Deliverable D4.1

Introducing CR elements into smart objects towards enhanced interconnectivity for Smart City applications

Editor:	Elias Tragos, FORTH
Deliverable nature:	Report (R)
Dissemination level: (Confidentiality)	Public (PU)
Contractual delivery date:	28 Feb 2015
Actual delivery date:	31 Mar 2015
Suggested readers:	Researchers, IERC, application developers, system administrators
Version:	1.0
Total number of pages:	150
Keywords:	Internet of Things, Reliability, Availability, Cognitive Radio, spectrum sensing, spectrum assignment, heterogeneous networks, Software-defined Radio, load balancing, frequency assignment, interface virtualisation, semantic representation, resource assignment, trusted routing, multipath forwarding

Abstract

This document presents the first results of Work Package 4 that deals with mechanisms for "Reliability, availability, robustness and scalability" of RERUM. Its goal is to present the techniques for optimizing the network reliability ensuring the always connectivity of the RERUM Devices (RDs) and the high availability of their data. The work presented here describes an integral functional entity of the RERUM System Architecture, namely the Communication and Networking Manager. Key techniques analysed are: (i) the adoption of Cognitive Radio (CR) technology on IoT devices with the development of lightweight and energy efficient mechanisms for spectrum sensing and spectrum assignment, (ii) the maximisation of spectral efficiency with the development of a technique for frequency reuse in clusters of RDs and with the bargaining of resources between neighbour clusters to resolve overloading issues, (iii) the virtualisation of network interfaces of RDs for the efficient assignment of Services to virtual interfaces, (iv) the optimal discovery and assignment of network resources in heterogeneous deployments with RDs that have both WiFi and 3G interfaces, (v) the calculation of trusted routing overlays for identifying trusted intermediate nodes in multipath scenarios and routing "sensitive" traffic only through trusted paths, and (vi) the lightweight multicast forwarding algorithm for RPL-based 6LoWPAN IoT deployments. These techniques will be used for the implementation of the RERUM use cases and will be tested in experiments and trials (if the implementation allows).

Disclaimer

The research leading to these results has received funding from the European Union's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 609094.

Impressum

Full project title	Reliable, resilient and secure IoT for smart city applications
Short project title	RERUM
Number and title of work-package	WP4 - Reliability, availability, robustness and scalability
Number and title of task	T4.1 – Cognitive radio based smart objects T4.2 – Network interconnectivity
Document title	Introducing CR elements into smart objects towards enhanced interconnectivity for Smart City applications
Editor: Name, company	Elias Tragos, FORTH
Work-package leader: Name, company	Elias Tragos, FORTH
Estimation of person months (PMs) spent on the Deliverable	

<this page is intentionally left blank>

Executive summary

This deliverable presents techniques developed within the RERUM project for optimizing the network interconnectivity of the RERUM Devices (RDs) and in general of IoT deployments. These techniques are part of the Devices/Communication layer of RERUM (as presented in deliverable D2.3 [RD23] Figure 27), which handles the communication of the devices and optimizes their networking for ensuring the high availability of their data to support any type of Smart City use cases. These techniques were developed based on the RERUM networking requirements (presented in deliverable D2.2 [RD22]) and aimed at guaranteeing the efficient interconnectivity of large networks of RDs, mitigating interference and ensuring the maximum Quality of Service (QoS).

Acknowledging that the devices are crucial components for the IoT world, RERUM develops techniques to ensure that the devices will always have network connectivity, so that whenever a Service requests some data, the respective device will be capable of providing them on time. This requirement for timely delivery is of outmost importance for many IoT applications that are related to emergencies or health issues and RERUM's mechanisms aim to address it efficiently.

In general, RERUM aims to address some of the key challenges in the wireless IoT world, namely interference, always connected capability, scarce spectrum and mobility. In this respect, the following techniques are analysed in this deliverable:

- Cognitive Radio (CR) technology has been adapted on the IoT world by designing a CR-inspired agent that can run on resource constrained devices,
- A lightweight and energy efficient framework for gathering spectrum occupancy measurements aiming to model the behaviour of licensed users,
- A mechanism for allowing the RDs to select the most appropriate spectrum fragment and the band width that it should access for meeting the service QoS requirements,
- A technique for frequency reuse in clusters of RDs that maximizes the spectral efficiency,
- A mechanism for negotiating and borrowing resources between neighbour clusters of RDs to resolve overloading issues,
- A technique for virtualisation of network interfaces of RDs for the efficient assignment of Services to those virtual interfaces,
- A mechanism for optimal discovery and assignment of network resources in heterogeneous deployments with RDs that have both WiFi and 3G interfaces,
- A framework for calculating trusted routing overlays that identify trusted intermediate nodes in multipath scenarios in order to route "sensitive" traffic only through trusted paths, and
- A lightweight multicast forwarding algorithm for RPL-based 6LoWPAN IoT deployments.

This document includes the results of 11 scientific papers, 6 of which are already published and presented in top tier IEEE conferences [TA13, LHAY14, STT14, SMP14, BLKS14, AAPY15], while 2 are already submitted [STT15, STT15b] and 3 are under submission (an extension of [STT15b] and the works presented in section 5.1 and section 5.2). The purely technical nature of this deliverable may create difficulties to the non-expert readers; thus, an introductory part is included at the beginning of each section describing briefly (i) the motivation for developing each technique that is presented at the specific section, (ii) the relation with the IoT and the RERUM System Architecture and (iii) the relation with the RERUM Use Cases and the practical problem the technique tries to solve.

