, R. Suyama, F. J. V. Zuben, J. Marcos, and T. Romano, "An evolutive algorithm for blind adaptive beamforming in gps applications," *Proceeding of the IEEE Signal Processing Society Workshop Machine Learning for Signal Processing, Sao Luis, Brazil*, pp. 715 – 724, May 2004, 10.1109/MLSP.2004.1423037.

[29] X.-H. Wang, X.-W. Shi, P. Li, Y.-F. Bai, B. Liu, R. Li, and H.-J. Lin, "Smart antenna design for gps/glonass anti-jamming using adaptive beamforming," *Proceeding of the International Conference on Microwave and Millimeter Wave Technology (ICMMT), Chengdu, China*, pp. 1149 – 1152, July 2010, 10.1109/ICMMT.2010.5525060.

[30] X. dong Yuan, J. wei Wan, Z. Cheng, B. bin Shi, and H. Fu, "A novel adaptive multi-beam gnss anti-interference processing system," *Proceeding of the International Conference on Wireless Communications and Signal Processing (WCSP), Nanjing, China*, pp. 1 – 4, Dec. 2011, 10.1109/WCSP.2011.6096831.

[31] M. Sahmoudi and M. G. Amin, "Unitary cyclic music for direction finding in gps receivers," *Proceeding of the 4-th IEEE Workshop on Sensor Array and Multichannel Processing, Waltham, MA, USA*, pp. 70 – 73, Aug. 2006, 10.1109/SAM.2006.1706093.

[32] Q. Li, W. Wang, D. Xu, and X. Wang, "A robust anti-jamming navigation receiver with antenna array and gps/sins," *Published in the IEEE Communications Letters*, vol. 18, pp. 467 – 470, Apr. 2014, 10.1109/LCOMM.2014.012314.132451.

[33] W. Lijun, Z. Huichang, and Y. Xiaoniu, "Adaptive array antenna for gps interference mitigation and its performance analysis," *Proceeding of the International Conference on Microwave and Millimeter Wave Technology (ICMMT), Builin, China*, pp. 1 – 4, Apr. 2007, 10.1109/ICMMT.2007.381315.

[34] F. Tokan and F. Gunes, "Interference suppression by optimising the positions of selected elements using generalised pattern search algorithm," *Published in the IET Microwaves, Antennas and Propagation Journal*, vol. 5, pp. 127 – 135, Jan. 2011, 10.1049/iet-map.2009.0374.

[35] F. Dovis, M. Pini, M. Spelat, and P. Mulassano, "Benefits of a reconfigurable software gnss receiver in multipath environment," *Published in the Journal of Global Positioning Systems*, vol. 3, pp. 49 – 56, Oct. 2004.

[36] P. C. Fenton, W. H. Falkenberg, T. J. Ford, K. K. Ng, and A. J. V. Dierendonck, "Novatel's gps receiver - the high performance oem sensor of the future," *Proceedings of the 4th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS)*, pp. 49 – 58, Sept. 1991.

[37] J. Ramos, M. Zoltowski, and M. Urgos, "Robust blind adaptive array. a prototype for gps," *Proceeding of the IEEE International Symposium on Phased Array Systems and Technology, Boston, MA, USA, USA*, pp. 406 – 410, Oct. 1996, doi:10.1109/PAST.1996.566125.

[38] J. Wang, D. H. S. Dae Heon LEE, W. H. KIM, H. Y. Yang, and E. Park, "Anti-jamming apparatus and method for compact array antenna," Patent US 2014/375 500 A1, 2014.

[39] C. E. McDowell, "Method and apparatus for reducing jamming by beam forming using navigational data," Patent US 5 952 968 A, 1999.

[40] N. Padros, J. I. Ortigosa, J. Baker, M. F. Iskander, and B. Thornberg, "Comparative study of high-performance gps receiving antenna designs," *Published in the IEEE Transactions on Antennas and Propagation*, vol. 45, no. 4, pp. 698 – 706, Apr. 1997, doi:10.1109/8.564096.

[41] C. A. Balanis, *Chapter 14: Microstrip Antennas*, June 2009, isbn:8126524227, 9788126524228.

[42] MS Windows NT toko dak datasheets. [Online]. Available: http://datasheet.octopart.com/DAK1575MS50T-Toko-datasheet-101207.pdf

[43] E. C. Niehenke, V. V. D. Marco, and A. Friedberg, "Linear analog hyperabrupt varactor diode phase shifters," *Proceeding of the International Microwave Symposium Digest (MTT-S), St. Louis, MO, USA, USA*, pp. 657 – 660, June 2010, 10.1109/MWSYM.1985.1132067.

[44] MS Windows NT varactors datasheets. [Online]. Available: http://www.skyworksinc.com/uploads/documents/SMVA1248_079LF_203203B.pdf

[45] MS Windows NT orbitron software. [Online]. Available: http://www.stoff.pl/

[46] J. Kennedy and R. Eberhart, "Particle swarm optimization," *Proceeding of the IEEE International Conference on Neural Networks, Perth, WA, Australia*, vol. 4, pp. 1942 – 1948, Aug. 1995, 10.1109/ICNN.1995.488968.

Appendix A

(Publications)

Open-source implementation of an ad-hoc IEEE802.11a/g/p software-defined radio on low-power and low-cost general purpose processors

Open-source implementation of an ad-hoc IEEE 802.11 a/g/p software-defined radio on low-power and low-cost general purpose processors

Simone CICCIA [1,2], *Giorgio GIORDANENGO* [2], *Giuseppe VECCHI* [1]

[1] Dept. of Electronics and Telecommunication (DET), Politecnico di Torino, Corso Duca degli Abruzzi, 24, 10029, Torino, Italy

[2] Istituto Superiore Mario Boella, Via Pier Carlo Boggio, 61, 10138, Torino, Italy

simone.ciccia, giuseppe.vecchi@polito.it; ciccia, giordanengo@ismb.it

Manuscript received 21 June 2017

Abstract. *This work proposes a low-cost and low-power software-defined radio open-source platform with IEEE 802.11 a/g/p wireless communication capability. A state-of-the-art version of the IEEE 802.11 a/g/p software for GNU Radio (a free and open-source software development framework) is available online, but we show here that its computational complexity prevents operations in low-power general purpose processors, even at throughputs below the standard. We therefore propose an evolution of this software that achieves a faster and lighter IEEE 802.11 a/g/p transmitter and receiver, suitable for low-power general purpose processors, for which GNU Radio provides very limited support; we discuss and describe the software radio processing structuring that is necessary to achieve the goal, providing a review of signal processing techniques. In particular, we emphasize the advanced reduced-instruction set (RISC) machine (ARM) study case, for which we also optimize some of the processing libraries. The presented software will remain open-source.*

Keywords

Software-defined radio (SDR), General Purpose Processor (GPP), Low-power wireless communications, Advanced RISC machine (ARM), Open-source software.

1. Introduction

Software Defined Radio (SDR) plays a key role in present and future wireless communications; its advantage over hardware realization is in providing flexible, maintainable, and low-cost radio equipment [1]. Since software can be modified, cancelled and upgraded, SDR has become popular both in products and in research, like mainstream cognitive radio [2] as well in more selected applications like reconfigurable antennas [3, 4]; SDR is also an optimal approach for researchers and students to experiment with radio communication and signal processing concepts [5]. It is also well accepted in military, satellite communication and Vehicular Ad Hoc Networks (VANET) primarily for safety, security and privacy issues which can be meet by means of the SDR flexibility [6, 7]. However, such flexibility comes at the expense of increased power consumption, that is the challenge to overcome today. This work targets the implementation of a low-power and low-cost IEEE 802.11 a/g/p SDR on General Purpose Platforms (GPP).

Since GPP implementation is not the only option addressed by the scientific community we first briefly review the current state-of-the-art of general SDR implementation; we will next focus on the state-of-the-art for GPP SDR implementation.

1.1 State of the Art: General implementation of SDR

Existing platforms to implement SDR include Field Programmable Gate Array (FPGA), Digital Signal Processors (DSP), GPP and combinations of these to distribute the computational complexity of the software radio to the most appropriate resource. Low-power and low-cost FPGA platforms are usually employed to accelerate some performance-critical functions of a software radio [8]. However, a full stack IEEE 802.11 a/g/p SDR implementation requires a costly hardware along with a large consumption as reported by the open-source SDR implementation proposed in [9]. Another disadvanges of FPGA is that, at present, they can be only partially reconfigured at run-time [10]. Performance-critical SDR tasks are also well assessed on DSP platforms, which are architectures specialized for signal processing. Atomix, a framework able to convert C code into DSP is presented in [11]. This publication illustrates a partial implementation of the IEEE 802.11a transceiver, with a reported consumption of 7W. On the contrary our implementation targets the full IEEE 802.11 a/g/p SDR (i.e. transmitter and receiver) on a

low-power GPP. Furthermore, [11] does not account for the consumption of the Front-End (FE); the Universal Software Radio Peripheral (USRP)2 series employed in the test has a consumption of about 14W, and a significant cost. Finally, [11] does not clarify when and if the framework (including the partial IEEE 802.11a implementation) will be released as a free and open-source. It is also not clear if this framework could be used with a specific DSP hardware only (i.e. TI6670). Therefore, this solution is interesting and promising, but at present it is still at too early a stage to be assessed as low-cost and/or low-power; in contrast, our proposed work provides a complete implementation of the system and its validation.

1.2 State of the Art: SDR implementation on GPP

GPPs are architectures designed to perform more general operation and they are commercially available at low-cost. This explains why they are so widespread. They are fully flexible and provide the advantages of dynamically controlling both hardware and software components of a radio. For example, a Linux shell script can schedule when the radio app and the FE should be active or off to further reduce energy consumption.

However, the development of a low-power and low-cost full stack IEEE 802.11 a/g/p SDR is still challenging on GPP. First, this implementation requires a low-cost and low-power FE for direct baseband conversion over the desired frequency range. Commercial Digital Video Broadcast-Terrestrial (DVB-T) hardware receivers, like Realtek-2832U chip [12], have been discovered to be programmable by the SDR community, and software to feed the In-phase and Quadrature components (I/Q) of the received signal into a host PC has been developed [13]. This device offers the cheapest way to play with SDR concepts, but its frequency coverage is limited to Ultra-High Frequency (UHF) band, therefore is useless for standards that rely on Industrial, Scientific and Medical (ISM) bands (e.g. 2.4 GHz and 5.8 GHz). Thus, a transceiver equipment like the HackRF FE is required to play with SDR at the mentioned bands [14].

Secondly, code performance is an issue. A straightforward software implementation of the IEEE 802.11 a/g/p standard, like the one offered by the GNU Radio community, can only run on desktop PC, due to its high computational power demand [15]. This limitation is immediately apparent when investigating the performance of the software into low-power boards; this is seen, e.g., in the snapshots reported in Fig. 1, that refer to the output of GNU Radio *flowgraph*. While the 802.11 standard would require a sampling rate of 20 MHz, when the available GNU Radio software is run at a sampling rate of 2 MHz, the integrated cores reach saturation and events like buffer underflow/overflow occur over the upcoming/incoming samples stream respectively. At this sampling rate ($F_{s,max}$ = 2MHz), the software receiver shows a high error rate of about 75% that comes with overflows ("O" in the output display) and core saturations, which are inter-related, and most of the frames are corrupted and discarded. This happens because the processing capability of the low-power board is insufficient for the existing state of the art GNU radio implementation.

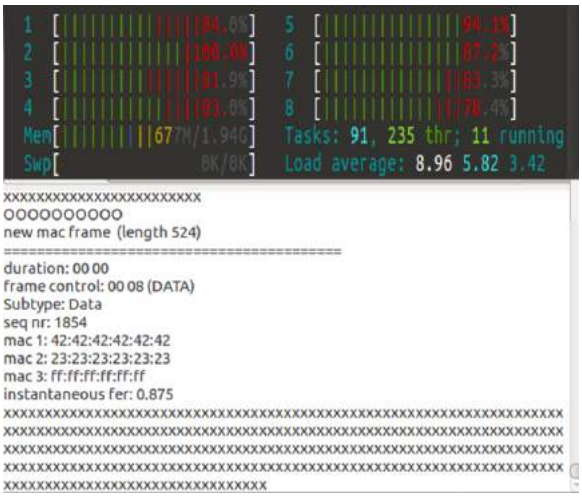

Fig. 1. Receiver performance at sampling rate $F_{s,max}$ = 2MHz, existing code on the targeted low-power board (see Sec. 2.2). The figure shows a snapshot of the GNU Radio *flowgraph* utility output for the existing IEEE 802.11 receiver, configured as described in Sec. 4.2. Top: Linux - htop program output;

Even more evidently, perhaps, the board goes in protection state after a few minutes of operation.

Some parts of this code have been recently optimized in [16]. The work in [16] was a source of inspiration to the present work; unfortunately, the code is not portable in architectures different from X86 (i.e. SSE2 instruction set) due to the lack of a generic implementation of a part of the software. This issue prevents installation on platforms different from X86, as reported in Figure 2, thus requiring modifications to run on the targeted low-power platforms. Furthermore, although the code has optimized for real-time performance, [16] does not state or discuss the amount of throughput achieved with the proposed improvements.

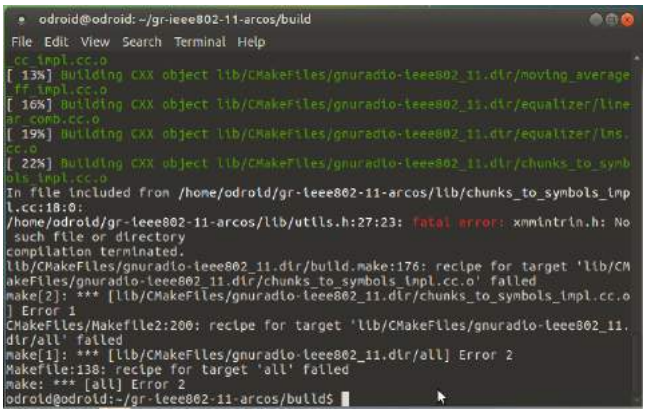

Fig. 2. Compiler output (make) building the work [16]. Issue preventing usability on platform different from X86 due to the lack of general implementation of part of the IEEE 802.11 software.

Among all possible framework for the development of SDR [17], we selected GNU Radio since it is free, open-source, comes with a wide availability of signal processing libraries and a vast community [18]. To the best of our knowledge, the GNU Radio framework (as well as other SDR implementations) is fully optimized for the X86 processor architecture, where the compiler toolchain can fully exploit specific architectural features to improve efficiency. For instance, in [19] the authors describe a SDR code optimized for the X86 architecture, targeting multicore acceleration; but it is not released as an open-source. Conversely, Advanced Reduced instruction set computing Machine (ARM)-based processors still suffer from a lack of good compiler support, and from the lack of specific features available in X86 counterparts [20, 21]. In fact, no low-power implementations have been reported so far, especially on the new generation of low-cost, power-efficient ARM platforms [22]. It is also worth noting that GNU Radio framework officially supports only X86 processors, thus providing no guarantees for its porting on ARM (or other) processors.
The open-source Sora project is, at present, a complete realtime SDR compliant with the IEEE 802.11 standard; however, this implementation requires a specific and costly hardware, and it needs to be employed on a PC [23]- as opposed to the application targeted here.
With the aim to provide low-cost and low-power SDR, the present work is intended to overcome the above limitations, and to yield SDR software requiring minimal hardware, i.e. a low-power and low-cost FE and GPP.

1.3 Proposed SDR implementation on GPP

In our intended application, there are no stringent requirements in terms of real-time performance; this allows to target low-power and low-cost GPP processor instead of DSP or FPGA implementations. As a result, our solution also maintains a significantly larger flexibility and code portability over other platforms. For example, in the case of a change of the GPP architecture, our solution only requires the code to be recompiled without the need for making changes in the source code. We start from a generic SDR implementation of the IEEE 802.11 a/g/p for desktop PC that at present is the only free and open-source software, and illustrate how to enhance code performance to achieve integration in low-power GPP while maintaining software portability, at minimal SDR cost.
To the best of our knowledge, there is no other implementations that provide SDR IEEE 802.11 a/g/p with the obtained performance on low-power and low-cost ARM-based boards. The code is available online at [24].
The paper is structured as follows. We start by reviewing requirements and available hardware for the FE (Sec. 2.1) and the GPP (Sec. 2.2); we then describe the realization of the SDR by describing the proposed method (Sec. 3) and we present the achieved results in Sec. 4; Sec. 5 collects the conclusions and perspectives.

2. Hardware Issues

We begin by discussing hardware issues, with specific attention to low power, and their impact and requirements on software development.

2.1 Front-End Issues

This section provides an overview of possible FEs suitable for the SDR implementation of the IEEE 802.11 a/g/p in GNU Radio and discusses the related issues.
To the best of our knowledge, valuable FE options for low-cost and low-power SDR which cover the IEEE 802.11 bands are summarized in the following with comments relevant to the present endeavor:

- Hack RF (299$), supported by GNU Radio; Half-duplex operations only.
- microSDR (expected 100$) according to [8]; at the time of this publication it is still under fabrication.
- PicoZeD SDR 1x1 (750$) [25]; Supported by Matlab and Simulink, but not guaranteed for GNU Radio.
- USRP-B series (more or less 800$ based on the model); supported by GNU Radio, Full-duplex operations.

Other FEs as Lime SDR are not considered in this discussion since they do not provide coverage at 5.8 GHz [26, 27]. In view of the above, our analysis and tests have considered the only two viable solutions, i.e. Hack RF and USRP-B. In the following we will discuss tests and findings that have guided our final selection.
A simplified circuit that shows analog and digital components of the receiver side is depicted in Figure 3.

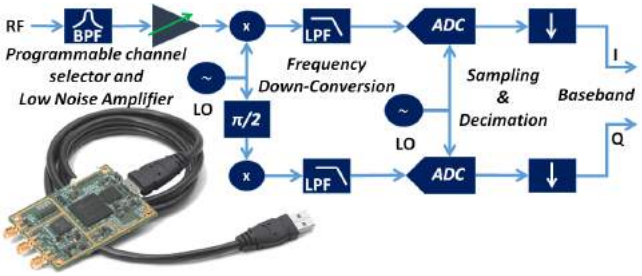

Fig. 3. Schematic circuit description of a typical Homodyne receiver.

The signal received in the bandwidth of interest is baseband transformed via a frequency down-conversion process. This analog elaboration produces a complex signal made up of I/Q components. Then Analog to Digital Converters (ADC)s map the analog values of the I/Q components to discrete levels. These converters have fixed sampling rate, thus a digital interpolation and decimation process is usually applied to get the sampling rate required by the intended communication standard. Finally, the digital I/Q sample stream is feed in

the GPP for digital signal processing tasks. The same chain (reversed), applies to the transmission process.
The first issue to address is the interface between the FE and the low-power GPP. In other worlds, the interface has to support the transfer rate R_T in bps as described by the following equation:

$$R_T = (I_{bits} + Q_{bits})F_{s,max} \quad (1)$$

where I_{bits} and Q_{bits} are the ADC bits dedicated to the In-phase and Quadrature component respectively, while $F_{s,max}$ is the sampling rate (in Hz) required by the intended communication standard.
The Hack RF uses the Universal Serial Bus (USB) v2.0 interface, which nominally supports 480 Mbps transfer rate; in this specific device the ADCs have 8 bits resolution for the I/Q components; inserting these values in Equation 1, together with the IEEE 802.11 maximum sampling rate of $F_{s,max} = 20$ MHz, one obtains $R_T = 320$ Mbps; thus, it appears to be fine for unidirectional (half-duplex) communication.
Ettus Research USRP-B operates via USB v3.0, which nominally supports 5 Gbps transfer rate. This FE provides 12 bits ADCs for I/Q, that coupled to a sampling rate of $F_{s,max} = 20$ MHz (IEEE 802.11) gives a transfer rate of 560 Mbps, i.e. a large margin with respect the USB v3.0 limit [28].
However, as part of the selection process, we performed standard transfer tests on both these front-ends; we found the behaviors reported in the snapshots in Figs. 4 and 5, and summarized by the data in Table 1.

Front-End	Samp.Rate[MHz]	Mode	Description
USRPB	8	full-duplex	Standalone connection
USRPB	0.5	full-duplex	Coupled with IEEE802.11 software
USRPB	17	half-duplex	Standalone connection
USRPB	1	half-duplex	Coupled with IEEE802.11 software
HACKRF	20	half-duplex	Standalone connection
HACKRF	0.5	half-duplex	Coupled with IEEE802.11 software

Tab. 1. The maximum sampling rate at which front-end devices passed the standard transfer tests on the low-power board in standalone, and coupled with existing IEEE 802.11 software.

It is apparent from these data that none is actually compliant with the IEEE 802.11 standard as both are unable to sustain the required sampling rate (i.e. 20 MHz). Hack RF is capable to transfer I/Q data into low-power GPP at such rate in standalone, while only in half-duplex mode. On the contrary, when associated to any IEEE 802.11 software the effective rate is very reduced, probably because the transceiver loses the execution priority with respect the tasks of the running software (i.e. availability of threads to perform its task on time). Conversely, USRP-B fails the transfer test at 20 MHz rate in full-duplex mode even in standalone. In view of the results in Table 1, we have selected USRP-B as front-end; it is apparent that (in a symmetrical link) there is no significant difference in throughput between half- and full-duplex; most of our results will refer to half-duplex operation for simplicity.

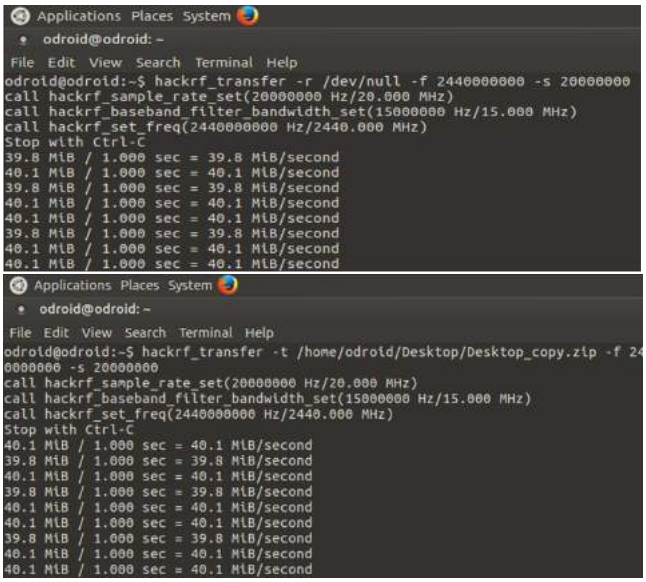

Fig. 4. Data rate logger output of the Hack RF standard transfer test at 20 MHz rate (Half Duplex). Receiver test (Up) and transmitter test (Bottom).

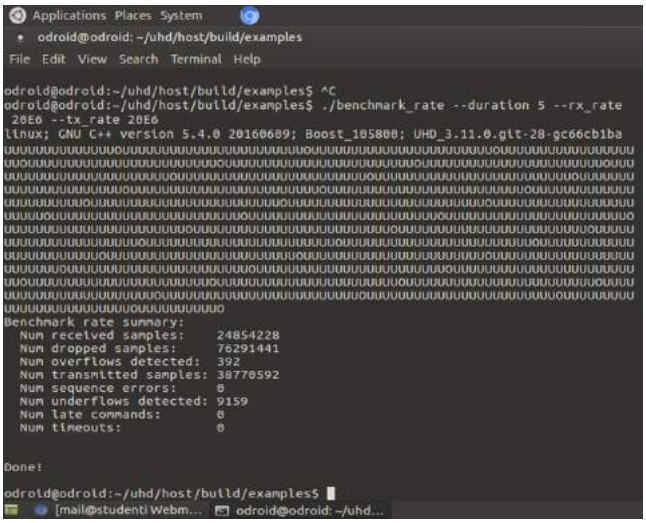

Fig. 5. Data rate logger output of the USRP-B205 mini standard transfer test at 20 MHz rate (Full Duplex).

2.2 GPP Platform

To satisfy the requirements of a more energy-efficient hardware, with low-cost property (74$) and small form factor, we targeted this project to the Odroid-XU4 [29]. This ARM-based board is made up by a selectable high performance quad-core Cortex-A15 and a lower-performance quad-core Cortex-A7 Cental Processing Unit (CPU)s. These processors include Neon-technology, an instruction set able to accelerate signal processing algorithm [30]. The drawback is that they come with a specific Linux distributions that in general does not perform fast enough for a SDR environment, e.g. slow USB buffer copies could prevent the interaction with the FE. A kernel modification is then required to fully

exploit the performance of such boards. The downgrade of the kernel is accomplished by recompiling with options full kernel-preemption, overclocking and Neon.

3. Operations, Algorithms and Optimizations

The processing required by the IEEE 802.11 to code and decode its waveform involves a massive number of arithmetic operations. Therefore, in the following we will analyze the most relevant such computational tasks, with a specific emphasis on those that can be substantially improved; this will constitute the starting point of the optimization.
In order to identify bottlenecks and heavy processing blocks of the IEEE 802.11 a/g/p *flowgraph* for GNU Radio we employed Control Port (CP) monitor performance, a tool recently added in GNU Radio by the work [31].
Figure 6 shows a simplified *flowgraph* description of the IEEE 802.11 a/g/p baseband processing showing with a lighter color the blocks we found critical according to CP, and that have been restructured by this work.

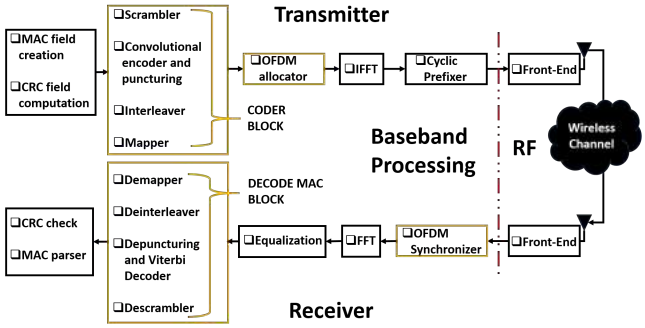

Fig. 6. Simplified GNU Radio *flowgraph* of the IEEE 802.11 a/g/p baseband processing chain.

The following sub-sections will discuss the restructuring of such *flowgraph* in the order of the proposed optimization methods.
A first note is about Fast Fourier Transform (FFT) and its inverse; they are basic functions that GNU Radio provides; accelerations of these functions are already available [32], and thus will be not discussed here.

3.1 Platform Generic Optimization

A ubiquitous source of inefficiency is the handling of shift register which are to perform a large number of operations. However, shift registers have a finite set of possible states and therefore their associated operation can be computed once at run-time and stored in Look Up Tables (LUT). This operation is beneficial in typical SDR platform since they are processor-based system with large amounts of available memory that can be exploited for reducing the runtime complexity of signal processing tasks. This approach is also discussed in [23] but never addressed from the implementation point of view.
Another significant enhancement in performance is obtained by operating on bytes instead of bits since the same operation is parallelized over 8 bits at the same time.
Some examples of operations to which this optimization process is applied are:

- Scrambler process
- Convolutional encoder
- Interleaving
- CRC computation

Significant improvements can also be reached by unrolling for–loops. This is the case of the OFDM Carrier Allocator. Its body was constituted by an initialization of a large output buffer onto which two sequential loops write information and pilot symbols, respectively. In our implementation of the OFDM Carrier Allocator the information, pilot symbols and guard intervals are allocated cyclically avoiding initialization and stressing loops. This effectively reduces the computational complexity on all architectures and, even more so, on ARM-based board since the compiler is not able to efficiently handle for-loops. The proposed implementation of the OFDM Carrier Allocator is reported in Appendix.
Regarding the demapping-decision process, we have implemented the same suboptimal method that has recently added online by the work [16]. We found that this method gives a $70x$ boost in the code speedup with respect the optimal method proposed by the IT++ library. Several authors have investigated the performance of this approach, and the works in [33, 34, 35] have demonstrated a loss of only 2 dB in the Bit Error Rate (BER).
Finally, frequency offset compensation and equalization are not modified since the CP analyses has not revealed critical aspects in performance.

3.2 A Practical Example of Platform Generic Optimization

To provide a concrete example, in the following we report a step-by-step demonstration to speedup the scrambling processing function, showing some tricks. The IEEE 802.11 a/g/p scrambler is defined by the equation 2, which shows the polynomial representation.

$$p(x) = x^7 + x^4 + 1 \tag{2}$$

Its flow graph is reported in Figure 7. The scrambler Shift Register (SR) starts with the all-ones state. The Pseudo-Random Binary Sequence (PRBS) is periodic and repeats after a period defined in the Equation 3, which shows the number of possible states.

$$PRBS\_PERIOD = 2^7 - 1 \tag{3}$$

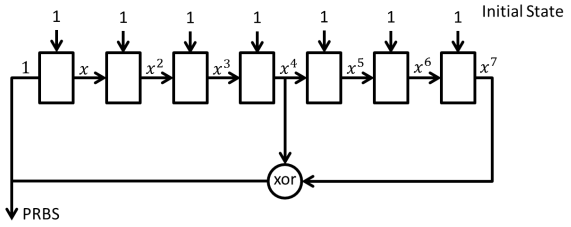

Fig. 7. Scrambler block diagram and illustration of the Pseudo-Random Binary Sequence (PRBS) generator.

The scrambler provided by the IEEE 802.11 a/g/p for GNU Radio computes the PRBS state for each input data bit (i.e. unpacked bytes) as shown in the following code:

```
0  /**General work**/
1  for (int i = 0; i < n_data_bits; i++)
2    {
3    PRBS = (!!(state & 64)) ^ (!!(state & 8));
4    out[i] = PRBS ^ in[i];
5    state = ((state << 1) & 0x7e) | PRSB;
6    if(state > 127) state = 1;
7    }
```

where the general work describes the part that reads inputs, processes, and writes outputs. This code offers poor performance since carry out many operations for each input bit. The complexity of signal processing can be reduced by exploiting memory availability. Since PRBS is periodic, it can be generated once in the constructor and stored in a LUT. In C/C++ programming language, the constructor is a section of code executed only once at run-time (i.e. when the object is created). We proposed the following restructuring of the PRBS:

```
0  /**Code constructor**/
1  SR = 0x7F;
2  for(int state = 0; state < PRBS_PERIOD; state++)
3    {
4    PRBS[state]  =  SR & 0x01;
5    PRBS[state] ^= (SR >> 3) & 0x01;
6    SR  = SR >> 1;
7    SR |= PRBS[state] << 6;
8    }
```

This optimization yields the following general work:

```
0  /**General work**/
1  for(int i = 0; i < n_data_bits; i++)
2    {
3    out[i] = in[i] ^ PRBS[state];
4        state = (state + 1) % PRBS_PERIOD;
5    }
```

As a consequence, the number of operation reduces to a memory access and an XOR operation for each bit. Furthermore, a second step optimization is still possible by operating on bytes instead of bits. In this case eight information bits are packed in one byte. This operation requires a different LUT that stores eight successive states of the PRBS, as reported in Table 2.

STATE	PRBS	PACKED_PRBS
0	0	11111110
1	0	11111100
2	0	11111000
3	0	11110000
4	1	11100001
5	1	11000011
6	1	10000111
7	0	00001110
8	1	00011101
9	1	00111011
10	1	01110111
11	1	11101111
12	0	11011110
..	.	.....
126	1	11111111

Tab. 2. Look Up Tables (LUT)s for the Scrambler processing function.

In this way, the new PRBS LUT is generated in the constructor as:

```
0  /**Code constructor**/
1  for(int state = 7; state < PRBS_PERIOD+7; state++)
2    for(int j = 7; j >= 0; j--)
3      PACKED_PRBS[state % PRBS_PERIOD] |=
         (PRBS[((state - j) % PRBS_PERIOD)] & 0x01) << j;
```

Then, assuming that input data are packed in a byte (i.e. 8 bit to process at once), the processing operation reduces to:

```
0  /**General work**/
1  for(int i = 0; i < n_data_bytes; i++)
2    {
3    packed_out[i] = packed_in[i] ^ PACKED_PRBS[state];
4    state = (state + 8) % PRBS_PERIOD;
5    }
```

where 8 bits are processed with one access to memory and an XOR operation. This optimization is architecture independent (i.e. the software is accelerated on any processor architectures). This method significantly boost the performance at the expense of using more memory. On complex architectures like ARM, the modulo division could be time consuming with respect to an IF condition due to low-level initialization issues. Thus, a further step of which all architecture benefits is restructuring the general work in the following way:

```
0  /**General work**/
1  for(int i = 0; i < n_data_bytes; i++)
2    {
3    packed_out[i] = packed_in[i] ^ PACKED_PRBS[state];
4    state += 8;
5        if( state > PRBS_PERIOD )
6              state %= PRBS_PERIOD;
7    }
```

In this way, the modulo division is performed only once every 16 byte operations.

3.3 Platform Specific Optimization

High speed Single Instruction Multi Data (SIMD) functions are very efficient and allow parallel operations, and their use was proposed in work [16, 19, 23, 36]. They are usually employed for the most critical task requiring acceleration. However, this type of optimization is strictly related to the architecture employed, because it exploits registers and instruction sets intrinsic to the processor. The code written with this method is therefore not portable; because of this, we used it only for the most critical tasks (e.g. to accelerate the Viterbi decoder), and always also keeping (i.e. providing) a generic version suitable for all processors.

At the receiver side, the synchronizer implements the maximum normalized correlation algorithm. This part involves many arithmetic operations at the signal sampling rate (i.e. before packet detector) and the use of accelerators is usually required. For this purpose we used the SIMD functions whose implementation and performance impact on ARM processors are well discussed in [36]. As these functions are architecture dependent, our proposed software provides also a platform-independent (generic) implementation to maintain the code portable. At the software building stage the compiler can choose the SIMD functions if supported on the target platform, otherwise select the generic-implementation counterpart.

We found out that the Viterbi decoder was a significant computational load, and in fact it turns out to be the decoding function requiring the highest computational power at the receiver side. At the beginning, the IEEE 802.11 a/g/p for GNU Radio employed the Viterbi decoder provided by the mathematical library of communication functions (ITpp). However, this straightforward implementation was not efficient for a SDR [37], and GNU Radio community provided a more performing version of the Viterbi which was actually employed in its DVB-T receiver project, usually referred to as the Karn-Ettus Viterbi [38]. This function has also an accelerated SIMD version, but unfortunately it can run on X86 machine only (i.e. supporting SSE2 instruction set). We adapted and tested the generic version of the Karn-Ettus Viterbi on our IEEE 802.11 a/g/p implementation, and we found that it is still too slow on ARM-based boards, i.e. packet are frequently missed due to overloading. Therefore, we implemented a SIMD accelerated version suitable for ARM processors, by exploiting Neon-technology. Since this optimization is architecture dependent we also maintained the generic version (i.e. Karn-Ettus Viterbi) for all architectures and the SSE2 optimization which can be exploited on X86 architectures.

4. Results

4.1 Software Performance

To assess the performance of single sub-blocks of code within a block – for instance scrambler, convolutional encoder, interleaver and mapping within the block MAPPER– we used the following metric:

$$F_{speedup} = \frac{\langle T_{proc,Ex} \rangle}{\langle T_{proc,Prop} \rangle} \tag{4}$$

where we define the speedup factor $F_{speedup}$ as the ratio between the average processing time $\langle T_{proc,Ex} \rangle$ required by the existing code and that, $\langle T_{proc,Prop} \rangle$ of the proposed sub-block of code. In other words, $\langle T_{proc} \rangle$ is the time required to process a frame in a specific task, averaged over N instances. In the performed test it is evaluated over the processing of 500-Byte frames, and averaged over $N = 100$ frames. Speedup factors have been obtained for the convolutional encoder, puncturing/depuncturing, and interleaver/deinterleiver, which are part of the blocks CODER and DECODE MAC, according to the method presented in Sec. 3.1. The improvements with respect the existing IEEE 802.11 a/g/p for GNU Radio have been reported in Table 3.

Function	Speedup factor
Convolutional Encoder	5x
Scrambler/Descrambler	7x
Interleaver/Deinterleaver	6x

Tab. 3. Improvements w.r.t. the existing IEEE 802.11 a/g/p for GNU Radio.

According to CP profiling of all blocks of Figure 6 we found the OFDM Carrier Allocator to be the block consuming the majority of the runtime. After the restructuring of the OFDM Carrier Allocator (Sec. 3.1) we obtained a large improvements as reported in Figure 8, which illustrates the run-time speedups based on Equation (4).

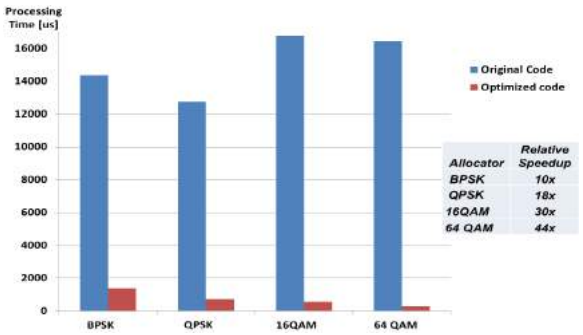

Fig. 8. Processing time comparison of proposed OFDM Allocator w.r.t GNU Radio OFDM Allocator.

As shown in Fig. 8 the acceleration increases as the modulation becomes more complex. This is justifiable from the fact that transmitting the same quantity of information, the highest the modulation order the lesser is the frame size (i.e. bytes). This achievement does not saturate buffers along the

run-time, as the original code does. This also explains why the latter code has a high execution time that results independent from the modulation employed.
The results of our proposed Viterbi implementation compared to the others supported by the target platform are reported in Figure 9, while a piece of code illustrating the most intensive computational part rewritten with ARM intrinsic functions is reported in appendix.

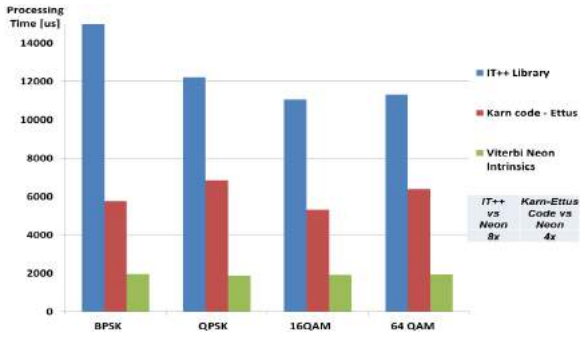

Fig. 9. Processing time comparison of proposed Viterbi implementation (Neon) against others on ARM-based processor.

The improvements obtained from the restructuring of the blocks CODER and DECODE MAC, depicted in Figures 6, have inspected with CP which allows to test the overall performance of blocks, instead of sub-blocks. Figure 10 and Figure 11 reports the results of this analysis for both existing and proposed blocks when running the software for different packet lengths and modulations. The bar chart shows the percentage of the total runtime (i.e. the number of CPU ticks) during the call to a specific block. We also report the runtime ratio between the existing and proposed blocks (i.e. speedup factor within the bar). The line chart, with secondary axis at the right side, reports the percentage of the buffers full allocated by the inspected block.
The CODER Block analysis, Figure 10, revealed a speedup factor around 30x for all modulation orders and packet lengths, while buffers are less busy of a factor of 15.
The DECODE MAC proves an overall speedup factor around 11x which in average persists for all modulation orders and packet lengths. Buffers result less full of a constant factor of 10.

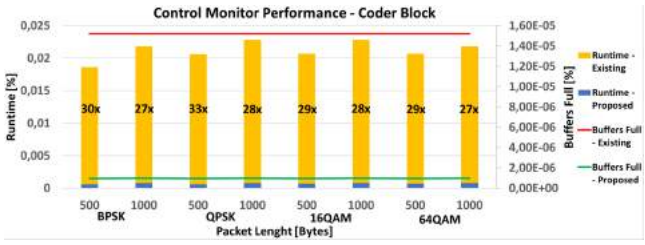

Fig. 10. Time and memory comparison, CODER block. The figure shows the percentage of runtime (bar plot) and buffers full (line plot) of the existing and proposed CODER block when running the GNU Radio *flowgraph* for different modulation orders and packet lengths.

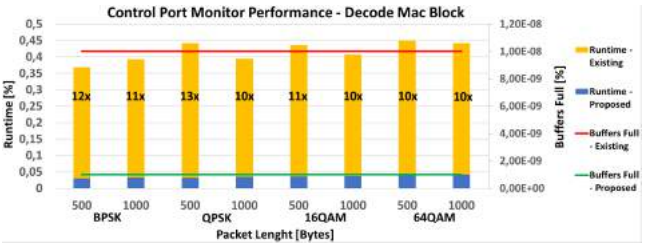

Fig. 11. Time and memory comparison, DECODE MAC block. The figure shows the percentage of runtime (bar plot) and buffers full (line plot) of the existing and proposed DECODE MAC block when running the GNU Radio *flowgraph* for different modulation order and packet lengths.

4.2 Wireless communication link performance

We begin by noting that the communication performance of IEEE 802.11 a/g/p software has been provided in the work [15]. Since there are no difference in coding/decoding chains which form the IEEE 802.11 waveform there also no significant change in communication performance (including the negligible loss of the demapper). Therefore, we will not address this issue further in our discussion.
The performed tests aim to evaluate the maximum throughput that the proposed implementation achieves in the target low-power platform. The analysis addresses a point-to-point wireless communication between two equal nodes spaced 5m apart. The characteristic of the single node (transmitter or receiver) is reported in Table 4.

Component	Type
CPU	Quad Cortex-A15 2.0GHz, Quad Cortex-A7 1.6GHz
Operating System	Ubuntu 15.04, Linux odroid 3.10.92-rt101
Environment	GNU Radio 3.7.10.1
Front-End	USRP B205 mini
Interface	USB v3.0
Antenna	Antenova Titanis AE030054-I 2.4GHz, peak gain 2.2dBi

Tab. 4. Overview of the most important components of the test bed.