List of authors

Company	Author	Contribution
FORTH	Elias Tragos George Stamatakis Apostolos Traganitis Alexandros Fragkiadakis Pavlos Charalampides Antonis Makrogiannakis Stefanos Papadakis Manolis Surligas Vasilios Siris	Overall editing of the document. Contribution in Section 1 (introduction), Section 2 (Cognitive Radio mechanisms), section 4 (the introduction part), Section 5.1 (trusted routing overlays) and Section 6.
UNIVBRIS	George Oikonomou Marcin Wójcik	Contribution in Section 5.2 (Bi-Directional Multicast Forwarding Algorithm (BMFA) for RPL-based 6LoWPANs). Reviewing the document.
LiU	Vangelis Angelakis Scott Fowler Ioannis Avgouleas David Gundlegård Niklass Danielsson	Contribution in Section 2.2 (CR-agent), Section 3 (optimization of wireless resources), and Section 4.2 (interface virtualisation)
Cyta	Athanasios Lioumpas	Contribution in Section 4.1.2 (naming/addressing) and Section 4.3 (discovery/assigning resources)
SSRL	Septimiu Nechifor George Moldovan	Contribution in Section 4.1 (semantic representation of Resources). Reviewing the document.

Table of Contents

<this page is intentionally left blank>

Table of Figures

<this page is intentionally left blank>

Table of Tables

<this page is intentionally left blank>

Abbreviations

AWGN	Additive White Gaussian Noise
6LoWPAN	IPv6 over Low power Wireless Personal Area Networks
AES	Advanced Encryption Standard
API	Application Programming Interface
BMFA	Bi-Directional Multicast Forwarding Algorithm
CoAP	Constrained Application Protocol
CR	Cognitive Radio
CRSO	Cognitive Radio Smart Object
CSMA	Carrier Sense Multiple Access
DAC	Digital to Analog Converter
DAG	Directed Acyclic Graph
DODAG	Destination-Oriented Directed Acyclic Graph
DSA	Dynamic Spectrum Access
DSRC	Dedicates Short Range Communications
FFR	Fractional Frequency Reuse
FPGA	Field Programmable Gate Array
HB	Home Box
HBHO	Hop-By-Hop Option
I2C	Inter-Integrated Circuit
ICMPv6	Internet Control Message Protocol version 6
IEEE	Institute of Electrical and Electronics Engineers
IERC	Internet of Things European Research Cluster
IETF	Internet Engineering Task Force
IMSI	International Mobile Subscriber Identity
INR	Interference to Noise Ratio
IoT	Internet of Things
IP	Integer Program
ISM	Industrial Scientific Medical
ITS	Intelligent Transportation Systems
LLN	Low-Power, Lossy Network
LLSEC	Link-Layer Security
LPN	Low Power Nodes
LTE	Long Term Evolution
M2M	Machine to Machine
MAC	Medium Access Control
MAPE	Mean Absolute Percentage Error
MDP	Markov Decision Process
MILP	Mixed Integer Linear Program
MLD	Multicast Listener Discovery
MLME	MAC Sublayer Management Entity
MOP	Mode of Operation
MPL	Multicast Protocol for Low power and Lossy Networks
MQTT	Message Queuing Telemetry Transport
MSPS	Mega Samples per Second
MTU	Maximum Transfer Unit

NAA	Network Aware Applications
NAPS	Naming, Addressing and Profile Server
NCO	Numerically Controlled Oscillator
ND	Neighbour Discovery
NRM	Network Resource Manager
OFDMA	Orthogonal Frequency Division Multiple Access
ONS	Object Name Service
OS	Operating System
PLCP	Physical Layer Convergence Procedure
PP	Partitioning Problem
PU	Primary User
QoE	Quality of Experience
QoL	Quality of Life
QoS	Quality of Service
RD	RERUM Device
RDC	Radio Duty Cycling
REST	Representational Transfer State
RFC	Request For Comments
RL	Reinforcement Learning
RPL	IPv6 Routing Protocol for Low Power and Lossy Networks
RSSI	Received Signal Strength Indicator
RSU	Road Side Units
RTSA	Real-Time Spectrum Analyzer
SDR	Software Defined Radio
SDU	Service Data Unit
SIA	Service to Interface Assignment
SIMD	Single Instruction, Multiple Data
SINR	Signal to Interference-plus-Noise Ratio
SMDP	Semi Markov Decision Process
SMRF	Stateless Multicast RPL Forwarding
SOAP	Simple Object Access protocol
SOS	Sensor Observations Service
SPI	Serial Peripheral Interface
SPS	Sensor Planning Service
SRC	Sample Rate Conversion
SSDF	Spectrum Sensing Data Falsification
SSN	Semantic Sensor Networks
TNDRP	Topology and Network Resource Discovery Protocol
UART	Universal Asynchronous Receiver Transmitter
UCB-E	Upper Confidence Bound Exploration
UE	User Equipment
URI	Uniform Request Identifier
USRP	Universal Software Radio Peripheral
VCAN	Virtualised Content-Aware Network
WSN	Wireless Sensor Network
WWRF	Wireless World Research Forum

Definitions

Term	Definition	Source
Acting element	An (embedded) device that has the capability to affect the condition of a Physical Entity, (like changing its state or moving it) by acting upon an electrical signal	RERUM/ IoT-A part of actuator [IOTA]
Actuator	A smart device that includes one or several acting elements and receives (IT-based) commands translating them to electrical signals for the acting elements. An actuator can also include a sensor so that there is knowledge on the Physical Entity it acts upon, in order to translate correctly the command into the electrical signal.	RERUM/ IoT-A
Application server	The point responsible for the end-user services (e.g., automation services, energy management, etc.). The Application server may reside either in the internet or in the RERUM domain and is responsible for accepting dynamic resource requests, executing the appropriate actions, and returning the results to the user.	RERUM/ IoT-A
Context	Context is any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves.	[AG99]
Cluster	A group of wireless (mainly sensor) nodes that work together for a more efficient and scalable organisation and management of the network.	RERUM, based on [AY07]
Cluster Head (CH)	The RERUM Device that plays the role of the Head of a Cluster within the RERUM network. The CH is responsible for routing the data from the members of the cluster to the rest of the network, as well as to take centralized networking decisions. The CH is either pre-assigned or can be selected by the RERUM Devices.	RERUM, based on [AY07]
Clustering	The process of splitting the network in clusters and electing CHs.	RERUM, based on [AY07]
Consent	Within RERUM the user consent is used for privacy purposes, when the system will ask the user if he allows to send his data to an application that requests them.	RERUM
Device	It can be a single or a combination of the following elements: • Sensors, which provide information about a Physical Entity • Tags, which are used to identify Physical Entities • Actuators, which can modify the physical state of a Physical Entity	IoT-A

Federation Head (FH)	A functional component that executes the process of the Federation of VRDs. It can be assigned to any powerful RD, the GW or a centralized server.	RERUM
Federation of Virtual RERUM Devices	Several Virtual RERUM Devices are forming a Federation if they cooperate to offer a joint service for a Virtual Entity (VE). The logic necessary to orchestrate the service is associated to the Virtual Entity that offers the service.	RERUM
Gateway (GW)	Network node equipped for interfacing with another network that uses different protocols.	Federal Standard 1037C [SF96]
Generic Virtual RERUM Object (GVO)	This is a software artefact that groups both virtualisations found in RERUM, namely the Virtual Entities and Virtual RERUM Devices that share properties like, that • they allow to be discovered, • they allow to be addressed, and they allow to be interacted with in a standardized manner.	RERUM
Internet Resources	These are sources of data/measurements that originate from outside of the RERUM domain and can be used as input for the applications.	RERUM
RERUM Middleware (MW)	Within RERUM, the Middleware is assumed to be a software layer or a group of functionalities that allows heterogeneous devices to be discovered, addressed and accessed by the applications in a seamless and unified way. The Middleware includes the virtualisation of devices to hide their heterogeneity.	RERUM
Physical Entity (PE)	A discrete, identifiable part of the physical environment which is of interest to the user for the completion of his goal. Physical Entities can be almost any object or environment.	Merriam-Webster dictionary[1]/ IOT-A
RERUM Aggregator	A RERUM Device can play the role of an Aggregator, when it collects, processes (aggregates, encrypts, filters, etc.) data/measurements from many other RERUM Devices and forwards them to the GW/Middleware/Application Server. A RERUM aggregator can be considered as an RD playing the role of a Federation Head and could be very helpful in terms of privacy, because this aggregation will avoid the leaking of personal information that may be contained in the data that are aggregated.	RERUM
RERUM Device (RD)	A RERUM Device (RD) is a piece of hardware and software (incl. the Operating System) that is equipped with intelligence. It has	RERUM

[1] Merriam-Webster Online: Dictionary and Thesaurus www.merriam-webster.com

or RERUM Smart Object	one or more Resources that the RERUM Device is able to either fill with interpreted and pre-processed sensory data or able to read and interpret the commands that are given. The RERUM Device has some Sensing, Tag or Acting elements directly attached to it.	
RERUM Deployment	The specific topology of software components on the physical layer, as well as the physical connections between these components.	IoT-A / RERUM
RERUM Gateway	A RERUM Gateway is a physical device that plays the role of a network gateway interconnecting different RERUM networks. Furthermore, the RERUM Gateway is responsible for managing the RDs that are connected to it. In this respect it can also include various Middleware functionalities.	RERUM
Resources	Resources are software components that provide some functionality. When associated with a Physical Entity, they either provide some information about or allow changing some aspects in the digital or physical world pertaining to one or more Physical Entities. In general, they are typically sensor Resources that provide sensing data or actuator Resources, e.g. a machine controller that effects some actuation in the physical world.	IOT-A On-device Resources
Sensing element	An (embedded) device that perceives certain characteristics of the real-world environment (Physical Entities), translating a change into an electrical signal.	RERUM
Sensor	A smart device that includes one or several sensing elements and is able to translate the electrical signal of the sensing elements to some type of information (digital representation) with specific value and semantic.	IoT-A
(IoT/RERUM) Service	Software component enabling interaction with resources through a well-defined interface, often via the Internet.	IoT-A, RERUM
Smart Object	See RERUM Device	RERUM
Virtual Entity (VE)	The digital synchronized representation of a Physical Entity.	IoT-A
Virtual RERUM Device (VRD)	A Virtual RERUM Device (RD) is a digital representation of a RERUM Device. The same one physical RERUM Device at one time is represented by one Virtual RERUM Device. This is a software artefact, like a Virtual Entity (VE), but represents a RERUM Device (RD).	RERUM
User	A Human or a software that interacts with a system for transferring information.	Based on IoT-A

<this page is intentionally left blank>

1 Introduction

1.1 Scope

This document presents the results of the EU-FP7-SMARTCITIES-2013 project RERUM [RERUM] with regards to technologies for enhancing the "Reliability, availability, robustness and scalability" of the system. More specifically, this document is a joint effort of two tasks that have been running for 16 months: Task 4.1 "Cognitive radio based smart objects" and Task 4.2 "Network interconnectivity".

Up until now, the IoT world focused mostly on enabling the interconnectivity of the devices through the virtualisation of the physical devices and objects and the centralized management of the virtual devices. However, developing only server-side or gateway-side mechanisms without any focus on the devices themselves does not solve availability and reliability issues, because it does not solve efficiently the networking problems of the devices. Within RERUM, the devices have a very important role in the system architecture and the goal is to embed intelligence on the devices so that their networking and their interconnectivity is being improved with regards to the availability of the devices to send and receive data. This will improve significantly the reliability of the overall system, increasing the overall availability of the devices so that their Resources can be delivered on-time whenever they are requested by the RERUM Middleware.