Proper Radio Frequency (RF) gains have been set to provide a reliable and stable communication for the data rate under test. In view of the results in Sec. 2.1, for the sake of simplicity we have chosen half-duplex communication in these tests. We conducted the tests by stressing the receiver transmitting continuous frames (e.g. broadcasting communication) to validate the tolerance to saturation. We successfully tested the code at $F_{s,max} = 2$ MHz without encountering issues related to the execution of the code, such as underflow at the transmitter and/or overflow at the receiver, and transmitting with the maximum order modulation (i.e. 64QAM). We also found that the decoder became the bottleneck after such wireless data rate.
For completeness, we performed the same test also with the HACK RF Front-End, however $F_{s,max}$ remained 0.5 MHz as with the initial code (in Table 1). With this data we can thus assess the associated maximum bit rate R of the wireless link,

$$R = (\frac{R}{B})F_{s,max} \quad (5)$$

where $\frac{R}{B}$ is the maximum spectral efficiency, that with the 64QAM modulation of IEEE 802.11 a/g/p is $\frac{R}{B} = 2.7$ bit/s/Hz [39].

On the transmitter side, we experimentally assessed that the transmitter supported up to 5MHz sampling rate with the existing software, and 9 MHz sampling rate with the proposed code. This corresponds to a wireless bit rate of 24.3 Mbps for the proposed code, and it is apparent that the link bit rate is limited by the lower sampling rate achievable by the receiver, due to decoding saturation; for our measured data of $F_{s,max} = 2$ MHz with USRP-B we obtain an overall maximum bit rate of $R = 5.4$ Mbps in the receiver.
These results can be different with other processors, dependent on the availability of parts of code that can exploit accelerators and the efficiency of the compiler.
Comparison of the existing and proposed code performance have been reported in Figure 1 and Figure 12, focusing only on the (more critical) receiver side. For the sake of readability, in generating these figures we have employed 16QAM modulation instead of 64QAM. At a difference with the results in Figure 1, our proposed solution is able to keep the sampling rate ($F_{s,max} = 2$ MHz) and correctly decode the information.

Fig. 12. Receiver performance at $F_{s,max} = 2$MHz, existing code. The figure shows a snapshot of GNU Radio *flowgraph* output for the proposed IEEE 802.11 receiver, configured as described in Table 4. (see Sec. 2.2). Top-left: Linux - htop program output; bottom-left: GNU Radio terminal output; right: GNU Radio Constellation GUI. (16QAM has been used instead of 64QAM only for ease of visualization.)

In conclusion, the wireless communication improvement in transmission rate is measured from the achievable sampling rate $F_{s,max}$, which is 1MHz for the existing code (Tab. 1), and 2 MHz for the proposed implementation; this corresponds to a two-fold increase in (maximum) transmission rate, from 2.7 to 5.4 Mbps in half-duplex for the 64QAM of the 802.11 standard. There is a likely increase in the transmitter performance, but this is less relevant due to the limitation coming from the receiver.

Consumption of both transmitter and receiver have been evaluated by measuring the current absorbed from the Odroid-XU4 through a digital ammeter. The estimation includes the FE which is self-powered through the USB and the operating system that together with the FE establish the power overhead of the idle state. The measurements are reported in Table 5. Usually the idle state is used for less time (i.e. when boot and before shutdown), since in a real application scenario, e.g. Wireless Sensor Networks (WSN), great energy efficiency is obtained by turning on the platform when data need to be transmitted and turning off as soon as the communication terminates.

Radio Platform State	**Typical Power Consumption** [W]
Idle	5
Transmitter	11
Receiver	16

Tab. 5. Consumption Profile Odroid-XU4 (Linux) and USRP B205 mini transmitting and receiving at the maximum achieved throughput of 6 Mbps.

5. Conclusions and Future Perspectives

This work described an extension of GNU Radio on ARM with a lighter implementation of the IEEE 802.11 a/g/p transmitter and receiver to allow operation on low-power and low-cost boards. To the best of our knowledge, there is no other instrumentation/implementation that provides a full SDR IEEE 802.11 a/g/p with these performance on low-cost ARM-based GPPs, without employing external hardware. Among the extensions of this work, a valuable option to further extend the bandwidth is to bring some computational complexity in the small FPGA of the FE as, for example, the receiver synchronizer or the Viterbi task. Another valuable option could be analysing the code at a very low level (i.e. assembly-ARM to further reduce the required number of operations).

Acknowledgments

The authors would like to thank the GNU Radio community for the valuable help.

References

[1] RAYCHAUDHURI, D., MANDAYAM, N. B. Frontiers of wireless and mobile communications. In *Proceedings of the IEEE*, 2014, vol. 100, no. 4, p. 824–840. DOI: 10.1109/JPROC.2011.2182095

[2] LLAMES, G. J. M., BANACIA, A. S. Spectrum sensing system in software-defined radio for determining spectrum availability. In *International Conference on Electronics, Information, and Communications - ICEIC*. Danang (Vietnam), 2016, p. 1–5. DOI: 10.1109/ELINFOCOM.2016.7562961

[3] HALL, P. S., GARDNER, P., FARAONE A. Antenna requirements for software defined and cognitive radios. In *Proceedings of the IEEE*, 2012, vol. 100, no. 7, p. 2262–2270. DOI: 10.1109/JPROC.2012.2182969

[4] CICCIA, S., VECCHI, G., GIORDANENGO, G., RENGA, F. Software-defined reconfigurable antenna for energy efficient wireless links. In *IEEE International Symposium on Antennas and Propagation*. Puerto Rico (USA), 2016, p. 1241–1242. DOI: 10.1109/APS.2016.7696328

[5] STEWART, R. W., CROCKETT, L., ATKINSON, D., et al. A low-cost desktop software defined radio design environment using matlab, simulink, and the rtlsdr. *IEEE Communications Magazine*, 2015, vol. 53, no. 9, p. 64–71. DOI: 10.1109/MCOM.2015.7263347

[6] BERGSTROM, C., CHUPRUN, S., GIFFORD, S., MAALOULI, G. Software defined radio (SDR) special military applications. In *Proceedings of MILCOM*. Anaheim, CA (USA), 2002, p. 383–388. DOI: 10.1109/MILCOM.2002.1180472.2016.7696328

[7] LO PRESTI, L., FALLETTI, E., NICOLA, M., TROGLIA GAMBA, M. Software Defined Radio technology for GNSS Receivers. In *IEEE International Workshop on Metrology for Aerospace*. Benevento (Italy), 2015, p. 314–319. DOI: 10.1109/MetroAeroSpace.2014.6865941

[8] KUO, Y., PANNUTO, P., SCHMID, T., DUTTA, P. Reconfiguring the software radio to improve power, price, and portability. In *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems - SenSys*. Toronto, Ontario (Canada), 2012, p. 267–280. ISBN: 978-1-4503-1169-4. DOI: 10.1145/2426656.2426683

[9] MURPHY, P., SABHARWAL, A., AAZHANG, B. Design of WARP: A wireless open-access research platform. In *14th European Signal Processing Conference - EUSIPCO*. Florence (Italy), 2006, p. 1–5. ISSN: 2219-5491. DOI: 10.5281/ZENODO.52692

[10] GAZZANO, J. D. Integrating FPGAs: A dynamically reconfigurable FPGA-based grid for high performance computing. In *International Conference on Advances in Electrical, Electronic and Systems Engineering - ICAEES*. Putrajaya (Malaysia), 2016, p. 1–4. DOI: 10.1109/ICAEES.2016.7887998

[11] BANSAL, M., SCHULMAN, A., KATTI, S. Atomix: A Framework for Deploying Signal Processing Applications on Wireless Infrastructure. In *Proceedings of the 12th USENIX Symposium on Networked Systems Design and Implementation - NSDI*. Oakland, California (USA), 2015, p. 173–188. ISBN: 978-1-931971-218. DOI: Not Available

[12] REALTEK, TAIWAN. *RTL 2832U*. [Online] Cited 2017-06-30. Available at: http://www.realtek.com.tw/

[13] LANG, J., OSMOCOM. *RTL-SDR*. [Online] Cited 2017-06-30. Available at: http://sdr.osmocom.org/

[14] OSSMANN, M. *Great Scott Gadgets - HackRF One*. [Online] Cited 2017-06-30. Available at: https://greatscottgadgets.com/

[15] BLOESSL, B., SEGATA, M., SOMMER, C., DRESSLER, F. Towards an open source ieee 802.11p stack: A full sdr-based transceiver in GNU Radio. In *Proceedings of 5th IEEE Vehicular Networking Conference - VNC*. Boston (USA), 2013, p. 143–149. ISSN: 2157-9865. DOI: 10.1109/VNC.2013.6737601

[16] ARCOS, G., FERRERI, R., RICHART, M., EZZATTI, P., et al. Accelerating an IEEE 802.11 a/g/p Transceiver in GNU Radio. In *Proceedings of the 9th Latin America Networking Conference - LANC*. Valparaiso (Chile), 2016, p. 13–19. ISBN: 978-1-4503-4591-0. DOI: 10.1145/2998373.2998443

[17] BANSAL, M., MEHLMAN, J., KATTI, S., LEVIS, P. DATASHVILI, L., BAIER, H. OpenRadio: A Programmable Wireless Dataplane. In *Proceedings of the first workshop on Hot topics in software defined networks - HotSDN*. Helsinki (Finland), 2012, p. 109–114. ISBN: 978-1-4503-1477-0. DOI: 10.1145/2342441.2342464

[18] BLOSSOM, E. GNU Radio: tools for exploring the radio frequency spectrum. *Linux Journal*, 2004, vol. 2004, no. 122, p. 4. ISSN: 1075-3583. DOI: Not Available

[19] BERGER, C. R., ARBATOV, V., VORONENKO, Y., et al. Real time software implementation of an ieee 802.11a baseband receiver on intel multicore. In *IEEE International Conference on Acoustics, Speech and Signal Processing - ICASSP*. Prague (Czech Republic), 2011, p. 1693–1696. DOI: 10.1109/ICASSP.2011.5946826

[20] MELNIK, D., BELEVANTSEV, A., PLOTNIKOV, D., LEE, S. DATASHVILI, L., BAIER, H. A case study: optimizing GCC on ARM for performance of libevas rasterization library. In *International Workshop on GCC Research Opportunities - GROW*. Pisa (Italy), 2010, p. 4–8. DOI: Not Available

[21] MAHALINGAM, P. R., ASOKAN S. A framework for optimizing GCC for ARM architecture. In *Proceedings of the International Conference on Advances in Computing, Communications and Informatics - ICACCI*. Chennai (India), 2012, p. 337–342. DOI: 10.1145/2345396.2345452

[22] SEGARS, S., UK. *ARM Holdings*. [Online] Cited 2017-06-30. Available at: https://www.arm.com/company

[23] TAN, K., LIU, H., ZHANG, J., ZHANG, Y., et al. Sora: High-performance software radio using general-purpose multicore processors. *Communications of the ACM*, 2011, vol. 54, no. 1, p. 99–107. DOI: 10.1145/1866739.1866760

[24] CICCIA, S., ITALY. *Low power 802.11a/g/p software*. [Online] Cited 2017-06-30 Available at: http://tinyurl.com/h3bglfk

[25] AVNET, USA. *PicoZed SDR 1X1*. [Online] Cited 2017-06-30. Available at: http://zedboard.org/

[26] Lime microsystems, UK. *LimeSDR*. [Online] Cited 2017-06-30. Available at: http://www.limemicro.com/

[27] RTL SDR. *Roundup of Software Defined Radios*. [Online] Cited 2017-06-30. Available at: http://www.rtl-sdr.com/roundup-software-defined-radios/

[28] ETTUS RESEARCH, USA. *USRP B205 mini*. [Online] Cited 2017-06-30. Available at: https://kb.ettus.com/

[29] HARDKERNEL, SOUTH KOREA. *ODROID-XU4*. [Online] Cited 2017-06-30. Available at: http://www.hardkernel.com/

[30] ARM. *NEON instruction set*. [Online] Cited 2017-06-30. Available at: https://www.arm.com/products/processors/technologies/neon.php

[31] ROUNDEAU, T. W., O'SHEA, T., GOERGEN, N. Inspecting GNU radio applications with controlport and performance counters. In *Proceedings of the second workshop on Software radio implementation forum - SRIF*. Hong Kong (China), 2013, p. 65–70. ISBN: 978-1-4503-2181-5 DOI: 10.1145/2491246.2491259

[32] FFTW3. *Multiplatform DFT*. [Online] Cited 2017-06-30. Available at: http://www.fftw.org/

[33] M ULLER, B., HOLTERS, M., Z OLZER, U., Low complexity Soft-Input Soft-Output Hamming Decoder. In *50th FITCE Congress - ICT: Bridging an Ever Shifting Digital Divide*. Palermo, Italy, 2011, p. 1–5. DOI: 10.1109/FITCE.2011.6133448

[34] TOSATO, F., BISAGLIA, P. Simplified soft-output demapper for binary interleaved COFDM with application to HIPERLAN/2. In *IEEE International Conference on Communications - ICC*. New York City, New York (USA), 2002, p. 664–668. DOI: 10.1109/ICC.2002.996940

[35] YEH, H., FUNG, W. DATASHVILI, L., BAIER, H. Performance and DSP implementation of soft bit-level demapper for M-QAM-OFDM systems. In *Proceedings of the Annual IEEE System Conference*. Vancouver (Canada), 2015, p. 810–815. DOI: 10.1109/SYSCON.2015.7116850

[36] WEST, N., GEIGER, D., SCHEETS, G. DATASHVILI, L., BAIER, H. Accelerating Software Radio on ARM: Adding NEON Support to VOLK. In *Proceedings of the IEEE Radio and Wireless Symposium - RWS*. San Diego, California (USA), 2015, p. 174–176. DOI: 10.1109/RWS.2015.7129727

[37] ITpp. *library of mathematical, signal processing and communication classes and functions*. [Online] Cited 2017-06-30. Available at: http://montecristo.co.it.pt/itpp/

[38] KARN, P. *Viterbi Decoder*. [Online] Cited 2017-06-30. Available at: http://www.ka9q.net/

[39] GAST, P. *802.11 Wireless Networks: The Definitive Guide*. 2nd ed., rev. Sebastopol, CA (USA): O'Reilly Media, Inc., 2005. ISBN: 0596100523, 9780596100520

Appendix

OFDM Allocator (Proposed implementation)

```
0  /**General work - Packet based**/
1  // Copy Sync word
2  for (unsigned i = 0; i < d_sync_words.size(); i++)
3    {
4    memcpy((void *) out, (void *) &d_sync_words[i][0],
                          sizeof(gr_complex) * d_fft_len);
5    out += d_fft_len;
6    }
7  // Allocates information/pilot symbols and guard intervals
8  long n_ofdm_symbols     = 0;
9  int symbols_to_allocate = d_occupied_carriers[0].size();
10 int symbols_allocated   = 0;
11 for (int i = 0; i < ninput_items[0]; i++)
12   {
13   out[(n_ofdm_symbols) * d_fft_len +
         d_occupied_carriers[0][symbols_allocated]] = in[i];
14   symbols_allocated++;
15   if(symbols_allocated==48)
16     {
17     unsigned int symoffset = n_ofdm_symbols * d_fft_len;
18     int piloffset          =
          n_ofdm_symbols % d_pilot_symbols.size();
19     out[symoffset + 0]     = 0;
20     out[symoffset + 1]     = 0;
21     out[symoffset + 2]     = 0;
22     out[symoffset + 3]     = 0;
23     out[symoffset + 4]     = 0;
24     out[symoffset + 5]     = 0;
25     out[symoffset + 11]    =
         d_pilot_symbols[piloffset][0];
26     out[symoffset + 25]    =
         d_pilot_symbols[piloffset][1];
27     out[symoffset + 39]    =
         d_pilot_symbols[piloffset][2];
28     out[symoffset + 53]    =
         d_pilot_symbols[piloffset][3];
29     out[symoffset + 32]    = 0;
30     out[symoffset + 59]    = 0;
31     out[symoffset + 60]    = 0;
32     out[symoffset + 61]    = 0;
33     out[symoffset + 62]    = 0;
34     out[symoffset + 63]    = 0;
35     symbols_allocated      = 0;
36     n_ofdm_symbols++;
37     }
38   }
```

VITERBI Decoder (Proposed implementation)

```
0  /** Trellis transition - main computation **/
1  // The code operates on 4 symbols at a time
2  void viterbi_butterfly2_neon(unsigned char *symbols,
     uint8x16_t *mm0, uint8x16_t *mm1,
       uint8x16_t *pp0, uint8x16_t *pp1)
3  {
4  // Definition of neon register for the computation
5  int i;
6  uint8x16_t *metric0, *metric1, *path0, *path1,
   decision0, decision1, survivor0, survivor1,
   m0, m1, m2, m3, metsv, metsvm, sym0v, sym1v, tmp0, tmp1;
7  uint16x8_t shift0, shift1;

8  // They are the paths and metric related to the
9  // previous and actual state
10 metric0 = mm0;
11 metric1 = mm1;
12 path0   = pp0;
13 path1   = pp1;

14 // Copy first two symbols on neon registers
15 sym0v = vdupq_n_u8(symbols[0]);
16 sym1v = vdupq_n_u8(symbols[1]);

17 // compute all possible brunch metric of
18 // trellis for the first two symbols, 0 and 1
19 for (i = 0; i < 2; i++)
20  {
21  if (symbols[0] == 2)
22   {
23   metsvm = veorq_u8(d_branchtab27_neon[1].v[i],sym1v);
24   metsv = vsubq_u8(vdupq_n_u8(1),metsvm);
25   }
26  else if (symbols[1] == 2)
27   {
28   metsvm = veorq_u8(d_branchtab27_neon[0].v[i],sym0v);
29   metsv = vsubq_u8(vdupq_n_u8(1),metsvm);
30   }
31  else
32   {
33   metsvm = vaddq_u8(
        veorq_u8(d_branchtab27_neon[0].v[i],sym0v),
          veorq_u8(d_branchtab27_neon[1].v[i],sym1v));
34   metsv = vsubq_u8(vdupq_n_u8(2),metsvm);
35   }
36
37  // Each symbol generates 2 output metrics
38  m0 = vaddq_u8(metric0[i], metsv);
39  m1 = vaddq_u8(metric0[i+2], metsvm);
40  m2 = vaddq_u8(metric0[i], metsvm);
41  m3 = vaddq_u8(metric0[i+2], metsv);

42  // Decision metric1 - metric0 > 0 and so on...
43  decision0 = vcgtq_s8(
       vsubq_s8(vreinterpretq_s8_u8(m0),
          vreinterpretq_s8_u8(m1)), vdupq_n_s8(0));
44  decision1 = vcgtq_s8(
       vsubq_s8(vreinterpretq_s8_u8(m2),
          vreinterpretq_s8_u8(m3)), vdupq_n_s8(0));

45  // Best metric selection m1 or m0...
46  survivor0 = vorrq_u8(
```

```
        vandq_u8(decision0,m0),
           vandq_u8( vmvnq_u8(decision0),m1));
47   survivor1 = vorrq_u8(
        vandq_u8(decision1,m2),
           vandq_u8( vmvnq_u8(decision1),m3));

48   // Update the metric and the path
49   // based on the new survivors
50        uint8x16x2_t interleave =
        vzipq_u8(survivor0,survivor1);
51   metric1[2*i]= interleave.val[0];
52   metric1[2*i+1]= interleave.val[1];

53   shift0 =  vshlq_n_u16(
        vreinterpretq_u16_u8(path0[i]), 1);
54   shift1 =  vshlq_n_u16(
        vreinterpretq_u16_u8(path0[2+i]), 1);
55   uint8x16_t temp_shift0 =
        vreinterpretq_u8_u16(shift0);
56   uint8x16_t temp_shift1 =
        vaddq_u8(vreinterpretq_u8_u16(shift1), vdupq_n_u8(1));

57   tmp0 = vorrq_u8( vandq_u8(decision0,temp_shift0),
        vandq_u8(vmvnq_u8(decision0),temp_shift1));
58   tmp1 = vorrq_u8( vandq_u8(decision1,temp_shift0),
        vandq_u8(vmvnq_u8(decision1),temp_shift1));

59   uint8x16x2_t interleave2 = vzipq_u8(tmp0, tmp1);
60   path1[2*i]= interleave2.val[0];
61   path1[2*i+1]= interleave2.val[1];
62   }

63   // Swap current and next states
64   metric0 = mm1;
65   metric1 = mm0;
66   path1   = pp0;
67   path0   = pp1;

68   // copy second two symbols on neon registers
69   sym0v = vdupq_n_u8(symbols[2]);
70   sym1v = vdupq_n_u8(symbols[3]);

71   // compute all possible brunch metric of
72   // trellis for the first two symbols, 2 and 3
73   for (i = 0; i < 2; i++)
74    {
75    if (symbols[2] == 2)
76     {
77     metsvm = veorq_u8(
           d_branchtab27_neon[1].v[i], sym1v);
78     metsv = vsubq_u8(vdupq_n_u8(1),metsvm);
79     }
80    else if (symbols[3] == 2)
81     {
82     metsvm =  veorq_u8(
           d_branchtab27_neon[0].v[i], sym0v);
83     metsv = vsubq_u8(vdupq_n_u8(1),metsvm);
84     }
85    else
86     {
87     metsvm = vaddq_u8(
           veorq_u8(d_branchtab27_neon[0].v[i],sym0v),
              veorq_u8(d_branchtab27_neon[1].v[i],sym1v));
88     metsv = vsubq_u8(vdupq_n_u8(2),metsvm);
89         }

90    // Each symbol generates 2 output metrics
91    m0 = vaddq_u8(metric0[i], metsv);
92    m1 = vaddq_u8(metric0[i+2], metsvm);
93    m2 = vaddq_u8(metric0[i], metsvm);
94    m3 = vaddq_u8(metric0[i+2], metsv);

95    // Decision metric1 - metric0 > 0 and so on...
96    decision0 = vcgtq_s8( vsubq_s8(vreinterpretq_s8_u8(m0),
        vreinterpretq_s8_u8(m1)), vdupq_n_s8(0));
97    decision1 = vcgtq_s8( vsubq_s8(vreinterpretq_s8_u8(m2),
        vreinterpretq_s8_u8(m3)), vdupq_n_s8(0));

98    // Best metric selection m1 or m0...
99    survivor0 = vorrq_u8( vandq_u8(decision0,m0),
        vandq_u8( vmvnq_u8(decision0),m1));
100   survivor1 = vorrq_u8( vandq_u8(decision1,m2),
        vandq_u8( vmvnq_u8(decision1),m3));

101   // Update the metric and the path
102   // based on the new survivors
103   uint8x16x2_t interleave =
        vzipq_u8(survivor0, survivor1);
104   metric1[2*i]=interleave.val[0];
105   metric1[2*i+1]=interleave.val[1];

106   shift0 =  vshlq_n_u16(
        vreinterpretq_u16_u8(path0[i]), 1);
107   shift1 =  vshlq_n_u16(
        vreinterpretq_u16_u8(path0[2+i]), 1);
108   uint8x16_t temp_shift0 = vreinterpretq_u8_u16(shift0);
109   uint8x16_t temp_shift1 = vaddq_u8(
       vreinterpretq_u8_u16(shift1), vdupq_n_u8(1));

110   tmp0 = vorrq_u8( vandq_u8(decision0,temp_shift0),
        vandq_u8(vmvnq_u8(decision0),temp_shift1));
111   tmp1 = vorrq_u8( vandq_u8(decision1,temp_shift0),
        vandq_u8(vmvnq_u8(decision1),temp_shift1));

112   uint8x16x2_t interleave2 = vzipq_u8(tmp0, tmp1);
113   path1[2*i]=interleave2.val[0];
114   path1[2*i+1]=interleave2.val[1];
115   }
116  return;
117  }
```

Reconfigurable antenna system for wireless applications

Reconfigurable Antenna System for Wireless Applications

S. Ciccia[2], G. Giordanengo[1,2], S. Arianos[1], F. Renga[1], P.Ruiu[1], A.Scionti[1], L.Mossucca[1], O. Terzo[1] and G. Vecchi[2]
1 Istituto Superiore Mario Boella (ISMB), Via Pier Carlo Boggio 61, 10138 Torino, Italy,
Email: {giordanengo, arianos, renga, ruiu, scionti, mossucca, terzo}@ismb.it,
2 Antenna and EMC Lab (LACE), Politecnico di Torino, C.so Duca degli Abruzzi 24, 10129 Torino, Italy
Email: {simone.ciccia, giuseppe.vecchi}@polito.it

***Abstract*— Nowadays wireless applications are widespread, and the demand for smart antenna technology grows exponentially. Although there are a large variety of effective algorithms to control antennas, they lack in the requirements of the next generation smart devices for industrial and societal applications which demand integration in compact, low cost and low power architectures. In this work we present a Wi-Fi device coupled with an antenna, where the receiver is able to adapt to a changing signal environment by providing a constant and reliable connectivity. Design criterion follows a strong low power approach, single front-end USB powered connected to an embedded low energy consumption platform. In addition, we come up in a low cost solution that does not require any external hardware. Decoding and antenna control algorithm are software-defined, while antenna beam steering is obtained by means of varactor diodes, voltage biased, used in a suitable way through hybrid couplers in phase shift configuration.**

***Index Terms*—Smart Antennas, Adaptive Antennas, Wi-Fi, Phased Array, Software Defined Radio**

I. INTRODUCTION

Reconfigurable Antennas (RA) represent promising communication systems able to ameliorate the reliability and quality of a radio link. In this context, antenna re-configurability is the key element since, the cutting-edge semiconductor technology (e.g. Micro Electro-Mechanical Systems MEMS), capable to eliminate the mechanical parts mainly used in the past, it allows the realization of low-cost and low-power systems. These types of antennas can realize frequency, polarization and/or pattern reconfigurability, enabling radio systems to accommodate different communication standards and handle channel conditions [1]-[5].

RA generally refers to a system composed by an antenna connected to a signal processor that attempts to adjust or adapt, for example its radiation pattern, in order to select or isolate the signal of interest. Such adaptive systems could be generalized in:

- Switched Beam systems, that provide fixed beam patterns, selected by means of the parameters to optimize (e.g. Signal to Noise Ratio SNR, Received Signal Strength Indication RSSI, to name a few) [6].

- Beam-Formed Adaptive systems in which the beam can be steered in any direction of interest while simultaneously nulling interfering signals [7]-[8].

In this work RA technology coupled with Software Defined Radio (SDR) context is introduced. The proposed receiver is able to reconfigure itself by searching the signal of interest, including packet decoding for the 802.11 standard [9]. However, it is worth noting that communication standard is not a limitation, since the controller system is fully independent from it.

SDR systems demand high processing capabilities, and modern embedded system can provide it. Among several alternatives, processor architectures from Advanced RISC Machine ARM are nowadays widely adopted as a basic-block for many electronic designs, thanks to their performance and small energy footprint. They can operate both with 32- and 64-bit word size, targeting different application scenarios (microcontroller-based, real-time, mobile and server). Applications requiring high performance can leverage many hardware optimizations, such as automatic load balancing among the processor cores, and core heterogeneity. All these features make the ARM-based computing boards suitable for implementing the computing module responsible for controlling the antenna and the base-band code of the receiver and transmitter units.

The paper is organized as follows: a general overview of the wireless communication system is provided, before describing in detail the antenna architecture, the front-end, the low-power computing board, the software environment, and the antenna controller. In section II, an overall overview of the system is given, before entering in the detailed description. Experimental results are discussed in section III, while section IV presents some conclusions and perspectives.

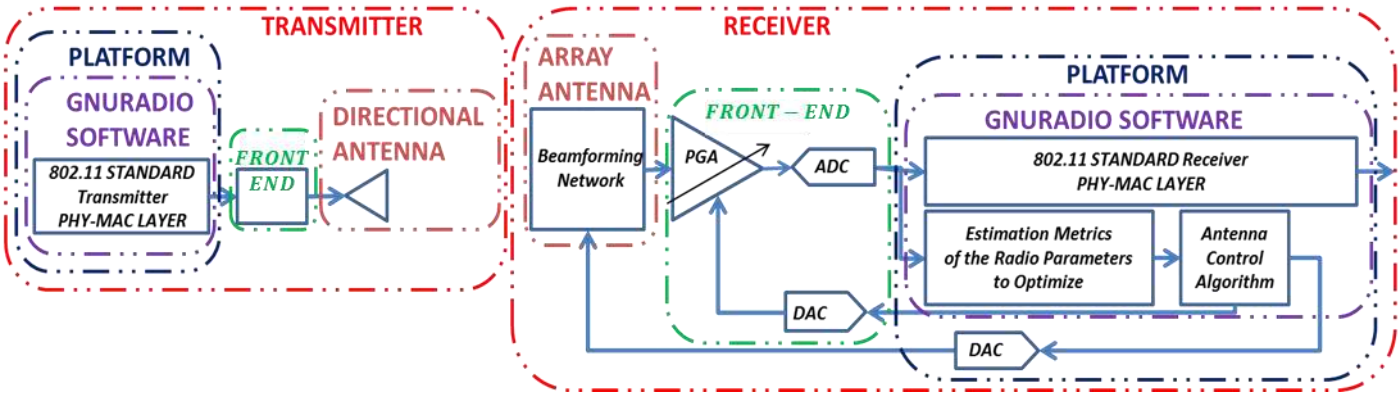

Fig. 1. Scheme of the proposed wireless communication system

II. System Overview

The system is depicted in Fig. 1. On the left side of the figure the transmitter block, with a directive antenna, while on the right side, the reconfigurable receiver whose antenna is an array with beam steering mechanism aimed at searching the signal of interest. Preliminary tests using commercial transceiver provided by Ettus Research [10], USRP – mod. N210, has been performed, while for the final design it has been selected the HACK RF One, a commercial transceiver from Great Scott Gadget [11]. The software (i.e. PHY - MAC layers, antenna control algorithm, transceiver control etc.) and the hardware needed to control the antenna are well integrated in a low power platform, ODROID-XU from HardKernel [12]. This adds simplicity to the system, without any further cost for external hardware.

A. Antenna Architecture

The transmitter antenna is a commercial Horn Antenna, R&S HF 906, operating in the frequency range 1 - 18 GHz, which provides a directive beam [13], while for the receiving part it has been used a planar 4x4 array antenna composed by dipoles, operating in the frequency range 2.4 – 2.48 GHz [14, 15]. By means of 90° hybrid branch couplers loaded with varactor diodes the system is capable to steer the beam in one plane, spanning from -55° to 50°.

B. Transceiver (Front-End)

The Front-End is the Radio Frequency device responsible for the up/down carrier frequency conversion and digitalization of the signal of interest.
As mentioned before, for preliminary tests, it has been used a commercial transceiver, the USRP N210. However, its application is limited to proof-of-concept because this device has the major impact on the overall costs and energy consumptions of the system. To overcome this aspect we changed the front-end with HACK RF One.
Compared to the previous device, this last offers lower power consumption (3W vs. 12W) and costs (300$ vs. 1500$).
The main disadvantage is represented by ADCs resolution that is 8 bits for HACK RF with respect to 14 bits provided by USRP.

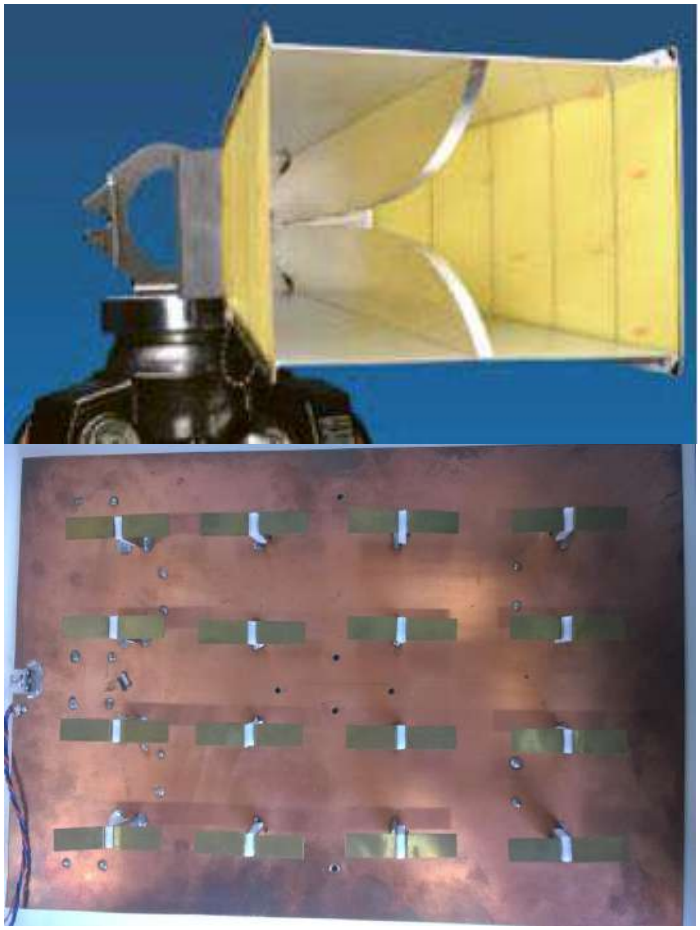

Fig. 2.
R&S HF 906 antenna - transmitter side (top),
Voltage Controlled Steerable Array - receiver side (bottom)

C. Board

Complex algorithms such as those used to control the antenna and the IEEE802.11 baseband code of receiver and transmitter require relatively high computing resources that are typical of a desktop computer. However, standard desktop computers have too high power requirements and thus they cannot be used as a building-block for the proposed system. On the contrary, embedded systems are extremely power efficient computing platform that are progressively evolving into platform with a high degree of parallelism and computing power [16]-[18].

For the implementation of the computing unit of the proposed wireless system, it has been chosen the HardKernel Odroid-XU board. It is a low-cost and low-power computer build upon the Samsung Exynos5 Octa-core. This is a System on Chip equipped with two clusters of 4 cores each, one containing a high-performance ARM Cortex-A15, while the other low power ARM Cortex-A7. The clusters are placed side by side with an embedded Graphical Processing Unit. On top of the processor, a 2 GByte low- power DRAM memory module is directly connected. The board also features a full set of I/O ports (e.g., Ethernet, Universal Serial Bus etc.) that ease the connection of the board to the rest of the communication platform.

Although the performance provided by the board running the antenna control algorithms and IEEE 802.11 base-band code are in line with those of a standard desktop computer, power consumption is far below that of the latter, since the board has a maximum power dissipation of 5W, while we experimentally measured an average consumption of 2,5W in the worst conditions when the software is running.

D. Software Environment

The software environment employed is GNURadio [19] since it is an open-source software development toolkit that provides signal processing blocks to implement software radios.

Furthermore, it provides several blocks, such as the one for the 802.11 signal processing that have been exploited to test the system with a real communication system [20, 21], while keeping the antenna controller independently from it.

E. Antenna Controller

The proposed system (i.e. the receiver, the controller and the antenna) is able to ameliorate the performance of the wireless link by steering the antenna main beam in the direction where the received signal power is stronger. In particular, the software receiver calculates the zero correlation of the incoming base band signal for a certain observation time and for each position of the scan plane (horizontal). Then the software defined antenna controller takes the decision by setting the beam in the angular position in which the received signal power present its maximum.

After the communication is established (i.e. the receiver is able to correctly decode the incoming packets), with the aim of increasing the quality of the received signal, the antenna controller switches to control the gain of the integrated Programmable Gain Amplifier (PGA) by means of an Automatic Gain Control (AGC) circuit. The receiver antenna loses the track of the transmitter when the normalized correlation signal does not overcome a predefined threshold within a certain time interval, meaning that packets are not decoded anymore. At this point, transmitter search restart.

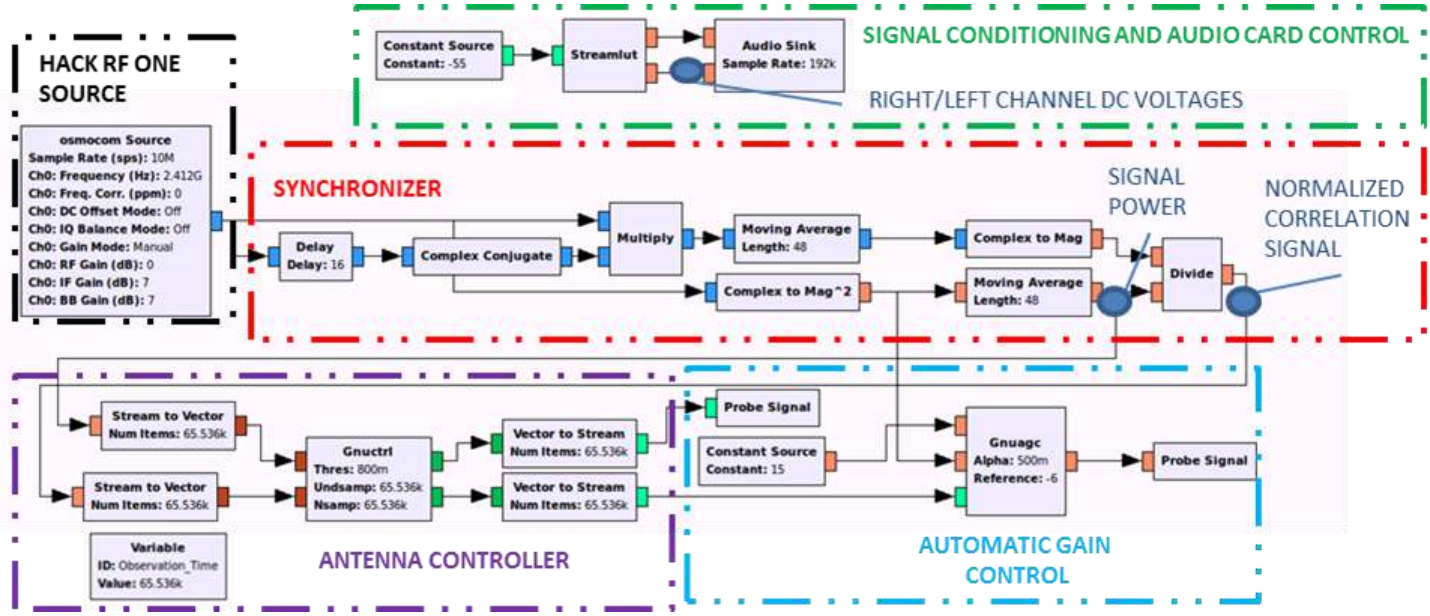

Fig. 3. GNURadio Flowgraph that implements the SDR system (receiver)

From uniform linear array antenna theory, the direction of the beam is controlled by the phase shift among radiating elements. In the case under test, this phase shift is provided by 90° hybrid branch couplers loaded with varactor diodes which require two DC bias voltages. These DC signals are generated by means of Digital to Analog Converters (DACs) integrated in the audio card of the Odroid platform (i.e. right and left audio channels). In this way, with a proper signal conditioning, we developed a fully software controller.
However, the majority of the audio cards available in the market are sold with a bypass capacitor to block the DC component of the signal. There are three methods to overcome this limitation:

1. Using an Odroid-XU board, since the integrated audio card does not have any bypass capacitor

2. Removing physically this capacitance from the hardware

3. Last and more elegant solution, since avoids any operation on the hardware, is the use of an envelope detector to extract the amplitude information of the sinusoidal signal generated by the audio card.

III. Test Results

Tests on the complete system have been performed in laboratory, evaluating the received packets during transmission. Fig.4 shows the test bed.
As a first experiment we alternated on-off transmissions by playing on the transmitted power, and we registered if the receiver system correctly woke up on the signal of interest. With a second experiment, the receiver and/or the transmitter were moved, to analyze the capability of the receiver to track the transmitter.

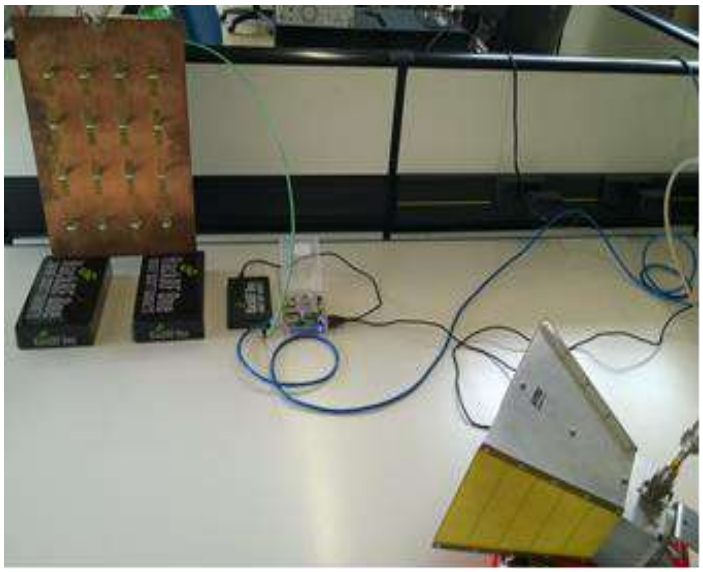

Fig. 4. Test Bed for the trials of the proposed wireless system

In particular, every time the normalized correlation signal does not overcame a predefined threshold within a certain time interval (e.g. two symbol periods), the receiver begins to reconfigure itself, thus optimizing the received power, converging in the direction of the transmitter. Empirical results showed in some cases that the antenna did not point the precise direction of the transmitter. This behavior is mainly due to the multipath environment in which we performed the test. However, we experimented that the longer the observation time the better the positioning, since power estimation in each possible space angle become more accurate.

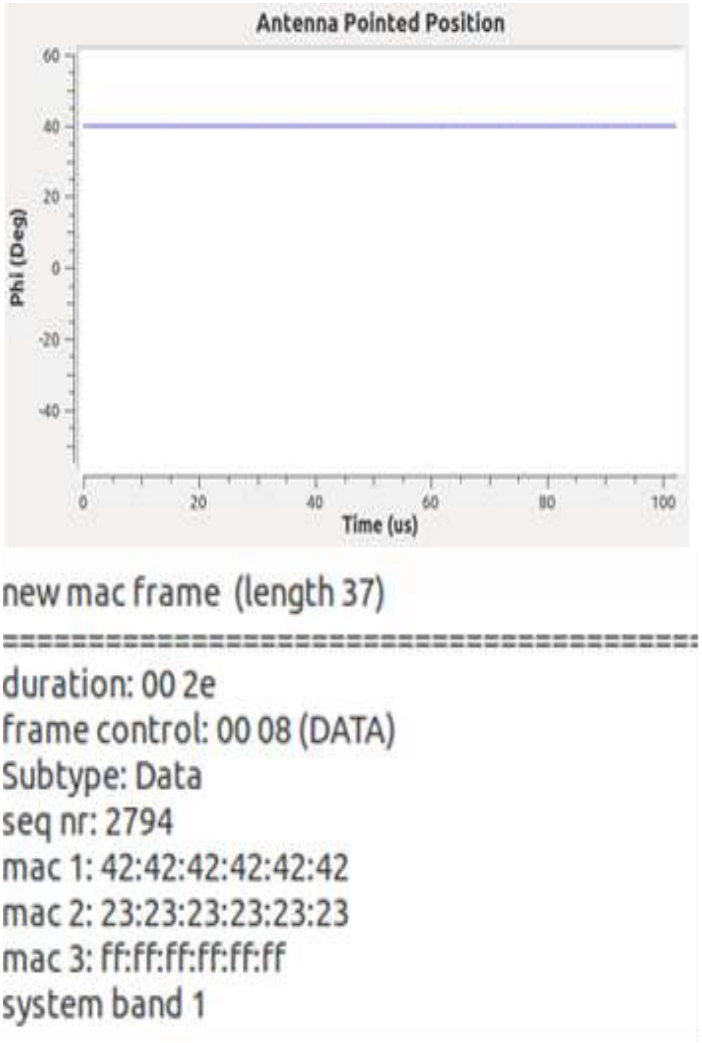

Fig. 5. A screenshot during simulation,
main beam position vs. Time (top),
decoded information (bottom)

In Fig.5 a screenshot of a simulation is reported. In the GUI the plot of the antenna position in the horizontal plane (i.e. by considering polar coordinates, $\vartheta = 90°$ and φ free) in the time domain is depicted. In the first case the transmitter is positioned as reported in Fig.4, while the receive antenna looks in the direction of the transmitter with an angle $\varphi = 40°$. The bottom of the figure shows decoded packets as they appear in the Linux console of the monitored machine.

Fig. 6 illustrates the results when moving the receiver in front of the transmitter, after the receiver antenna get the transmitter track (i.e. φ = 0°), at the top, while mac address and received information at the bottom. These results demonstrate that the proposed communication platform is able to correctly work with standard wireless signals, while keeping low the overall power dissipation and cost. This makes the use of this platform as a good candidate for implementing the communication system in different application scenarios.

IV. Conclusions and Future Works

In this paper we presented a first attempt to integrate a reconfigurable antenna and a software defined radio system, into a single flexible platform. The proposed solution evaluates the quality of the received signal by attempting to maintain a reliable connection by means of antenna beam steering when signal quality degrades. A low-power, low-cost approach have been maintained with the aim of satisfying today requirements. The choice of a conventional array antenna at the receiver side is due to the fact that it was available in laboratory. However, we are approaching challenges more attractive. A low profile dual band antenna with beam steering is currently in the design phase. Future works will focus on a system capable to receives signals over the two 802.11 bands (i.e. 2.4GHz and 5GHz), as well as integrating a software-defined spectrum sensing controller block in the receiver.

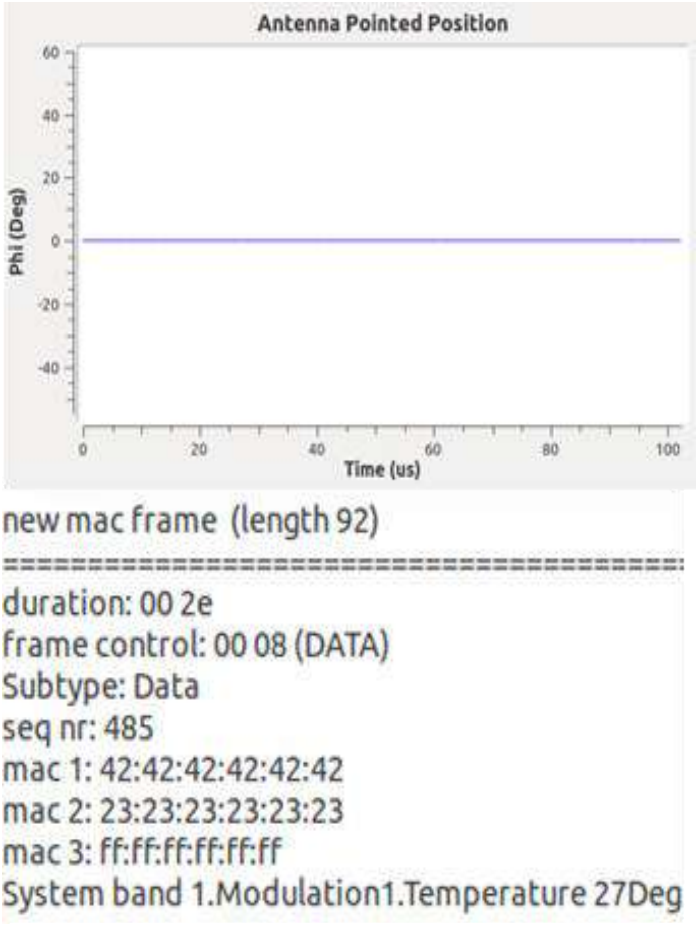

Fig. 6. A second screenshot during simulation, main beam position vs. Time (top), decoded information (bottom)

Acknowledgment

The authors would like to acknowledge the invaluable help of Prof. Mario Orefice and Eng. Gianluca Dassano, from Politecnico di Torino, who provided the steerable antenna.

References

[1] Lima, A.G.M., de Menezes, L.R.A.X., "A multiband fractal antenna array to software defined radio," Microwave and Optoelectronics, SBMO/IEEE MTT-S International Conference, pp. 256 – 258, 2005

[2] Jafargholi, A. ; Emadi, M. ; Esmi, N. ; Mousavi, M. ; Nayebi, M.M., "VHF-LB SDR Phased Array Jamming Resistant Antenna," ITS Telecommunications Proceedings, 6th International Conference, pp. 442 – 445, 2006

[3] Orefice, M. ; Dassano, G.L. ; Matekovits, L. ; Pirinoli, P. ; Vecchi, G. ; Shurvinton, B., "A Wide Coverage Scanning Array for Smart Antennas Applications," Microwave Conference, 31st European, 2001

[4] Arianos, S. ; Quijano, J.L.A. ; Vipiana, F. ; Dassano, G. ; Vecchi, G. ; Orefice, M., "Optimization procedures for the design of reconfigurable compact multi-band antennas," Antennas and Propagation (EUCAP), 6th European Conference, 2012

[5] Araque Quijano, J.L. ; Vecchi, G., "Optimization of a Compact Frequency- and Environment-Reconfigurable Antenna Antennas and Propagation," IEEE Transactions on Volume: 60 , Issue: 6, 2012

[6] Donelli, M. ; Sacchi, C., "Implementation of a low-cost reconfigurable antenna array for SDR-based communication systems," Aerospace Conference, 2012 IEEE, pp. 1 – 7, 2012

[7] Elhag, N.A.A. ; Osman, I.M. ; Yassin, A.A. ; Ahmed, T.B., "Angle of arrival estimation in smart antenna using MUSIC method for wideband wireless communication," Computing, Electrical and Electronics Engineering (ICCEEE), International Conference, pp. 69 – 73, 2013

[8] H. K. Hwang, Zekeriya Aliyazicioglu, Marshall Grice, Anatoly Yakovlev, "Direction of Arrival Estimation using a Root-MUSIC Algorithm," Proceedings of the International MultiConference of Engineers and Computer Scientists 2008 Vol II IMECS 2008, 19-21 March, 2008, Hong Kong

[9] Wireless LAN Medium Access Control (MAC) and Phisical Layer(PHI) Specications. Std 802.11-2012, IEEE, 2012

[10] Ettus Research USRP N210 transceiver, 30/07/2015 http://www.ettus.com/product/details/UN210-KIT

[11] Great Scott Gadgets HACK RF One transceiver, 30/07/2015 https://greatscottgadgets.com/hackrf/

[12] HardKernel ODROID board, 30/07/2015 http://www.hardkernel.com/main/main.php

[13] R&S®HF906 Antenna - Brief description, 26.08.2008 http://cdn.rohde-schwarz.com/pws/dl_downloads/dl_common_library/dl_brochures_and_datasheets/pdf_1/HF906_brief_e.pdf

[14] Orefice, M. ; Dassano, G.L. ; Matekovits, L. ; Pirinoli, P. ; Vecchi, G. ; Shurvinton, B., "A Wide Coverage Scanning Array for Smart Antennas Applications," Microwave Conference, 31st European, 2001

[15] G. Dassano ; M. Orefice, "Voltage Controlled Steerable Array for Wireless Sensors Networks," In: EUCAP 2007 2nd European Conference on Antennas and Propagation (EuCAP 2007), Edinburgh, U.K 11-16 Nov. 2007

[16] Gschwind, Michael and Hofstee, H. Peter and Flachs, Brian and Hopkins, Martin and Watanabe, Synergistic Processing in Cell's Multicore Architecture, IEEE Micro, March 2006

[17] A. Branover, D. Foley, and M. Steinman, AMD Fusion APU: Llano, IEEE Micro, vol. 32, no. 2, 2012, pp. 28–37.

[18] D. Bouvier et al., Kabini: An AMD accelerated processing unit system on a chip, IEEE Micro, April 2014

[19] GNURadio Project, 30/07/2015 http://gnuradio.org/redmine/projects/gnuradio/wiki

[20] Bastian Bloessl, Michele Segata, Christoph Sommer and Falko Dressler, "Decoding IEEE 802.11a/g/p OFDM in Software using GNU Radio," Proceedings of 19th ACM International Conference on Mobile Computing and Networking (MobiCom 2013), Demo Session, Miami, FL, October 2013, pp. 159-161

[21] Bastian Bloessl, Michele Segata, Christoph Sommer and Falko Dressler, "An IEEE 802.11a/g/p OFDM Receiver for GNU Radio," Proceedings of ACM SIGCOMM 2013, 2nd ACM SIGCOMM Workshop of Software Radio Implementation Forum (SRIF 2013), Hong Kong, China, August 2013, pp. 9-16

Low power advanced wireless communication exploiting reconfigurable antennas

Low Power Advanced Wireless Communication Exploiting Reconfigurable Antennas

S. Ciccia[2], G. Giordanengo[1,2], F. Renga[1], G. Vecchi[2]
[1]Istituto Superiore Mario Boella (ISMB), Via Pier Carlo Boggio 61, 10138 Torino, Italy,
Email: $\{giordanengo, renga\}$@ismb.it
[2]Antenna and EMC Lab (LACE), Politecnico di Torino, C.so Duca degli Abruzzi 24, 10129 Torino, Italy,
Email: $\{simone.ciccia, giuseppe.vecchi\}$@polito.it

***Abstract*—In this work antenna re-configurability is realized via Software Defined Radio (SDR) by implementing a phase controller for array antenna in a flexible low power General Purpose Platform (GPP) with single Analog to Digital Converter (ADC). The high software flexibility approach of the system represents a promising technique for the newer communication standards and could be applied to Wireless Sensor Network (WSN).**

***Index Terms*—Reconfigurable Antennas, Wireless, Phased Array, Software Defined Radio.**

I. Introduction

Concurrently to the evolution of wireless standards 802.11b/g/n (Wi-Fi), reconfigurable antennas have attracted attentions in optimizing Signal to Noise Ratio (SNR), reducing interferers by directing the beam towards terminals and/or access points, with the purpose of increasing link capacity, saving battery life and reducing terminals costs [1-4]. Recently, Baseband Beamforming (BB) and Multi Input Multi Output (MIMO) are the most widely used technologies, but in low power communication there is no optimal choice since these systems rely on multiple front-end and powerful Digital Signal Processors (DSPs) [5-8]. This work would like to demonstrate that high quality communication link is possible with a simple array antenna controlled by software defined radio receiver able to direct the main beam in the direction of the transmitter, ensuring that the global energy consumption of the link is still lower than the classical approach realized with a traditional on-chip Wi-Fi module. Although SDR receiver implementation shows a higher energy consumption with respect these commercial modules, the introduced system is able to search and direct the antenna main beam toward the transmitter via software and then switch to the standard receiving mode with the purpose of reducing energy consumption. This is possible because the higher SNR provided by the array antenna allows to reduce the transmitted power, parameter that weighs on the global energy consumption of the link. This application could improve battery life and transmission speed, crucial parameters in Wireless Sensor Networks [9,10]. The same discussion could achieve benefits on the sensors distribution. By establishing the transmitted power, the proposed system could increase the distance among sensors without any signal repeater. With respect MIMO, it could be employed in low multipath environment, where the attenuation provided by secondary lobes in non-Line of Sight (LOS) direction falls under acceptable limits. The rest of the paper is organized as follows: section II exposes a general overview of the designed system, section III includes simulation results of the idea, section IV concludes discussing the test on the real system.

II. System Overview

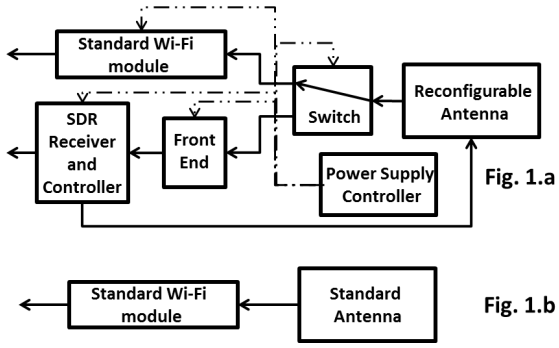

Fig. 1. a) Proposed system, b) Conventional Wi-Fi module flow-graphs

In Fig.1.a the reconfigurable antenna is a planar $4x4$ array composed by dipoles, operating in the frequency range $2.42.48GHz$ with $17dBi$ gain [11,12]. By means of $90°$ hybrid branch couplers loaded with varactor diodes the system is capable to steer the main beam in the horizontal plane, spanning from $-55°$ to $50°$. A switch isolates the software receiver from the generic Wi-Fi module, refer to Fig. 2.b, in order to select the path of interest. In steering mode the antenna is connected to the SDR receiver composed by the commercial transceiver Hack RF One and the software module that run on a GPP, an ODROID-XU for the tests, refer to Fig. 2.a, which controls the main beam by providing voltages to varactor diodes through the integrated audio card, and runs the software receiver to decode the received information [13-15]. When the software module locks the transmitter an external power supply controller enables the standard Wi-Fi reception, by turning off the SDR module. If a new steering is needed, power controller turns off Wi-Fi module, turns on the SDR module and so on. As a reference system, the same standard

Wi-Fi module is used, equipped with a half wave dipole, $2.3dBi$ gain, refer to Fig. 2.b.

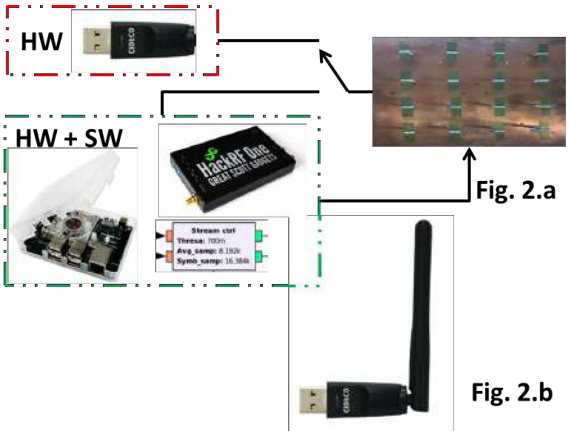

Fig. 2. a) Proposed system, b) Conventional Wi-Fi module flow-graphs

III. Simulation Results

A first simulation has been performed to evaluates the situation of the link by employing the reference system and the proposed system at parity of link distance ($4700m$) and with a required SNR of $20dB$, as reported in Fig. 3. In LOS, thanks to the difference in antenna gains, our system needs a much lower transmitted power with respect to the reference module.

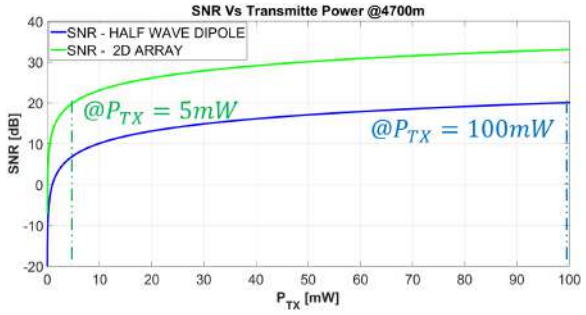

Fig. 3. SNR at the input terminal (antenna noise temperature is included in the computation)

Becomes clear that the transmitted power reduction could compensate the energy imbalance caused by SDR steering. In Fig. 4 an estimation of the energy consumption in one day for the reference and target systems are presented, considering that the Standard Wi-Fi module has a consumption of about $1.8W$ [16], while the SDR module has an estimated value of $8W$ [14,15].

An analysis on this figure shows that energy consumption of the proposed system increase linearly with how many times it searches the transmitter. Depending on the scan metrics (e.g. number of averaged power samples per steered position), in this test the time needed for the search was set to $17ms$.

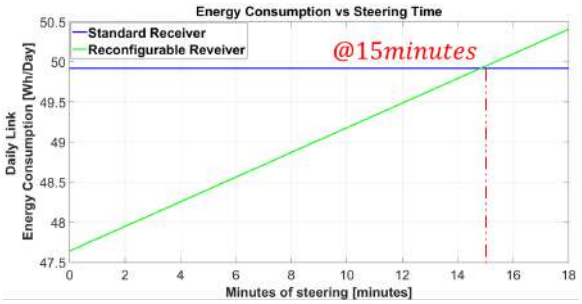

Fig. 4. Energy consumption in one day in function of the minutes spent in steering

The break-even point for this configuration happens when steering is applied for $15min$. This means that in one day our system could perform 52941 searches before reaching to the same energy consumption of a standard Wi-Fi module. This information is set on the power supply controller that switches configurations with a predefined timing.

A. Realized System and Conclusions

The author wants to conclude the paper showing a real test on antenna re-configurability. Fig. 5 presents the scan and lock mechanism used for the tests. When SDR module is enabled, the antenna controller algorithm computes for some angles of the horizontal plane the average power of the received signal, maintaining the position in which the mean power is maximum. The lock condition indicates to the power supply controller to change the actual configuration. There are still possibility to reduce the energy consumption of the system. First of all the computational cost of the SDR receiver could be optimized, second the time needed for the transmitter search could be reduced. However improvements go out from the scope of this paper that want to leave a proof-of-concept. Concluding this section, the proposed system shows an ad-hoc solution able to provide high capacity link at a reasonable energy consumption.

Acknowledgment

The authors would like to acknowledge the invaluable help of Prof. Mario Orefice and Eng. Gianluca Dassano, from Politecnico di Torino, who provided the steerable antenna.

References

[1] Guo, Y.J., Pei-Yuan Qin, “Advances in reconfigurable antennas for wireless communications”, Antennas and Propagation (EuCAP), 2015

[2] Dau-Chyrh Chang, Cheng-Nan Hu, “Smart Antennas for Advanced Communication Systems”, Proceedings of the IEEE, Year: 2012, Volume: 100, Issue: 7, Pages: 2233 - 2249

[3] Shiann Shiun Jeng, Chen Wan Tsung, “Performance Evaluation of IEEE 802.11g with Smart Antenna System in the Presence of Bluetooth Interference Environment”, Vehicular Technology Conference (VCT), 2007

[4] Wenjiang Wang, S. Krishnan, Khoon Seong Lim, Aigang Feng, Boonpoh Ng, “A simple beamforming network for 802.11b/g WLAN systems”, 11th IEEE Singapore International Conference on Communication Systems (ICCS), 2008

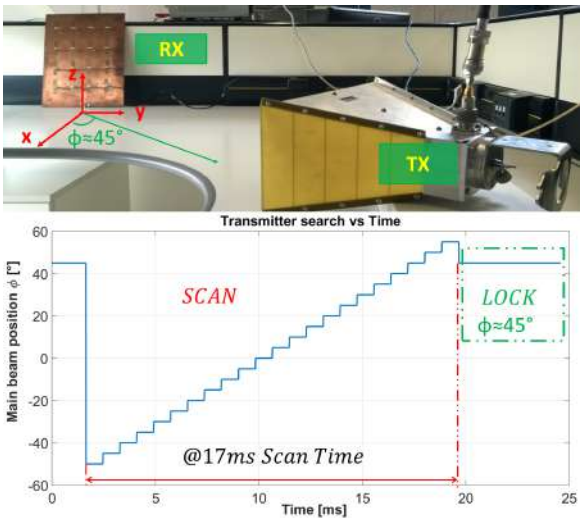

Fig. 5. Re-configurability Test and Transmitter Lock

[5] M. Uthansakul, P. Uthansakul, "Experiments with a Low-Profile Beamforming MIMO System for WLAN Applications", Antennas and Propagation Magazine, IEEE Year: 2011, Volume: 53, Issue: 6, Pages: 56 - 69

[6] A. Hakkarainen, J. Werner, K.R. Dandekar, M. Valkama, "Widely-linear beamforming and RF impairment suppression in massive antenna arrays", Journal of Communications and Networks, IEEE Year: 2013, Volume: 15, Issue: 4, Pages: 383 - 397,

[7] A.S. Prasad, S. Vasudevan, R. Selvalakshmi, K.S. Ram, G. Subhashini, S. Sujitha, B.S. Narayanan, "Analysis of adaptive algorithms for digital beamforming in Smart Antennas", International Conference on Recent Trends in Information Technology (ICRTIT), 2011

[8] Jo Kyung-Tae, Ko Young-chai, Yang Hong-Chuan, "RF beamforming considering RF characteristics in MIMO system", International Conference on Information and Communication Technology Convergence (ICTC), 2010

[9] R. Shelke, G. Kulkarni, R. Sutar, P. Bhore, D. Nilesh, S. Belsare, "Energy Management in Wireless Sensor Network", 15th International Conference on Computer Modelling and Simulation (UKSim), 2013 ,

[10] David Tolar, Milan Stork,"Acknowledged data transmissions in wireless sensor network with optimized sensor energy consumption", International Conference on Applied Electronics (AE), 2015

[11] M. Orefice, G.L. Dassano, L. Matekovits, P. Pirinoli, G. Vecchi, B. Shurvinton, "A Wide Coverage Scanning Array for Smart Antennas Applications", Microwave Conference, 31st European, 2001

[12] G. Dassano ; M. Orefice, "Voltage Controlled Steerable Array for Wireless Sensors Networks", 2nd European Conference on Antennas and Propagation (EuCAP 2007), Edinburgh, U.K 11-16 Nov. 2007

[13] Bastian Bloessl, Michele Segata, Christoph Sommer and Falko Dressler, "Decoding IEEE 802.11a/g/p OFDM in Software using GNU Radio", Proceedings of 19th ACM International Conference on Mobile Computing and Networking (MobiCom 2013), Demo Session, Miami, FL, October 2013, pp. 159-161

[14] *https://greatscottgadgets.com/hackrf/*, Great Scott Gadgets HACK RF One transceiver, 12/01/2016

[15] *http://www.hardkernel.com/main/main.php*, HardKernel ODROID board, 12/01/2016

[16] *http://www.datasheet4u.com/datasheet/*, RTL8188CUS-GR Wi-Fi module, 12/01/2016

Software-defined reconfigurable antenna for energy efficient wireless links

Software-defined Reconfigurable Antenna for Energy Efficient Wireless Links

Simone Ciccia, Giuseppe Vecchi
Antenna and EMC Lab (LACE), Politecnico di Torino,
C.so Duca degli Abruzzi 24, 10129 Torino, Italy
Email: {$simone.ciccia, giuseppe.vecchi$}@polito.it

Giorgio Giordanengo, Flavio Renga
Istituto Superiore Mario Boella (ISMB),
Via Pier Carlo Boggio 61, 10138 Torino, Italy,
Email: {$giordanengo, renga$}@ismb.it

***Abstract*—The ability to reconfigure an antenna can also allow energy saving in autonomous system as, e.g. the nodes of a deployable wireless sensor network. This option appears not yet fully exploited. Software-defined radio (SDR) plays an important role in this scenario. This paper shows how a simple antenna with reconfigurable beam can reduce the energy consumption of wireless links (WI-FI), while preserving channel capacity and reducing interferers.**

I. Introduction

The advent of reconfigurable antennas has stimulated a vastness of studies and researches on the employments of a single antenna to perform a multitude of tasks. These types of antennas can realize frequency, polarization and/or pattern re-configurability, making radio systems able to accommodate different communication standards and handle channel conditions [1, 2]. Currently, Baseband Beamforming is the most widely used technique employed at the base stations to generate customized patterns for the users that require service. Although this method provides high performance in interference suppression and channel capacity it relies on multi front-end structures (i.e. one transceiver for each radiating element), with an array processing method loaded into powerful Digital Signal Processors and Field Programmable Gate Array, leading to high computational complexity and unacceptable energy consumption if applied to user terminal [3-5]. This work would like to demonstrate that taking advantage of a steerable beam antenna, which permits to radiate the overall useful energy in the Line Of Sight LOS direction of a communication link; it is possible to realize a high capacity channel with interferences attenuated to a level defined by the secondary lobes of the array. This concept is applied to the user terminal, making it accessible and feasible for civilian/general applications. The rest of the paper is organized as follows: section II exposes a general overview of the designed system highlighting the method used to achieve minimum energy consumption; section III includes the results obtained by means of simulation, while section IV concludes discussing the re-configurability tests.

II. System Overview

The reference system is composed by a standard half wave dipole, $2.3dBi$ gain, connected to a hardware decoder chip [6], as illustrated in figure 1.a. Instead, the proposed system is designed to work with a reconfigurable antenna made up of $4x4$ dipoles, operating in the frequency range $2.4 - 2.48GHz$ with $17dBi$ gain. By means of $90°$ hybrid couplers loaded with varactor diodes, controlled by means of voltages, the antenna main beam can be steered in the horizontal plane, spanning from $-55°$ to $50°$ [7].

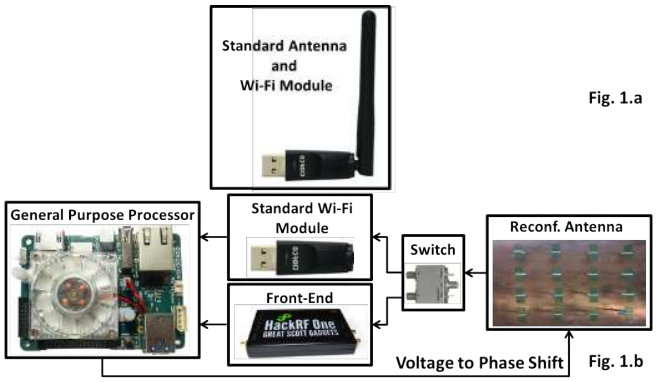

Fig. 1. a) Conventional Wi-Fi module; b) Proposed system

As depicted in Fig. 1.b, the array is connected by means of an RF switch that isolates the Software Defined Radio SDR (i.e. transceiver and the software receiver loaded in the General Purpose Processor GPP) from the conventional hardware receiver, due to the fact that the SDR has an energy consumption greater than the conventional module, therefore it is only used for steering the main beam of the array antenna toward the transmitter, and then is shut down. The SDR is composed by the commercial transceiver Hack RF One connected to the GPP, an ODROID-XU4, which runs the software program to decode the received information and steer the main beam of the antenna, by providing the bias voltages to the beam forming network through the integrated audio card [8, 9]. When the SDR detects the transmitter position, call it lock condition, an external power supply controller enables the standard Wi-Fi reception, by turning off the SDR module. If a new steering is needed, power controller turns off Wi-Fi module, turns on the SDR module and so on.

III. Energy Consumption - Simulation Results

A first simulation has been performed to evaluate the situation of the link by employing the reference and the

proposed system at parity of link distance ($4700m$), with a required SNR of $20dB$. In LOS, thanks to the difference in the gains of the antennas, the reconfigurable system needs a lower transmitted power with respect to the conventional module. The link budget computation reveals that a standard wireless link requires a transmitted power of $100mW$, while with the proposed array only $5mW$ are neededd. Becomes clear that the transmitted power reduction could compensate the energy imbalance caused by SDR steering. In Fig.2 an estimation of the energy consumption in one day for the reference and target systems are presented, considering that the Standard Wi-Fi module has a consumption of about $1.8W$ [7], while the SDR module has an estimated value of $8W$ [9, 10].

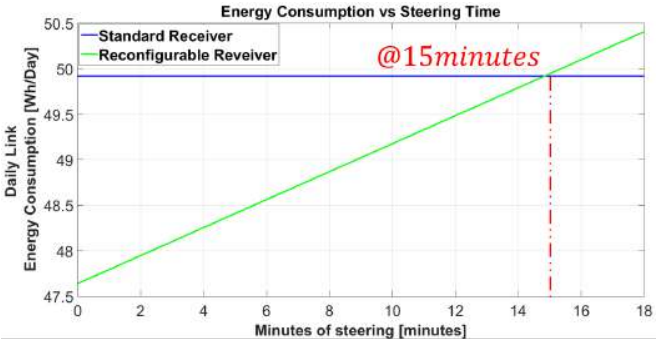

Fig. 2. Energy consumption in one day in function of the minutes spent in steering

As Fig.2 suggests, the energy consumption of the proposed system increase linearly with how many times the scan is applied. Depending on the metric (e.g. number of averaged power samples per steered position), this configuration requires $17ms$ to lock the transmitter. The break-even point happens when steering is applied for $15min$. This means that in one day the SDR could perform 52941 scan before reaching to the same energy consumption of a standard Wi-Fi module. This information is set on the power supply controller that switches configurations with a predefined timing. A note of interest, imagine the advantages on the energy consumption that could be achieved in the case of stationary link, where transmitter search is applied one time per day.

IV. Realized System and Conclusions

This section shows the re-configurability tests through Fig.3 that presents the scan and lock mechanism employed. When the SDR is enabled, the antenna controller algorithm computes for some angles of the horizontal plane the average power of the received signal, maintaining the position in which the mean power is maximum. The lock condition indicates to the power supply controller to change the actual configuration. There are still possibility to reduce the energy consumption of the reconfigurable system. First of all the computational cost of the software-defined receiver could be optimized, second the time needed for the transmitter search could be reduced. However improvements go out from the scope of this paper that want to leave a proof-of-concept. Concluding this section, the proposed system shows an ad-hoc solution able to provide high capacity link at a reasonable energy consumption.

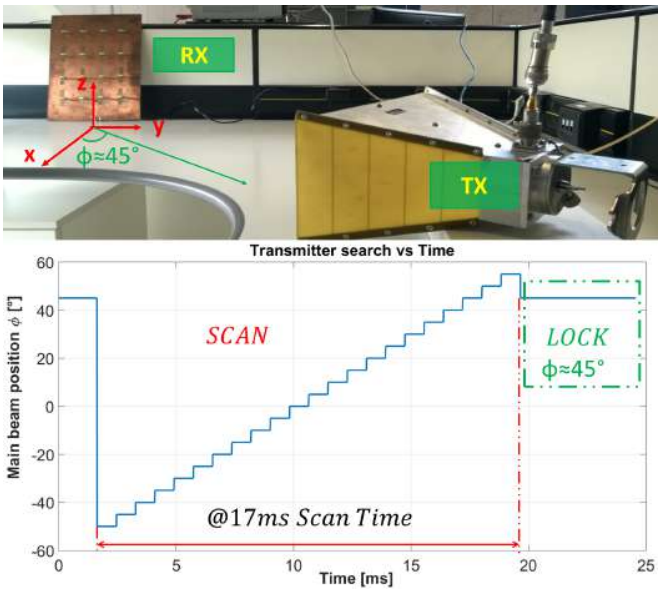

Fig. 3. Re-configurability Test and Transmitter Lock

ACKNOWLEDGEMENT

The authors would like to acknowledge the invaluable help of Prof. Mario Orefice and Eng. Gianluca Dassano, from Politecnico di Torino, who provided the steerable antenna.

References

[1] M. Orefice, G.L. Dassano, L. Matekovits, P. Pirinoli, G. Vecchi, B. Shurvinton, "A Wide Coverage Scanning Array for Smart Antennas Applications," Microwave Conference, 31st European, 2001

[2] J.L. Araque Quijano, G. Vecchi, "Optimization of a Compact Frequency- and Environment-Reconfigurable Antenna," IEEE Transactions on Antennas and Propagation, 2012

[3] M. Uthansakul, P. Uthansakul, "Experiments with a Low-Profile Beamforming MIMO System for WLAN Applications," Antennas and Propagation Magazine, IEEE Journals and Magazines, 2011

[4] A. Hakkarainen, J. Werner, K.R. Dandekar, M. Valkama, "Widely-linear beamforming and RF impairment suppression in massive antenna arrays," Journal of Communications and Networks, IEEE Journals and Magazines, 2013

[5] A.S. Prasad, S. Vasudevan, R. Selvalakshmi, K.S. Ram, G. Subhashini, S. Sujitha, B.S. Narayanan, "Analysis of adaptive algorithms for digital beamforming in Smart Antennas," International Conference on Recent Trends in Information Technology (ICRTIT), 2011

[6] *http://www.datasheet4u.com/datasheet/*, RTL8188CUS-GR Wi-Fi module, 18/01/2016

[7] G. Dassano, M. Orefice, "Voltage Controlled Steerable Array for Wireless Sensors Networks," 2nd European Conference on Antennas and Propagation EuCAP, Edinburgh, U.K, Nov. 2007

[8] *https://greatscottgadgets.com/hackrf/*, Great Scott Gadgets HACK RF One transceiver, 18/01/2016

[9] *http://www.hardkernel.com/main/main.php*, HardKernel ODROID board, 18/01/2016

Low power computning and communication system for critical environments

Low Power Computing and Communication System for Critical Environments

Luca Pilosu*, Lorenzo Mossucca*, Alberto Scionti*, Giorgio Giordanengo*,
Flavio Renga*, Pietro Ruiu*, Olivier Terzo*, Simone Ciccia†, Giuseppe Vecchi†
*Istituto Superiore Mario Boella (ISMB), Torino, Italy
†Department of Electronics and Telecommunication (DET), Politecnico di Torino, Torino, Italy
E-mails: *{pilosu, mossucca, scionti, giordanengo, renga, ruiu, terzo}@ismb.it †{simone.ciccia, giuseppe.vecchi}@polito.it

***Abstract*—The necessity of managing acquisition instruments installed in remote areas (e.g., polar regions), far away from the main buildings of the permanent observatories, provides the perfect test-case for exploiting the use of low power computing and communication systems. Such systems are powered by renewable energy sources and coupled with reconfigurable antennas that allow radio-communication capabilities with low energy requirements. The antenna reconfiguration is performed via Software Defined Radio (SDR) functionalities by implementing a phase controller for the array antenna in a flexible low power General Purpose Platform (GPP), with a single Front-End (FE). The high software flexibility approach of the system represents a promising technique for the newer communication standards and could be also applied to Wireless Sensor Networks (WSNs).**

This paper presents the prototype that is devoted to ionospheric analysis and that will be tested in Antarctica, in the Italian base called Mario Zucchelli Station, during summer campaigns. The system, developed to guarantee its functionality in critical environmental conditions, is composed of three main blocks: Communication, Computing and Power supply.
Furthermore, the computing and communication system has been designed to take into account the harsh environmental conditions of the deployment site.

***Keywords*—Low power, Reconfigurable Antennas, Wi-Fi, Software Defined Radio, Linux Container**

I. Introduction

Low power computing and communication technologies have recently gained a key role for helping the scientific community to improve the capability of processing and exchanging data, while keeping low the impact on power consumption. Low power computing inherits most of the challenges faced in managing large computing infrastructures (i.e., optimizing the resource allocation and their utilization, scalability and ease of use), putting the stress on a coordinated operation of several low power units aiming at obtaining the highest processing capabilities. From this perspective, embedded computer domain (which comprises low power computing) is moving from traditional low performance platforms to more powerful, highly parallel, general purpose systems. The availability of small computer boards equipped with full chipsets and I/O interfaces represents an inexpensive opportunity to explore parallel computer systems in many critical application scenarios. Heterogeneity of hardware modules largely contributes to the adoption of such platforms, since heterogeneity in the form of hardware accelerators (e.g., GPUs embedded in the same System-on-Chip along with CPUs) allow to take advantage from specialization without defeating flexibility and adaptability of the whole system. Finally, the relative low price of such platforms further extends their attractiveness.

On the other hand, low power communication aims at obtaining excellent data transmission features while remaining compliant to the standard protocols (e.g., IEEE 802.11) and keeping the power consumption at acceptable limits.

Concurrently to the evolution of wireless standards 802.11 b/g/n, reconfigurable antennas have gained a lot of attention for optimizing Signal to Noise Ratio (SNR), and reducing interferences by directing the beam towards terminals and/or access points, with the purpose of increasing link capacity, extending battery life and reducing terminals costs [1], [2], [3], [4]. Recently, base-band beamforming and Multi-Input Multi-Output (MIMO) are the most widely used technologies; however, in low power communication this does not represent the optimal choice, since MIMO systems rely on multi front-end receivers and power-hungry Digital Signal Processors (DSPs) [5], [6], [7], [8].

This work would like to demonstrate that high quality communication link can be achieved by means of terminals equipped with reconfigurable antenna - single Front-End and a software defined radio controller to steer the main beam in the direction of the base station, thus ensuring that the global energy consumption of the link is still lower when compared to the traditional approach implemented with a general on-chip IEEE802.11 wireless module.

These technological improvements, i.e., low power processing and wireless communication, have been integrated together for the purpose of a case study in a critical environment: the Antarctic region. Antarctica is a valuable resource for scientific research in several domains: today the continent hosts stable research stations from 28 countries. It is also a unique scenario for stressing various technologies, thanks to its extreme environmental and isolation conditions. The scientific application addressed by the system described in this paper is based on the acquisition of GNSS satellite data for the study of radio propagation in the ionosphere [9], [10].

The prototype architecture, which is depicted in Fig.1, is split into two parts that will be installed in the Italian Mario Zucchelli Station (i.e., the permanent observatory):

- the *base station* will be placed inside the main building of the observatory, meaning the module does not have

particular constraints regarding power supply and temperature;
- the *measurement station* is located 500 meters far away from the base station, and it needs to take into account some crucial environmental aspects, such as climatic constraints, power supply and management, distances and Line of Sight (LoS) conditions for wireless communications.

Fig. 1. Main blocks of the prototype architecture

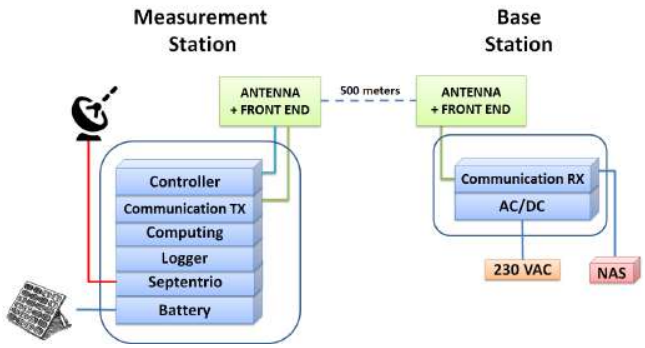

GNSS data are currently acquired in the observatory but, so far, this study has consisted of three mandatory phases:

1) raw data acquisition on a fixed measurement station, located into a container outside the base station;
2) processing of the received data on a dedicated PC;
3) data storage on a networked hard disk, exploiting a wired link between the measurement and the base stations.

This configuration implies that two main constraints emerge: *(i)* the GNSS receiver can only be used in the measurement station that has power supply and it is reached by the network infrastructure; *(ii)* expensive hardware must be dedicated to a single task.

The goal of the activity presented here is to extend the existing architecture, by implementing a ground-disposable system powered by renewable sources, and equipped with low power processing and wireless communication capabilities, in such a way it can overcome the limitations of the current fixed installations.

The rest of the paper is organized as follows: Section II provides a description of the scientific and technological background. In Section III the design is described, highlighting the main constraints and the motivation of the chosen approach and architecture. Section IV presents the experimental results of the preliminary tests on the building blocks of the prototype. Section V draws conclusions and gives some ideas for future works.

II. Background and Motivation

Nowadays, the continuous progresses in the miniaturization of electronic chips allow to create powerful computing platforms which can easily fit the space of a very small board. Such types of computing platforms, also known as Single Board Computers (SBCs), integrate all the components of a traditional desktop-like computer (e.g., CPU, main memory, I/O interconnections, storage, etc.) in a compact boards. Raspberry Pi [11] and Arduino [12] are two popular projects proposing SBCs with the size of a credit card, with the latter more oriented to the creation of systems able to interact with the outside environment, thanks to a large set of sensors and actuators available. Although initially designed to cover applications in the embedded domain (this is mainly due to the adoption of low power embedded processor or micro-controllers), nowadays, SBCs sport high-performance Systems-on-Chip (SoCs) coupling a powerful CPU and a GPU within the same silicon die. Such platforms are able to run also complex parallel applications. These improved features make them more attractive for implementing advanced cyber-physical applications. Furthermore, inexpensive hybrid solutions mixing general purpose computational capabilities and reconfigurable logic (i.e., Field Programmable Gate Arrays – FPGAs) [13] are becoming popular. It is worth to note that, albeit the performance exhibited by such computing modules are continuously growing [14], their power (and mostly energy) consumption remains within an acceptable range, so that it is possible to effectively power them through a battery or using renewable energy sources.

The growing demand for processing and storage capabilities lead in the past computer designers to implement parallel platforms in order to crunch as much as possible data. Historically, computer clusters have been used to process large amount of data recurring to standard commercial off-the-shelf computers instead of more expensive dedicated ones, like those available in a supercomputer system. Today, the large availability of high-performance SBCs allow the creation of clusters hosted in a standard server chassis. HPE Moonshot [15] is an example of a commercial high-density server system, which is able to pack up to 180 server nodes in a single enclosure. Similarly, several research projects are investigating on the design and use of such systems both in the high-performance and cyber-physical contexts. For instance, the FIPS project [16] developed a heterogeneous datacenter-in-a-box equipped with up to 72 ARM-based low power processors, and mixing also standard X86 processors and FPGA acceleration modules. The 1U server box is mainly designed to run complex simulations coming from the life-science domain. Conversely, the AXIOM project [17] proposes a modular design based on the UDOO boards [18] for accelerating video-processing applications that are common in the embedded domain.