Task 4.1 focused on developing a framework that will allow the adaptation of the Cognitive Radio (CR) technology on the IoT devices. The motivation for that was to increase the availability of the IoT communications, enabling the devices to mitigate efficiently the wireless interference in the overcrowded free spectrum bands (Industrial Scientific and Medical – ISM bands). In this respect, as also presented briefly in the System Architecture deliverable (D2.3) [RD23], a lightweight CR-agent has been designed. The CR-agent will be running on the RERUM Devices (RDs) and will be capable of accessing the wireless spectrum in an opportunistic manner, not being limited by standard protocols like IEEE 802.11 or 802.15.4 that are widely used in the IoT world. Furthermore, for the optimization of the wireless interconnectivity of the devices, mechanisms for dynamic frequency assignment and for negotiation of the assignment of wireless resources to devices have been developed and adapted from standard wireless solutions to the specific requirements of IoT devices.

Task 4.2 focused on optimizing the interconnectivity of heterogeneous wireless IoT devices. Virtualisation of the network interfaces of the RDs is a key achievement within RERUM, progressing significantly the state of the art against standard existing solutions. Its goal is to assign in a more efficient and dynamic way the network resources to services considering their Quality of Service (QoS) requirements. In the same area, RERUM developed a mechanism for discovering and allocating network resources (that can be equipped with metadata to allow their semantic representation) to RDs aiming to increase the spectral efficiency of wireless deployments of RDs within city environments. For optimizing the use of virtual resources, algorithms for routing and forwarding have been developed. The former focuses on identifying trusted routes in overlay networks, without requiring knowledge of the underlying network connections, thus is suitable for virtualised environments. The latter focuses on both optimizing the network bandwidth and minimizing the energy consumption of the nodes during message exchanges.

The document, apart from the theoretic evaluations of the proposed techniques, also includes a partial prototype implementation of an Software Defined Radio (SDR) based gateway that is able to interconnect heterogeneous wireless IoT networks. The specifications of the prototype implementation are included in Section 2, and a video recording of its performance as well as a short report will be available in https://ict-rerum.eu/sdr-prototype/.

1.2 Intended audience

This document presents pure technical solutions for improving the reliability and availability of IoT networks, so it has a very narrow targeted audience. It aims mainly for researchers that are working in the areas of Cognitive Radios, Software Defined Radios, network optimization, network virtualisation, modelling of network resources and wireless sensor networks. The solutions presented in the document are explained in detail so that the respective readers can easily implement them and test them on their systems. The document also aims at other IoT related projects and the IERC members to provide them with the RERUM solutions on improving the networking reliability and the data availability of their systems. In this respect, a dialogue with other projects for integrating the RERUM device-oriented solutions with the middleware-oriented solutions of other projects can start, in order to develop jointly an optimised IoT framework with emphasis (but not exclusively focused) on smart city applications.

1.3 Position within the project

1.3.1 Relation with other tasks and WPs

As mentioned above, this document (D4.1) is the output of the activities of Tasks 4.1 and 4.2 of Work Package 4 (WP4). In Figure 1, the relation of the tasks of WP4 with the rest of the tasks/WPs of the RERUM project is depicted. D4.1 used as input the results of Tasks T2.2 and T2.3 (which were described in D2.2 [RD22]) and the results of Task 2.4 (that were described in D2.3 [RD23]). From Task 2.2 the input relates to the adaptation of the CR-agent model for the RDs. From Task 2.3 the input is related to the networking, system and device requirements that were used as basis for the design and the development of the respective technologies that are developed within Tasks 4.1 and 4.2 and are detailed in this document. From Task 2.4, the input to this deliverable (and the rest of WP4) is related to the specific modules of the system architecture that are developed within WP4.

The output of D4.1 is used within Task 4.3 (D4.2) for optimizing the developed mechanisms for their energy efficiency, as well as within Task 4.4 (D4.3) for assessing their scalability. Furthermore, D4.1 provides output to Task 5.2 for the implementation of the mechanisms for running the use cases within experiments and trials in Tasks 5.3, 5.4 and 5.5. The early results from the experiments and trials will be used for refining and optimizing the developed mechanisms within WP4 that would provide results for the optimization of the implemented mechanisms within WP5 (as it is depicted in the feedback loop shown with the red lines in the figure). D4.1 will also provide feedback to Task 2.4 for the refinement of the RERUM architecture within the second year of the project.