The usage of small computing clusters made of inexpensive boards is becoming attractive for the scientific community [19], [20], [21], since they can provide autonomous computing and wireless communication capabilities, which are of great interest to collect data from several sensors or to act on the environment through actuators. In many applications, such systems can be powered by a battery pack or by external renewable sources. In these contexts, the adoption of a multi-board solution becomes almost mandatory, since tasks regarding monitoring the power source or controlling the wireless interface require dedicated computing resources. In particular, adaptable wireless communication systems based on smart

antennas result to be high computationally demanding, thus requiring fast computation capabilities. This becomes even more true when wireless communications are managed through more flexible software-defined functionalities, since all the low level communication tasks are run as dedicated software components. However, to our best knowledge there is not yet a well defined architecture that provides the computational capabilities of a small cluster, which can be powered through a battery pack, and that is able to communicate through a wireless channel through a reconfigurable smart antenna.

Given this premise, our solution represents a first attempt to pack cluster computing and reconfigurable wireless communication capabilities into a complete autonomous system, which is able to operate in critical environmental conditions, as those of the Antarctic region.

A. *Environmental constraints*

Extreme climatic conditions are the main characteristic of the polar environment, and their impact on electronic equipment is crucial. For an effective selection of the appropriate devices and precautions, data logged from a weather station close to the site where the prototype will be installed have been downloaded and analyzed [22], both for taking into account the main constraints and for evaluating the available resources in the area. Temperature, shown in Fig.2, is the first aspect that must be considered, and in this area it spans from a minimum of $-35\,^{\circ}\mathrm{C}$ to a maximum of $+8\,^{\circ}\mathrm{C}$. This means that all the electronic equipment must operate in extended temperature range, being the standard range from $0\,^{\circ}\mathrm{C}$ to $+40\,^{\circ}\mathrm{C}$.

Fig. 2. Hourly temperature over one year [°C]

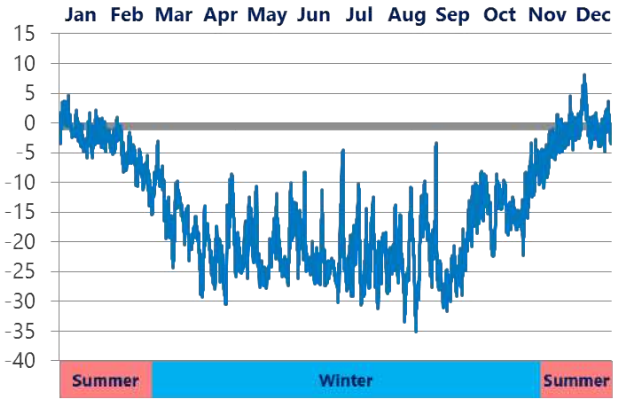

Wind, as can be seen in Fig.3, is always very strong in that area, being a very promising resource for wind power, but there is also a tangible risk of prejudicing the equipment's functionality. In this specific case, it was not feasible to develop a wind-power supply, because of the presence of a magnetic observatory nearby that could possibly suffer from interference issues.

Also solar energy is an important source of supply for Antarctica, and it is commonly adopted because, despite having a dark period of some months in winter (see Fig.4), the solar radiation values are quite high during summer time.

Fig. 3. Hourly wind speed over one year [knots]

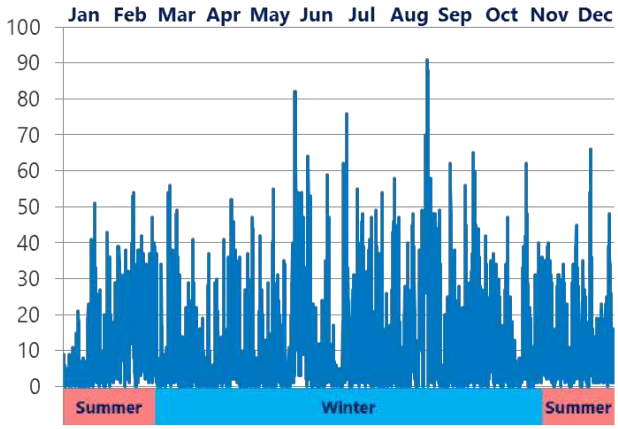

Therefore, this is the renewable power source that has been chosen for this specific application, specifically addressing a summer campaign.

Fig. 4. Hourly solar radiation over one year [W/m^2]

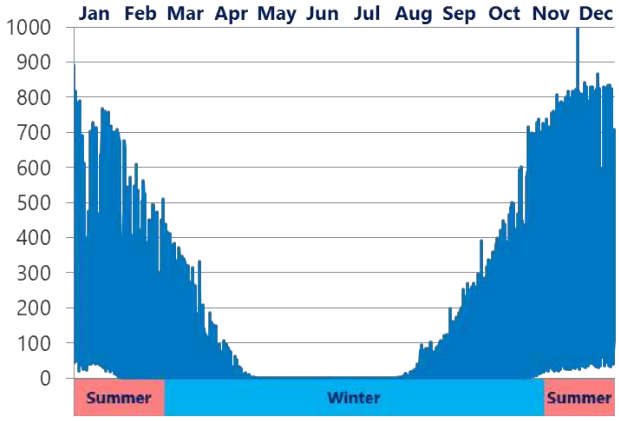

III. Design and Architecture

As depicted in Fig.1, the prototype is composed by the following subsystems that will be described in next paragraphs:

- base station side: Reconfigurable Antenna, Front-End, Communication Receiver, Antenna controller, Power supply (provided by 230 VAC), Network Attached Storage (NAS);
- measurement station side: Reconfigurable Antenna, Front-End, Communication Transmitter, Antenna controller, Power supply (provided by Photovoltaic panels), Computing board, Logger board, Septentrio PolaRxS receiver and GNSS Antenna.

Power efficiency is the main concern of this ground disposable prototype so that, in the following paragraphs, the main logical components of the prototype architecture are described and the specific issues for the design of each building block are detailed. The final prototype implementation of the measurement station is shown in Fig.5.

Fig. 5. Integrated prototype of the measurement station

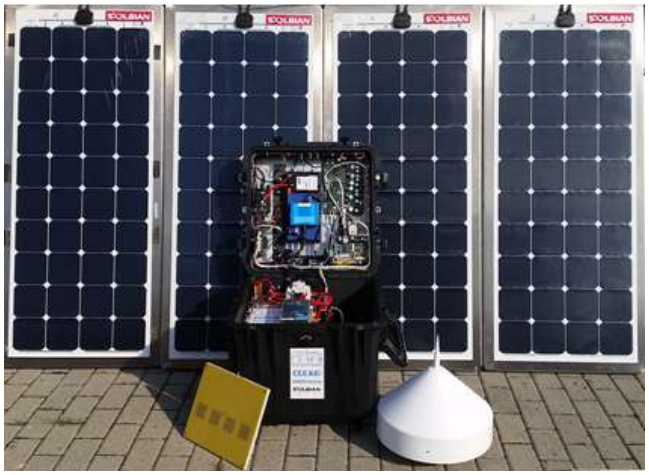

A. Power management

A fine power management is of crucial importance for the measurement station, which is intended as a completely autonomous device (in particular from the energy provisioning point of view), powered by renewable sources (i.e. solar energy). Solar power availability depends from many factors: weather conditions, latitude of the installation, yearly seasons and time of the day. Therefore, care should be taken in order to optimize the overall power efficiency of the entire system. In this paragraph, the architecture and behavior of Power Management Unit (PMU) are described in details. The PMU has been designed and developed by *CLEAR Srl*

Fig. 6. Power supply unit

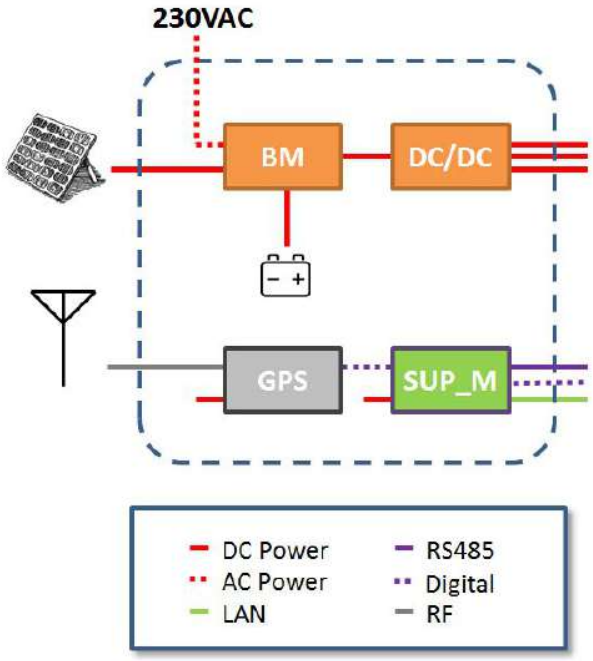

Algorithm 1 Procedure to safely power off the slave modules

```
1: procedure SWITCH OFF PROCEDURE
2:     ShutdownCompleted ← 0                    ▷ Initial value for shutdown status flag
3:     SendCommand (SHUTDOWN)                   ▷ Send the SHUTDOWN command
4:     ResetTimeoutTimer(t)
5:     while ShutdownCompleted = 0 do
6:         GetCommand (c)                       ▷ Check for command answer
7:         GetPowerConsumption (p)              ▷ Monitor the slave power consumption
8:         GetEvent (e)                         ▷ Check for timeout event
9:         if ((c == CLEAR_TO_SHUTDOWN) or
10:            (p < MIN_POWER_THRESHOLD) or
11:            (e == EVT_TIMEOUT)) then
12:            ShutdownCompleted ← 1            ▷ Set shutdown status flag to true
13:        end if
14:    end while
15:    PowerOff()                               ▷ Safely power off the slave module
16: end procedure
```

In Fig.6 the power management architecture is depicted. The Battery Manager (BM) is in charge of extracting all the available electrical energy coming out from the photovoltaic panels and storing it into the battery stack. The BM implements also a Maximum Power Point Tracking (MPPT) algorithm in order to maximize the power extracted as the external conditions change (e.g., panel temperature, solar radiation, presence of shades on a subset of photovoltaic cells). The DC/DC converter module is in charge of extracting the power available from the battery, generating the voltage needed to power the systems. The GPS module in Fig.6 provides the absolute timing information to every component of the system architecture. An input power supply $230VAC$ is also included in order to be able to test the device in laboratory, without the need to connect photovoltaic modules.

The supervisor SUP_M is the module in charge of further optimizing the power consumption of the whole system. Such module feeds the electrical power to all the other components of the system (e.g., GNSS receiver, computing board, antenna) using energy saving policies. In fact, according to a series of preset operating profiles, the modules that are not used are safely switched off using the procedure in Algorithm 1.

Various operating profiles (duty cycles) can be set on each power line, in order to implement specific energy saving behaviors (e.g., the TX module of the measurement station can be powered on only when the acquisition and pre-processing of new GNSS data is completed). These profiles can be changed and stored on the fly using remote commands. In Table I the main electrical components used in PMU are reported.

B. Communications

The ability to reconfigure an antenna may allow saving energy in autonomous systems as, for instance, the nodes of a deployable wireless sensor network [23]. The fundamental concept relies on saving energy radiated in unwanted directions. This is commonly achieved by means of directional antennas, but the real advantage of reconfigurable antennas is their capability to focus their energy toward the intended direction dynamically, without requiring a mechanical movement. Generally, this is accomplished by

TABLE I
ELECTRICAL COMPONENTS USED IN PMU

Component	Manufacturer	Description
AD8221	*Analog Devices*	*Precision Instrumentation Amplifier*
L76	*Quectel*	*GNSS Module*
STM32F407	*STM*	*ARM Cortex-M4 32b MCU+FPU, 210DMIPS, up to 1MB Flash/192+4KB RAM, USB, OTG HS/FS, Ethernet, 17 TIMs, 3 ADCs, 15 comm. interfaces & camera*
TEN40	*Traco Power*	*DC/DC converter, TEN 40N Series, 40 Watt*
THL20WI	*Traco Power*	*DC/DC converter, THL 20WI Series, 20 Watt*
THN30	*Traco Power*	*DC/DC converter, THN 30 Series, 30 Watt*
TMR9	*Traco Power*	*DC/DC converter, TMR 9 Series, 9 Watt*
IK0505SA	*XP POWER*	*DC/DC converter, IK Series, 0.25 Watt*

varying the voltage of internal components of the antenna (e.g., the beamformer). Since less transmission power will be required, this leads to a more energy-efficient development [24].

Fig. 7. Reconfigurable Antenna System

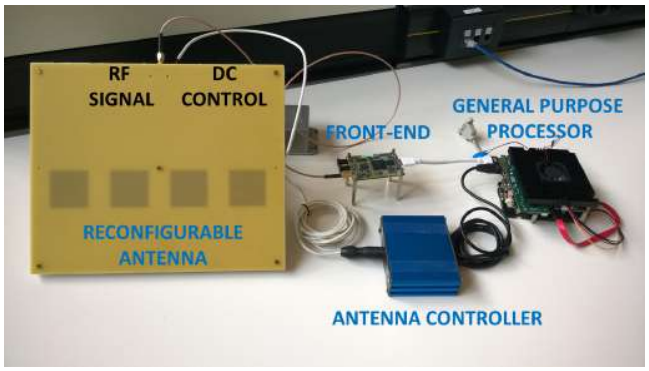

An overview of the communication system is reported in Fig.7, which presents the basic components employed on both base and measurement stations. The realized antenna is a 4x1 array connected to a Radio-Frequency (RF) beamformer that, by means of voltage controllable phase, is able to steer the main beam. These four signals are summed within the RF beamformer and the resulting signal is downconverted to baseband and discretized by a single FE. The digital signal is then processed by a GPP which executes the software implementation of the employed communication standard and the algorithm controlling the antenna, that scans the space to search for the LoS. Both software are fully optimized to be fitted in low power boards. The controller is a PC-to-Antenna beamformer interface which essentially converts the digital input in to a voltage with the purpose of controlling the position of the main beam.

The link has been optimized for low power point-to-point communication. It derives from the idea of placing the measurement station in an arbitrary position without knowing a-priori the base station location. When the measurement station needs to transmit the collected data, by means of the reconfigurable antenna controlled via SDR, it aligns the main beam autonomously with the one of base station. Furthermore, thanks to the modular structure, a power supply controller switches off the transmitter module of the measurement station when it has finished to send data. This allows a fine-grain control on the power consumption of the communication module which requires peak power when transmitting and zero consumption in the idle time. A different situation applies to the base station: in fact the availability of AC power supply allows this side to be always on. The Base station is responsible to transmit a carrier signal to be found by the measurement station and waiting to receive the data transmitted from the latter.

C. Computing

The computing module is devoted to the pre-processing of raw data collected from a Septentrio PolaRxS receiver, and it is activated by the logger module whenever a given amount of raw data are received, then it is powered off immediately after the elaboration in order to optimize power consumption.

This module allows to execute several analyses based on the type of data retrieved. One of these analyses is surely related to the ionosphere which is the single largest contributor to the GNSS (Global Navigation Satellite System) error budget, and ionospheric scintillation in particular is one of its most harmful effects. The Ground Based Scintillation Climatology (GBSC), developed by INGV [25], can ingest data from high sampling rate GNSS receivers for scintillation monitoring like the widely used GISTM (GPS for Ionospheric Scintillation and PolaRxS). Each analysis needs a predefined elaboration environment with particular software and libraries, so in order to separate each environment Docker tool [26] has been implemented. Docker is a new high-level tool which automates the deployment of applications inside software containers and provides an additional layer of abstraction and automation of operating system-level virtualization. It is designed to help delivering applications faster by using a lightweight container virtualization platform surrounded by a set of tools and workflows that help developers in deploying and managing applications easier. The idea is to associate each application with a container, where these containers can be loaded on to the main board and then upload and download to different locations. Containers are basically self-contained executable applications. To this end, the developer of the application is in charge of collecting all the required dependency components and libraries, and pack them within the application itself. Docker is built on top of LXC Linux containers API, and its main feature is that it offers an environment (in terms of functionalities and performance) as close as possible to a traditional Virtual Machina (VM), but without the overhead of running a whole separate kernel. In addition, it allows to expose kernel features used to simulate the presence of dedicated hardware resources, as well as a full OS (guest OS) with its own memory management installed with virtual device drivers, and controlled by a hypervisor, which works on the top of the host OS. By adding low overhead over host machines, containers perform better than other virtualization methods based on traditional hypervisors, such as KVM [27]

and Xen [28]. Unlike a VM, which runs the full operating system, a container can be even only a single process.

IV. Results

Several tests have been done on each building block of this project. First, each part has been tested separately from the others, then they have been tested together to define an internal communication protocol able to manage all the modules composing the complete prototype. The following sections illustrate the results obtained from the functional tests on:

A. sizing of the renewable energy components, taking into account the prototype consumption;
B. tests in the environmental chamber, bringing the instrumentation to extreme conditions;
C. communication module, in particular antenna reconfigurability and LoS data transmission.

A. Sizing of the renewable energy components

In order to correctly carry out the sizing of the PMU, the available data logged from the above mentioned weather station [22] has been considered. As a first step, the local radiance has been measured all over one year (Fig.4). This information is important to achieve the correct dimensioning of the photovoltaic panels to be installed on the site.

The system should be dimensioned in the worst case period, therefore the months of November and March have been selected as the dimensioning temporal points. It is worth noting that, in the energy budget evaluation process, an optimized operating profile has been considered in order to achieve the correct dimensioning in the real operating conditions (i.e., suitable power on-off duty cycle has been considered for the slave components that don't need to be always powered on).

Fig. 8. Daily solar radiation (March 1st)

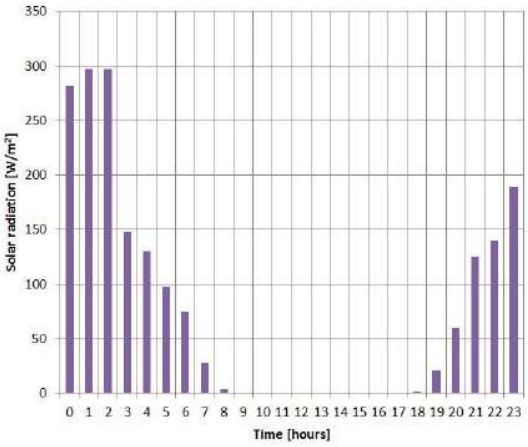

Considering the above mentioned periods, the typical daily radiance trend is depicted in Fig.8. The overall daily radiation for the March 1st is $1706\,W/m^2$ per day, which corresponds to 1.7 equivalent hours[1]. A single photovoltaic module[2] is characterized by $125\,Wp$ peak power, therefore it produces:

$$E_{day} = h_{eq} \cdot P_{STC} = 212.5\,Wh$$

The mean electrical load power requested by the selected operating profile is about $30\,W$, with a peak value of $50\,W$. This mean power corresponds to $720\,Wh$ of requested energy per day. Consequently, it is possible to calculate the minimum number of photovoltaic modules requested for the system:

$$N_{panels} > \frac{E_{eff}}{E_{load}} = 3.38$$

Rounding this value to the next larger integer gives the minimum number of required panels. This rounding operation also takes into account the losses due to the DC/DC converters and the efficiency of the energy storage unit (i.e., the battery stack). We finally selected 4 panels of the above mentioned photovoltaic module.

For the sizing of the battery stack, an autonomy of 3 days without renewable sources has been considered: this corresponds to $2160\,Wh$ of stored energy ($30\,W$ per $72\,h$). The LiFeYPO4 battery from GWL Power [3] is characterized by a $320\,Wh$ capacity, it is therefore possible to calculate the minimum number of battery cells requested for the system, which is:

$$N_{batteries} > P_{load} \cdot \left(\frac{h}{C}\right) = 30\,W \cdot \frac{72\,h}{320\,Wh} = 6.75$$

Considering the losses due to the energy storage unit, we selected 8 cells of the above mentioned battery model.

B. Tests in the environmental chamber

The equipment selected for the prototype has been tested in an environmental chamber with a twofold goal:

1) a preliminary validation of the full functionality of the single building blocks of the architecture when brought to low temperatures;
2) an evaluation of the internal temperature of the enclosure where the hardware is placed, when external conditions are varied.

As for the first objective, this was assumed to be intrinsically guaranteed by the specification of the single devices selected. On the other hand, the coexistence of several devices, switched on and running in a single box, could drive to higher temperatures within the box, and possibly to overheating conditions. As can be seen in Fig. 9, the internal temperature of the box retraces the external condition, with a considerable and almost constant gap of $+30\,°C$.

This is a positive remark for the polar application specifically addressed, but should be taken into consideration for possible future extensions in different environmental conditions.

[1]Considering $1000\,W/m^2$ radiation in STC conditions.
[2]SOLBIAN FLEX SP125 model.
[3]The battery family based on LiFeYPO4 technology has been selected because it is particularly suitable for harsh environmental conditions.

Fig. 9. Tests of the equipment in the environmental chamber: temperature comparison between chamber and equipment box [°C]

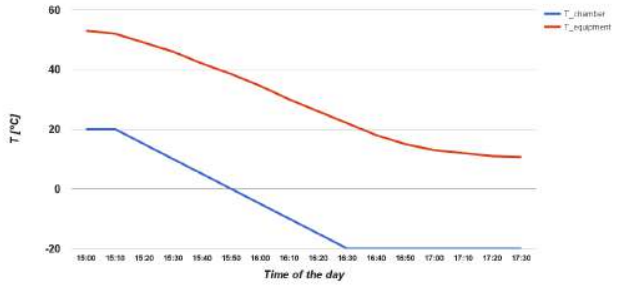

Considering the output of the previous test, where the equipment was switched on for the whole duration, a further assessment has been done: the environmental chamber has been set to the minimum temperature of $-20\,°C$ for 3 hours, while the devices under test were switched off. After this period of time, the equipment was switched on in order to see if every device was able to start working at low temperature without being warmed by other ones nearby. Also in this case, the test was successful, giving a positive feedback on the real working conditions in the final test-bed.

C. Base station detection and data transmission

The settings for the reconfigurability test are illustrated in Fig. 10. Base station is always listening and waiting for data, while transmitting a carrier signal with the purpose of being revealed.

Fig. 10. Reconfigurability Test Settings

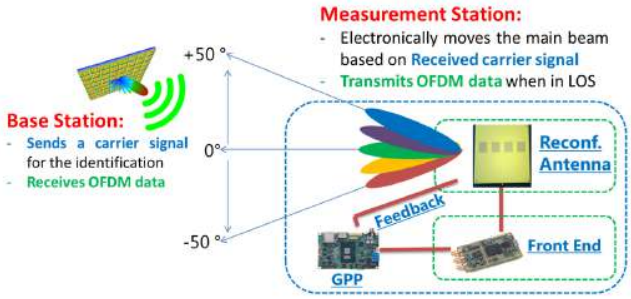

When measurement station is ready to transmit collected data, its communication module is switched on, and before starting the transmission the antenna controller algorithm computes the received power for each position in which the main beam can be steered (e.g., Fig. 11, executes a scan that ranges from $-50\,°C$ to $+50\,°C$ with a step of $5\,°C$).

After the best position is detected (signal is the strongest), the LoS condition is established and the measurement station sends the pre-processed GNSS data to the base station. The scan and lock mechanism reported in Fig. 11, is repeated every time a transmission is required. Scan time can also be reduced, since it mainly depends on sampling frequency, observation windows for the power estimation and scan step. Thanks to software flexibility, they can be adjusted based on needs.

Fig. 11. Scan and Lock Algorithm

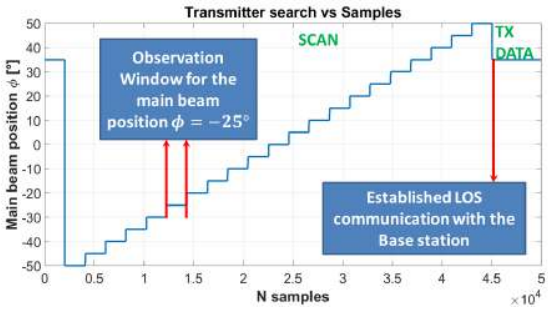

V. Conclusions and Future Work

An advanced prototype of a low power computing and communication system, self-sustained by means of renewable energy and advancements in communication, has been designed and realized for being deployed in the Antarctic continent within the scope of the Italian National Research Program for Antarctic Research (PNRA).

The system is meant to support the collection and transmission of scientific data in field, being easily movable and suitable to extreme environmental conditions. All the constituting components of the prototype have been selected to be robust to low temperatures, and tested in an environmental chamber before the on-site campaign; this has allowed to recreate the temperature conditions that characterize the Antarctic summer time, when the system is intended to be operative. Communication between measurement station and base station has been successfully tested outdoors, at various transmission rates, through the approximation of the positions of the two stations in a scenario as close as possible to the field set-up.

Next steps include the field tests of the system in the Antarctic region, which is also a crucial step to collect feedback that will allow a fine tuning of the single building blocks, as well as possible improvements of the whole system. The system's flexibility will also be enhanced, leveraging its modularity to work on different scenarios (e.g., changing the deployment site will drive to different constraints on hardware and renewable sources), preserving the current features already implemented for this experimentation.

VI. Acknowledgements

The authors are grateful to the PRNA-Programma Nazionale di Ricerche in Antartide for supporting the project "Upper atmosphere observations and Space Weather" within PNRA D.C.D. 393 del 17/02/2015 PNRA14_00110 - Linea A1 and to all contributors: CLEAR elettronica Srl, SOLBIAN energie alternative Srl, INGV and ENEA.

References

[1] Y.J. Guo and Pei-Yuan Qin. Advances in reconfigurable antennas for wireless communications. In *9th European Conference on Antenna and Propagation*, April 2015.

[2] S. S. Jeng and C. W. Tsung. Performance evaluation of ieee 802.11g with smart antenna system in the presence of bluetooth interference environment. In *2007 IEEE 65th Vehicular Technology Conference - VTC2007-Spring*, pages 569–573, April 2007.

[3] Wenjiang Wang, Sivanand Krishnan, Khoon Seong Lim, Aigang Feng, and Boonpoh Ng. A simple beamforming network for 802.11b/g wlan systems. In *Communication Systems, 2008. ICCS 2008. 11th IEEE Singapore International Conference on*, pages 809–812, Nov 2008.

[4] D. C. Chang and C. N. Hu. Smart antennas for advanced communication systems. *Proceedings of the IEEE*, 100(7):2233–2249, July 2012.

[5] M. Uthansakul and P. Uthansakul. Experiments with a low-profile beamforming mimo system for wlan applications. *IEEE Antennas and Propagation Magazine*, 53(6):56–69, Dec 2011.

[6] A. Hakkarainen, J. Werner, K. R. Dandekar, and M. Valkama. Widely-linear beamforming and rf impairment suppression in massive antenna arrays. *Journal of Communications and Networks*, 15(4):383–397, Aug 2013.

[7] K. T. Jo, Y. C. Ko, and Hong-Chuan Yang. Rf beamforming considering rf characteristics in mimo system. In *2010 International Conference on Information and Communication Technology Convergence (ICTC)*, pages 403–408, Nov 2010.

[8] A. S. Prasad, S. Vasudevan, R. Selvalakshmi, K. S. Ram, G. Subhashini, S. Sujitha, and B. S. Narayanan. Analysis of adaptive algorithms for digital beamforming in smart antennas. In *Recent Trends in Information Technology (ICRTIT), 2011 International Conference on*, pages 64–68, June 2011.

[9] et al. Prikryl P. Gps phase scintillation at high latitudes during geomagnetic storms of 7-17 march 2012 - part 2: Interhemispheric comparison. *Annales Geophysicae*, 2015.

[10] et al. Prikryl P. An interhemispheric comparison of gps phase scintillation with auroral emission observed at the south pole and from the dmsp satellite. *Annals of Geophysics*, 2013.

[11] Raspberry pi. https://www.raspberrypi.org. Accessed: 2016-09-07.

[12] Arduino single board computer. https://www.arduino.cc. Accessed: 2016-09-07.

[13] The parallella board. https://www.parallella.org. Accessed: 2016-09-07.

[14] D. Richie, J. Ross, S. Park, and D. Shires. Threaded mpi programming model for the epiphany risc array processor. *Journal of Computational Science*, 9:94 – 100, 2015.

[15] Hpe moonshot system. https://www.hpe.com/us/en/servers/moonshot.html. Accessed: 2016-09-07.

[16] Fips project. https://www.fips-project.eu/wordpress/. Accessed: 2016-09-07.

[17] Roberto Giorgi. Scalable embedded systems: Towards the convergence of high-performance and embedded computing. In *EUC-2015*, 2015.

[18] E. Palazzetti. *Getting Started with UDOO*. Packt Publishing, 2015.

[19] Simon J. Cox, James T. Cox, Richard P. Boardman, Steven J. Johnston, Mark Scott, and Neil S. O'Brien. Iridis-pi: a low-cost, compact demonstration cluster. *Cluster Computing*, 17(2):349–358, 2014.

[20] Yiran Zhao, Shen Li, Shaohan Hu, Hongwei Wang, Shuochao Yao, Huajie Shao, and Tarek Abdelzaher. An experimental evaluation of datacenter workloads on low-power embedded micro servers. *Proc. VLDB Endow.*, 9(9):696–707, May 2016.

[21] Sheikh Ferdoush and Xinrong Li. Wireless sensor network system design using raspberry pi and arduino for environmental monitoring applications. *Procedia Computer Science*, 34:103 – 110, 2014.

[22] Climantartide web site. http://www.climantartide.it/. Accessed: 2016-09-07.

[23] M. Orefice G. Dassano. Voltage controlled steerable array for wireless sensors networks. In *2nd European Conference on Antennas and Propagation EuCAP*, November 2007.

[24] Constantine A. Balanis. *Antenna Theory - Analysis and Design*. Wiley, third edition, 2005.

[25] O. Terzo L. Spogli L. Alfonsi V. Romano A. Scionti, P. Ruiu. Demograpе: Managing scientific applications in a cloud federated environment. In *CISIS-2016*, 2016.

[26] Docker. https://www.docker.com/. Accessed: 2016-09-07.

[27] Kvm - kernel virtual machine. http://www.linux-kvm.org/page/Main_Page. Accessed: 2016-09-07.

[28] Xen. https://www.xenproject.org/. Accessed: 2016-09-07.

GreenLab: autonomous low power system extending multi-constellation GNSS acquisition in Antarctica

32nd URSI GASS, Montreal, 19-26 August 2017

GreenLab: autonomous low power system extending multi-constellation GNSS acquisition in Antarctica

L. Mossucca (1), L. Pilosu* (1), P. Ruiu (1,2), G. Giordanengo(1), S. Ciccia (1,2), G. Vecchi (2), O. Terzo (1)
V. Romano (3,4), L.Spogli (3,4), C. Cesaroni (3), I. Hunstad (3) and A. Serratore (3).
(1) Istituto Superiore Mario Boella, Torino, Italy
(2) Politecnico di Torino, Torino, Italy
(3) Istituto Nazionale di Geofisica e Vulcanologia, Roma, Italy
(4) SpacEarth Technology, Roma, Italy

Abstract

The opportunity to extend monitoring campaigns in Antarctica, overcoming the limitations of the current fixed installations and getting data directly on-field, becomes increasingly attractive for researchers and scientists. The GreenLab is an energy-efficient and self-sufficient system specifically conceived for critical environments. It is equipped with a multi-constellation GNSS receiver and low power systems for computing and communications. Exploiting an automatically reconfigurable antenna, the GreenLab enables a real-time sharing of received data with the base station while optimizing power consumption. This paper presents the deployment and validation of the first GreenLab prototype, during the XXXII summer campaign in Antarctica. Acquired data have been compared to existing GNSS receivers already installed in Antarctica in fixed observatories, with encouraging results.

1. Introduction

The Antarctic observatories are a great resource for scientific communities of different fields: astrophysics, magnetism, biology adaptation, evolution of the earth's crust and climate change are only some of them. Antarctica is one of the best locations on the Earth for investigating the electromagnetic propagation in the ionosphere and studying phenomena related to the interactions between the Sun and Earth.
The work presented in this paper is focused on the analysis of the ionosphere, which is the largest contributor to the Global Navigation Satellite System (GNSS) error budget. In particular, ionospheric scintillation is one of the most harmful and interesting effects.
At present, critical environmental conditions and logistic difficulties in polar regions (e.g., power supply and network outside the bases) are such that the experimental acquisitions are almost exclusively carried out within the base stations. In the case where remote stations are available, the acquired data are not shared in real time with the base station.
The GreenLab aims to be an alternative to this scenario, acquiring and pre-processing data directly on field (remote stations), and sending them to the base station.

2. GreenLab architecture

The GreenLab is an autonomous, disposable system powered by renewable sources, equipped with low power processing and wireless communication technologies. It is composed of four main blocks, as shown in **Figure 1**: power supply by renewable sources (i.e. photovoltaic panels); GNSS receiver (PolaRxS), low power computing for acquiring and pre-processing data; low power communications for transferring data to the base station.
Solar energy has been chosen as the main power source because, even if in winter the Sun never rises for several months in Antarctica, the solar radiation values are quite high during summer time (24 sun hours per day), which is the addressed season for the on-field experiments.
The first prototype of the GreenLab has been installed and configured in the Italian station "Mario Zucchelli", in Terra Nova Bay, as an extension of the existing observatory [1]. This is the first attempt to test such type of solution in a harsh environment like Antarctica and **Figure 2** shows the first installation.
The installation is composed by two parts: *(i)* the *receiver module*, composed by the communication and storage devices, placed inside the base station and *(ii)* the GreenLab, hosted in a measurement station, 500 m far away from the base station. While *(i)* does not have particular constraints, *(ii)* is affected by some critical aspects, such as climatic constraints, power supply and management, distances and Line of Sight conditions for wireless communications.
In this area, temperature spans from a minimum of -35°C to a maximum of +8° C: this entails that all involved hardware in the prototype is able to operate in this temperature range, while commercial equipment usually works in standard temperature range from 0°C to +40°C.
The GreenLab is realized to be suitable for the critical environmental conditions of Antarctica, and it aims at being simply deployable on field, leveraging a new antenna technology for automatically discovering the receiver's position without the need of an accurate pointing. The process flow of the GreenLab is composed by three sequential macro-tasks: (a) 24/7 continuous acquisition of raw data from a Septentrio PolaRxS GNSS receiver, (b) processing of the received data on the low

power computing block, and (c) wireless transmission of pre-processed data to the receiver module.

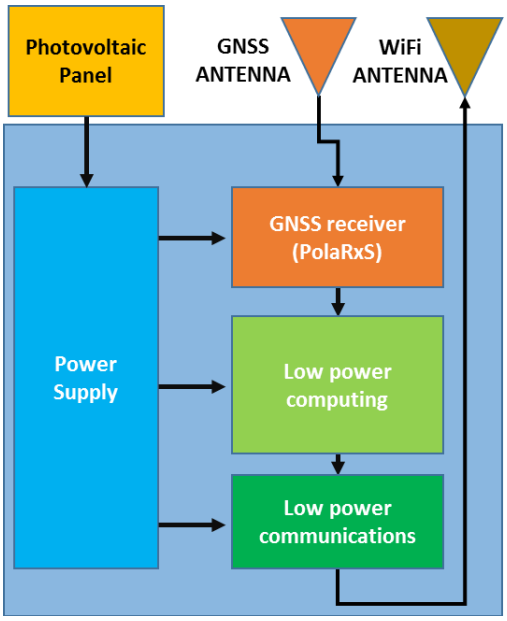

Figure 1. Block diagram of the GreenLab

After the acquisition of raw data from the GNSS receiver, the low power computing block runs a pre-processing algorithm to extract only the most significant information that will be sent to the base station. For this purpose, the this processing block is activated hourly, and powered off immediately after the elaboration in order to optimize power consumption.

The power supply module aims at optimizing power management, both gathering electrical energy from the photovoltaic panels and orchestrating the ON-OFF duty cycles of all the other modules. An *ad-hoc* firmware has been developed to minimize the power consumption of the whole system, using energy saving policies which allow to switch off modules as soon as they finish their tasks.

The low power communications block is devoted to transmit efficiently the pre-processed data coming from the low power computing block. A reconfigurable antenna, controlled by a Software Defined Radio algorithm has been chosen to improve communication link and to reduce power consumption [2].

Figure 3 shows the core components of the first GreenLab prototype.

3. Data acquisition and management

The Septentrio PolaRxS is a receiver capable to sample the GNSS L-band signals from GPS (L1, L2 and L5), GLONASS (L1 and L2) and GALILEO (E1, E5a and E5b) constellations [3]. Starting from raw measurements of the phase and amplitude of the signals taken at 50 Hz sampling rate, the PolaRxS is able to calculate every minute the slant values of the phase (σ_Φ) and amplitude (S4) scintillation indices on all available frequencies from all satellites in view [4, 5]. PolaRxS is able also to provide 1-minute measurements of the uncalibrated Total Electron Content (TEC) and of the rate of TEC change (ROT).

Figure 2. GreenLab installed in Mario Zucchelli Station

After the acquisition of GNSS signals, data are stored on a local hard drive and organized in two main sets: (*i*) raw data and (*ii*) pre-processed data.

Raw data, being of a significant size (up to 1.5GB/day), are only stored locally into the GreenLab. Pre-processed data, whose size is almost 30 MB/day, are forwarded to the wireless transmitter and sent directly to the receiver module, that is always active, waiting for incoming data. These data, being of smaller size and thus easier to be transferred with the limited resources available in the base, are shared almost in real time with the scientific community.

The building blocks of the low power communications system are the same for transmission and reception: a classical phased-array (formed by 4 patch antennas) coupled with a Beam-Forming Network based on phase-shifter, able to steer the antenna beam of about ±50° in the horizontal plane.

The signal acquired by the antenna is down-converted to baseband and digitized by exploiting a Universal Software Radio Peripheral (USRP) B200mini-i. It is a transceiver with the dimensions of a business card and can be employed in a frequency range from 70 MHz to 6 GHz. At the end of this processing chain, the signal is digitized, through an Analog to Digital Converter (ADC), and it can be elaborated by the communication board.

The digital stream is elaborated with a fully software implementation of a standard WiFi 802.11 receiver, which exploits the signal information to control the antenna and

to optimize the direction and the power of the signal to be transmitted/received. The automatic pointing has been performed by a control board developed by ISMB through an interface that converts the digital signal coming from the SW to a voltage that it is able to steer the main beam.

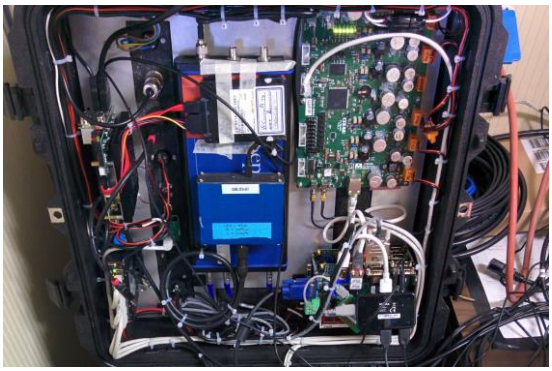

Figure 3. Core components of the GreenLab

For the specific test case, a point-to-point communication has been implemented and, to perform the link in a low-power perspective, the data collected and processed by the computing board are sent to the main station once an hour. In this way the power consumption is optimized since the board and all peripheral devices (e.g. antenna controller, USRP, etc.) are switched-on every hour only for the time to perform the transmission, allowing to further reduce the power consumption of the entire system.

4. Validation results

Data acquired from the PolaRxS receiver have been analyzed and tested against the acquisition of GSV4004 receiver [4]. A GSV4004 is present in Mario Zucchelli Station since 2006 and, similarly to PolaRxS, it is able to provide 1-minute values of σ_Φ, S_4, TEC and ROT starting from GPS L1 and L2 signals data. To perform such comparison, we focused on the data acquired during 24 December 2016, which was characterized by the presence of phase scintillation events in the second half of the day, as shown in **Figure 4**. To minimize the impact of the multipath, an elevation angle mask of 30° is applied to the data of both receivers. Moreover, to minimize the impact of the geometry of the satellite-receiver configuration, the slant values of both indices are projected to the vertical, as described in [6, 7], assuming the ionosphere as a single thin irregular layer located at 350 km. The criticalities of the index verticalization are critically reviewed in [8].

Hereafter we refer to the vertical values of the phase and amplitude scintillation indices as σ_Φ and S_4, respectively. **Figure 4** reports Time profiles of ROT (top), σ_Φ (middle) and S4 (bottom) as recorded by the GSV4004 (panels a) and the PolaRxS (panels b) receivers on 24 December 2016. Colors represents different satellites in view. Both receivers are able to catch the phase scintillation events, highlighted with blue shaded rectangles. No meaningful enhancement of S_4 is found, confirming again that the phase scintillation is most probable to occur at high latitude with respect to amplitude one [6]. In correspondence with the bulk of the scintillation events, the ROT values enhance, confirming the presence of the meaningful electron density gradients, leading to scintillation [7].

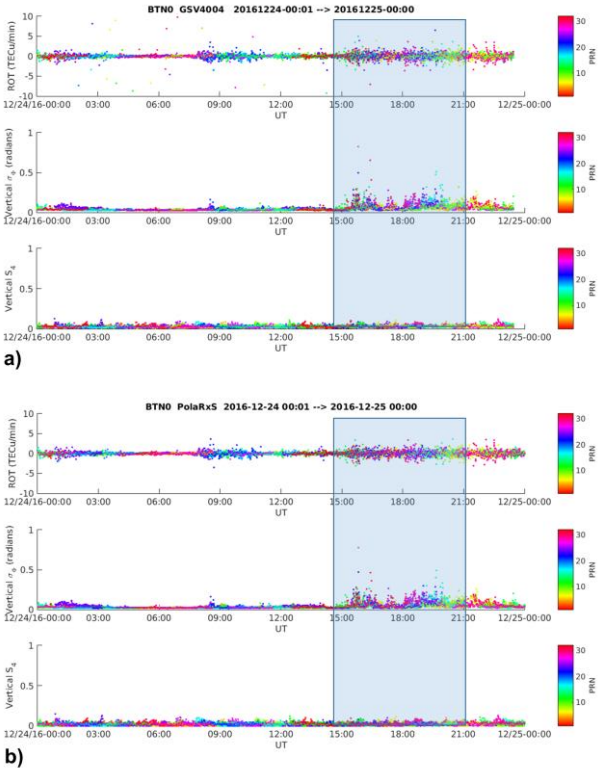

Figure 4. Time profiles of ROT (top), vertical σ_Φ (middle) and vertical S_4 (bottom) as recorded by the GSV4004 (panels a) and the PolaRxS (panels b) receivers.

To better highlight the capability of the GreenLab in catching the scintillation events, a scatter plot reporting the values of σ_Φ by both receivers on 24 December 2016 is reported in **Figure 5**.

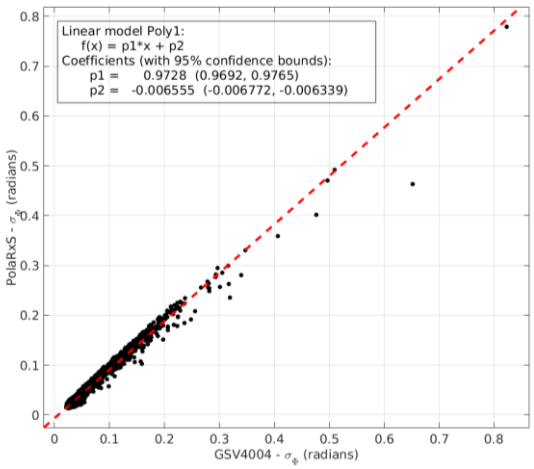

Figure 5. Scatter plot of vertical σ_Φ from GSV4004 (x axis) and PolaRxS (y axis). Red dashed lines indicates the linear fit, whose results are reported in the box.

The agreement between the two independent measurements is demonstrated by the linear fit (red dashed line in the figure), whose angular coefficient and intercept are 0.97 and -0.01 radians, respectively.

5. Final remarks

The provision of real time data acquired in very hard conditions, like the Antarctica ones, is very challenging for many reasons, both environmental and logistic. In fact, outdoor installations are exposed to extreme climate, while fixed observatories are tightly bound to the base where they are deployed.
The GreenLab is an all-in-one, self-powered solution integrating a GNSS receiver, a data-logger with a pre-processing cluster and a wireless communication system with an automatic pointing system. It allows to extend the measurements campaigns to remote sites and to make data available to the base station in real time. The technological innovations implemented in this prototype allow to move it in different locations while it automatically reconfigures the wireless antenna in order to point exactly towards the base station optimizing the transmitted power.
This solution has been realized and tested on field in Antarctica, during the XXXII summer campaign at the Italian base "Mario Zucchelli", and the recorded data are being saved continuously.
The goodness of data recorded from the GreenLab has been validated, compared to an existing fixed receiver.

6. Acknowledgements

The authors are grateful to the PRNA-Programma Nazionale di Ricerche in Antartide for supporting the project "Upper atmosphere observations and Space Weather" within PNRA D.C.D. 393 del 17/02/2015 PNRA14_00110 - Linea A1 and to all contributors: CLEAR elettronica Srl, SOLBIAN energie alternative Srl and ENEA.

7. References

[1] L. Pilosu, L. Mossucca, A. Scionti, G. Giordanengo, F. Renga, P. Ruiu, O. Terzo, S. Ciccia, and G. Vecchi, "Low Power Computing and Communication System for Critical Environments," Series Lecture Notes on Data Engineering and Communications Technologies, Springer, Vol. 1, 2016, pp. 221-232, doi: 10.1007/978-3-319-49109-7_21.

[2] S. Ciccia, G. Vecchi, G. Giordanengo and F. Renga, "Software-defined reconfigurable antenna for energy efficient wireless links," *IEEE International Symposium on Antennas and Propagation (APSURSI)*, Fajardo, 2016, doi: 10.1109/APS.2016.7696328.

[3] B. Bougard, J. M. Sleewaegen, L. Spogli, S. V. Veettil, and J. F. Monico, "CIGALA: Challenging the solar maximum in Brazil with PolaRxS," *Proceedings of the 24th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS*, 2011, pp. 2572-2579, doi: 10.1016/j.jastp.2013.10.003

[4] A. J. Van Dierendonck, J. Klobuchar, and Q. Hua, "Ionospheric scintillation monitoring using commercial single frequency C/A code receivers," *Proceedings of ION GPS*, Vol. 93, September 1993, pp. 1333-1342.

[5] T. L. Beach, and P. M. Kintner, "Development and use of a GPS ionospheric scintillation monitor," *IEEE Transactions on Geoscience and Remote Sensing*, **39**, 5, May 2001, pp. 918-928, doi: 10.1007/BF02830312.

[6] L. Spogli, L. Alfonsi, G. De Franceschi, V. Romano, M. H. O. Aquino, and A. Dodson, "Climatology of GPS ionospheric scintillations over high and mid-latitude European regions," *Annales Geophysicae*, 27, September 2009, pp. 3429-3437, doi: 10.5194/angeo-27-3429-2009.

[7] L. Alfonsi, L. Spogli, G. De Franceschi, V. Romano, M. Aquino, A. Dodson, and C. N. Mitchell, "Bipolar climatology of GPS ionospheric scintillation at solar minimum," *Radio Science*, **46**,3, June 2011, doi: 10.1029/2010RS004571.

[8] L. Spogli, L. Alfonsi, V. Romano, G. De Franceschi, G. M. J. Francisco, M. H. Shimabukuro, B. Bougard, and M. Aquino, "Assessing the GNSS scintillation climate over Brazil under increasing solar activity," *Journal of Atmospheric and Solar-Terrestrial Physics*, 105, December 2013, pp. 199-206, doi: 10.1016/j.jastp.2013.10.003.

Energy efficient system for environment observation

Energy Efficient System for Environment Observation

Giorgio Giordanengo, Luca Pilosu, Lorenzo Mossucca, Flavio Renga, Simone Ciccia, Olivier Terzo, Giuseppe Vecchi, Vincenzo Romano, Ingrid Hunstad

Abstract Environment observations provide a unique source of consistent information about the natural environment and they provide resource managers the information to assess the current state of the environment, weight the requirements of different uses by multiple stakeholders, and manage the natural resources and ecosystems in a sustainable manner. Most of the observations are based on satellites, but remote-sensing technologies alone cannot guarantee observations at the spatio-temporal resolution and with the accuracy requested for monitoring and modeling applications targeting, like weather and climate extremes and the complex feedback processes between the natural environment and human activities. Dense networks of standard and in-situ weather related sensors are present in EU and US, but it may happen that their data are not always available in real-time or updated with the required scale for various weather and climate applications. Then, high-resolution, (near) real-time on field monitoring systems are needed to satisfy the demand to sample environmental data, both in dense populated regions and in less developed and getting more populated regions, where essential in-situ observational capabilities can be lacking or deteriorating. The paper would demonstrate the possibility to have energy efficient computing and communication systems that can be employed for environment observation and that can enrich traditional in-situ and remote sensing environmental data, to enable a significant step forward in the environment monitoring of a wide range of weather and climate data. The paper will present an approach going in this direction (computing/communication everywhere with low-power constrains), tested in a harsh environment, by exploiting low-power boards to perform data pre-processing and reconfigurable antennas to send data in a more energetically convenient way applied to a real case as it may be the monitoring of ionospheric scintillation in Antarctica.

1 Introduction

Since 2003 a specially modified GPS receiver has been used by INGV, Upper Atmosphere Research Unit, for recording the phase and amplitude of the L1 sig-

Giorgio Giordanengo, Luca Pilosu, Lorenzo Mossucca, Flavio Renga and Olivier Terzo
Istituto Superiore Mario Boella (ISMB), Torino, Italy, e-mail: {giordanengo,pilosu,mossucca,renga,terzo}@ismb.it

Simone Ciccia and Giuseppe Vecchi
Department of Electronics and Telecommunication (DET), Politecnico di Torino, Torino, Italy, e-mail: {simone.ciccia,giuseppe.vecchi}@polito.it

Vincenzo Romano and Ingrid Hunstad
Istituto Nazionale di Geofisica e Vulcanologia (INGV), Rome, Italy, e-mail: {vincenzo.romano,ingrid.hunstad}@ingv.it

nals during events of ionospheric scintillation at high latitude [1]. In Antarctica the first Global Navigation Satellite Systems (GNSS) receiver, for ionospheric purposes, was installed in 2006 at the Italian station "Mario Zucchelli", Terra Nova Bay (74.7°S, 164.1°E), in the frame of the Italian National Program for Antarctic Research (PNRA). This site became part of the wider Ionospheric Scintillations Arctic Campaign Coordinated Observation (ISACCO), the Italian project for monitoring ionospheric scintillations and total electron content (TEC) in polar regions with modified GNSS receivers. Since then, three more installations have been done in Antarctica in the Italian-French station of Concordia (75.1°S, 123.4°E) on the Antarctic plateau, at 1200 km from Mario Zucchelli base.

GNSS receivers constitute a very useful tool to investigate the Earth's ionosphere. A radio wave, such as GNSS radio signals, that travels through the ionosphere is affected by a rapid amplitude and phase fluctuations called ionospheric scintillations [2]. Amplitude scintillation can create deep signal fades while phase scintillation are characterized by rapid carrier-phase changes that can produce cycle slips. Both effects can prevent a GNSS receiver from locking on to the signal and make it impossible to estimate a position [3]. Following the success of the first ionospheric scintillation event on Galileo, recorded in Antarctica in the frame of the DemoGRAPE Project [4], a multi-constellation GNSS receiver has been installed in the Mario Zucchelli station during the last summer campaign, in December 2016. The Italian Antarctic Ionospheric Observatory is nowadays covering almost one solar cycle of continuous GNSS data. The new facility to record also Galileo and GLONASS data make the observing facility unique and relevant to investigate the solar-earth physics and the ionosphere dynamics. Maintaining a continuous and unmanned observatory, needs an appropriate infrastructure in terms of computing, power, and connectivity. To fulfill this need the receiver has been integrated by an ad-hoc energy-efficient and self-sufficient GreenLab system realised by Istituto Superiore Mario Boella (ISMB) in the frame of "Upper Atmosphere Observations and Space Weather" PNRA Project [5]. Data management system is also crucial to make data available and useful to scientific and application community. For this reason in the last years a lot of efforts have been devoted to develop data management platforms [6] and ICT infrastructures [7].

In this paper, the authors present a system able to satisfy the demand of more energetically efficient systems able to perform processing and (wireless) communications, with high quality link. To reach this result, the properties of reconfigurable antenna coupled to a software defined radio controller able to perform a real-time control of the antenna itself (e.g. to steer the main beam in the direction of the main station, or to vary the RF power needed to perform the communication), will be exploited with low power processing boards in order to ensure a global energy consumption sustainable with a system powered by (limited) renewable energies [8, 9].

The paper is organized as follows: Section 2 introduces the scientific and technological background and the chosen approach and architecture, Section 3 presents the experimental results of the tests performed in Antarctic Region, while Section 4 draws conclusions and perspectives.

2 System Design and Architecture

The system, formed by low power computing and communication technologies and powered by renewable sources, is sketched in Figure 2. This system has been installed in the Italian Antarctic station *Mario Zucchelli* [10]. As shown in Figure 2, the system is composed by four blocks:

Fig. 1 Panoramic view of the test site in which the Tx and Rx are highlighted.

1. energy supply unit, formed by solar panel and batteries to ensure energy in case of bad weather;
2. data acquisition, by the GNSS antenna and receiver for this specific application;
3. computing section, where collected data are pre-processed;
4. communication part, where pre-processed data are sent to the base-station.

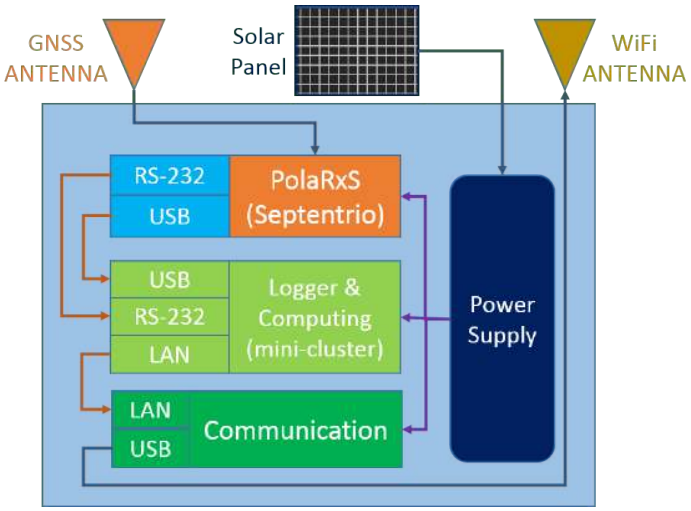

Fig. 2 Pictorial representation of the designed system.

Solar energy is commonly employed in this Region because the sun, during the Antarctic summer, is always present and this is the season selected for the on-field experiments. The power supply module exploits this type of energy source through an optimized power management, and orchestrating the duty cycles of all other modules by switching off modules as soon as they finish their tasks. During winter periods, when the sun never rises for several months, other renewable sources can be exploited (e.g. wind), but this is out of the scope of this demonstrator.

The core of the system is formed by the computing and communication parts which allow, together, to exchange data between the observation point and the base station, where all data are collected, keeping low the overall power consumption.

The computing mini-cluster performs raw-data pre-processing from the GNSS receiver (in the specific case a Septentrio PolaRxS [11]), to extract only a resume of all information collected to be sent in (near) real-time to the base-station. The communication module reads the reduced set of information coming from the computing part and, through a reconfigurable antenna coupled to Software Defined Radio (SDR) algorithms controlling the antenna itself, sends the data to the base-station. This solution allows to automatically discover the receiver's position without the need of an accurate pointing and, in the meantime, to reduce the transmitted power since the Signal to Noise Ratio (SNR) is reduced with this technology.

All components have been selected and developed taking into account the harsh environmental conditions of the deployment Region; in the installation area, the temperature ranges are between $-35°$C to $8°$C. This means that all the materials employed need to be able to operate in these climate conditions and then components working in extended temperature range have to be used (since commercial equipment usually works in a range from $0°$C to $40°$C). Finally, the whole system has also a receiver-side part located in the base station, which includes a receiver board and a network storage server. Even if important, this part is out of the scope of this work since it does not have constraints regarding power supply and temperature and then it will not be treated in the rest part of the paper.

2.1 Power management

In order to implement a completely autonomous device, the system described in this paper is powered by renewable sources (i.e. solar energy). Depending on where the system is installed, solar energy is not always available in site. This is determined by many factors: weather conditions, latitude of the installation, yearly seasons and time of the day. Therefore, care should be taken in order to optimize the overall power efficiency of the entire system. In this paragraph, the adopted power management architecture and behavior is described in detail. Figure 3 shows the

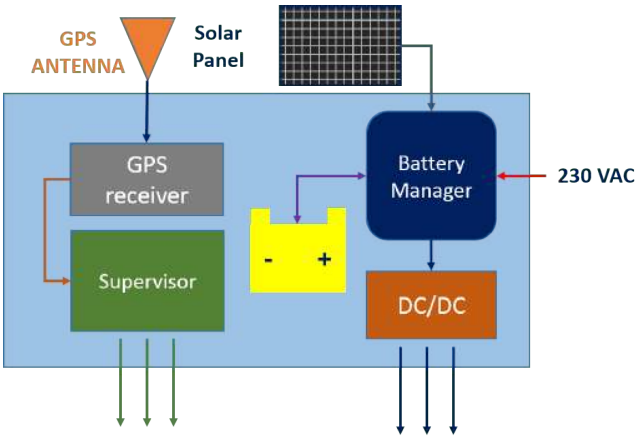

Fig. 3 Architecture of the power management unit.

power management architecture. The Battery Manager (BM) is in charge to ex-

tract all available electrical energy generated by photovoltaic panels (see Figure 4) and to store it into the battery module. The BM integrates also a Maximum Power Point Tracking (MPPT) algorithm, which is able to maximize the extracted power, as the external conditions change (e.g., panel temperature, solar radiation, presence of shades on a subset of photovoltaic cells), etc.). The DC/DC converter module, directly connected to the battery manager, is in charge to extract the available power from the battery, by generating all voltage needed to power the entire system. A GPS module is also included in the power management unit, with the main function to provide the absolute timing information to every component of the system architecture. An additional input power supply at 230 VAC is also included in the system architecture, to allow testing of the device without the need of photovoltaic modules.

In order to further optimize the power consumptions of the whole system, the power management unit integrates also a system supervisor. Such module feeds the electrical power to all other subcomponents of the system (i.e. computation, receiver, storage and antenna unit) using energy saving policies. Indeed, according to a series of preset operating profiles, the modules that are not used are safely switched off to save power. Various operating profiles (duty cycles) can be adopted on each power line, in order to implement specific energy saving behavior (e.g. the TX module can be temporarily powered on only once the acquisition and computation of new data is available). These profiles can be changed and stored on the fly using remote commands.

Fig. 4 Solar panels used as primary source of energy from the system.

2.2 Data Management

To acquire and process data from a multi-constellation GNSS Receiver (i.e. Septentrio PolaRxS), as well as to enhance data analysis by allowing pre-processing and providing it to users in near real time, data collection from Low Energy Boards has

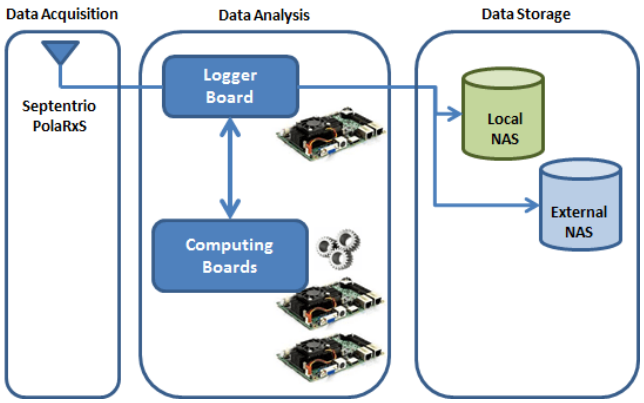

Fig. 5 Data Management System.

been exploited. The Data Management System, depicted in Figure 5, consists of three main modules:

1. Data Acquisition
2. Data Analysis
3. Data Storage

Data acquisition is based on the Septentrio PolaRxS receiver which runs over multiple frequencies and constellations for ionospheric monitoring and space weather observations. It is equipped with low noise control, and its features are high-quality, simultaneous Galileo, GLONASS and GPS signal tracking. This kind of receiver is particularly valuable for tracking rapid signal dynamics like those found in scintillation events. Thanks to its technology for analyzing interference and notch filtering for mitigation, it can be installed in crucial environments for radio broadcasting. GNSS receiver and logger board work together and are always active, 24 hours a day. They are fitted out of automatic procedures for auto-configuration in case of unexpected reboot.

The Data Analysis module is composed of two boards, i.e. logger and computing. It is devoted to data acquisition from the receiver and its pre-processing. The logger is a Single Board Computer (SBC) which integrates all the components of a traditional desktop computer (e.g., CPU, main memory, I/O interconnections, storage, etc.) in a compact board. The Computing stage is formed by several SBCs that have been devoted to process data coming from the logger board. Periodically, the computing boards are switched on as soon as a given amount of raw data are received and they allow to execute several analysis based on the type of data retrieved. One of these analysis is related to the ionosphere which is the single largest contributor to the GNSS error budget, and ionospheric scintillation in particular is one of its most harmful effects. The Ground Based Scintillation Climatology can ingest data from high sampling rate GNSS receivers for scintillation monitoring like the widely used GPS for Ionospheric Scintillation and PolaRxS (GISTM). As mention before, the logger board is always active, 24 hours a day, while, in a energy-saving perspective, the computing boards can be activated only when they are really necessary, based on actual workloads.

Finally, data storage allows to manage and store data and associated metadata. It consists of two NAS servers: the first one, namely Local NAS, is used for saving raw data; the second one, namely External NAS, stores elaborated data and makes them

available for external users outside the base. In particular, the Septentrio PolaRxS receiver generates about 2 GB of raw data each day, that are reduced to about 100 MB a day after the pre-elaboration. The data storage module maintains data provenance and provides the necessary tools and interfaces to allow scientists to quickly identify outdated simulations using chronological traces. Data stored in Local NAS server allow to rerun simulations at any time, such as when a algorithm is updated.

2.3 Communications

The communication system is composed mainly by three parts:

1. Antenna;
2. Front-End;
3. Low Power PC with communication protocol fully software-defined.

The system is configured, by default, in "bidirectional" mode, i.e. radio nodes are able to both transmit and receive data. Neverthless, the communication system can also be set up in unidirectional mode. For instance, the typical one-way data-flow situation from the observatory point (where the data are acquired) to the base-station (where the data are stored). For brevity, the description of the communication system building blocks is addressed from the receiving point of view. However, the transmission side is identical.

The antenna is constituted by a classical phased-array, as depicted in Figure 6, and it is formed by 4 patch antenna (Fig. 6a), coupled with a Beam-Forming Network based on phase-shifter (Fig. 6b) able to steer the beam of the antenna itself of about $\pm 50^{\circ}$ in the horizontal plane [12, 13]. This configuration allows a raw pointing since the receiver position is automatically discovered.

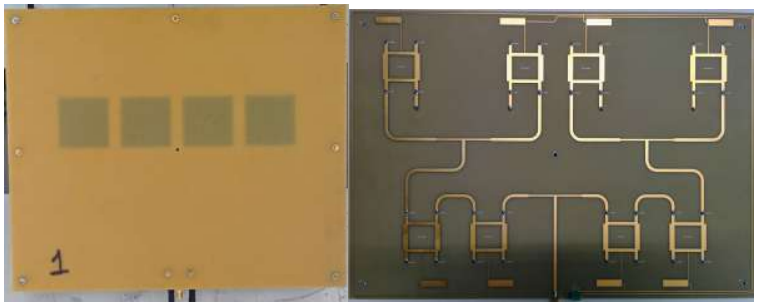

(a) (b)

Fig. 6 Antenna used: phase-array formed by 4 patches (a) coupled with a Beam-Forming Network based on phase-shifter (b).

The Radio Frequency (RF) signal received by the antenna it is down-converted to base-band and digitized by exploiting a Universal Software Radio Peripheral (USRP) B200mini-i [14]. This is the front-end that, through an Analog to Digital (ADC) converter with 56 MHz of instantaneous bandwidth allows the interface between the RF and the digital world. It is a transceiver with the dimensions of a card (see Figure 7), which can be employed in a frequency range from 70 MHz to

6 GHz. Furthermore, it is provided by a user-programmable Xilinx Spartan-6 Field Programmable Gate Array (FPGA). Such devices, coupled with the reconfigurable

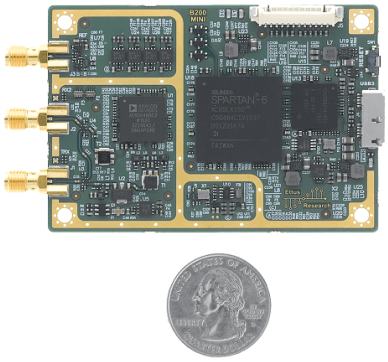

Fig. 7 USRP B200mini-i SDR/cognitive radio.

antenna, constitute the RF part and it is coupled to the board, through a high-speed USB 3.0 connection. Now, downstream of this device, the signal is digitized and it can be elaborate by the communication board, which is in charge to transmit/receive data coming from the computing part. For this purpose it has been exploited a Pico-ITX SBC board (see Figure 8); this embedded board mount the latest generation of Ultra-Low Power Intel Core (i.e. 6th generation 14nm) processors and it is equipped with one DDR4 SO-DIMM with 4 GB memory capacity [15]. Finally, the PICO500 is built to withstand wide temperature conditions, ranging from -20 to +70 Celsius.

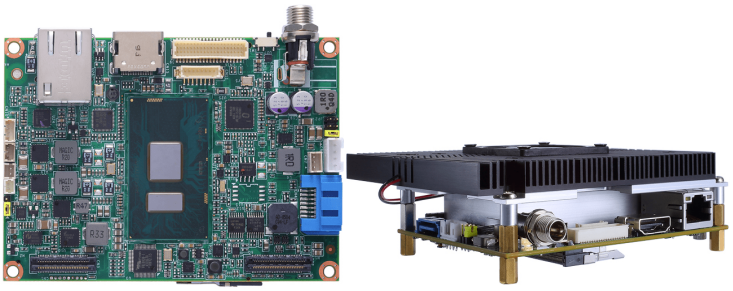

(a) (b)

Fig. 8 Pico-ITX SBC embedded board: top view (a) and coupled with an with active thermal solution (b).

The board elaborates and treats the digital stream with a fully software implementation of a standard WiFi 802.11 receiver. The software algorithm exploits the signal information to control the antenna and to optimize the direction and the power of the signal to be transmitted/received. The IEEE802.11a/g/p wireless communication protocol has been implemented by exploiting GNURadio [16], a free and open-source software development toolkit that provides signal processing blocks to

implement software radios. GnuRadio itself provides a baseline open-source version of the communication standard in question. Unfortunately, its computational complexity impedes operation in low power General Purpose Processors; thus, starting from this version, we developed a faster and lighter implementation of the transmitter and receiver, customized for Low-Power General Purpose Processor. The last link in the chain is the control of the antenna starting from the software; to do this an interface between the board and the antenna has been developed. This interface it is in charge to convert the digital signal coming from the software to a voltage that, applied to the beam-forming network of the designed antenna, it is able to steer the main beam, allowing to focus towards the base-station, where the receiver is waiting for the data.

The implemented test case is a point-to-point communication between the acquisition and the base-station; to perform the communication in a low-power perspective, the data collected and processed by the computing board are sent to the main station once an hour. This allows to reduce the power consumption, since the board and all stuffs connected to that board (i.e. antenna controller and USRP) are switched-on only to perform the data transmission, allowing to reduce the power consumption of the entire system.

3 Experimental Results

In this Section the results obtained on field will be depicted to demonstrate the validity of this type of system to perform low power processing and wireless communication, integrated together to be deployed in a critical environment like the Antarctic region.

3.1 Communication Performance

Prior showing the tabulated results of the wireless communication tests, it is essential to illustrate the working principle of the realized communication link. The link implemented for this application is a point-to-point topology composed by two radio nodes:

- The Observatory Point Radio (OPR), i.e. the radio in charge to transmit data, in this case the node presenting the highest consuming. Thus, the energy minimization process has been applied on this side;
- The Base Station Radio (BSR), i.e. the radio in charge to receive the data and to identify the radio process.

These nodes employ the following high level ad-hoc protocol in order to communicate efficiently:

1. Observatory Point side, the energy supply unit turn on, with a programmable timer, the OPR board to load and transmit the stored information;
2. OPR board disseminates a request of transmission and poses to listening mode later;
3. BSR, at the reception of the request, sends a training sequence with the aim of revealing its spatial position;
4. OPR elaborates the training sequence by means of a spatial scan, i.e. electronically moves the antenna main beam. For each discovered position, the optimization algorithm keeps track about the training sequence power. At the end

of the spatial scan, the processor directs the antenna main beam toward the direction of maximum received power. This configuration is maintained for all the permanence of the data transmission;
5. After data transfer success, OPR advices the energy supply unit that shutdown. The energy supply unit cut off power supply to the radio.

The reconfigurable antenna approach in a wireless communication link establishes two important advancement: first, by focusing all radiated energy toward the base station, it reduce in considerable manner the amount of power emitted by the transmitter. The latter is directly proportional to the radio instantaneous power consumption. In other words, the reduction in the transmitted power results in a decreased consumption of the radio. Second, since the antenna main beam is strongly focused toward the intended direction, noise and interferer are attenuated accordingly. This allows the transmitter to employ the highest communication rates to transfer the information. The greater is the communication rate the faster is the data transfer. Consequently, less time the radio is in active mode. These two factors critically affect the daily energy consumption of the proposed link Figure 9 reports the out-

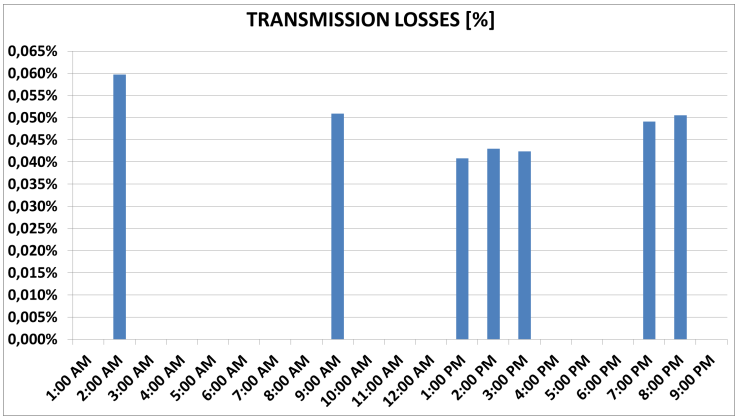

Fig. 9 Automatic Transmission Outcomes.

comes of one-day transmissions. Each data transfer consists in a merge of 4 files of variable size, collected from the observatory system in one hour of monitoring. An example is reported in Table 1. The information about the total size of the data to be transmitted is included in the merged file. This way, the BSR system knows in advance how much data expects from the OPR system. At the end of the data transfer, BSR counts the amounts of received data. Then, the comparison of this number with the one retrieved from the header of the merged file it gives an estimation on the amounts of loosen bytes.

4 Conclusions and Perspectives

An advanced prototype of a low power computing and communication system, self-sustained by means of renewable energy and advancements in communication, has been designed, realized and tested in Antarctica, allowing to test an autonomous system able to manage, pre-process and send data for a real test-case. All components,

Table 1 Example of one hour data collection

Merged Files	Transmitted[Bytes]	Received[Bytes]	Losses[%]
BTN0P_352o00.16_.ismr	152.765	152.465	0,20
BTN0P_352o15.16_.ismr	148.123	148.123	0,00
BTN0P_352o30.16_.ismr	145.168	145.168	0,00
BTN0P_352o45.16_.ismr	143.902	143.902	0,00

designed to allow the collection and transmission of scientific data in field, have been selected to be able to work in really harsh conditions; this allows to overcome unscathed the Antarctic summer campaign. Communication between measurement station and base station has been successfully tested and compared with data collected and sent in more standard way (i.e. wired transmission).

Future development of this systems include a fine tuning of its single building blocks, as well as possible improvements in terms of energy consumption and transmission extension by exploiting other low-power small-form-factor boards.

5 Acknowledgements

The authors are grateful to the Programma Nazionale di Ricerche in Antartide (PNRA) for supporting the project "Upper atmosphere observations and Space Weather" within PNRA D.C.D. 393 del 17/02/2015 PNRA14_00110 - Linea A1. The authors would like also to acknowledge the invaluable help of CLEAR elettronica srl, which provide the control board for the scheduling and power management and SOLBIAN energie alternative srl, which provide the solar panels. Finally we are grateful to ENEA for the logistic and the support during the campaign.

References

1. Giorgiana De Franceschi, Lucilla Alfonsi, and Vincenzo Romano. Isacco: an italian project to monitor the high latitudes ionosphere by means of gps receivers. *GPS Solutions*, 10(4):263–267, 2006.
2. C. N. Mitchell, L. Alfonsi, G. De Franceschi, M. Lester, V. Romano, and A. W. Wernik. Gps tec and scintillation measurements from the polar ionosphere during the october 2003 storm. *Geophysical Research Letters*, 32(12), 2005.
3. Marcio Aquino, Terry Moore, Alan Dodson, Sam Waugh, Jock Souter, and Fabiano S. Rodrigues. Implications of ionospheric scintillation for gnss users in northern europe. *Journal of Navigation*, 58(2):241âĂŞ256, 2005.
4. L. Alfonsi, P. J. Cilliers, V. Romano, I. Hunstad, E. Correia, N. Linty, F. Dovis, O. Terzo, P. Ruiu, J. Ward, and P. Riley. First observations of gnss ionospheric scintillations from demograpе project. *Space Weather*, 14(10):704–709, 2016.
5. Upper atmosphere observations and space weather. PNRA D.C.D. 393 del 17/02/2015 PNRA14_00110 - Linea A1.
6. V. Romano, S. Pau, M. Pezzopane, L. Spogli, E. Zuccheretti, M. Aquino, and C. M. Hancock. eswua: A tool to manage and access gnss ionospheric data from mid-to-high latitudes. *Annals of Geophysics*, 56(2), 2013.
7. A. Scionti, P. Ruiu, O. Terzo, L. Spogli, L. Alfonsi, and V. Romano. Demograpе: Managing scientific applications in a cloud-federated environment. In *2016 10th International Conference on Complex, Intelligent, and Software Intensive Systems (CISIS)*, pages 426–431, July 2016.

8. S. Ciccia, G. Giordanengo, S. Arianos, F. Renga, P. Ruiu, A. Scionti, L. Mossucca, O. Terzo, and G. Vecchi. Reconfigurable antenna system for wireless applications. In *2015 IEEE 1st International Forum on Research and Technologies for Society and Industry Leveraging a better tomorrow (RTSI)*, pages 111–116, Sept 2015.
9. S. Ciccia, G. Vecchi, G. Giordanengo, and F. Renga. Software-defined reconfigurable antenna for energy efficient wireless links. In *2016 IEEE International Symposium on Antennas and Propagation (APSURSI)*, pages 1241–1242, June 2016.
10. The mario zucchelli station. `http://www.italiantartide.it/stazione-mario-zucchelli`.
11. Polarxs. `http://www.navtechgps.com/septentrio_polarxs_packaged_receivers/`.
12. M. Orefice, G. L. Dassano, L. Matekovits, P. Pirinoli, G. Vecchi, and B. Shurvinton. A wide coverage scanning array for smart antennas applications. In *2001 31st European Microwave Conference*, pages 1–4, Sept 2001.
13. G. Dassano and M. Orefice. Voltage controlled steerable array for wireless sensors networks. In *The Second European Conference on Antennas and Propagation, EuCAP 2007*, pages 1–4, Nov 2007.
14. Usrp b200mini-i. `https://www.ettus.com/product/details/USRP-B200mini-i`.
15. Pico-itx sbc. `http://www.axiomtek.it/`.
16. Gnu radio, free and open-source software development toolkit. `http://gnuradio.org/`.

Integration of Reconfigurable Antennas in Ultra Low-power Radio Platforms Based on System-on-chip for Wireless Sensor Network

Integration of Reconfigurable Antennas in Ultra Low-power Radio Platforms Based on System-on-chip for Wireless Sensor Network

Simone Ciccia[1,2], Giorgio Giordanengo[2], Giuseppe Vecchi[1], *Fellow, IEEE*
[1]Department of Electronics and Telecommunications, Politecnico di Torino, 10129 Torino, ITALY
[2]Istituto Superiore Mario Boella (ISMB), 10138 Torino, ITALY

This work deals with energy consumption reduction in wireless sensor network nodes by means of antennas with reconfigurability at radio-frequency. A first novelty with respect to existing literature is that reconfiguration is effected with a platform comprised of a commercial ultra-low power microcontroller and a commercial system-on-chip radio, with specific reference to the IEEE 802.11 b/g standard. We consider explicitly the energy overhead involved in the antenna configuration process and its impact on the global energy efficiency of the system, which was not done so far. We discuss two complementary ways of reducing energy consumption, which is also new. We employ an antenna array with steerable beam, controlled by an algorithm embedded in the firmware of the ultra-low power radio platform; it is designed for prevalently line-of-sight conditions, with beam steering aiming at maximizing the received power. With a common IEEE 802.11 b/g system-on-chip module and a moderate antenna maximum gain (12dB) a power saving factor in excess of 27% can be obtained by reducing the transmit power, at unchanged data rate; additionally, by keeping transmit power constant but increasing the data rate, an energy saving from 15% to 95% can be achieved, depending on link distance and achievable data rate. Our experimental tests confirmed global power saving due to transmit power reduction. The employed hardware showed limitations in throughput that did not allow the constant-power option, and we report the analysis leading to attribute the fault to the hardware interface between micro-controller and the radio.

Index Terms—**Reconfigurable antennas, Ultra-low power communications, Wireless sensor networks.**

I. Introduction

ONE of the most critical issues in Wireless Sensor Network (WSN) is energy efficiency. The main reason of this concern is the fact that they are mostly supplied by renewable sources and/or small batteries, with the need of guaranteeing a constant monitoring operativity [1], [2].

Several techniques to improve the lifetime of a node in a WSN have presented in the literature, especially to lower the consumption of the radio part [3], [4]. The following are a non-exhaustive review of the proposed approaches. A recently approach is the network layer solution proposed in [5], [6]. By means of a modification of the communication protocol it adjusts the transmit power based on some receive parameters. In other words, the link quality is adaptively analyzed in order to select an appropriate transmit power level. In this way, the consumption is reduced by saving the unnecessary amount of transmit power. Another option is scheduling the transmissions toward the base station. It consists in a sleep/awake mechanism which activates the transmit radio only when a data transfer is required [7]. Where applicable (e.g. environmental and traffic monitoring systems) it allows to significantly improve energy efficiency since the power spent for data transmissions is spread along large lowest-power idle intervals.

The use of a directive antenna would allow a higher received power with all other factors unchanged, as the latter is proportional to the transmit power and the product of Receive (RX) and Transmit (TX) antennas gains. Therefore, it can provide energy efficiency alone, or in synergy with any other energy saving approach [8]. It is well known that higher antenna gain can only be achieved by a non-omnidirectional antenna, thus requiring a steerable beam that can be adapted to different link directions, e.g. as in [9], [10]. Since the beam direction needs to be changed dynamically, or at least automatically in the initial deployment, the necessary antenna systems are termed "reconfigurable".

Manuscript received August 29, 2017; revised Month DD, 201Y. Corresponding author: S. Ciccia (email: simone.ciccia@polito.it, ciccia@ismb.it).

A. Reconfigurable antennas for energy saving: State of the art

Despite the good prospects of reconfigurable antennas, this energy saving option appears not yet fully exploited in WSN for different reasons; the main such reason can be attributed to the difficulty of effectively integrating these antenna systems in low-power architectures. In the following, we review the approaches in the literature; we will then address our approach and its innovation with respect to such a state of the art in the next section.

An example of integration of antenna reconfigurability in WSN can be found in [11], where the authors attempt to integrate a switched beam antenna into a radio based on System on Chip (SoC); this SoC, however, is driven by an external PC. The main drawback to this approach is that the solution relies on a Linux Operative System (OS) which entails a consistent power consumption overhead to run.

A wake-up method along with reconfigurable antennas is presented in [12], where two reconfigurable antennas are employed, for receiving and transmitting. The first radio apparatus continuously switches the main beam of the antenna to find the location of the base station, as well as pay attention when the latter ask for data. The second reconfigurable antenna and the radio are activated only when a request of data is detected. The data transfer is performed with the main beam oriented towards the base station since the first radio apparatus is already aware

for the configuration of the antenna. In this interesting solution, although, the sensor has the receive circuitry constrained to be always on, with associated power consumption.

The interesting work in [13] discusses how controllable directive antennas achieve transmission power saving as a result of minimizing the collision rate at base station. In that work a switchable beam antenna and the related consumption for its configuration were analyzed. However, the impact of energy required for the antenna configuration phase has not been considered in the *global* energy efficiency of the system. This initial energy overhead is negligible in a static scenario with a single, initial configuration phase; in a dynamic (time dependent) scenario where node positions can change over time, this overhead has to be considered, as will be done here.

A number of other works have addressed the issue of reconfigurable antennas for energy saving in WSN [8], [10], [12], [14]–[17]. Some of the accounts are of purely simulative nature [8], [10], [15]–[17], while some demonstrated actual antenna reconfiguration either in stand-alone [14] or in association to a system [12], [13], implemented in the short range and low rate IEEE 802.15 SoC. A common feature of all these interesting works is the neglect of the consumption impact of the antenna configurations in the total energy budget (i.e. time model); this important issue will be addressed in the present work.

B. Proposed approach

The goal of our work is the integration of the reconfigurable antenna in an Ultra Low-power (ULP) SoC-based system. The ULP performance comes from the ULP microcontroller that manages the SoC radio module; antenna reconfiguration is effected by a firmware algorithm running on the ULP microcontroller. We demonstrate the approach on a system composed of low-cost commercial devices, the ULP microcontroller and the IEEE 802.11 b/g SoC radio.

We observe that antenna (re)configuration does produce the described advantages, but this comes at an energy cost that has not been considered so far [8], [10]–[13], [15]–[17]. In fact, any search and any processing implies an on-time, and the associated overall energy cost, spent without transmitting useful data; this configuration overhead needs to be accounted for in order to achieve a true energy gain. To the best of the authors' knowledge, such an analysis on the energy efficiency provided by reconfigurable antenna still lacks in literature and is addressed in this paper for the first time.

In contrast to the approach in [11], [14], our solution does not require an external (non low-power) PC; it relies on a small firmware: this allows to avoid the (large) power consumption related to the OS. For example, we have employed a common micro-controller (STM32 Nucleo) in Fig.2a and [18], which consumes a power of about $200mW$ as opposed to the Watts required by a board to run an OS, even if it were a low-power board.

In order to overcome the discussed limitations of the interesting approach in [12], we propose an energy efficient link where only the base station is in charge of continuously scanning the space and finding which sensor node requests to transmit, while the latter are activated by their own radio awakening system and transmit only when the collected data need to be transferred (with the adopted policy). Since usually the base station has less stringent energy requirements than the sensor nodes, this approach is advantageous. This notwithstanding, the number of antenna search (i.e. the "configuration process") must be carefully considered at the latter side, since this number influences the overall lifetime of the node; this will be discussed later.

Another advantage of the proposed solution is that we consider a network architecture that need not be of the star type (i.e. single hop) [13]; in fact, in any (re-)configuration initialization phase we map the Service Set IDentifier (SSID) and associated Receive Signal Strength Indication (RSSI) of all nodes in the scan range (both single-hop and multi-hop); this allows to the system to configure the beam in the optimal way for any connection topology in a manner that it is tied to the employed network protocol. Furthermore, our solution does not affect the MAC protocol. An important point of departure from the current literature is also the application to a dynamic scenario where true re-configuration may be necessary due to changes in the network geometry.