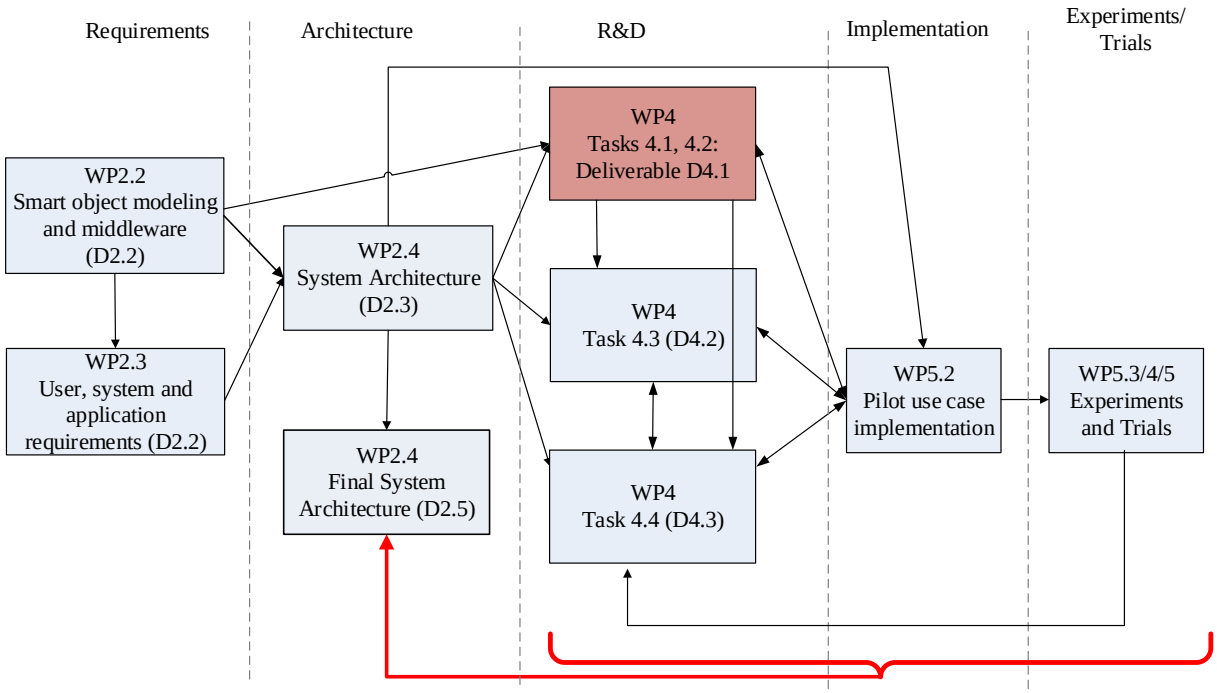

Figure 1. Position of D4.1 within the project activities

1.3.2 Relation with the architecture

D4.1 is the main technical deliverable of WP4 that aims to develop mechanisms for enabling the seamless and efficient communications (virtualised or not) of the RERUM Devices (RDs). Here we will briefly describe the position of this deliverable with regards to the functional entities of the RERUM system architecture that were presented in Deliverable D2.3 [RD23].

The contents of this deliverable are related with the functional entity “Communication and Network Manager” that is depicted in Figure 2. This entity consists of two entities, namely the “Communication Manager” and the “Network Manager”.

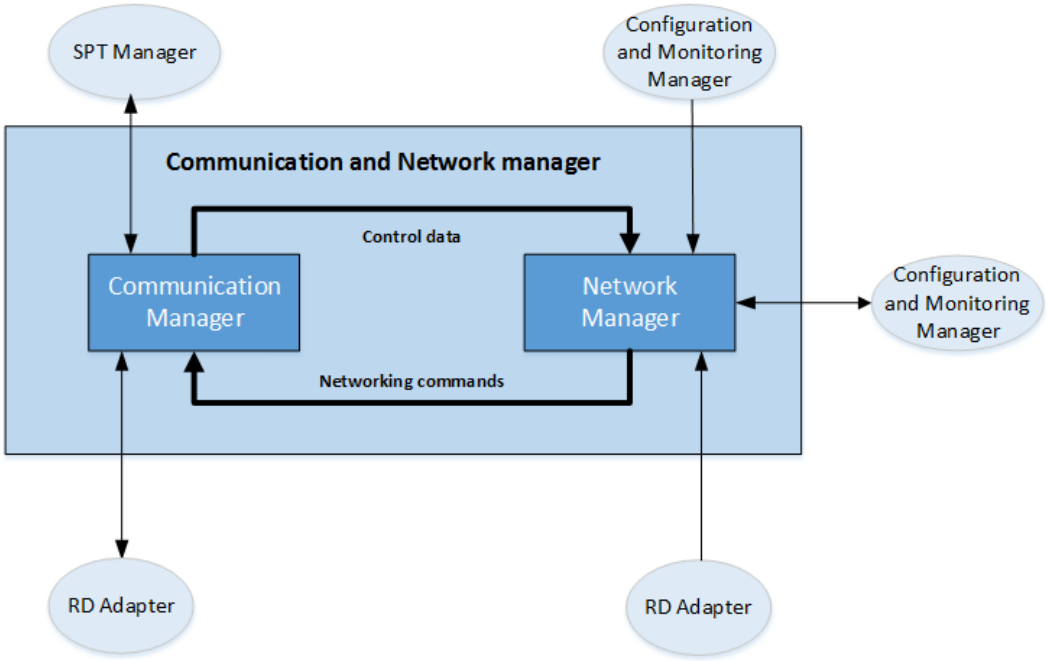

Figure 2. Communication and Network Manager internal components

The **Communication Manager** consists of four functional components as seen in Figure 3. The figure presents the conceptual model of the functional entity, which does not necessarily mean that all the depicted modules will be developed within RERUM. For several parts that are not of major interest of the project, existing solutions can and will be applied. In this deliverable, the focus is given on the following functional components:

- **Interface selection,** parts of which are described in Section 2 and parts in Section 4 in this deliverable,
- **Routing,** parts of which are described in Section 5 in this deliverable that presents the multicast forwarding,
- **Protocol translation,** described in Section 2 in this deliverable where the CR-based gateway is presented.

Figure 3. RERUM Communication Manager functional components

The module for **Communication Security** is the responsibility of WP3 and its internal functionalities are described in detail within Deliverable D3.1 [RD31], so the reader is advised to refer to this deliverable for more information on this module. The functional component for **Scheduling** is not a standalone module, but parts of it are included in other functionalities, i.e. in the interface selection, when the Spectrum Assignment is performed, the flows are assigned to spectrum fragments according

to their priority and their transmission takes place in a prioritized manner and similarly in the Routing module.

The **Network Manager** consists of eight functional components as depicted in Figure 4. It is responsible for managing and configuring the components that deal with the network connectivity of the RDs either in a centralized or in a distributed approach (mainly depending on the design choice of the network administrator). This figure presents the conceptual model and not all components are developed and implemented within RERUM. The following modules are described within D4.1:

- **Spectrum management**: the CR-related functionalities are described together with the interface selection in Section 2. However, the interference management functionalities are described in Section 3,
- **Local QoS manager**: basically it is inherent in all sections, because all developed functionalities take into account the QoS requirements of the services. However, specific QoS-related functionalities are described in Section 4,
- **Routing management**, described in Section 5 that presents the trusted routing overlay mechanism, and
- **Clustering**, parts of which are described within Section 3.