We also observe that a wireless sensor system equipped with such a reconfigurable directive antenna has two complementary advantages over one with a standard (ominidirectional) antenna:

- At the same link distance, one has a given received power with a lower transmit power. As a result one can maintain the same data rate as with a standard system, but lower the transmit power with unchanged data link performances. In this case, one has a power saving, as power consumption of the wireless device is proportional to the transmitted power. Then, a given quantity of information is transferred with a lower energy because of the lower transmit power. This is the traditional motivation of higher antenna gains.
- Alternatively, with the same transmit power, one has a higher received power, which allows to transmit data at higher rates (e.g. as in IEEE 802.11 b/g standards). This implies that a given quantity of information can be transferred in a shorter time, i.e. with a smaller energy. This alternative route to energy saving does not appear to have been considered so far in the literature.

It is apparent that both result in a lower energy-per-bit cost; of course the two advantages can be combined, and/or the distance can be increased with all the rest unchanged. We will address both routes in the following.

Among all possible reconfigurable antenna design for WSN, we have chosen a standard beam steerable antenna array to prove the concept [19]–[22]. This antenna has a coverage of about 90° in the scan plane and offers good spatial selectivity.

We addressed the integration of the reconfigurable antenna in the IEEE 802.11 b/g ST SoC-radio in Fig.2b and [23], since this communication standard provides a wide availability of data rates along with a robust coding, but this choice is neither restrictive nor limited to other actuation forms [24]. In fact, the same ULP microcontroller is compatible

with other commercially available SoC-radios, as for instance, LoRaWAN and Bluetooth Low Energy [25].

The proposed wireless system is suitable for WSN applications requiring long range communication link, avoiding signal repeaters. In particular, thanks to the electronically scanned beam of the antenna, the wireless communication link reaches distances not accessible with the state-of-the-art wireless systems, and can be exploited in application like:

- Transport (road traffic monitoring)
- Environmental and disaster monitoring (agrifood, fire prevention, rivers levels and floods, air pollution and other phenomena)
- Security (surveillance and intrusion detection)
- Smart grid and energy control system
- Industrial (structural monitoring)
- Remote sensing (whether forecasting)

As an example, the proposed system has been applied to road traffic monitoring. In this application, the proposed ULP radio platform is coupled with a ULP computing module (ST - SecSoC) that perform local video processing in order to detect road traffic events. Some features of the ST - SecSoC processing tasks are traffic congestion, wrong way vehicle detection and cycle detection. The average distance between nodes is of 500-1000m. The result of the local video-processing translates in the transmission of some images (alarms) to a traffic road management and elaboration center. The same ULP radio platform is used to bridge the receive data to a local router that is in charge to deliver the alarms to the supervisor.

C. Structure of the paper

The rest of the paper is organized as follows. Sec. (II) presents the employed method in order to analyze the performance of the proposed solution with respect the standard solution as depicted in Fig. 1. Sec. (III) introduces the hardware and software components involved in the realization of the proposed system, with special attention to the ULP SoC-based radio platform, the designed antenna, the developed firmware respectively and the reconfiguration tests under real environmental conditions. Sec.(IV) details the consumption profile of the system in possible scenarios. Finally, Sec.(V) analyses the energy saving by changing the data rate and then some conclusions are drawn.

II. Background and Analysis Method

A. Link Budget

The analyzed scenario is reported in Fig. 1 which illustrates the State-of-the-Art Wireless (SAW) system (above) and the proposed Reconfigurable Antenna Wireless (RAW) system (bottom). Concerning the SAW system, it is composed by the SoC-based radio sold with an integrated antenna that has an omnidirectional-like coverage [26]. Instead, the RAW system exploits the SoC module sold with the U.FL connector and integrates the proposed reconfigurable antenna which provides a directive beam towards the intended direction. The first quantity of interest is the received power that has evaluated with the standard link budget formulation described by (1).

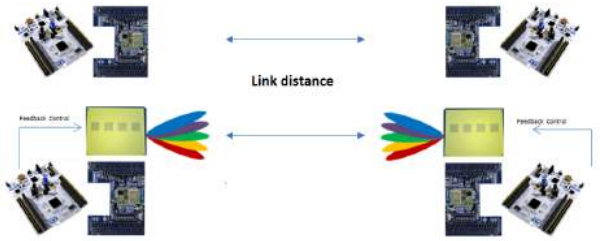

Fig. 1. State-of-the-Art system (above) vs reconfigurable antenna configuration (bottom)

$$P_{rx} = P_{tx} + G_{tx} + G_{rx} - L_{att}\ [dBm] \tag{1}$$

where P_{rx} is the received power in dBm, P_{tx} is the transmit power in dBm, G_{tx} and G_{rx} are the transmit and receive antenna gains respectively and L_{att} is the overall attenuation defined in the following way:

$$L_{att} = L_{path} + L_{add}\ [dB] \tag{2}$$

where L_{add} is an additional loss that accounts for the multipath propagation and cable losses at the transmitter and receiver side respectively, while L_{path} is the path loss which is in turn described as:

$$L_{path} = 20\log_{10}(\frac{4\pi d f_{max}}{c})\ [dB] \tag{3}$$

where d is the link distance, c is speed of light and f_{max} is the maximum operating frequency.

B. Power Saving

A significant point of the link budget comparison analysis is that the RAW link can lower the transmit power and/or increase the data rate with respect the SAW link; the related saving in transmit power is obtained by equating the receive power, which yields

$$\begin{aligned} PW_{saving} &= P_{tx,SAW} - P_{tx,RAW} = \\ &= (G_{tx,RAW} - G_{tx,SAW}) + (G_{rx,RAW} - G_{rx,SAW})\ [dB] \end{aligned} \tag{4}$$

It appears to be independent from the link distance d but directly related to the difference between the antenna gains of the RAW and the SAW links. The power saving enabled by the antenna gain difference PW_{saving} is effected via control of the Programmable Gain Amplifier (PGA), and this power saving could thus be limited by the minimum PGA gain that can be set in the device. This will be further discussed in Sec. IV-B.

C. Energy Saving for equal Data Rate

Lowering the transmit power is one of the way to save energy. Because this relies on antenna directivity, it requires a beam alignment, i.e. a *configuration phase* in which power is transmitted and received without transfer of useful information. We will now investigate the impact of this energy consumption of this antenna configuration phase on the overall operation. It is apparent that the RAW system always begins in energy debit with respect to the SAW system. Thus, we

examine how long it takes to the RAW system to reach the energy break even point, i.e. the time at which the RAW system power saving has equaled the energy spent in the configuration phase; if the scenario does not change, after that time the system always saves energy. This energy breakeven time is an important parameter to judge the effective advantage of the RAW approach with respect to intended system scenarios: energy saving happens if system scenario changes (relocation of nodes, e.g.) are slower than the time scale set by the breakeven time.

In view of this analysis, the energy consumption of the non-reconfigurable (SAW) system is given by:

$$E_{SAW}(\Delta t) = \int_{to}^{to+\Delta t} P_{a,SAW}\ dt\ [J] \tag{5}$$

where to is an initial time, $P_{a,SAW}$ is the power absorbed by the SAW system during data transfer, and Δt is the considered data transfer duration. In order to transmit the same amount of data, the reconfigurable system RAW uses the same data rate, and thus the same transmit time Δt but requiring a lower power $P_{a,RAW}$; hence the associated energy cost is:

$$E_{RAW}(\Delta t) = E_{scan} + \int_{to}^{to+\Delta t} P_{a,RAW}\ dt\ [J] \tag{6}$$

where E_{scan} is the energy required in the antenna configuration phase. In this phase, for each searched direction, one has to point the beam in the given direction, and then transmit a pilot sequence and acquire the result to form a decision on the searched direction; this per-direction cost has to be multiplied by the number of necessary search directions N_p. The beam pointing phase will be labeled with "BP"; in our realization the power absorption is mainly due to the Digital to Analog Converter (DAC) needed to drive the analog phase shifters, and this cost is obviously dependent on the specific hardware realization of the antenna and its driver. The acquisition phase corresponds to the time t_{ACQ} required to perform a SSID/RSSI search, and overall the energy cost can be written as

$$E_{scan} = N_p\left(\int_0^{t_{BP}} P_{standby}\ dt + \int_{t_{BP}}^{t_{BP}+t_{ACQ}} P_{a,search}\ dt\right)\ [J] \tag{7}$$

where $P_{standby}$ is the power absorbed during beam pointing, and $P_{a,search}$ the power required during the SSID/RSSI search. On the other hand, obviously initial time at which useful data transmission starts is now $to = t_{BP} + t_{ACQ}$.

The breakeven achieves after a time $\Delta t = t_{be}$ for which $E_{SAW}(t_{be}) = E_{RAW}(t_{be})$; the involved powers are typically constant during the related time phases, and the integrals simplify to a product, thus giving

$$t_{be} = \frac{E_{scan}}{P_{a,SAW} - P_{a,RAW}}\ [s] \tag{8}$$

Otherwise said, t_{be} gives an estimation of how long the RAW system needs to wait in order to perform another antenna configuration without exceeding the consumption of the SAW.

D. Energy Saving: Data rate impact

A second way to save energy is through increase of data rate. Indeed, as long as the power absorbed by the radio module does not depend on data rate, the energy required to transmit a frame is reduced if the transmit time is lowered, which happens increasing the data rate. This invariance of absorbed power with data rate will be indeed verified in the appropriate Sec. V.

In view of this discussion the energy consumption is expressed as:

$$E_{SAW}(N_f, R) = N_f \int_{to}^{to+t_f(R_{SAW})} P_{a,SAW}\ dt\ [J] \tag{9}$$

$$E_{RAW}(N_f, R) = E_{scan} + N_f \int_{to}^{to+t_f(R_{RAW})} P_{a,RAW}\ dt\ [J] \tag{10}$$

where $t_f(R)$ represents the frame duration that is a function of the achieved data rate. In turn, the data rate is a function of the receiver sensitivity which defines the required receive power level in order to use a specific modulation. As in the present case, this characteristics is usually declared in the receiver datasheet.

The energy break even can be found equating (9) with (10):

$$N_f^{be}(R) = \frac{E_{scan}}{(P_{a,SAW}\ t_f(R_{SAW})) - (P_{a,RAW}\ t_f(R_{RAW}))}\ [s] \tag{11}$$

where $N_f^{be}(R)$ is the number of transmitted frames that achieves breakeven, while $t_f(R_{RAW})$ and $t_f(R_{SAW})$ are the frame duration related to the data rate employed by the RAW and the SAW system respectively. Therefore, Equation (11) states the quantity of information (number of frames) that has to be transmitted before another antenna configuration can be performed in order to break even with configuration energy cost.

III. Implementation and Description of the Proposed System

A. STM32 Nucleo

The STM32 Nucleo platform is a ULP Advanced Reduced instruction set computing Machine (ARM)-based microcontroller that can be programmed on-the-fly by means of writing the firmware in a small flash memory (of about 1MB). The platform is depicted in Fig. 2a. This solution provides fast prototyping and flexibility since come with a comprehensive software library that can be easiliy included in a standard ANSI C program and provides the standard General Purpose Input/Output (GPIO)s to interface external radio modules, as the IEEE 802.11 b/g, LoRaWAN and other SoC radios [25]. It also offers Input/Output pins that can be configured as DACs, Pulse Width Modulator (PWM)s and switches, that are optimum control signals for reconfigurable antennas.

B. SoC-based radio

The radio module, depicted in Fig. 2b, is based on a low-power IEEE 802.11 b/g SoC which expands the STM32

Nucleo boards with wireless capability [23]. This module integrates a PGA and a power manager which control the active, sleep and standby states of the radio. The module can operates as a base station (server) and sensor node (client) in all possible wireless mode, like access point, ad-hoc and infrastructure mode. As base station the module can accept up to 8 connection simultaneously.

(a) (b)

Fig. 2. This figure shows the main components of the ULP radio platform (STM32 Nucleo Platform and SoC-based radio): ULP microcontroller (a) and Low-power IEEE 802.11 b/g SoC-based (b).

C. Reconfigurable antenna

The antenna exploited for the test is made by two parts, i.e. the radiating elements, depicted in Fig. 3a and the Beam Forming Network (BFN), Fig. 3b. Together they form a conventional phased array, based on micro-strip patch antennas. Such antenna topology consists by a superposition of both parts, in a sandwich-like form with the radiating elements facing the internal part of the structure, avoiding deterioration and external contacts. Radiating elements are aperture coupled microstrip patches built over a substrate stack composed of IS400 and Rohacell (e_r = 4.46 and 1.05, h = 1.55 and 4 *mm* respectively). The aperture feed avoids reflection due to soldering while the stack helps to improve the patch bandwidth [27]. Inter-elements distance has computed by meas of the following equation,

$$d \leqslant \frac{1 - 1/N}{1 + \mid cos(\alpha_{max}) \mid} \lambda_0 \tag{12}$$

where N is the number of radiant elements, α_{max} is the maximum scan angle while λ_0 the free space wavelength. Equation (12) guarantees that the grating lobe does not enter in the visible range of the radiation pattern when the main beam is steered at the maximum scan angle α_{max}. With the selected 4 elements array and a desired scan angle of 45°, (12) gives and inter-elements distance of 65 *mm*. The entire structure has the dimensions of 230*mm* x 260*mm* x 7*mm*. Radiating elements are distributed in such a way to form a directive beam, with a maximum gain of 12*dB*, a narrow beam in the scan plane, approximately 30° of Half Power Beam Width (HPBW), and a wider beam in the vertical plane (HPBW of about 60°). In Fig. 4, measured radiation pattern for horizontal and vertical plane are shown. The BFN, which if formed by one input and four outputs, is in charge to perform the beam steering through a phase shifting between the different branches (while there is an uniform distribution in amplitude) in a constant way, by continuously varying the DC voltages. To achieve this result, several quadrature coupler have been exploited [28]. Loaded by two varactors, the coupler allows to easily change the phase continuously simply by varying the voltage applied to the diodes. When no voltage (or the same voltage) is applied to the BFN branches, the array is pointing broadside, since the phase difference among all output ports is zero. When different voltages are applied to the BFN branches, the main beam is moving in the scan plane. For the realized array, a scan in the range ±45° is achieved and in Fig. 5 some example of measured radiation patterns are depicted. The antenna operates in the Industrial, Scientific and Medical (ISM) frequency band 2.4 - 2.48*GHz* and it comes with an SMA to U.FL connector in order to directly interface the radio module described in Subsec. III-B. The measured reflection coefficients for the antenna prototypes are reported in Fig. 6, which shows a good matching (return loss lower than -15*dB*) in the overall bandwidth.

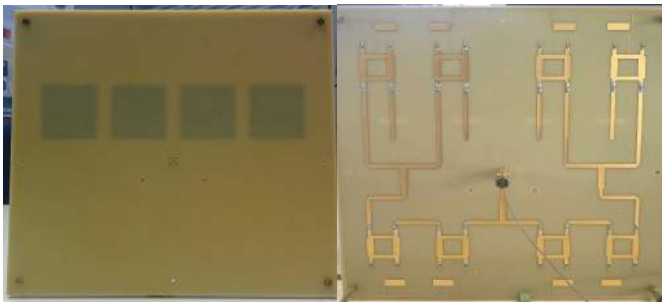

(a) (b)

Fig. 3. Reconfigurable antenna: Radiating elements (a) and Beam Forming Network (b).

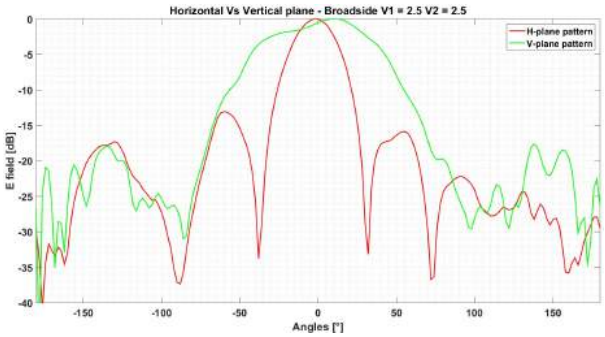

Fig. 4. Reconfigurable antenna: Measured radiation pattern (horizontal and vertical plane)

D. Developed Firmware

The advanced wireless communication capabilities are managed by a firmware running on the STM32 Nucleo board previously described. The developed firmware is *generic* in the sense that can be coupled with any sensor (e.g environmental data loggers etc.) which has a peripheral as well as Serial Peripheral Interface (SPI) or Universal Serial Bus (USB) in order to interface with the proposed radio. For the intended

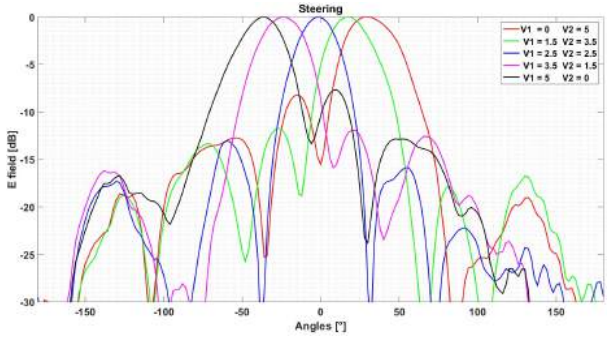

Fig. 5. Reconfigurable antenna: Measured radiation pattern (steering capability in the scan plane)

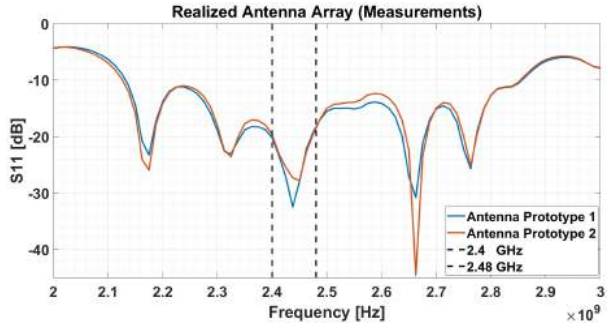

Fig. 6. Reconfigurable antenna - Measured reflection coefficient (S11)

application, the firmware for the STM32 Nucleo has configured with the minimum number of peripherals needed:

- A DAC to control the antenna (i.e. to perform steering).
- GPIO interface for the Nucleo-SoC radio communication.
- SPI to transmits and receive information with the external world, as the ST-SecSoC board for the sensor node.

The developed firmware manages the aforementioned peripherals in the following way:

- An optimization algorithm controls the radiation of the antenna by providing the "right voltage value" to steer the antenna main beam in the best direction (by exploiting the DAC).
- At sensor node side, the firmware on STM32 Nucleo retrieves data from the SPI buffer and transfers it to the SoC radio in order to create the IEEE 802.11 b/g waveform. The STM32 Nucleo and its SoC radio are supplied by the ST-SecSoC board when the latter has events to transmit.
- At base station, the IEEE 802.11 b/g signal is decoded and sent to an external device via SPI, for storage or further processing. In our application we transferred data to a router in order to inform the elaboration center of the event occurrence.

On both sides (i.e. sensor and base station), the communication software creates an ad-hoc network protected by a shared key. In order to communicate, two devices must know the key, otherwise the connection attempt is automatically rejected. The key is generated via software and it is flashed in memory together with the firmware.

The working principle of the energy efficient wireless link is the following. The base station continuously steers the main beam in order to find a sensor node that requires the connection. The radio part of the sensor node is powered only when data need to be transmitted. In this way, the radio will be off for most of the time. As soon as the sensor node supply the radio, the base station directs his beam toward it. Soon after, the sensor node applies the same process to direct the beam toward the base station. At that time the main beam of both sensor node and base station are aligned and the link can advantage of energy savings discussed in Sec.(II).

Finally, with the sole purpose of giving a practical application example, Fig. 7 shows the prototype for surveillance monitoring alluded to in the Introduction.

(a) (b)

Fig. 7. Example of application of proposed system: prototype for road traffic monitoring. (a) The ST - SecSoC module and its camera is visible in the top part; below that, the ULP radio platform (i.e. STM32 and IEEE 802.11 b/g SoC) on the right, a the board with conditioning circuitry on the left (conditioning circuitry is employed to interface the ST - SecSoC board with the ULP radio platform). (b) The reconfigurable antenna (backside), shown in Fig. 3a and Fig. 3b.

E. Antenna reconfiguration tests

In order to validate the steering capability of the proposed system we chosen our laboratory that is an optimal scenario for testing against interferers. Fig. 8 illustrates the base station and the wireless sensor node (i.e. $OPERAVAL1$) slightly at the right side. The most interfering contributions come from wireless routers present in the same room, $LACE-2G$ on the left side and TP_LINK at about $2m$ behind the base station. As soon as the wireless sensor node is powered, the base station detects the $OPERAVAL1$ network (i.e. the desired SSID) and scan the visible space. A step of 10° has been chosen to enhance the precision over the received RSSI. For each visited position the system maps the SSID/RSSI of all devices operating at $2.4GHz$, generating the matrix reported in Tab. I. At this time, the base station algorithm process the matrix in order to find the best compromise between the maximum RSSI received from $OPERAVAL1$ against the RSSIs of the most interfering signals (i.e. $LACE-2G$ and $TP-LINK$). As validated in the proposed tests, the system is able to spread the overall antenna radiation in the intended direction (i.e. 15°) enhancing the communication capability with the target. In this case, $LACE-2G$ results attenuated of $8dB$ and TP_LINK of $6dB$ with respect the maximum

TABLE I
BASE STATION (RX) ANTENNA CONFIGURATIONS - THE TABLE'S CONTENT REPRESENTS THE ELABORATED RSSI[dBm] FOR EACH POINTING DIRECTION OF THE ANTENNA MAIN BEAM

	Link directions $\theta[°]$									
SSDI	$-45°$	$-35°$	$-25°$	$-15°$	$-5°$	$5°$	$15°$	$25°$	$35°$	$45°$
OPERAVAL1	-45	-46	-46	-48	-46	-44	-41	-43	-44	-45
LACE − 2G	-39	-40	-40	-43	-45	-50	-49	-52	-54	-54
polito	-84	-92	-84	-88	-78	-81	-89	-90	-85	-89
eduroam	-80	-91	-92	-89	-84	-91	-83	-85	-85	-90
PiTEAC	-83	-84	-83	-85	-86	-86	-90	-91	-84	-85
Almanac	-86	-90	-89	-90	-81	-90	-82	-81	-90	-92
CDI	-76	-78	-70	-69	-65	-64	-68	-73	-72	-74
TP − LINK	-49	-49	-49	-48	-47	-47	-47	-47	-47	-48

TABLE II
SENSOR NODE (TX) ANTENNA CONFIGURATIONS - THE TABLE'S CONTENT REPRESENTS THE ELABORATED RSSI[dBm] FOR EACH POINTING DIRECTION OF THE ANTENNA MAIN BEAM

	Link directions $\theta[°]$									
SSDI	$-45°$	$-35°$	$-25°$	$-15°$	$-5°$	$5°$	$15°$	$25°$	$35°$	$45°$
OPERAVAL1	-49	-47	-45	-43	-41	-44	-47	-51	-53	-54
LACE − 2G	-60	-58	-56	-54	-51	-50	-46	-42	-42	-42
polito	-90	-88	-90	-90	-86	-83	-83	-83	-85	-86
eduroam	-87	-91	-90	-90	-87	-83	-89	-85	-87	-86
PiTEAC	-87	-89	-91	-91	-90	-88	-85	-83	-86	-90
Almanac	-90	-85	-83	-87	-93	-93	-91	-89	-87	-85
CDI	-73	-75	-74	-71	-71	-69	-72	-75	-76	-80
TP − LINK	-50	-51	-50	-48	-48	-47	-49	-48	-48	-50

value of $OPERAVAL1$. This is not the case, but moving the beam toward the angle proving the maximum RSSI value of $OPERAVAL1$ could be not always the optimal choice. In fact, in a worst case scenario with an interferer placed between the sensor node and the base station it could be better choosing a different position in which, although the $OPERAVAL1$ RSSI is lower the interferer is reduced. This discussion is more important for the base station than the sensor node because the former receives a lot of signals in the same bandwidth, thus working in a very noisy scenario. For the sensor node is instead more important to be aligned with the base station antenna since the maximum directivity has to be exploited to reduce the transmit power. In reason of this the sensor node perform consequently a scan to align the main beam of the antenna with the base station. The result of the scan is reported in Tab. II. Fig. 9 proposes the polar representation of Tab. I and Tab. II illustrating the final direction the antenna main beams are positioned.

IV. POWER CONSUMPTION AND PERFORMANCE

A. Wireless radio module: power consumption analysis

The ULP radio platform has been analyzed in laboratory in order to validate the relationship between the transmit power and instantaneous power consumption as reported in Fig. 10 and Fig. 11. According to Fig. 11 the measurements has been performed providing $V_{cc} = 3.3V$ to the wireless module. A digital oscilloscope is added in series to the supply voltage and the current absorption were measured for different configurable gains of the integrated PGA. The latter provides a transmit power P_{tx} dynamics that span from 0 to 18dBm, as

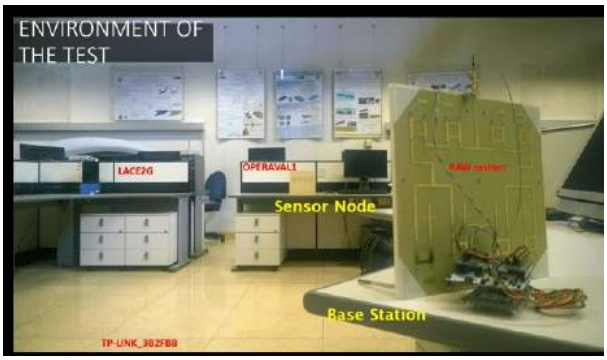

Fig. 8. Antenna configurations for the line of sight communication - test scenario

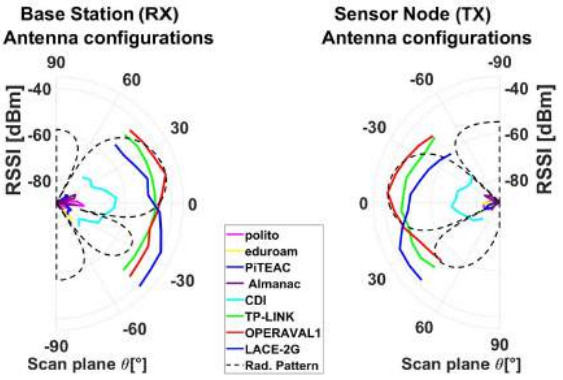

Fig. 9. Elaborated RSSI[dBm] for each pointing direction of the antenna main beam. The beam (Rad. pattern) indicates the final position of the scan process: Base Station (RX) left side - Sensor node (TX) right side

declared in [23]. The overall consumption of the ULP radio platform is reported in Tab. III. In addition to this investigation we also measured the consumption of the standby state which is a timed programmable low power state that cuts off the power to the SoC-radio, while the STM32 Nucleo remains active. This state is used at base station in order to lower its consumption. For instance, the base station scans the space and if there are no sensors that attempt to communicate then it goes in standby for several seconds before resuming with another scan. The sleep state is instead a condition in which the SoC-radio lower its consumption but is always listening in order to wake-up until a configurable number of beacon frames are received. This state has not been employed at sensor node side since better energy efficiency is achieved by supplying the ULP radio platform only when data need to be transmitted. It is the sensor itself that has to manage the power supply to the ULP radio platform (i.e. the ST-SecSoC board) since acts as the master of the communication.

The frame duration, which depends on the employed data rate, was also investigated. The IEEE 802.11 b/g allows a maximum frame length of 1500B of information. In contrast to this we discovered that the module can transmits a maximum of 1000B per frame, thus the estimation of the frame duration is based on this packet length and shown in Tab. IV. The frame duration consists of a header composed by 24B always transmitted at 1$Mbps$, thus providing a fixed timing of 192μs.

TABLE III
OVERALL SOC-BASED RADIO AND STM32 NUCLEO STATES

Radio platform State	**Abs. Current** I_a *[mA]*	**Abs. Power** P_a *[mW]*
Transmit (P_{tx} = 0dBm)	296	977
Transmit (P_{tx} = 10dBm)	310	1023
Transmit (P_{tx} = 18dBm)	404	1333
Receive	152	502
Standby	96	317
Sleep	142	467

TABLE IV
FRAME DURATION FOR DIFFERENT ACHIEVABLE DATA RATES DEFINED BY THE STANDARD IEEE802.11B/G

IEEE-STD	**Rate** R *[Mbps]*	**Required sensitivity** *[dBm]*	**Frame duration** t_f *[µs]*
802.11b	1	-96	8192
802.11b	2	-93	4192
802.11b	5,5	-91	1646.54
802.11b	9	-89.5	1080.89
802.11b	11	-87	919.27
802.11g	18	-86	638.22
802.11g	36	-80	415.11
802.11g	54	-74.5	340.74

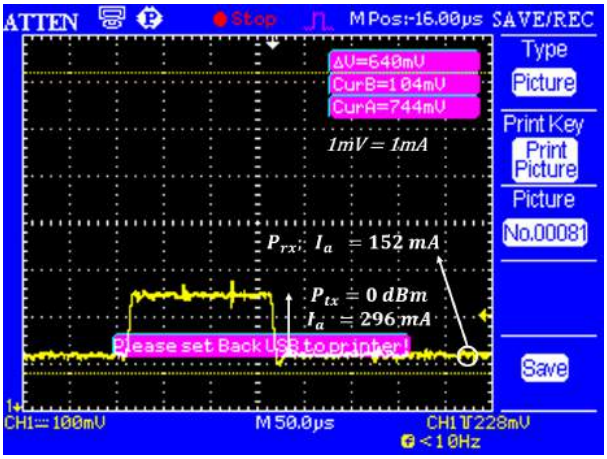

Fig. 10. Current absorption vs frame transmission at P_{tx} = 0dBm

A second header is added to the highest data rate only (i.e. the IEEE 802.11g) and consist of 4B transmitted at the selected rate, however its impact on the overall frame is not so relevant. The rest of the frame duration is given by the quantity of information bytes divided the selected rate.

B. Energy savings due to transmit power reduction

The estimate of the power saving, as per Equation (4) is readily done with the system parameters. In our case the integrated antenna of the SAW has a maximum gain $G_{tx,SAW} = G_{rx,SAW} \approx 2dB$; the RAW system has a maximum gain of $G_{tx,RAW} = G_{rx,RAW} = 12dB$ in the broadside setup, with a gain decrease of 2dB at the maximum scanning direction of 45°. As a result, Equation (4) yields a power saving PW_{saving} ranging from 20dB to 16dB, depending on whether the link direction is in the broadside direction of both antennas (best case), or at maximum scanning angle of both

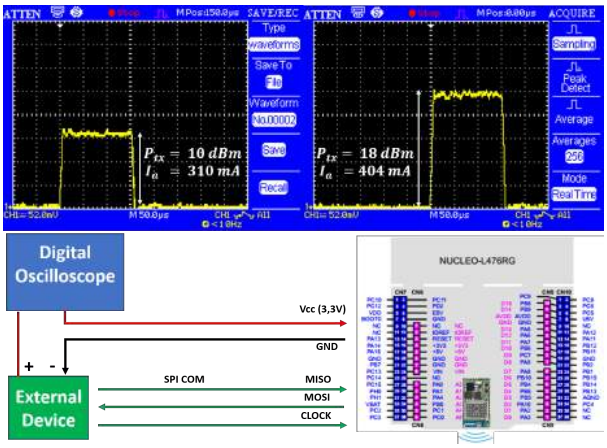

Fig. 11. Measurement test-bed and current absorption transmitting frames at P_{tx} = 10dBm and 18dBm

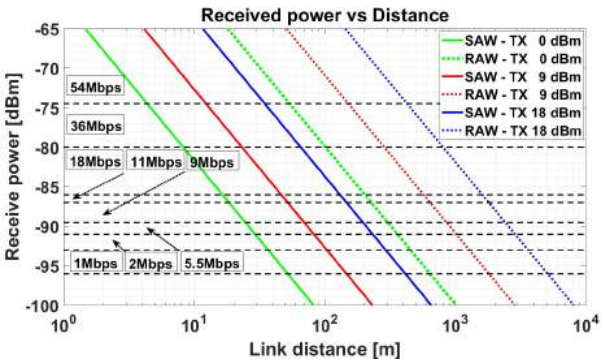

Fig. 12. Link budget: Comparison between state-of-the-art and reconfigurable antenna wireless link for different transmit power

(worst case). Fig. 12 gives a graphical synopsis of the power savings also reporting data rates as per the standard.

With the available hardware the programmable gain range of the amplifier is limited to 18dB, this means that expected power saving will range from 16 to 18dB, depending on link alignment. While the hardware dynamic range limits the achievable power saving, this extra gain (above 18dB) can be used to increase the transmission data rate, which in turn is also a way of reducing energy requirements (see Sec.(II-D). These expected power savings have indeed verified by measuring current absorption as described, resulting in a net power saving from 313mW (worst alignment) to 356mW (gain saturation).

C. Energy budget for the antenna configuration phase

As we have seen, the energy dedicated to the antenna configuration phase has to be accounted for in order to estimate an accurate energy saving. We found that 900ms are necessary to move the beam and perform the SSID/RSSI evaluation. The 40% of such interval is dedicated to set the DAC with the radio module set in the standby state, while for the rest of the time is mostly in the receive state. The overall consumption of the antenna configuration phase depends on the required RSSI accuracy in the scan plane. In other words, the proposed

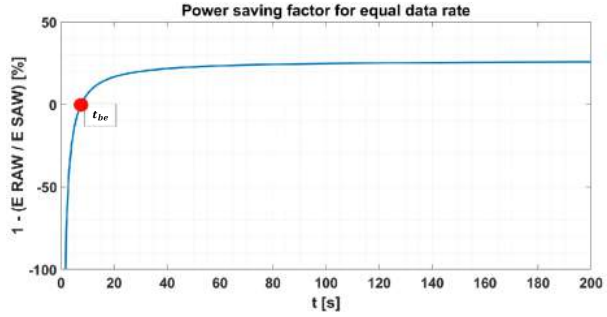

Fig. 13. Power saving factor for equal data rate and breakeven point t_{be}

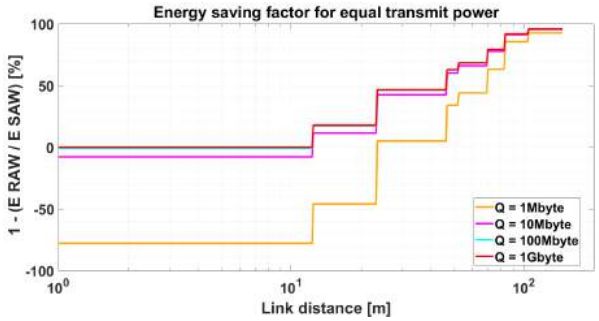

Fig. 14. Energy saving factor for equal transmit power and variable data rate. The saving is shown as a function of link distance d and quantity of transmitted information Q; the discontinuities are due changes in data rate along with the distance, according to the IEEE 802.11 b/g

antenna has a HPBW of 30°, thus an appreciable change of $3dB$ in RSSI is sensed by steering the main beam with steps of 15°. As a result, only 7 points (i.e. −45°, −30°, −15°, 0°, 15°, 30°, 45°) are necessary to analyze the entire scan plane when a precision of $3dB$ is enough. Other solutions are possible, however an increment of accuracy impacts the consumption. With the aforementioned configuration the energy required to scan the whole plane (Eq.(7)) results to be $2.7J$.

D. Overall energy Saving for equal data rate

The RAW and SAW energy consumption has compared for equal data rate. Fig. 13 reports the energy saving ratio in the form $1 - E_{SAW}(\Delta t)/E_{RAW}(\Delta t)$, as a function of the data transmission duration Δt after antenna configuration; the figure is relative to $PW_{saving} = 18dB$. As we already discussed, the RAW system obviously starts in energy debit due to the scan phase, with a negative saving factor. In this case, the breakeven point $E_{RAW}/E_{SAW} = 1$ is obtained at $t_{be} = 7.57s$, as also deduced by (8). After breakeven, the gain quickly increases with transmit time until saturation to $1-(E_{RAW}/E_{SAW}) = 0.27$ (i.e. 27% less energy). We recall that the energy saving in this case is independent of the link distance. However, at $P_{tx} = 9dBm$ the SAW link becomes unable to operate at $150m$, while with the same transmit power the RAW system can achieve a link distance of around $2km$.

V. Energy saving for equal transmit power and variable data rate

In this section the case the two system operate at different data rate is analyzed according to the method presented in Sec.(II-D). Looking at Fig. 12 for equal transmit power (e.g. $9dBm$), the SAW link achieves maximum rate (i.e. $54Mbps$) for few meters, then significantly drops to the minimum (i.e. $1Mbps$) at around $100m$ to consequently being unable to communicate. On the other hand, the RAW is able to maintain the maximum rate along with the SAW operation range. Fig. 14 shows the energy saving factor by comparing (10) with respect (9) for different quantities of transmitted information and in the aforementioned conditions. As intuitive, for a link distance lower than $10m$ there is no benefit in using the RAW system since both wireless links operate at equal rate and transmitted power. In these conditions, and for small quantity of transmitted data (e.g. $Q = 1MB$) the energy required in the scan phase is not compensated by the ensuing power gain, resulting in a negative energy saving factor; however, even at short distances, the initial overhead is fully recovered for large quantities of transmitted data (i.e. greater than $100MB$), with the energy saving factor approaching zero (i.e. breakeven). Significant saving factors are then obtained at larger distances when the SAW system needs to lower the data rate because of the distance (i.e. received power), while the RAW allows to keep it constant. In this case, saving factors that span from 18% to 95% (i.e. $54Mbps$ vs $36Mbps$ and $54Mbps$ vs $1Mbps$ by looking at the curve $Q = 1GB$), can be obtained. Overall, the RAW link achieves best savings on longer distances, especially when the quantity of data to be transmitted is of the order of few MB.

As for the experimental testing, we first verified power absorption as a function of data rate; this has been done by measuring the absorbed current (also as a function of time, with a scope) with different data rates. We observed that power absorption remains unchanged as data rate increases, while absorbed power does drop at times fully coherent with frame duration, in accordance with Tab. IV.

On the other hand, we run into a data rate limitation of the currently employed hardware. We experienced a long latency in between the transmission of consecutive frames, and this latency is constant with respect to employed data rate. Several tests led to establish that this latency is caused by an interface bottleneck, i.e. a slow data transfer between the STM32 Nucleo platform and the IEEE 802.11 b/g SoC radio module; this prevents the latter to transmit a frame until a complete packet of data has been transferred.

This (unexpected) hardware limitation has made it impossible to experimentally verify the saving mode based on variable data rate at present.

VI. Conclusion

The paper has addressed the implementation of reconfigurable antennas on SoC radio demonstrating its feasibility on new generation ULP radio platforms. Our research has demonstrated that the reconfigurable antenna reduces the overall energy consumption of the wireless link, as well as providing a more efficient communication in interfering conditions. The efficacy of the system has been validated and

tested in a critical interfering scenario. We have verified that further enhancement in energy saving is presently hindered by current ULP hardware limitations, with good prospect for future systems.

Work currently under way concerns a limitation of the present beam-steering; with a standard steerable beam antenna the base station can direct its beam toward the direction of only one sensor node at a time. An improved version of this work will show a novel antenna array concept which is currently under investigation to remove this limitation. In fact, an optimization algorithm elaborates the pattern based on the SSID/RSSI matrix that in an adaptive way form a custom multi-beam pattern to host more than one sensor nodes at a time as well as reducing the interference.

ACKNOWLEDGMENT

This work was supported by the OPERA project, which has received funding from the European Union H2020 call ICT4-2015 programme under grant agreement No. 688386.

REFERENCES

[1] G. Giordanengo, L. Pilosu, L. Mossucca, F. Renga, S. Ciccia, O. Terzo, G. Vecchi, V. Romano, and I. Hunstad, "Energy efficient system for environment observation," in *Proc. of the Int. Conf. on Complex, Intelligent, and Software Intensive Systems (CISIS), Torino, Italy*, July 2017, pp. 987 – 999, doi:10.1007/978-3-319-61566-0-93.

[2] D. Ye, D. Gong, and W. Wang, "Application of wireless sensor networks in environmental monitoring," in *Proc. of the Int. Conf. on Power Electronics and Intelligent Transportation System (PEITS), Shenzhen, China*, vol. 1, Dec. 2009, pp. 205 – 208, doi:10.1109/PEITS.2009.5407035.

[3] Z. Zhou, J. Xu, Z. Zhang, F. Lei, and W. Fang, "Energy-efficient optimization for concurrent compositions of wsn services," *IEEE Access*, vol. PP, pp. 657 – 660, Sept. 2017, doi:10.1109/ACCESS.2017.2752756.

[4] H. H. Kenchannavar, S. Beedakar, and U. P. Kulkarni, "Optimization techniques to improve lifetime of wireless sensor networks: A review," in *Proc. of the Int. Conf. on Energy Systems and Applications, Pune, India*, July 2015, pp. 446 – 450, doi:10.1109/ICESA.2015.7503389.

[5] V. Rao, P. P. P, and S. Kar, "Adaptive transmission power protocol for heterogeneous wireless sensor networks," in *Proc. of the Nat. Conf. on Communications (NCC), Mumbai, India*. IEEE, Apr. 2015, pp. 1 – 5, doi:10.1109/NCC.2015.7084898.

[6] D. Basu, G. S. Gupta, G. Moretti, and X. Gui, "Energy efficiency comparison of a state based adaptive transmission protocol with fixed power transmission for mobile wireless sensors," *Journal of Telecommunications System and Management*, vol. 6, 2017, doi:10.4172/2167-0919.1000149.

[7] P. Le-Huy and S. Roy, "Low-power 2.4 ghz wake-up radio for wireless sensor networks," in *Proc. of the IEEE Int. Conf. on Wireless and Mobile Computing, Networking and Communications, Avignon, France*, Oct. 2008, pp. 13 – 18, doi:10.1109/WiMob.2008.54.

[8] G. Manes, R. Fantacci, F. Chiti, M. Ciabatti, G. Collodi, D. D. Palma, I. Nelli, and A. Manes, "Energy efficient mac protocols for wireless sensor networks endowed with directive antennas: a cross-layer solution," in *Proc. of the IEEE Radio and Wireless Symposium, Orlando, FL, USA*, Mar. 2008, pp. 239 – 242, doi:10.1109/RWS.2008.4463473.