The **Network Monitoring** module is directly related to the security and the reconfiguration of the system and thus is described in detail in the deliverable D3.1, so the reader is advised to refer to this deliverable for more information. The rest of the modules are included in the conceptual model of the Network Manager in terms of completeness, but no new solutions are developed within RERUM. These modules are for **Neighbour Discovery, Opportunistic Communications** and **Mobility Management**.

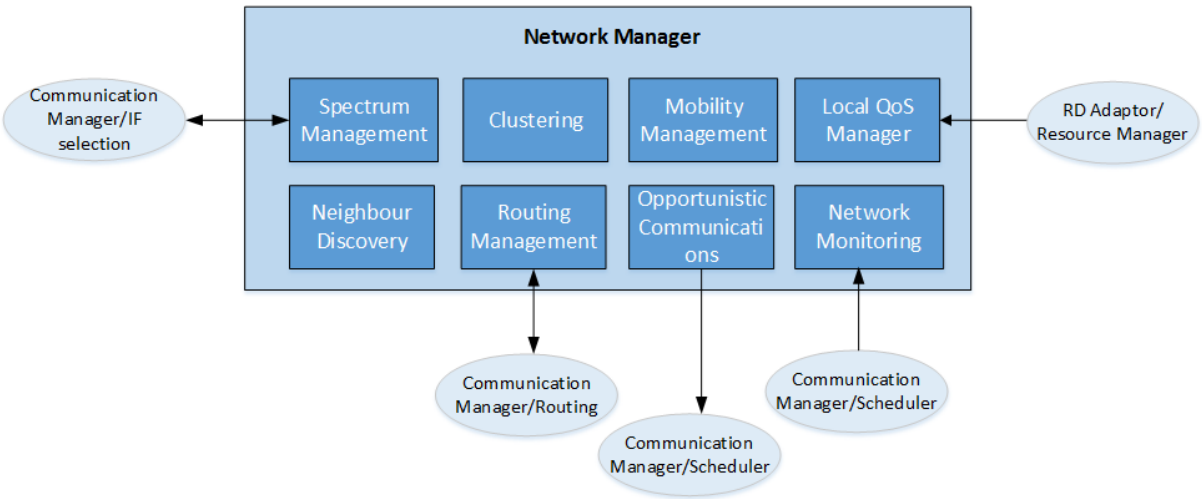

Figure 4. Network manager internal components

It is evident from the above that D4.1 plays an important role within RERUM since it describes the key project innovations for networking and communication. More details regarding the interfaces for the interconnectivity of these modules will be given in the final version of the system architecture to be published in D2.5 due end of August 2015.

1.3.3 Relation with the use cases

The results of this deliverable will be used for the implementation of the use cases. The efficient networking of the devices is of paramount importance for the performance of any IoT device in order to increase the data availability in the overall system. No matter how effective, useful and advanced an appplication is, it cannot provide any actual benefit if the devices are not able to transmit their data. Especially in situations when the timely delivery of the latest data is important (i.e. in emergency situations where there is a fire or a health incident) the data delivery should be prioritized and delivered as quickly as possible to the destination (e.g. te emergency response teams). For this, the devices should be "always connected" to the system in an optimized way. The results described in this deliverable are mostly applicable to three of the four use cases of RERUM. With the exception of the smart transportation use case, all other use cases are dealing with RDs that are basically fixed devices connected through their wireless interfaces that are either IEEE 802.11 or 802.15.4. Thus, all solutions presented in this deliverable can be applied to these three use cases. The smart transportation use case considered within RERUM, basically utilizes mobile phones carried by citizens. The results presented here can also be adapted to this use case in the future when the mobile phones will have reconfigurable radio interfaces. However, currently RERUM has identified a key challenge for mobile devices, which is related related with the optimal selection of interfaces when both mobile and wireless networks are available and such a solution is described in Section 4. More details for the implementation of the described mechanisms in the use cases and their testing in experiments and trials are presented in deliverable D5.1 [RD51], therefore, the reader is advised to refer to that deliverable.

1.4 Structure of the document

The document is structured as follows:

- Section 2 presents the adaptation of the Cognitive Radio (CR) technology to the IoT world, with the description of a CR-inspired agent that allows the development of lightweight opportunistic spectrum access technologies within the RDs. Respective mechanisms for spectrum sensing and spectrum assignment are also described together with the prototype implementation of a CR-inspired gateway that interconnects heterogeneous RDs using standard Software-Defined-Radio (SDR) implementations of IEEE 802.11 and 802.15.4.
- Section 3 presents techniques for mitigating interference in wireless IoT environments optimizing the networking connectivity of networks consisting of RDs. Focusing on large RERUM deployments where RDs are grouped in clusters we present an intelligent technique for optimizing the frequency reuse among clusters.
- Section 4 presents techniques for optimizing the resource access in virtualised environments. A novel technique for virtualising the RD's network interfaces and for efficient assigning of Services in virtual interfaces is presented. Solutions for modelling and naming Resources and virtual RDs, as well as for discovering and allocating resources to virtual RDs are also described.
- Section 5 presents routing and forwarding solutions that aim to optimize the assignment of virtual resources and their efficient delivery from the RDs to the rest of the RERUM system. The focus is on developing trusted routing overlays and on bi-directional multicast forwarding.
- Section 6 concludes the deliverable with an overview of the key results.

2 Enhancing RDs with CR capabilities

2.1 Introduction

2.1.1 Motivation

In the last decade, there is a significant increase in the usage of wireless networking technologies in urban environments, with an estimation of reaching almost 4 billion mobile and wireless users in a few years. These numbers only consider humans being the users of the wireless networks that is basically optimized for human-oriented communications. However, with the recent growth of the Internet of Things (IoT), the number of devices that are connected to the Internet (mostly wirelessly) in the next decade is expected to increase dramatically. There are various estimations about the number of IoT devices that will be interconnected to the global internet, starting from the least optimistic ones, like the report from Analysys Mason (in 2010) that talks about 2.1 billion devices by 2020 and going to the most optimistic one from WWRF that estimates that 7 trillion devices will serve 7 billion people by the year 2017 [W++12,SS09]. The continuous growth of IoT is proved also by the inaccuracy of any such predictions. For example, the prediction of Analysys Mason in 2010 for 2.1 billion devices [MAS11] has now been refined and there is a new report that talks about 20+ billion devices by the year 2023 [MAS14].