[9] G. Dassano and M. Orefice, "Voltage controlled steerable array for wireless sensors networks," in *Proc. of the Eur. Conf. on Antennas and Propagation (EuCAP), Edinburgh, UK*, Feb. 2007, pp. 1 – 4, doi:10.1049/ic.2007.1269.

[10] M. Hanaoui and M. Rifi, "Directional communications with smart antenna system to improve energy efficiency in wireless sensor networks," in *Proc. of the Int. Conf. on Electrical and Information Technologies (ICEIT), Tangiers, Morocco*, May 2016, pp. 203 – 207, doi:10.1109/EITech.2016.7519590.

[11] M. Mofolo and A. A. Lysko, "Incorporating antenna beamswitching technique into drivers for ieee802.11wlan devices," in *Proc. of the IEEE Int. Conf. on Communications (MICC), Kuching, Sarawak, Malaysia*. IEEE Conference Publications, Oct 2015, pp. 96 – 101, doi:10.1109/MICC.2015.7725415.

[12] L. Catarinucci, S. Guglielmi, R. Colella, and L. Tarricone, "Pattern-reconfigurable antennas and smart wake-up circuits to decrease power consumption in wsn nodes," *IEEE Sensors Journal*, vol. 14, no. 12, pp. 4323 – 4324, 2014, doi:10.1109/JSEN.2014.2360939.

[13] T. N. Le, A. Pegatoquet, T. L. Huy, L. Lizzi, and F. Ferrero, "Improving energy efficiency of mobile wsn using reconfigurabledirectional antennas," *IEEE Communications Letters*, vol. 20, pp. 1243 – 1246, 2016, doi:10.1109/LCOMM.2016.2554544.

[14] A. A. Lysko, "Towards an ultra-low-power electronically controllable array antenna for wsn," in *Proc. of the IEEE-APS Topical Conf. on Antennas and Propagation in Wireless Communications (APWC), Cape Town, South Africa*, Oct 2012, pp. 642 – 645, doi:10.1109/APWC.2012.6324932.

[15] E. D. Skiani, S. A. Mitilineos, and S. C. A. Thomopoulos, "A study of the performance of wireless sensor networks operating with smart antennas," *IEEE Antennas and Propagation Magazine*, vol. 54, no. 3, pp. 50 – 67, Aug. 2012, doi:10.1109/MAP.2012.6293950.

[16] R. R. Rout, S. Ghosh, and S. K. Ghosh, "Efficient data collection with directional antenna and network coding in wireless sensor networks," in *Proc. of the Conf. on Advanced Networks and Telecommuncations Systems (ANTS), Bangalore, India*, June 2013, pp. 81 – 86, doi:10.1109/ANTS.2012.6524233.

[17] K. Kucuk, A. Kavak, and H. Yigit, "A smart antenna module using omnet++ for wireless sensor network simulation," in *Proc. of the Int. Symp. on Wireless Communication Systems (ISWCS), Trondheim, Norway*, Dec. 2007, pp. 747 – 751, doi:10.1109/ISWCS.2007.4392440.

[18] STM, "Stm32 32-bit arm cortex mcus." [Online]. Available: http://www.st.com/en/microcontrollers/stm32-32-bit-arm-cortex-mcus.html

[19] K. S. Hwang, J. Ahn, K.-J. Kim, H. K. Yoon, and Y. J. Yoon, "Pattern reconfigurable antenna for a wireless sensor network sink node," in *Proc. of the Asia-Pacific Microwave Conference, Yokohama, Japan*, Mar. 2010, pp. 2021 – 2024.

[20] A. Dihissou, A. Diallo, P. L. Thuc, and R. Staraj, "Technique to increase directivity of a reconfigurable array antenna for wireless sensor network," in *Proc. of the Eur. Conf. on Antennas and Propagation (EUCAP), Paris, France, France*, May 2017, pp. 606 – 610, doi:10.23919/EuCAP.2017.7928636.

[21] L. Lizzi, F. Ferrero, J. M. Ribero, R. Staraj, T. N. Le, A. Pegatoquet, and L. H. Trinh, "Differential pattern-reconfigurable antenna prototype for efficient wireless sensor networks," in *Proc. of the IEEE Int. Symp. on Antennas and Propagation (APSURSI), Fajardo, Puerto Rico*, Oct. 2016, pp. 1239 – 1240, doi:10.1109/APS.2016.7696327.

[22] M. Rzymowski and L. Kulas, "Design, realization and measurements of enhanced performance 2.4 ghz espar antenna for localization in wireless sensor networks, zagreb, croatia," in *Proc. of the IEEE Eurocon*, Oct. 2013, pp. 207 – 211, doi:10.1109/EUROCON.2013.6624988.

[23] STM, "St serial-to-wi-fi modules." [Online]. Available: http://www.st.com/content/st_com/en/products/wireless-connectivity/wi-fi/spwf01sc.html

[24] W. San-Um, P. Lekbunyasin, M. Kodyoo, W. Wongsuwan, J. Makfak, and J. Kerdsri, "A long-range low-power wireless sensor network based on u-lora technology for tactical troops tracking systems," in *Proc. of the Conf. on Defence Technology (ACDT), Phuket, Thailand*, March 2017, pp. 32 – 35, doi:10.1109/ACDT.2017.7886152.

[25] STM, "St wireless connectivity." [Online]. Available: http://www.st.com/en/wireless-connectivity.html

[26] Antenova, "Antenova m2m, 2.4ghz antennas." [Online]. Available: http://www.antenova-m2m.com/product/2-4-ghz-rufa-right-smd-antennas/#

[27] T. A. Milligam, "Chapter 6: Aperture-coupled stacked patches," in *Modern Antenna Design, 2nd ed (John Wiley and Sons)*, July 2005, isbn:978-0-471-45776-3.

[28] E. C. Niehenke, V. V. DiMarco, and A. Friedberg, "Linear analog hyperabrupt varactor diode phase shifters," in *Proc. of the IEEE Int. Microwave Symposium Digest (MTT-S), St. Louis, MO, USA, USA*, Feb. 1985, pp. 657 – 660, doi:10.1109/MWSYM.1985.1132067.

The Green computing continuum: the OPERA perspective

The Green Computing Continuum: the OPERA Perspective

A. Scionti, O. Terzo, P. Ruiu, G. Giordanengo, S. Ciccia, G. Urlini, J. Nider, M. Rapoport, C. Petrie, R. Chamberlain, G. Renaud, D. Tsafrir, I. Yaniv, and D. Harryvan

Abstract Cloud computing is an emerging paradigm in which users access to a shared pool of computing resources dynamically allocated (i.e., ubiquitous computing service), depending on their specific needs. Such paradigm exploits the infrastructural capabilities of modern data centers to provide computational power and storage space required to satisfy modern application demands. The seamless integration of Cyber-Physical Systems (CPS) and Cloud infrastructures allows the effective processing of the huge amount of data collected by smart embedded systems, towards the creation of new services for the end users. However, trying to continuously increase data center capabilities comes at the cost of an increased energy consumption. The OPERA project aims at bringing innovative solutions to increase the energy efficiency of Cloud infrastructures, by leveraging on modular, high-density, heterogeneous and low power computing systems, spanning data center servers and remote CPS. The effectiveness of the proposed solutions is demonstrated with key scenarios: a road traffic monitoring application, the deployment of a virtual desktop infrastructure, and the deployment of a compact data center on a truck.

A. Scionti · O. Terzo · P. Ruiu · G. Giordanengo
ISMB, Torino, IT. e-mail: {scionti,terzo,ruiu,giordanengo}@ismb.it

S. Ciccia
Politecnico di Torino, Torino, IT. e-mail: simone.ciccia@polito.it

G. Urlini
STMicroelectronics, Milano, IT. e-mail: giulio.urlini@st.com

J. Nider · M. Rapoport
IBM Research - Haifa, IL. e-mail: {joeln,rapoport}@il.ibm.com

C. Petrie · R. Chamberlain
Nallatech Ltd, Glasgow, UK. e-mail: {c.petrie,r.chamberlain}@nallatech.com

G. Renaud
HPE, Grenoble, FR. e-mail: gallig.renaud@hpe.com

D. Tsafrir · I. Yaniv
Technion – Israel Institute of Technology, IL. e-mail: {dan,idanyani}@cs.technion.ac.il

D. Harryvan
Certios, NL. e-mail: dirk.harryvan@certios.nl

1 Introduction

More powerful and smart computing systems are made possible by the continuous advancements in silicon technology. Embedded systems evolved into modern Cyber-Physical Systems (CPS): smart connected system that are enough powerful to enable (near) real-time interactions with the surrounding environment. For this reason, CPS are at the basis of implementing new services, albeit they generate an enormous amount of data, thanks to their capability of sensing/acting on the environment where they are deployed. Cloud computing, or simply Cloud, is the set of hardware and software technologies used to process and analyse such amount of data, so that it becomes possible to respond to the societal and industrial needs through innovative services. Also, Cloud technologies enables CPS to retrieve useful information to react to the changes in the environment where they operate. However, such welcome capabilities are today counterbalanced by the high power consumption and energy inefficiencies of the processing technologies that traditionally power data center (DC) servers. Generally, DCs have thousands of running servers arranged in multiple racks. Maintaining these infrastructures has high costs, and Cloud providers pay 40% to 50% of their total expenses as power bills [2]. The critical point of these large scale infrastructure is represented by their large inefficiency: the average server utilisation remains between 10% to 50% in most of the DCs [3, 4] and further, idle servers can consume up to 65% of their peak power [3, 5, 6]. New architectural solutions are thus required to provide more responsiveness, scalability, and energy efficiency.

One way of ameliorating the situation is to use virtualization techniques [7], which increase server utilisation by replicating many times the behaviour of a physical machine, by means of the abstraction of the main hardware features. Thanks to virtualization, servers with different resources can dynamically host different "virtual machines" (VMs). Hypervisors, i.e., software devoted to control VMs during their lifetime, play a key role in scheduling the allocation of VMs to the physical servers. Also, hypervisors consume computing resources for achieving their goals, which in turn translates into power consumption. In addition, traditional virtualization systems put large pressure on the servers' memory subsystem, thus further contributing to increase power consumption, and limiting the capability of the hypervisors to pack a large number of VMs on the servers. Cloud paradigm allows acquiring and releasing resources dynamically over time, making the problem of allocating computing and storage resources even more complex. The stochastic nature of Cloud workloads, along with blind resource allocation policies (i.e., resources are usually provisioned to their peak usage) employed by most DCs, lead to poor resource utilisation [8, 9]. Thus, there is a demand for more efficient resource allocation mechanisms, capable of exploiting the large architectural heterogeneity available in modern DCs, as well as capable of leveraging on less memory hungry virtualization systems. Furthermore, such allocation strategy should be helpful in reducing the workload imbalance and also to optimise resource utilisation.

Tackling these challenges, the OPERA project [1] aims at investigating on the integration of high-performance, low power processing elements within highly mod-

ular and scalable architectures. OPERA realises that heterogeneity is a key element for achieving new levels of energy efficiency and computational capabilities. To this end, solutions delivered in the project widely leverage on reconfigurable devices (field programmable gate arrays – FPGAs) and specialised hardware modules to maximise performance and reduce power consumption. Also, improving the efficiency of the virtualization layer represents one of the main objectives of the OPERA project. Improving the way memory is consumed by (virtualized) Cloud applications, adopting more lightweight technologies (e.g., Linux containerisation), and better supporting allocation of heterogeneous computing resources, OPERA aims at greatly reducing the overall energy consumption of next generation Cloud infrastructures, as well as of remotely connected CPS. Specifically, the following objectives have been identified:

- Exploiting heterogeneity through the adoption of different multicore processor architectures (i.e., ARM, X86_64, and POWER), along with specialised accelerators based on reconfigurable devices (FPGAs);
- Automatically splitting workloads in to a set of independent software modules which can be executed as "microservices" on specific target devices, in order to improve the overall energy-efficiency of the Cloud infrastructure;
- Leveraging on direct optical links and the Coherent Accelerator Processor Interface (CAPI) to increase scalability and maximise the performance of the Cloud infrastructure;
- Designing a scalable, modular, and standardised server enclosure (also referred to as "Data Center-in-a-Box") supporting tens to hundreds of processing cores (even with different architectures) in a very compact form-factor;
- Defining and integrating appropriate metrics with a more holistic approach to continuously monitor the efficiency of the Cloud computing platform;
- Exploiting ultra-low power multi-/manycore processors with a dedicated energy-aware instruction set architecture (ISA) to improve processing CPS capabilities in the context of computer vision applications;
- Exploiting reconfiguration and adaptability of the radio communication subsystem to enhance the CPS efficiency;
- Introducing offloading services (possibly accelerated through FPGA boards) to enable remote CPS accessing larger computing capabilities for critical tasks.

The project, spanning over 3 years, is coordinated by STMicroelectronics, along with the support of Istituto Superiore Mario Boella – ISMB (technical coordinator). The consortium involves also industrial partners, as well as public bodies and academic institutions (HPE, Nallatech, IBM, Teseo–Clemessy, Certios, CSI Piemonte, Le Département de Isère, Neavia technologies, Technion) to successfully achieve the above mentioned objectives.

In this chapter, we discuss the set of technologies, architectural elements, and their integration in the OPERA solution, which are the results of the research activity carried out in the project. Specifically, the design, and integration of a high-density, highly scalable, modular, server equipped with acceleration boards is presented, as well as the capability of integrating, at the DC level, HPC-oriented machines

equipped with POWER processors and FPGA accelerators. The integration of an energy-aware Cloud orchestration software is analysed, along with the analysis of the application impact on the memory subsystem, which is at the basis of a more energy-aware optimisation of the running applications. Finally, the extension of the Cloud infrastructure to include new smart and energy efficient CPS is presented: thanks to network connections, a mechanism for offloading some of the computation on the Cloud back-end is discussed as well.

The rest of the chapter is organised as follows. Section 2 introduces state-of-the-art related works; Section 3 gives an overview of Cloud paradigm evolution. In Section 4 and Section 5, we detail OPERA infrastructure resources, spanning from DC servers to remote CPS. Next, in Section 6 we describes the three real-life application scenarios where OPERA solutions is successfully applied. Finally, we conclude the presented chapter in Section 7.

2 Related works

The OPERA project exploits many different technologies with the ambition of integrating them into a more (energy) efficient platform, serving as the basis for creating the next generation Cloud infrastructures.

Improving energy efficiency in Cloud domain has seen several approaches being proposed in recent years [10]. Most of the proposed approaches aim at improving resource allocation strategies, targeting a more efficient use of the computing and storage resource, which in turn translates in reducing energy costs. Various algorithms have been used to implement smarter resource allocation strategies, including: best fit [11], first fit [12], constraint satisfaction problem (CSP) [13], control theory [14], game theory [15], genetic algorithms (GA) and their variants [16], queuing model [17], and also various other heuristics [18, 19]. It is interesting to note that greedy algorithms (e.g., best fit and first fit) are employed in commercial products, such as the distributed resource scheduler (DRS) tool from VMware [20]). As an emerging architectural style in developing Cloud applications, splitting them in to a set of independent microservices, it becomes possible to control their allocation on the Cloud resources with a fine-grain resolution with regards to traditional monolithic services, as well as it becomes possible to predict their incoming.

Resource allocation can be conveniently formalised as a Bin Packing Problem (BPP), which is a NP-hard optimisation problem. Generally, solving large instances of BPP requires the adoption of heuristics. Among the various, evolutionary-based algorithms, such as Particle Swarm Optimization (PSO), provide very promising results. Zhan et al. [31] proposed a hybrid PSO algorithm by combining the standard PSO technique with the simulated annealing (SA) to improve the iterations of the PSO algorithm. In [33], authors introduced a mutation mechanism and a self-adapting inertia weight method to the standard PSO to improve convergence speed and efficiency. Zhang et al. [32] describe a hierarchical PSO-based applica-

tion scheduling algorithm for Cloud, to minimise both the load imbalance and the inter-network communication cost.

Time series based prediction has been also used as a mean for foreseeing Cloud workloads in advance, so that it is possible to reserve and optimally allocate DC resources. A comprehensive set of literature is available in this context. A simplified method for modelling univariate time series is given by autoregressive (AR) model. Combined with the moving average (MA), it leads to the well-known ARMA model. In [21], authors proposed a ARMA-based framework for characterising and forecasting the daily access patterns of YouTube videos. Tirado et al. [22] present a workload forecasting AR-based model that captures trends and seasonal patterns. Prediction was used to improve the server utilisation and load imbalance by replicating and consolidating resources. Other works proposed AR-derived models to dynamically optimise the use of server resources and improve energy efficiency [24]. Similarly, the autoregressive integrated moving average (ARIMA) and autoregressive fractionally integrated moving average (ARFIMA) can be used to capture non-stationary behaviours in the time series (i.e., it performs better in describing mean, variance, autocorrelation of time series, which in turn change over time), and to extract statistical self-similar (or long memory) patterns embedded in the data sets. To this end, several research works have been presented [25, 26]. On the other hand, it is also a common practice for the Cloud providers to automatically provisioning virtual machines (VMs) via a dynamic scaling algorithm, without using any predictive method [27, 28].

Heterogeneity (including GPGPUs, FPGAs, Intel XeonPhi, many-cores) have been popularised in HPC domain to accelerate computations and to lower power consumption costs of computing infrastructures. Thanks to different architectural features of the processing elements, heterogeneous solutions allow to better adapt to the workload characteristics. However, they still are not common in the Cloud computing domain [34], where architectural homogeneity helps to reduce the overhead and costs of managing large infrastructures. Although, Cloud providers started offering instances running on powerful GPGPUs (e.g., Amazon AWS G2 and P2 instances, Microsoft Azure N-type instances), the access to FPGA-based images is still complex and very limited, since the difficulties in programming such devices. Some steps towards easing the programming of FPGAs has been done with language extensions and frameworks supporting the translation of code written with popular high-level languages (C/C++) into synthesisable ones (VHDL/Verilog). Examples of such compilation frameworks are represented by LegUp [41], ROCCC [42], and OpenCL [43]. However, the lack of precise control of the placement and routing of FPGA resources, generally still demands for manual optimisation of the final design. Despite many challenges to address, some works tried to reverse the situation, explicitly targeting Cloud workloads [35, 36, 37]. However, such works only provide a glimpse of what future evolution of programming languages and compilers will enable (in fact, most of this works strongly rely on an hand-made design optimisation phase).

Cloud-connected smart sensors (also referred to as Cyber-Physical Systems) are often demanded for processing intensive tasks to avoid large data transfer to the

Cloud back-end. To this end, CPS are equipped with high-performance processors, which run in an ultra-low power envelop. Examples of such design style are PULPino [38], an ultra-low power architecture targeting IoT domain, and the ReISC core architecture [39]. To achieve high-performance within an ultra-low power envelope, CPS processors exploits a mix of energy-optimised technologies, including an energy-aware instruction set architecture (ISA). Despite manufacturing technology and micro-architectural improvements, many application scenarios require computational capabilities (e.g., real-time elaboration of high-resolution streaming videos) that are far from that offered by an CPS. In that case, offloading computational intensive task to a remote Cloud back-end becomes the only feasible solution. RAPID EU-funded project [44] aims at accelerating the capabilities of low power embedded devices by taking advantage of high-performance accelerators and high-bandwidth networks. Specifically, compute or storage intensive tasks can be seamlessly offloaded from the low-power devices to more powerful heterogeneous accelerators. For instance, such mechanism has been successfully used to enable Android devices with no GPU support to run GPU-aware kernels, by migrating their execution on remote servers equipped with high-end GPGPUs [45]. Communication capabilities (both wired and wireless) are at the basis of the implementation of effective offloading services, and remain largely exploited in CPS deployed in unmanned contexts. Adapting the communication subsystem to the physical channel characteristics [40] greatly help the deployment of CPS also in rural areas, as well as greatly contributes to save energy on the battery.

3 The green computing continuum

Cloud computing (or simply Cloud) provides a way to access computing resources without hosting them on premise. The level of abstraction at which computing resources are accessed (i.e., the way Cloud users interact with the infrastructure) determines the adopted Cloud service model. Three main service models exist: Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastructure-as-a-Service (IaaS). OPERA aims at improving the way services and resources are managed within a DC, thus it mostly considers the IaaS model. At this level of abstraction (i.e., users may deal with the low level infrastructure, by requesting virtual machines and managing their interconnections) resource allocation problem emerges as one of the main challenges to address. In addition, recently we witnessed to the broadening of Cloud infrastructures, beyond their traditional concept. With the fast growth of the number of small embedded devices connected to Internet [46], Cloud providers need to support them anywhere at any time. The infrastructural Cloud support helps to integrate Cloudlets, Fog computing and the “Cloudification” of the Internet of Things (IoT) within traditional infrastructures (see Figure 1). Supporting such growing set of connected devices exacerbate the need for a more efficient way of managing and controlling resources within large-scale computing systems.

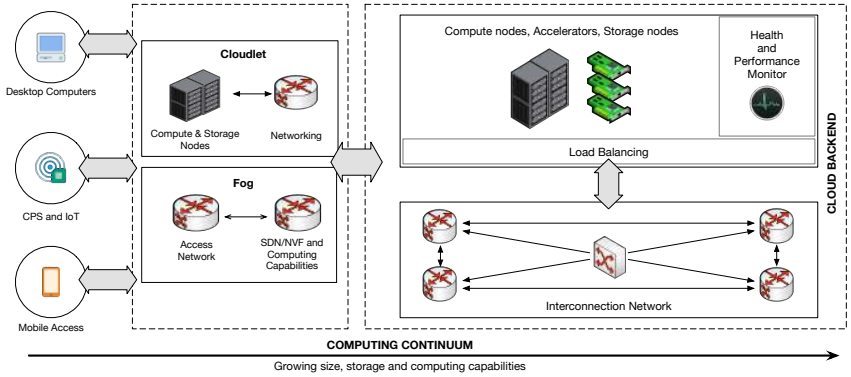

Fig. 1 A visual representation of the computing continuum: from left (smart embedded devices) to right (public cloud data centers), computing and storage capabilities grow. Along with compute and storage, energy consumption grows as well.

Nowadays, mobile computing has gained a momentum thanks to the progress in lowering power consumption of embedded hardware platforms. The higher the computing capabilities are, the more complex the applications running become. Thus, although the improvements in the capabilities of such platforms, Cloud becomes a choice of worth in supporting the remote execution of computing-intensive jobs every time the mobile device is not able to provide enough computational resources. To this purpose, Cloud infrastructures began to reduce the latency for processing jobs and to support (near) real time services. The adopted solutions, termed as Cloudlets [47], are trusted, capable systems co-located with the point-of-presence, and equipped with a fixed number of powerful servers. Using Cloudlets, users can instantiate, on-demand, custom VMs on a cluster of commodity machines, without the need for accessing traditional Cloud services. Fog computing is another emerging paradigm [48]. Here, the idea is to extend the Cloud infrastructure perimeter, by transferring and processing jobs within the network, aiming at reducing the access latency and improving QoS. Fog computing exploits the "virtualization" of the network equipment: router and switch functionalities are melded with more general purpose processing capabilities exposed through specific services [49]. The result of such transformation is a set of new technologies, such as software-defined networking (SDN) and network function virtualization (NFV) [50], which are emerging as a front runner to increase network scalability and to optimise infrastructural costs. SDN and NVF are also two examples of Cloud infrastructure elements that can strongly benefits from more efficient, high-density, computing machines (e.g., FPGA devices represents an optimal substrate for accelerating latency-sensitive network operations). Specifically, SDN provides a centralised, programmable network framework that can dynamically provision a virtualized server and storage infrastructure in a DC. SDN consists of SDN controllers (for a centralised view of the overall network), southbound APIs (for communicating with the switches and routers) and northbound APIs (to communicate with the applications and the busi-

ness logic). NFV helps to run network functions, such as network address translation (NAT) and firewall services as a piece of software. Another extension of the traditional Cloud infrastructures is represented by the Mobile Edge Computing (MEC) [51] paradigm. Here, the aim is to provide a dynamic content optimisation and also to improve the quality of user experience, by leveraging radio access networks (RANs) to reduce latency and increasing bandwidth.

The lowering cost of manufacturing technology led to integrate large computing capabilities, sensors and actuators within the same systems, unleashing the full potential of embedded systems as "cyber-physical systems" (CPS). The large availability of smart sensors/actuators and the corresponding emerging of the Internet-of-Things (IoT) paradigm [52], makes Cloud infrastructures necessary to store and process the enormous amount of collected data. Cloud-IoT applications are quite different compared to the traditional Cloud applications and services (due to a diverse set of network protocols, and the integration of specialised hardware and software solutions). From this standpoint, Cloud-based IoT platforms are a conglomeration of APIs and machine-to-machine communication protocols, such as REST, WebSockets, and IoT-specific ones (e.g., message queuing telemetry transport – MQTT, constrained application protocol – CoAP).

3.1 Energy efficiency perspective

Energy efficiency began an important topic for continuing to sustain the adoption of Cloud and IT technologies. Europe, as other countries, defined the objectives that must be achieved to make Cloud and IT technologies "green". To drive energy efficiency in the EU member states, targets have been detailed in the energy efficiency directive, EED [53]. This energy efficiency directive establishes a set of binding measures to help the EU reaching its energy efficiency target by 2020. Under this directive, all EU countries are required to use energy more efficiently at all stages of the energy chain, from its production to its final consumption. Interestingly, energy intensity in EU industry decreased by almost 19% between 2001 and 2011, while the increased use of IT and Cloud services by both individuals as well as organisations has resulted in a marked increase in datacenter energy use [54]. It is important to note that to decrease final energy use, either total production (output) must decrease and/or efficiency increases must outpace the increase in production.

Although the term "energy efficiency" (E_e) is used often in everyday life, the term warrants careful definition. In mathematical form all energy efficiency metrics are expressed as:

$$E_e = \frac{F_u}{E_{out}} \tag{1}$$

where F_u defines the functional unit (i.e., work done), and E_{out} represents the energy used to produce the output results. For the sake of correctly modelling energy efficiency, there are however many possibilities for describing *i*) energy consump-

tion, *ii*) system boundaries, and *iii*) functional units (this is mostly represented by the type of workload executed) and test conditions.

Looking at the energy consumption of DC equipment, to get an estimate, one must turn to one or more publically available data sources. These data sources are collected as LCI and LCA databases. LCI databases provide Life Cycle Inventory data-sets, while LCA databases include in addition Life Cycle Impact Assessments methods. In environmental impact studies, energy is most often expressed as primary energy (PE). Aside from the grid conversion, LCA studies add other external factors, such as datacenter cooling, into the PE calculation.

System boundaries are necessary to define in order to correctly calculate the energy efficiency of the system under evaluation. It implies to consider only the elements that actually influence the energy behaviour of the system. Further, workload composition becomes crucial to analyse the efficiency of DC equipment. To this standpoint, three main hierarchical levels can be defined. *System boundary*: this level incorporates all the equipment used in delivery of a service, including end-user equipment and transport networks; *Equipment boundary*: this level is a limitation of the system level, focused on a single machine or a tightly knit cluster of machines delivering certain functionalities (both the user and networking are excluded in this view); *Component boundary*: this level is a limitation of the equipment level focused on an element or component inside a machine or the cluster that performs a distinct function (a component view is not limited to a hardware component but can also be a software component).

Workload greatly influence the behaviour of the hardware units, which in turn determines their consumed energy. To demonstrate the effect of workload, it is useful to examine the results coming from the SPECpower benchmark. Interestingly, low server utilisation is common, even in the case of high load conditions. Similar outcomes remain valid for the Amazon EC2 service [55] where large fluctuations are present, but the long term average utilisation is in the range of 10% to 15%. In this range, even with an aggressive power management, servers do not perform efficiently. Not only workload volume influences the efficiency of computations, but also its composition. Different software components (possibly using different technologies, such as Java virtual machines, compiled languages, etc.) perform differently on a given machine, depending on the other software concurrently running. The influence of workload composition is harder to quantify than that of the workload volume. The issue being that what is measurable in a server is the power drawn along with the application output. Varying the workload composition modifies the application output, making comparison of results impossible. However, such differentiation can be turned into an advantage if heterogeneous hardware elements are used. In this case, it becomes possible to search for an advantage in bindings different workload components on different hardware elements, for better performance and energy reduction.

OPERA aims at providing a substantial improvement of the energy efficiency of a DC, by considering workloads made of software components used to deliver end-user services (e.g., access to remote virtual desktops, SaaS-based applications, offloading services, etc.). The influence of the workload volume and composition on

the computing elements is also investigated to better tune the architectural design, as well as the adopted technologies. The boundary is represented by the set of machines that perform computations, including also networking equipment and remote CPS (Cloud end-nodes). Thanks to a wide use of heterogeneous platforms, OPERA aims at greatly improving overall energy efficiency.

4 Heterogeneous data centers

Cloud computing paradigm is based on the availability of a large set of compute, storage and networking resources, which are in general heterogeneous in nature.

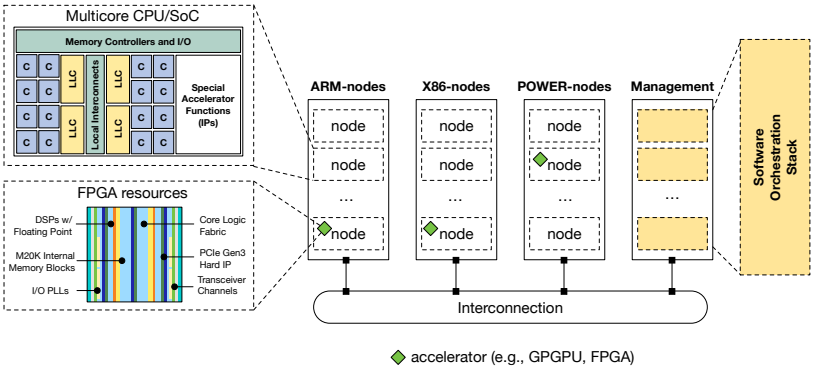

Fig. 2 Conceptual representation of a modern heterogeneous data center. Nodes have different instruction sets and micro-architectures, and can be accelerated through GPGPUs or FPGAs. Internal organisation of multicores and FPGAs is provided (left side).

Unlike in the past, the growing demand for more energy-aware systems, is leading data centers to rapidly embrace of different processor architectures and dedicated accelerators. The former are well represented by ARM-based systems, while the latter are represented by GPGPUs, with a smaller presence of FPGA devices. In OPERA, this practice is further extended by introducing also POWER-based systems. This is a platform originally intended for high-performance computing machines; however, with the emerging of Cloud services supporting HPC-oriented and scientific applications, the availability of dedicated computing nodes becomes more valuable. It is worth to highlight that heterogeneity in data centers extends in several dimensions, not only across architectures: machines configured in the same way and using the same CPU architecture can be still different each other since CPUs families change the features over generations (e.g., the most recent Intel Xeon processors support the AVX-512 instruction set extension which was not available on previous models), as well as the supported clock frequencies.

Figure 2 depicts the conceptual representation of such modern "heterogeneous" data center, highlighting the presence of software stack devoted to management. Looking at the internal organisation of modern chips, one can see the presence of multiple computing cores, along with dedicated interconnections and accelerating IPs (i.e., circuits providing specific functionalities in hardware). One of the main task in governing such kind of infrastructures regards the allocation of resources for the different applications. In order to be "efficient", such allocation should consider several factors to get the optimal allocation decision. To this purpose, in OPERA, power consumption of each platform and the relative load are considered good estimators of the relative energy consumption.

4.1 Accelerated servers

Figure 3 depicts the organisation of the heterogeneous DC as envisioned by OPERA. Hardware differentiation provides the proper computational power required by complex tasks running on top of diverse frameworks (e.g., Apache Hadoop, Apache Spark, MPI, etc.).

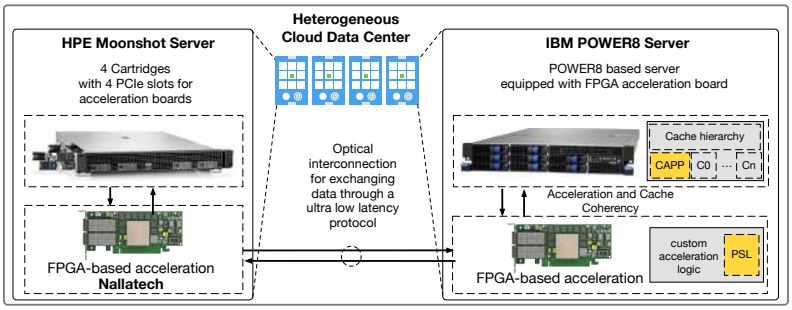

Fig. 3 OPERA envisioned heterogeneous Cloud data center infrastructure: high density, high performance and low power server nodes are linked through optical interconnect (via FPGA cards) with POWER-based nodes. FPGA provides also acceleration resources for computational heavy tasks.

To this end, OPERA integrates low power and high-performance processor architectures (ARM, X86_64, and POWER) into highly dense and interconnected server modules. Developed enclosures (HPE Moonshot) makes it possible to integrate hundreds of different processing elements by exploiting a microserver design: a single cartridge contains the specific processing element (i.e., CPU or DSP), the main memory and the storage. Each cartridge can be coupled with a dedicated accelerator (i.e., GPGPU, FPGA-based board, manycore device). Besides Cloud-specific applications, other algorithms may largely benefit from FPGA acceleration, such as network communication functions and protocols.

Effective usage of hardware components depends on the easiness in accessing their functionalities at the software level. To this end, in OPERA, FPGA accelerators are wrapped by an optimized Board Support Package (BSP) that deals with the low-level details of the FPGA architecture and peripherals (e.g., PCIe, CAPI, SerDes I/O, SDRAM memory, etc.), and that fits into the high-level OpenCL toolflow [57]. This allows the application (or a portion) to be represented as highly portable kernels, that Cloud orchestrator can decide at the last minute eventually to off-load from software running on the host processor to silicon gates. Furthermore, such designed accelerators furnish the DC servers with PCIe and Coherent Accelerator Processor Interface (CAPI) attached programmable logic. With the CAPI [59], the FPGA accelerator appears as a coherent CPU peer over the I/O physical interface, thus accessing to a homogeneous virtual address space spanning the CPU and the accelerator, as well as a hardware-managed caching system. The advantages are clear: a shorter software path length is required compared to the traditional I/O model. The interface is implemented through two hardware sub-blocks: the Power Service Layer (PSL), and the Accelerator Functional Unit (AFU), i.e., the silicon block implementing the acceleration logic. The PSL block contains the hardware resources that maintain cache coherency, acting as a bridge between the AFU and the main CPU. An Interrupt Source Layer (ISL) is available in order to create an access point to the AFU for the software layer. On the CPU side, the Coherent Attached Processor Proxy (CAPP) acts as a gateway for serving the requests coming from and directed to the external AFU. Although the implementation of PSL and ISL units consumes resources on the FPGA (i.e., Flip-Flops, LUTs, RAM blocks, etc.), the adoption of high-end reconfigurable devices leaves enough space to implement complex hardware logic modules. In this context, OPERA leverages on the last Intel products (Arria 10 System-on-Chip – SoC [56]), which supports IEEE-754 Floating Point arithmetic through newly integrated DSPs. Beside pure reconfigurable logic, the Intel Arria 10 SoC features the second-generation dual-core ARM Cortex-A9 MP-Core processor, which is integrated into the hard processor system (HPS) to obtain more performance, flexibility, and security with respect to the previous generation or equivalent soft-cores. On-die ARM cores allow the seamless integration of the reconfigurable logic with a general purpose elaboration pipeline, and in a broader perspective its integration in the complex DC architecture.

Scalability of the platform is further ensured by the adoption of on-board optical links. OPERA opted for standard Quad Small Form-factor Pluggable (QSFP) modules that permit up to 40 Gb/s of bandwidth, physically configured either as a single 40 Gb/s link or split into four independent 10 Gb/s links. This unprecedented level of flexibility allows for supporting a wide range of topologies. The hardware design issues and constraints are abstracted away and automatically handled by the Intel OpenCL compiler, leaving the software programmer to deal only with specific algorithms of interest. The compiler allows optimising high-level code (e.g., C-based code) enabling OpenCL channels (an OpenCL language construct) to be used for kernel-to-kernel or I/O-to-kernel data transfers at high bandwidth and low latency. Channels are also used to implement an application program interface

(API) intended for the host to communicate with the hardware accelerator, generally mapping the PCI Express interconnect.

4.2 Workload decomposition

To take full advantage of hardware specialisation, a mechanism to automatically assign tasks must be put in place. To this purpose, OPERA mostly exploits the *microservice* model. It has recently emerged in the Cloud community as a development style which allows building applications composed by several small independent but interconnected software modules [58]. Each module runs its own processes and communicating with others by means of a lightweight mechanism, typically consisting of an HTTP-based REST API, resulting in an asynchronous, shared-nothing, highly scalable software architecture. In order to automatise the deployment phase of such microservices-based Cloud applications, an ad-hoc "descriptor" is used. It allows the abstraction of the different software components and their relationships. To this end, OPERA leverages on the OASIS TOSCA [60] standard (Topology and Orchestration Specification for Cloud Applications), which enables the portability of Cloud applications and services across different platforms, and that has recently been extended to support Linux containers. TOSCA provides meta-model expressed with an XML-based language. It consists of two main parts: *i*) a topology template – a graph in which typed nodes represent service's components and typed relationships connect nodes following a specific topology; and *ii*) plans – work-flows used to describe managing concerns.

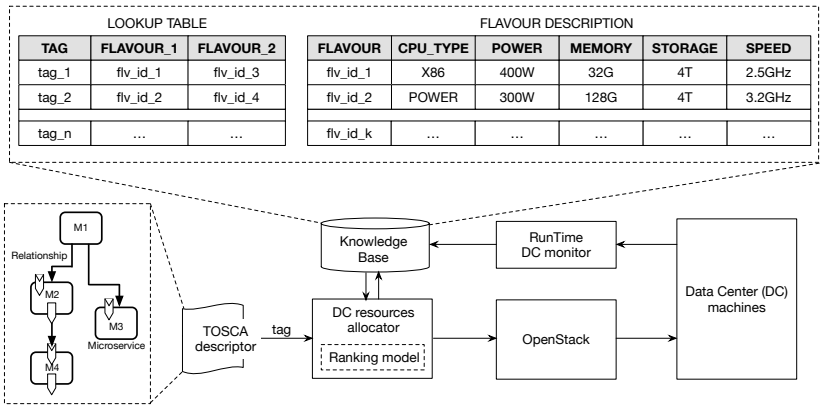

Fig. 4 The deployment chain used in OPERA to allocate resources within the heterogeneous data center. On top, the knowledge base – KB is depicted (table with list of nodes is not represented).

In OPERA, the flexibility offered by the TOSCA application descriptor is exploited to create hooks used to correctly assign microservices to the most suitable

hardware resource for their execution. Although, it is our intention to design and develop a solution that can be used in different context (i.e., not locked to a specific vendor existing platform), OPERA selected one popular platform as a reference: the OpenStack Cloud orchestration system. With regards to the integration of additional components to the OpenStack system, our solution takes into consideration a 2-steps allocation strategy:

- *Phase–1*: deploy the application components on the most suitable platform (i.e., by choosing the better processor architecture, amount of memory, storage space, etc.), with the aim of maximising the energy efficiency;
- *Phase–2*: periodically rescheduling (i.e., migrating) the application components on the most suitable platform if different load/efficiency conditions arise. For instance, a web front-end previously running on an ARM-based server can be moved on a X86_64 machine if the load of the ARM machine exceeded a threshold and/or the number of requests to the front-end increased.

In this regard, we refer to phase–1 as a *static deployment action*, while we talk of *dynamic (re-)scheduling* of the application components in phase–2. To perform such actions, the Cloud orchestrator needs to match microservices with the most suitable hardware resources based on the indication collected in a Knowledge Base (KB), and to monitor the status of the infrastructure.

Since energy consumption ultimately depends on the power consumption of the host machines in the DC, the enhanced Cloud orchestrator (we can refer to it as the Efficient Cloud Resources Allocation Engine – ECRAE) exploits a simple but rather effective power-model to select the host for execution. The power-based model becomes necessary to implement a greedy allocation strategy (Phase–1). Greedy allocation strategy does not ensure an optimal allocation for the whole set of microservices, thus, a further optimisation process is required (Phase–2). Here, a global optimisation algorithms is used to re-schedule all the allocated components with the objective of globally reducing the power (energy) consumption.

4.2.1 Efficiency model

The most critical element in the selection of the actual resource for executing a microservice is the model used to rank the machines belonging to the DC infrastructure. One point to keep into consideration is the relation between energy (E) and power consumption (P). The power P refers to the "instantaneous" energy consumed by a system and generally varies over the time. Based on that, the energy consumed by a system can be computed as the integral of the power consumption function on a given period of time: $E = \int_{t_0}^{t_1} P(t)dt$. The power consumed by a server machine depends on several factors; however, it can be assumed that it is mostly influenced by the consumption of the main hardware components (i.e., CPU, memory, storage and network interface). Power consumption of the CPU and the memory depends on how much the software running on that node stresses these components. Since the load generally changes over time, thus also the power consumed by the

CPU and memory (as well as other components) changes. Also, power consumption in idle state is critical. In literature [3, 6, 62] has been well documented that a conventional server machine, especially if not properly designed, may consume up to 65% of the maximum power consumption in the idle state. It is also worth noting that given the nature of a service, it becomes difficult to foreseen for how much time it will last. For instance, it is not possible to define the amount of time for which a database should run, since it is expected to be accessible unless a failure arises.