This explosion in the number of wireless devices raises significant requirements for new communication technologies that sshould be developed to support the IoT interconnectivity in both new and existing wireless deployments in urban environments. The increased network usage, the new connectivity requirements, the heterogeneous communication technologies involved and the huge amounts of heterogeneous types of traffic that these devices will flow through the Internet are key issues to be addressed by IoT technologies, since the current Internet is not fully capable of supporting this type of communication.

In such a complex networking landscape, the requirements for anywhere, anytime and anyplace connectivity seem to be valid not only for humans but also for devices, since many smart applications can have very strict Quality of Service requirements with respect to timely data delivery, i.e. medical applications or emergency/alarms. Current wireless networking technologies (e.g. Wireless Sensor Networks using the IEEE 802.15.4 protocol), which are also used for IoT communications, are severely impacted by interference, which can degrade significantly the performance and the connectivity of the devices.

As proved extensively in the literature [TFAS11, AP++08], interference is one of the main performance limiting factors in wireless, ad hoc and sensor networks. Controlling interference is essential to achieve maximum performance, since it affects directly the reception capabilities of the clients [QZ++07]. Interference is a key factor that can lead to reduced capacity and performance, since it can reduce the achievable transmission rate of wireless interfaces, increase the frame loss ratio and reduce the utilization of the wireless resources, due to contention in distributed channel access protocols, such as IEEE 802.11 or 802.15.4. Interference can be between links belonging to the same network or can originate from external sources.

Interference is directly related with the utilization of the radio spectrum by many users. The radio spectrum is a natural resource regulated by governmental and international agencies. It is assigned to license holders on a long term basis using fixed assignment policy, which affects the spectrum usage as measurements have shown [A++09]. This shows that for large portions of spectrum the utilization is quite low, while for the ISM bands the utilization is quite high leading to significant interference. Cognitive Radio (CR), proposed by J. Mitola in [MM9], has emerged the last decade as a promising technology able to exploit the unused portions of spectrum in an opportunistic manner. CR was first introduced for opportunistic radio access in niche applications, such as wireless microphones, later

adopted on traditional ad-hoc networks [A++09] for improving spectrum access. Lately the interest in being moved to Wireless Sensor Networks (WSNs) and the smart grid as presented in [LDP12].

Within RERUM, the goal is to utilize the CR technology to optimize the network performance enabling efficient IoT communications between the heterogeneous types of RERUM Devices (RDs). For efficiently interconnecting wirelessly numerous smart objects (that can be co-existing with other wireless technologies) several issues have been identified [V++11]:

- Hardware heterogeneity,
- Different communication protocols,
- Different communication technologies,
- Interference in WSN frequencies,
- Single-radio devices,
- Limited energy resources that enable power saving modes on the devices,
- No IoT-tailored QoS requirements,
- Low trustworthiness of IoT communications with regards to both security and performance,
- Dense deployments of IoT devices in small areas.

As part of the preparatory work for the RERUM project, a position paper for CR-inspired Smart Objects (CRSO) was published [TA13]. The platform proposed in this paper is adapted within RERUM in order to develop the Cognitive-Radio agent of the RDs (described in the following subsection). The RDs have increased radio intelligence and an ability to adapt to environmental conditions using the CR technology. The basic idea is that an RD should be able to exploit the unused spectrum bands opportunistically, but only in an energy efficient way (since one of the key aspects of RERUM is energy efficiency as mentioned in deliverable D2.2). In this respect, the cognitive cycle introduced by Mitola [MM99] can be changed according to Figure 5, in order to include the “energy efficiency” aspect. Here two new modules are included in the cycle: (i) **battery information module** and (ii) **energy efficiency module**. The battery information module is responsible for maintaining information regarding the remaining capacity of the battery and to keep track record of the average energy consumption at each spectrum band. The energy efficiency module is directly connected with all other modules of the cognitive cycle to consult them (and command when in critical battery state) for taking decisions and acting with regards to: (i) consume less energy when sensing the spectrum (i.e. sense the spectrum less frequently), (ii) analyse the spectrum with regards to the required energy to transmit in the band under examination and (iii) decide to use a spectrum band with low required energy consumption.