To tackle the above mentioned challenges, we elaborated a simple but still effective model for ranking nodes in the infrastructure. By profiling the behaviour of various microservices, we assume that each software component increases the CPU and memory load for a given quantity. Such quantity (C_l – represents the CPU load increase expressed as a percentage, M_l – represents the memory load increase expressed as a percentage) is measured as the average increase generated by the execution of that software component using the host machine in different working conditions. The following equation allows to emit a score value (R) for a given node:

$$R = \{\alpha\, C_l + (1-\alpha)\, M_l\}\; P_{avg} \tag{2}$$

The score R is the weighted measure of the current power consumption P_{avg} of the node (the power weighted value is biased by the power consumption of the node in idle state, so that the P_{avg} value is given by the power consumption in idle state incremented by the fraction due to the machine load); where the weight is expressed by a linear combination of the current CPU load (C_l) and the memory load (M_l). The α parameter allows to tune the power weight, considering the eventual imbalance between CPU and memory load factors. For instance, setting $\alpha = 0.25$, the load on the memory would be equal to 75%. Power consumption (P_{avg}) is obtained as a measure of the average power consumption of the host platform in different working conditions. Averaging the power consumption allows to capture the typical power profile of the host system. Such value is read directly from the Knowledge Base, thus it must be periodically updated to better reflect real machine behaviour. Such mechanism requires the availability of a hardware power monitor and an interface to query it. As part of the deployment mechanism, the envisioned high-density server enclosures are equipped with such monitoring infrastructure, to ease the action of the ECRAE system.

4.2.2 Static allocation strategy

TOSCA provides a hierarchical description of a generic Cloud application, which is at the basis of the mechanism used to trigger the “static” allocation strategy. For each element of the hierarchy, the set of scripts to manage the installation and the main functionalities exposed by the component are provided. The last element in the hierarchy is represented by host features. In OPERA we propose to use a tag to describe the affinity of the software components and the host features. Such affinity is representative of a possible configuration that is evaluated as the most suitable for

Algorithm 1: Static allocation strategy

Input: Knowledge Base (KB)
Output: DC machine where to execute the microservice

$T_c, T_m, \alpha \leftarrow get\_next\_task()$
$tag \leftarrow get\_tag()$
$a_1, a_2 \leftarrow get\_affinity()$
$N_1[\,] \leftarrow get\_node\_list(a_1)$
$N_2[\,] \leftarrow get\_node\_list(a_2)$
$score[\,] \leftarrow \emptyset$
for ***each*** n_{id} ***in*** N_1 **do**
    $C_l, M_l, P_{avg} \leftarrow get\_node\_status(n_{id})$
    $R \leftarrow \{\alpha\, C_l + (1-\alpha)\, M_l\}\ P_{avg}$
    $score[\,] \leftarrow add\_entry(R, n_{id})$
$score[\,] \leftarrow sort(score[\,])$
$sel\_node\_id\_best_1 \leftarrow select\_min\_score(score[\,], R)$
$go \leftarrow$ **True**
while $go =$ ***True or*** $score[\,] = \emptyset$
do
    $sel\_node\_id_1 \leftarrow select\_min\_score(score[\,], R)$
    if $sel\_node\_id_1(C_l) + T_c < 0.98$ ***and***
    $sel\_node\_id_1(M_l) + T_m < 0.98$ **then**
        $go \leftarrow$ **False**
    else
        $score[\,] \leftarrow remove\_entry(R, sel\_node\_id_1)$
if $go =$ ***False*** **then**
    $sel\_node\_id\_best_1 \leftarrow sel\_node\_id_1$
for ***each*** n_{id} ***in*** N_2 **do**
    $C_l, M_l, P_{avg} \leftarrow get\_node\_status(n_{id})$
    $R \leftarrow \{\alpha\, C_l + (1-\alpha)\, M_l\}\ P_{avg}$
    $score[\,] \leftarrow add\_entry(R, n_{id})$
$score[\,] \leftarrow sort(score[\,])$
$sel\_node\_id\_best_2 \leftarrow select\_min\_score(score[\,], R)$
$go \leftarrow$ **True**
while $go =$ ***True or*** $score[\,] = \emptyset$
do
    $sel\_node\_id_2 \leftarrow select\_min\_score(score[\,], R)$
    if $sel\_node\_id_2(C_l) + T_c < 0.98$ ***and***
    $sel\_node\_id_2(M_l) + T_m < 0.98$ **then**
        $go \leftarrow$ **False**
    else
        $score[\,] \leftarrow remove\_entry(R, sel\_node\_id_2)$
if $go =$ ***False*** **then**
    $sel\_node\_id\_best_2 \leftarrow sel\_node\_id_2$
if $score[sel\_node\_id\_best_2] < score[sel\_node\_id\_best_1]$ **then**
    return $sel\_node\_id\_best_2$
else
    return $sel\_node\_id\_best_1$

the execution of the specific component (e.g., a big-memory tag could be used to represent the configuration of a big memory machine, which is well suited for in-memory database operations). The correspondence between the affinity expressed in the TOSCA descriptor and the node configurations is provided by the Knowledge Base. Here, for each configuration, the list of nodes which provide that configuration is also available, along with the actual CPU load, memory load and the average power dissipation. Once the node that better fits with the ECRAE policies has been selected, the corresponding full configuration is used to replace the affinity element in the TOSCA descriptor. The result of such selection process for each application component is a fully-compliant TOSCA description file, which can be transformed into a OpenStack compliant description. Interestingly, although the affinity with a specific host configuration is a static information provided in the TOSCA description, OPERA aims at finding and integrating a mechanism able to allow the orchestration system to "automatically learning" which is the best affinity mapping, as a future investigation direction. To the purpose of fast retrieving information, the KB is organised as a (relational) database. Specifically, three tables provide the required information. One table is used to map tags with a list of possible machine configurations. A second table provides the information of each node matching the specific configuration. Finally, a third table provides the current CPU and memory load for each node. This information is updated by an external component (e.g., Ceilometer module in the OpenStack, Carbon/Graphite, etc.). All these information then are combined into a simple ranking model (see equation 2).

The algorithm used by ECRAE is provided as a pseudo-code (see Algorithm–1). The first step is to extract the information related to the application component to allocate (lines 1–2). This information regards the increment in terms of CPU and memory loads (as a percentage), the tunable α parameter, and the affinity tag. Given the affinity tag, in line 3 the corresponding configurations (affinities) are extracted. We assume that the first configuration (a_1) represents the best match with the requirements of the application component; however, an alternative configuration can be exploited (a_2). Given the effective configurations, the algorithm extracts the list of nodes in the data center that have those configurations (lines 4–5), then it creates an empty list to associate to each node a corresponding R score. In lines 7—10, a loop is used to create such list (the ranking model described in Section 4.2.1), and in line 11 the list is sorted. The first element (line 12) is the candidate node providing the lowest power increase, but it is needed to verify if the corresponding increase in the CPU and memory loads are acceptable. To this end, the algorithm analyses all the available nodes (lines 13–21). If no one of the nodes can be further loaded (a threshold of 98% is assumed), the algorithm maintains the initial candidate solution. Similarly, in lines 22–40, a candidate solution is searched in the list of alternative configurations. Finally, in lines 41–44, the best ever node is returned. Given this node, the corresponding exact configuration is substituted in the TOSCA descriptor, and the application is allocated through the OpenStack environment (i.e., the Linux container or the required virtual machine are instantiated on the selected node).

For instance, lets consider two nodes, each of them belonging to one of the flavours associated to a given affinity tag. For the sake of simplicity we can as-

sume two X86_64 nodes, each in the idle state, but with different average power consumption: we assume $node_1$ consuming up to 100 W (i.e., assuming 65% of idle power consumption that is equal to 65W), $node_2$ consuming up to 130 W (i.e., assuming 65% of idle power consumption that is equal to 84.5 W). Lets assume to schedule two tasks loading the nodes by 45% each (i.e., the CPU load and memory load are assumed equals to 0.45, using $\alpha = 0.5$). Given this premise, the basic allocation policy (i.e., assigning the task to the less loaded node) leads to a higher power consumption, as reported in figure 5.

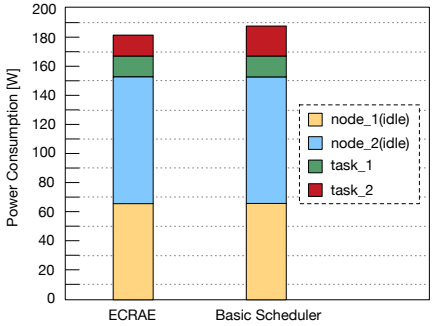

Fig. 5 Example of the power (energy) saving provided by the static allocation strategy: two different nodes are considered respectively with their idle power consumption.

In fact, when the first task is selected, the two nodes are in the idle state and both the strategies (basic and the one implemented by ECRAE) allocate the task to $node_1$. At this point, the power consumption of $node_1$ increases up to 80.75 W, with an overall power consumption, for the two nodes, equals to 165.25 W. On the other hand, the second task is allocated differently. ECRAE ranks the node depending on their weighted power consumption, thus selecting $node_1$ also for the second task (although $node_2$ is less loaded). This provides further 15.75 W of power consumption (increasing the load up to 90%). Conversely, basic allocation strategy selects the node with the lowest load, leading $node_2$ to be selected. In that case, the execution of second task on such node provides 20.475 W of power consumption, leading to an overall power consumption of 185.72 W. Although the power saving for the two nodes is around 5 W, if we consider such saving on a large slice of the DC servers, we obtain a huge improvement in terms of energy saving.

4.2.3 Dynamic allocation strategy

Static workload allocation exploits an aggressive greedy strategy that does not ensure global optimal allocation of the resources. To guarantee such optimality, an instance of a global scheduling problem (i.e., Bin Packing Problem – BPP) must be solved. Workload scheduling is a well-known optimisation problem that fits in the

complexity class of NP-Hard problems. It can be formulated as follows: *Given a set of different objects (VMs or containers), each with a volume S_i (i.e., the amount of CPU and memory used), the objective is to assign as much as possible objects to a bin (i.e., a server machine) that as a finite volume V_j (i.e., a certain amount of CPU and memory offer).* In the context of workload scheduling, the problem requires the minimisation of the number of running machines (i.e., the number of used bins), and also the whole power consumption of the data center (i.e., an heuristic should try to consolidate as much as possible the workload on the minimum number of active hosts). Among various algorithmic solutions, evolutionary-based (and in particular the Particle Swarm Optimisation) techniques allow to quickly solve large instances of this problem.

Particle Swarm Optimisation (PSO) is a population-based stochastic meta-heuristic developed by Kennedy and Eberhart in 1995 to optimise multi-modal continuous problems. In PSO, a group of independent solutions (termed as particles) are used to sample the search space and discover the optimal solution. The state of such group of particles is evolved over time, by updating particles' position within the multi-variable search space. Passing from one position in a given instant of time to another is made by taking into account particles' velocity. The velocity and the position of the particles are taken care by two components, which are described as two factors incorporating a form of distributed intelligence:

- *Cognitive factor*: encodes the information regarding the history of the best position assumed by the particles at certain moment in time.
- *Social factor*: encodes the information relating to the history of the best position assumed by the neighbourhood of the particle at certain moment in time.

These two factors are used to adapt the velocity of the particles in such way it can steer the position towards the optimal solution. PSO does not make use of operators devoted to combining solutions belonging to the same population. On the contrary, the social factor allows to incorporate the knowledge collected by other particles. The topology of the neighbourhood influences the behaviour of the heuristic, although the entire set of particles is used as the neighbourhood (i.e., lattice model). The lattice model also has the advantage of keeping the number of operations used to determine the absolute best position low.

The PSO algorithm is a key component of the ECRAE, since it allows to periodically redistribute the workload on the DC resources, towards their efficient exploitation. At the basis of this periodic imbalance reduction strategy, there is the possibility of migrating application components from one machine to another. Traditional virtual machines can be transparently migrated, but their associated overheads are not acceptable for an effective implementation of the microservices model. To overcome this limitations, Linux containers must be put in place. Linux containers are not designed to be migrated, but such feature becomes important to fully unleash their potential. In the following Section we discuss the way adopted in OPERA to solve this challenge.

4.3 Workload Migration

If microservices are in use (i.e., containerised application components), and the microservices have been designed to be scalable and resilient, it is possible that scaling up (i.e., creating more instances) or scaling down (i.e., destroying instances) of a given microservice is the best way to balance the load. In other cases, creating or destroying containers is not possible, and we must resort to migration. Virtual machines are more heavy-weight example of an execution context, and the price for instantiating new VMs may be considerably higher than migrating an existing one. On the contrary, Linux containers are an example of a light-weight execution context for the microservices model. In the OPERA project, not only we look at how to implement container migration, but also how to perform that in a heterogeneous data center. First, let us list and explain the options that do not make sense. We can then deal with the remaining options as viable, and look at which ones we are focusing on.

In the case of machines with different ISAs (but still general purpose CPUs) VM migration between two machines may have merit. Today, it is possible to migrate a VM to a server exposing a different ISA by employing a binary translation mechanism (such as QEMU), which can emulate the target ISA and translate from the source ISA through software functions. This kind of emulation is at least an order of magnitude slower than running native code, which means it should not be used in performance-critical situations. Due to the amount of software running, it also consumes more energy than the native equivalent, and is therefore not efficient from several perspectives.

Migrating workloads to accelerators (such as GPGPUs) is a possible avenue that may show results in the future. However, nowadays, due to many restrictions and architectural specialisation, accelerator are not able to manage the execution of a whole VM or container. In fact, such technologies relies on features available on general purpose processors, and they are mostly developed around the widely adopted X86_64 architecture. Although in OPERA the objective is to seamlessly access to accelerators (e.g., GPGPUs or even to FPGAs), their usage is limited to the offloading of computational-heavy functions. Here, migration in an heterogeneous context means focusing on multiple ISAs and different types of compute resources, but still relying on general purpose processors. For these reasons it does not make sense to migrate a virtual machine or container to such accelerators. It is likely that much more efficient methods of moving a workload to such compute resources will exist. Furthermore, in the case of FPGAs, the acceleration is mainly a static mapping between a computational-heavy function and a dedicated circuit. Even if we are going to consider the case of "high level synthesis" languages such as OpenCL, there is not yet a clear path for migrating general applications (or portions thereof) to such wildly different compute resources. Further, although techniques to dynamically reconfigure the FPGA devices exist and can be used, their overhead remains too high to justify their implementation on a large scale.

4.3.1 Container Migration

Container migration (and process migration as the generalised case) is a more relevant problem for which we can develop a solution. Containers gain importance due to the advent of microservices, in which application are decomposed, and each component is isolated from the others by means of containers. Some containers run services that cannot easily be replicated or scaled (such as in-memory databases), and the cost to migrate such containers is likely to be less than attempting replication. Moreover, container migration is performed at the system level, which means the application does not need to be aware of how to migrate itself, nor that the migration is even taking place. Thus we may provide migration support for such containerised applications that do not contain support for scaling or replication.

There are several popular implementations of containers on the Linux operating system, such as Docker, LXC/LXD, and Runc. Those container implementations rely on Checkpoint-Restore in Userspace (CRIU) tool for checkpointing, restoring and migrating the containers. Although CRIU relies heavily on advanced features found in the Linux kernel, it does not require any modifications of the kernel itself and it is able to perform checkpoint and restore operations entirely in userspace. At the basic level, CRIU allows freezing a running application and checkpointing it as a collection of files. These files can be later used for restoring the application and continuing execution from the exact point where it was frozen. This basic checkpoint-restore functionality enables several features such as application live migration, application snapshots, remote application analysis and remote debugging. Any flavor of Linux containers can be abstracted as a process tree along with additional properties required for process isolation and fine grained resource management. These processes may have access to various virtual and pseudo devices. CRIU is capable to snapshot the state of the entire process tree as well as the state of the virtual and pseudo devices the processes in the process tree are using. In addition, the properties required for process isolation and fine grained resource management are saved and become an integral part of the container state snapshot.

4.3.2 Comparing post-copy and pre-copy migration techniques

The container state snapshot contains several components that describe process state, open file descriptors, sockets, Linux namespaces, state of virtual and pseudo devices. Yet, all these objects are small and can be easily migrated between different hosts with negligible latency. The part of the container state requiring most of the storage capacity for a snapshot or most of the network bandwidth for a migration is the memory dump of the processes that run in a container. For the case of the container migration, the amount of the memory used by the applications running inside the container defines the time required to migrate the container, as well as the downtime of the application.

The simplest and naive implementation of container migration is as follows: *i*) freeze all the processes inside the container; *ii*) create a snapshot of the entire con-

tainer state, including complete memory dumps of all the processes; *iii*) transfer the container state snapshot to the destination node; and *iv*) restore the container from the snapshot. In this case, the time required to migrate the container and the downtime of the application running inside it are equal and both these times are proportional to the amount of memory used by the processes comprising the container. The application downtime during migration may be decreased with one or more round of memory *pre-copy* before freezing the container. With iterative memory pre-copy, container migration time is slightly longer that in the simple case, but the actual downtime of the application is significantly smaller in most cases. However, such approach may not work for applications with rapidly changing memory working set. For such applications, the amount of modified memory will always be higher than the desired threshold, and therefore the iterative pre-copy algorithm will never converge.

An alternative approach for reducing the application downtime is called *post-copy migration* (this is also the approach adopted in OPERA). With post-copy migration, the memory dump is not created and memory contents is transferred after the application is resumed on the destination node. The primary advantage of post-copy migration is its guaranteed convergence. Additionally, post-copy migration requires less network bandwidth than iterative pre-copy migration, since the memory contents are transferred exactly once. The migration time in this case is small because only the minimal container state snapshot is transferred before the execution is resumed on the destination node. The application downtime is almost as small as the migration time, however, immediately after migration the application will be less responsive because of the increased latency for memory access (this initial low responsiveness is generally tolerated in Cloud applications).

4.4 Optimizing virtual memory management

Also, finding the correlation between using different compute resources and run-time help the development of a methodology for performance estimation, in terms of run-time and energy consumption. Accurate and fast methods for performance estimation will be beneficial, for example, in workload management and dynamic allocation of resources. As the number of huge memory pages that can be allocated and the number of threads are limited in increasing performance perspective, then allocating these limited resources between applications that run on the same system requires to have some model to find the best allocation of these resources, to get the best performance. In modern computing platforms and data centers, the RAM size is not the main performance bottleneck (and energy consumption source), and modern computing platforms can support terabytes of RAM. But, increasing only the RAM size does not increase address translation table (TLB), which is used in modern CPUs to quickly access memory locations. Because TLB capacities cannot scale at the same rate as DRAM, TLB misses and address translation can incur crippling performance penalties for large memory workloads. TLB misses might de-

grade performance substantially and might account for up to 50% of the application run-time [61]. So, increasing only RAM size will not improve the performance for some applications (especially memory intensive ones), that suffer from TLB misses. Therefore, using huge memory pages can save some of these penalties by the fact that using TLB entries of huge pages covers much more memory space than when using the same number of TLB entries of base pages.

There are two main challenges in running the profiling work on modern computing platforms. The first is that these platforms are designed to run multiple tasks on multiple cores, and then we should profile only the running work without being affected by other tasks that run on the same system, or on the same core. The second challenge, is that developing an estimation model requires few different samples of run-time for different page-walks or threads, but getting diversity in TLB page walks for the same workload is more challenging, in terms of controlling the page walks or the allocated huge pages.

Hardware performance counters are set of special-purpose registers built into modern microprocessors to store the counts of hardware-related activities within computer systems. Advanced users often rely on those counters to conduct low-level performance analysis or tuning. The main counters and hardware events we are interested in using for analysing applications' behaviour and drawing an evaluation model are: *i*) the DTLB_LOAD_MISSES:WALK_DURATION, and *ii*) DTLB_STORE_MISSES:WALK_DURATION (this are available on X86_64 processors, but similar can be exploited on different architectures); which respectively count the total cycles Page Miss Handler (PMH) is busy with page walks for memory load and store accesses. The majority of the processors also implement performance counters collecting information regarding the power and energy events, which are of interest to find a correlation between the application performance and the energy efficiency of the hardware in use. An example of such performance counters is represented by the RAPL interface available on the Intel X86_64 processors. RAPL is not an analog power meter, but rather uses a software power model. This software power model estimates energy usage by using hardware performance counters and I/O models. Based on the analysis of the GUPS benchmark (this is representative of memory intensive workloads found in DCs) running on two reference machines, the impact on the memory subsystem can be analysed. Generally, plotting the application run-time as a function of DTLB load page walks provides a simple linear model. Such simple model offers more space for understanding the application behaviour and for optimising the memory access pattern (we can assume that page walks have a linear overhead on run-time). Such conclusion is confirmed by the analysis of the energy consumption as a function of DTLB load page walks. Again a linear model is enough to explain collected data, showing large opportunity to optimise the application energy impact. Although enough accurate to capture application behaviour in most of the cases, more complex model (e.g., quadratic one) should be used to limit the evaluation error.

5 Ultra-low power connected cyber-physical systems

Cyber-Physical Systems (CPS) are becoming an essential part of modern large scale infrastructures, such as in the case of Cloud data centers. CPS provide enough computing and storage capabilities to pre-process captured data, before streaming them in the Cloud back-end. Also, they embed different sensors and actuators, so that they make easy to remotely interact with the surrounding physical environment. However, to further enlarge CPS adoption, more energy efficient technologies must be put in place, as well as a more effective way of exploiting back-end capabilities for processing captured data. To this end, OPERA provides, not only a highly integrated design, but also a mechanism to effectively offload more computational intensive tasks on high-performance accelerated servers.

5.1 Accelerating smart vision applications

Video surveillance is one of the most interesting application fields for smart connected devices. In OPERA, "intelligent" cameras are used to monitor urban traffic aiming at recognising potential situations of risk. Such kind of application covers multidisciplinary fields, spanning computer vision, pattern recognition, signal processing, and communication. The complexity of the application is high, and when it has to be performed in a time-constrained manner, the situation is exacerbated. Meeting the requirements for such applications means implementing advanced hardware systems, albeit generally with a low energy efficiency. From this viewpoint, integration of functionalities in the form of dedicated hardware modules is considered as a technological key feature to increase CPS efficiency. Video surveillance also exposes large parallelism to the hardware: processing functions are applied to (groups of) pixels in parallel.

For this reason, OPERA design adopts a highly-parallel architecture based on an ultra-low power (ULP) manycore solution designed to operate with very low supply voltage and currents. Specifically, it exploits energy efficient cores (EE-cores) which can accelerate several operations in hardware, as well as exploit an energy-optimised instruction set architecture. With such features, the envisioned computing layer can perform operations only requiring few pJ of energy. In addition, a dedicated image processing unit allows performing complex operations, such as moving object detection and image/video compression, at a low energy consumption when compared with standard software implementations. Also, OPERA aims at further improving the performance/power ratio by integrating HW/SW components to accelerate convolutional neural networks (CNNs). CNNs allow, with a limited increase in the used resources, to improve the identification of classes of objects on the scene (this task is the basis for any video monitoring application). However, every time the scene to analyse requires more computational capabilities with respect to that available on the CPS, the analysis task can be effectively offloaded on a remote high-performance server. Thorough a dedicated API, the CPS can access to a Cloud

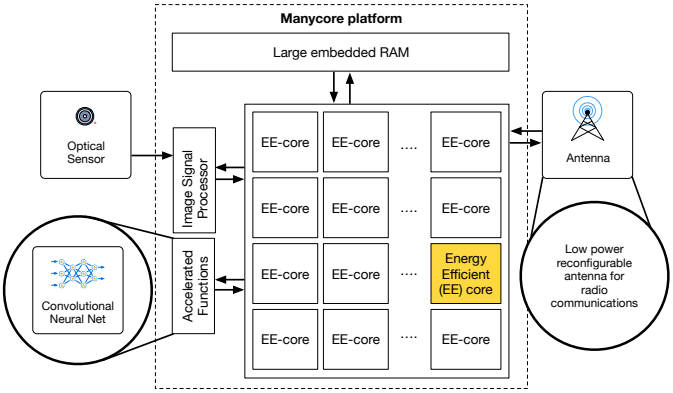

Fig. 6 OPERA Cyber-Physical System architecture: an ultra-low power manycore processor with acceleration function for CNNs is directly attached to the camera sensor and to the reconfigurable communication antenna.

(micro)service: here, more complex algorithm can be run and efficiently accelerated using GPGPUs or FPGA boards.

The complete CPS design (see Figure 6) envisages a low power reconfigurable radio-frequency (RF) communication interface. This kind of communication interface is of particular interest for video surveillance and monitoring applications where the environmental context is particularly critical, such as in the case of mountain roads where connectivity is not reliable. To deal with this issue, our reconfigurable RF module will be capable of adapting the transmission to the best channel/protocol features. Since RF transmission is generally power-hungry, we will design the RF interface with very low power components.

6 The real-life application context

OPERA aims at providing technology demonstration on improved energy efficiency, scalability, and computational performance by resorting to three real-world applications.

6.1 Virtual desktop infrastructure

The purpose of this use case is to demonstrate how the set of technologies taken into consideration by OPERA, allows to make data centers more scalable, and energy efficient. The idea is to exploit the "as-a-service" model to provide a virtual desktop infrastructure (VDI), i.e., a remotely accessible desktop environment. Users are increasingly demanding access to their applications and data anywhere and from

any device. The rapid growth of "nomadic" workers who roam from one computer/device to another leads organisations to provide access to the users' desktop experience at any computer in the workplace, effectively detaching the user from the physical machine. Virtualization is at the basis of this process: employees can access their applications and data very safely over a network, and the risk of data loss is minimised. On the other hand, such practice allows IT departments also to save costs by consolidating services on physical resources (server machines hosted in private or public DCs).

To address this challenge, OPERA implements a solution based on the open-source framework OpenStack. Specifically, OpenStack components such as Cinder and Ceph file system, are used to cover block storage needs for virtual machines and containers. Given the low-latency requirements of this use case, network management is also concerned. To this end, OPERA exploits the flexibility furnished by Neutron. In addition, network latency will be kept low by leveraging on a more powerful remote desktop protocol w.r.t. the traditional protocols (e.g., VNC, RDP, etc.). Finally, to keep as low as possible the overhead of the software virtualization layer, a mechanism based on the KVM hypervisor and containerisation is put in place to run lightweight virtual machines on low power servers.

6.2 Mobile data center on truck

OPERA intends to deliver mobile IT services for the Italian agency called Protezione Civile. IT services, such as forecasting and risk prevention, contrasting and overcoming emergencies, are delivered through a truck (operated by partner CSI Piemonte) equipped with electronic instruments which allow: *i*) creating a satellite communication link; *ii*) acquiring images and videos of area surrounding the truck; and *iii*) processing, temporary archiving and transferring acquired data (videos and images). Images and videos are captured using a drone equipped with cameras and wireless connectivity. The use of a drone is helpful also in the case of dangerous or difficult access to the target site. Once acquired, images and videos are then processed on the truck. Data processing consists of these two steps: *i*) the arrangement of the videos by deleting not useful parts, adding comments, etc; and *ii*) the creation of *orthoimages* for their subsequent comparison with others archived. An orthoimage is a normalised image with respect to a given reference framework, and the creation of orthoimages generally represent a compute-intensive task. For instance, 20 minutes of flight yield 300 photos that require approximately 15 hours to be processed with a standard X86_64 machine (having a resolution of 10^{-2} m. Moving from a standard computer to a high-density but low power server equipped with an FPGA accelerator enables OPERA to greatly speedup the processing task: the elaboration of the same set of photos can be completed in roughly 30 minutes, while keeping low the impact on the power source of the truck (currently, a gasoline power generator).

6.3 Road traffic monitoring

OPERA foresees a growing interest in using remotely-controlled CPS for traffic monitoring purposes in urban and rural contexts. Deploying ultra-low power CPS equipped with effective video processing and wireless communication capabilities, makes possible to monitor large geographic areas in a more energy-efficient manner. For instance, it becomes possible to quickly detect accidents or any situation of risk and communicate alerts (eventually also to vehicles). To this end, collected and pre-processed data are transferred to low power servers located in a remote data centers for further analysis and proactive actions intended to reduce such risk situations.

7 Conclusion

In this chapter, we have presented hardware and software technologies, as well as their integration into a fully functional infrastructure covering the whole computing continuum. Such mix of technological solution are still under development, within the context of the OPERA H2020 European project. This project aims at improving the energy-efficiency of current Cloud computing infrastructures by two order of magnitude when compared with current state-of-the-art systems. To accomplish with this challenging objective, the integration of several advancements on the data center side, as well as on the end nodes of the Cloud is envisioned. Particular attention to the integration of (ultra-)low power high performance technologies is of primary interest for the project.

Specifically, OPERA foresees to gain efficiency on the data center side, by proposing a modular, high-density server enclosure equipped with small low power server boards, and accelerator cards. To maximise the efficiency, FPGA devices will be used in the accelerator boards to provide acceleration for specific kernels as well as low latency connectivity towards POWER nodes. In addition, to fully exploit this wider heterogeneity, applications leverage on a modular architectural style (microservices) that allows to better scale. On the other hand, Cloud end-nodes (i.e., CPS) are made less power-hungry by integrating manycore processing elements with dedicated hardware accelerated functions and reconfigurable wireless communication interfaces. To assess the feasibility of the envisioned platform, OPERA aims at testing its solution on three real-world applications, albeit the results carried out in the project are of interest for a broader community.

Acknowledgement

This work is supported by the European Union H2020 program through the OPERA project (grant no. 688386).

References

1. http://www.operaproject.eu
2. Filani, David and He, Jackson and Gao, Sam and Rajappa, Murali and Kumar, Anil and Shah, Pinkesh and Nagappan, Ram, "Dynamic Data Center Power Management: Trends, Issues, and Solutions.", *Intel Technology Journal*, vol. 12, issue 1, 2008.
3. Barroso, Luiz André and Hölzle, Urs, "The Case for Energy-Proportional Computing", *Computer*, vol. 40, pp.33–37, 2007.
4. Barroso, Luiz André and Clidaras, Jimmy and Hölzle, Urs, "The datacenter as a computer: An introduction to the design of warehouse-scale machines", *Synthesis lectures on computer architecture*, vol. 8, no. 3, pp. 1–154, Morgan & Claypool Publishers, 2013.
5. Greenberg, Albert and Hamilton, James and Maltz, David A and Patel, Parveen, "The cost of a cloud: research problems in data center networks", *ACM SIGCOMM computer communication review*, vol. 39, no. 1, pp.68–73, ACM, 2008.
6. Fan, Xiaobo and Weber, Wolf-Dietrich and Barroso, Luiz Andre, "Power provisioning for a warehouse-sized computer", *ACM SIGARCH Computer Architecture News*, vol. 35, pp. 13–23, ACM, 2007.
7. Pearce, Michael and Zeadally, Sherali and Hunt, Ray, "Virtualization: Issues, security threats, and solutions", *ACM Computing Surveys (CSUR)*, vol. 45, no. 2, pp. 17, ACM, 2013.
8. Srikantaiah, Shekhar and Kansal, Aman and Zhao, Feng, "Energy aware consolidation for cloud computing", *Proceedings of the 2008 conference on Power aware computing and systems*, vol. 10, 20008.
9. Vogels, Werner, "Beyond server consolidation", *Queue*, vol. 6, no. 1, pp. 20–26, ACM, 2008.
10. Kaur, Tarandeep and Chana, Inderveer, "Energy efficiency techniques in cloud computing: A survey and taxonomy", *ACM Computing Surveys (CSUR)*, vol. 48, no. 2, pp. 22, ACM, 2015.
11. Beloglazov, Anton and Abawajy, Jemal and Buyya, Rajkumar, "Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing", *Future generation computer systems (FGCS)*, vol. 28, no. 5, pp. 755–768, Elsevier, 2012.
12. Murtazaev, Aziz and Oh, Sangyoon, "Sercon: Server consolidation algorithm using live migration of virtual machines for green computing", *IETE Technical Review*, vol. 28, no. 3, pp. 212–231, Taylor & Francis, 2011.
13. Van, Hien Nguyen and Tran, Frédéric Dang and Menaud, Jean-Marc, "Performance and power management for cloud infrastructures", *Cloud Computing (CLOUD), IEEE 3rd International Conference on*, pp. 329–336, IEEE, 2010.
14. Zhang, Qi and Zhu, Quanyan and Boutaba, Raouf, "Dynamic resource allocation for spot markets in cloud computing environments", *Utility and Cloud Computing (UCC), Fourth IEEE International Conference on*, pp. 178–185, IEEE, 2011.
15. Ardagna, Danilo and Panicucci, Barbara and Passacantando, Mauro, "A game theoretic formulation of the service provisioning problem in cloud systems", *Proceedings of the 20th international conference on World wide web*, pp. 177–186, ACM, 2011.
16. Quang-Hung, Nguyen and Nien, Pham Dac and Nam, Nguyen Hoai and Tuong, Nguyen Huynh and Thoai, Nam, "A genetic algorithm for power-aware virtual machine allocation in private cloud", *Information and Communication Technology*, pp. 183–191, Springer, 2013.
17. Li, Luqun, "An optimistic differentiated service job scheduling system for cloud computing service users and providers", *Multimedia and Ubiquitous Engineering, MUE'09. Third International Conference on*, pp. 295–299, IEEE, 2009.
18. Li, Kenli and Tang, Xiaoyong and Li, Keqin, "Energy-efficient stochastic task scheduling on heterogeneous computing systems", *Parallel and Distributed Systems, IEEE Transactions on*, vol. 25, no. 11, pp. 2867–2876, IEEE, 2014.
19. Ghribi, Chaima and Hadji, Makhlouf and Zeghlache, Djamal, "Energy efficient vm scheduling for cloud data centers: Exact allocation and migration algorithms", *Cluster, Cloud and Grid Computing (CCGrid), 13th IEEE/ACM International Symposium on*, pp. 671–678, IEEE, 2013.

20. Infrastructure – VMware, “Resource management with VMware DRS”, *VMware Whitepaper*, 2006.
21. Gürsun, Gonca and Crovella, Mark and Matta, Ibrahim, “Describing and forecasting video access patterns”, *INFOCOM, Proceedings IEEE*, pp. 16–20, IEEE, 2011.
22. Tirado, Juan M and Higuero, Daniel and Isaila, Florin and Carretero, Jesus, “Predictive data grouping and placement for cloud-based elastic server infrastructures”, *Proceedings of the 11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, pp. 285–294, IEEE Computer Society, 2011.
23. Roy, Nilabja and Dubey, Abhishek and Gokhale, Aniruddha, “Efficient autoscaling in the cloud using predictive models for workload forecasting”, *Cloud Computing (CLOUD), IEEE International Conference on*, pp. 500–507, IEEE, 2011.
24. Chandra, Abhishek and Gong, Weibo and Shenoy, Prashant, “Dynamic resource allocation for shared data centers using online measurements”, *International Workshop on Quality of Service*, pp. 381–398, Springer, 2003.
25. Kumar, Anoop S. and Mazumdar, Somnath, “Forecasting HPC Workload Using ARMA Models and SSA”, *Proceedings of the 15th IEEE Conference on Information Technology (ICIT)*, pp. 1–4, IEEE, 2016.
26. Calheiros, Rodrigo N and Masoumi, Enayat and Ranjan, Rajiv and Buyya, Rajkumar, “Workload prediction using arima model and its impact on cloud applications’ qos”, *IEEE Transactions on Cloud Computing*, vol. 3, no. 4, pp. 449–458, IEEE, 2015.
27. Iqbal, Waheed and Dailey, Matthew N and Carrera, David and Janecek, Paul, “Adaptive resource provisioning for read intensive multi-tier applications in the cloud”, *Future Generation Computer Systems*, vol. 27, no. 6, pp. 871–879, Elsevier, 2011.
28. Beloglazov, Anton and Buyya, Rajkumar, “Adaptive threshold-based approach for energy-efficient consolidation of virtual machines in cloud data centers”, *Proceedings of the 8th International Workshop on Middleware for Grids, Clouds and e-Science*, vol. 4, ACM, 2010.
29. Chieu, Trieu C and Mohindra, Ajay and Karve, Alexei A and Segal, Alla, “Dynamic scaling of web applications in a virtualized cloud computing environment”, *e-Business Engineering, ICEBE’09. IEEE International Conference on*, pp. 281–286, IEEE, 2009.
30. Lim, Harold C and Babu, Shivnath and Chase, Jeffrey S and Parekh, Sujay S, “Automated control in cloud computing: challenges and opportunities”, *Proceedings of the 1st workshop on Automated control for data centers and clouds*, pp. 13–18, ACM, 2009.
31. Zhan, Shaobin and Huo, Hongying, “Improved PSO-based task scheduling algorithm in cloud computing, *Journal of Information & Computational Science*, vol. 9, no. 13, pp. 3821–3829, 2012.
32. Zhang, Hongli and Li, Panpan and Zhou, Zhigang and Yu, Xiangzhan, “A PSO-Based Hierarchical Resource Scheduling Strategy on Cloud Computing, *International Conference on Trustworthy Computing and Services*, pp. 325–332, Springer, 2012.
33. Liu, Zhanghui and Wang, Xiaoli, “A PSO-based algorithm for load balancing in virtual machines of cloud computing environment”, *International Conference in Swarm Intelligence*, pp. 142–147, Springer, 2012.
34. Crago, S.P. and Walters, J.P., ”Heterogeneous Cloud Computing: The Way Forward”, *IEEE Computer*, vol. 48, no. 1, pp.59–61, 2015.
35. M. Lavasani, H. Angepat and D. Chiou, ”An FPGA-based In-Line Accelerator for Memcached,” in *IEEE Computer Architecture Letters*, vol. 13, no. 2, July-Dec. 15 2014.
36. A. Putnam et al., ”A reconfigurable fabric for accelerating large-scale datacenter services,” *2014 ACM/IEEE 41st International Symposium on Computer Architecture* (ISCA), Minneapolis, MN, 2014.
37. A. Becher, F. Bauer, D. Ziener and J. Teich, ”Energy-aware SQL query acceleration through FPGA-based dynamic partial reconfiguration,” *2014 24th International Conference on Field Programmable Logic and Applications* (FPL), Munich, 2014.
38. A. Traber, et al., ”PULPino: A small single-core RISC-V SoC”, *RISC-V Workshop*, 2016.
39. N. Ickes, et al., ”A 10 pJ/cycle ultra-low-voltage 32-bit microprocessor system-on-chip”, *Proceedings of the ESSCIRC*, Helsinki, 2011.

40. S. Ciccia et al., "Reconfigurable antenna system for wireless applications," *Research and Technologies for Society and Industry Leveraging a Better Tomorrow* (RTSI), IEEE 1st International Forum on, Turin, 2015, pp. 111-116.
41. Andrew Canis, Jongsok Choi, Mark Aldham, Victor Zhang, Ahmed Kammoona, Tomasz Czajkowski, Stephen D. Brown, and Jason H. Anderson, "LegUp: An open-source high-level synthesis tool for FPGA-based processor/accelerator systems", *ACM Trans. Embed. Comput. Syst.*, vol. 13, no. 2, ACM, 2013.
42. J. Villarreal, A. Park, W. Najjar and R. Halstead, "Designing Modular Hardware Accelerators in C with ROCCC 2.0", *18th IEEE Annual International Symposium on Field-Programmable Custom Computing Machines*, pp. 127–134, IEEE, 2010.
43. A. Munshi, "The OpenCL specification", *IEEE Hot Chips 21 Symposium (HCS)*, pp. 1–314, 2009.
44. http://www.rapid-project.eu
45. Montella R., Ferraro C., Kosta S., Pelliccia V., Giunta G., "Enabling Android-Based Devices to High-End GPGPUs", *In: Algorithms and Architectures for Parallel Processing (ICA3PP) – Lecture Notes in Computer Science*, vol. 10048, Springer, 2016.
46. Evans, D., "The internet of things how the next evolution of the internet is changing everything", *CISCO White papers*, 2011.
47. Satyanarayanan, Mahadev and Bahl, Paramvir and Caceres, Ramón and Davies, Nigel, "The case for vm-based cloudlets in mobile computing", *IEEE pervasive Computing*, vol. 8, no. 4, pp. 14–23, IEEE, 2009.
48. Vaquero, Luis M and Rodero-Merino, Luis, "Finding your way in the fog: Towards a comprehensive definition of fog computing", *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 5, pp. 27–32, ACM, 2014.
49. Willis, Dale F and Dasgupta, Arkodeb and Banerjee, Suman, "Paradrop: a multi-tenant platform for dynamically installed third party services on home gateways", *Proceedings of the 2014 ACM SIGCOMM workshop on Distributed cloud computing*, pp. 43–44, ACM, 2014.
50. Martins, Joao and Ahmed, Mohamed and Raiciu, Costin and Olteanu, Vladimir and Honda, Michio and Bifulco, Roberto and Huici, Felipe, "ClickOS and the art of network function virtualization", *Proceedings of the 11th USENIX Conference on Networked Systems Design and Implementation*, pp. 459–473, USENIX Association, 2014.
51. Patel, M and Naughton, B and Chan, C and Sprecher, N and Abeta, S and Neal, A and others, "Mobile-edge computing introductory technical white paper", *White Paper, Mobile-edge Computing (MEC) industry initiative*, 2014.
52. Hwang, Kai and Dongarra, Jack and Fox, Geoffrey C., "Distributed and cloud computing: from parallel processing to the internet of things", Morgan Kaufmann, 2013.
53. European-Commission, "Energy efficiency directive" https://ec.europa.eu/energy/en/topics/energy-efficiency/energy-efficiency-directive
54. Afman M., "Energiegebruik Nederlandse commerciele datacenters". http://www.cedelft.eu/publicatie/energy_consumption_of_dutch_commercial_datacentres%2C_2014-2017/1606
55. Huan L., "Host server CPU utilization in Amazon EC2 cloud", https://huanliu.wordpress.com/2012/02/17/host-server-cpu-utilization-in-amazon-ec2-cloud/
56. Altera Arria 10 FPGAs. https://www.altera.com/products/fpga/arria-series/arria-10/overview.html
57. Khronos Group, "The open standard for parallel programming of heterogeneous systems", https://www.khronos.org/opencl/
58. Thones J., "Microservices", *IEEE Software*, vol. 32, no. 1, 2015.
59. J. Stuecheli, B. Blaner, C. R. Johns and M. S. Siegel, "CAPI: A Coherent Accelerator Processor Interface" in *IBM Journal of Research and Development*, vol. 59, no. 1, 2015.
60. Organization for the Advancement of Structured Information Standards, "OASIS Topology and Orchestration Specification for Cloud Applications (TOSCA)", 2015.
61. Yaniv, Idan, and Dan Tsafrir, "Hash, Don't Cache (the Page Table)", *SIGMETRICS*, 2016.
62. C. Lefurgy, X. Wang, and M. Ware, "Server-level power control", *in Proc. of the IEEE International Conference on Autonomic Computing*, IEEE, 2007.