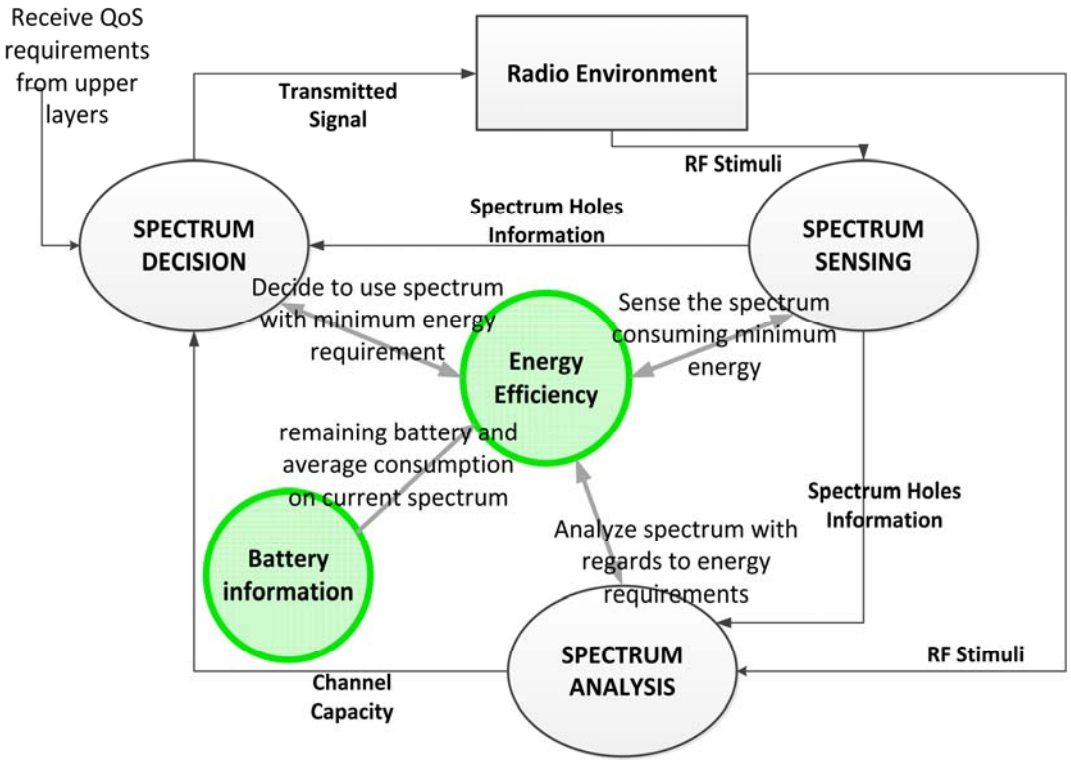

Figure 5. Energy efficiency-based cognitive cycle

2.1.2 Relationship with smart city applications and RERUM use cases

Due to the inherent intelligence of CR technology there are many potential smart city applications that can benefit greatly from the use of CR-based IoT devices. The idea is to allow the CR devices (and basically the RDs) to be interconnected with each other and with the gateways in a seamless and efficient way. A summary of the application areas that can benefit from CR-inspired RDs is discussed below:

Smart Buildings (including the RERUM indoor use cases for comfort quality monitoring and home energy management)

In the smart city domain, buildings are becoming more intelligent embracing a wide range of technologies to improve the everyday life of the inhabitants. The most common smart building applications are: (a) building automation and (b) building energy management. The wide deployment of wireless networks within buildings in the last decade resulted in severe interference in the ISM bands that are utilized by standard Wi-Fi Access Points. For implementing the RERUM indoor use cases, a large number of RDs have to be installed for monitoring the energy consumption of devices (i.e. TVs, air conditioners, lights) and for home automation (sensors and actuators in doors, windows). Equipping these devices with radio cognition can allow for a distributed self-configuration of the RDs in non-ISM frequencies. Normally, IoT devices do not require transmitting large amounts of data, therefore opportunistic usage of short timeslots of the ISM bands (even when they are congested by the Wi-Fi users) will be sufficient for these applications. When more resource-demanding applications are requested (i.e. transferring video for surveillance) then free spectrum in other bands (e.g. TV-bands) can be utilized without disturbing the Wi-Fi users. Furthermore, RDs accessing wide band widths in underutilized TV-bands would be ideal for in-house multimedia distribution and sharing between devices such as TVs, PCs and Hi-Fi systems. For this type of applications, high data rates are required (9-30Mbps) which cannot be achieved in congested Wi-Fi environment, but RDs can perform spectrum

aggregation in unused frequencies to utilize wider bandwidth, i.e. in TV bands for achieving high data rates.

Vehicular Communications (including the RERUM outdoor use case for smart transportation)

Intelligent Transportation Systems (ITS) [DD10] utilize road-based and vehicle-based sensors for several applications, like electronic tolling and traffic monitoring. In the 5.8 – 5.9 GHz band, 75MHz and 30MHz have been assigned in the US and in Europe respectively for Dedicated Short Range Communication (DSRC) communications with a channel width of 10MHz. This is only sufficient for delivering short amount of data in short distances to vehicles for alarms and emergencies. This requires, though, a wide deployment of many Road Side Units (RSUs) in order to provide full coverage. RDs installed on vehicles or playing the role of road-side ITS sensors can make current DSRC networks able to support huge amounts of data that will need to be exchanged for vehicular applications like traffic monitoring and management and smart parking in congested roads. CR technology can greatly assist in this direction, by allowing the opportunistic utilization of more spectrum frequencies to enable higher data rates and longer range with less transmission power. That way, a minimum deployment (or even no deployment) of RSUs would be feasible and vehicles will be able to communicate in long distances (in lower unused frequencies) so that they will know the traffic long ahead either by short messages or even by video multicast. Specifically for the RERUM smart transportation use case, the CR-based RDs can play both the role of clients (installed on vehicles) and RSUs. The RSUs play the role of the gateway, implementing all the required mechanisms for centralized handling of the client-RDs (spectrum sensing and assignment, routing, etc.). Speed measurements that are required for the smart transportation use case could be sent directly from the vehicles to the RSUs in long distances, so the deployment costs of such an application can be minimized.

Environmental Monitoring (one of the RERUM outdoor use cases)

For the second outdoor use case (monitoring the environmental conditions in a smart city), a large number of heterogeneous devices have to be installed in several key locations (i.e. squares) monitoring humidity, air quality, noise, temperature, CO and CO2 emissions. Another option would be to have dedicated vehicles (i.e. buses or garbage collector vehicles) moving on the city roads with sensors devices performing measurements in many areas. For the second approach, location determination devices (i.e. GPS) should also exist on the vehicle. To ensure the reliable and energy-efficient operation of the large number of devices, the CR technology should be adopted to exploit dynamic spectrum management techniques in an overcrowded spectrum that is sensed in current large cities as identified in [LDP12, T++11]. Furthermore, the heterogeneity of the devices will be concealed due to the reconfiguration capabilities of the RDs, so that their interconnectivity of the devices will be easier and seamless, allowing the transferring of data in long distances and in distant areas with no network coverage.

Smart Grid (not directly a RERUM use case, but it relates with the indoor use case for Home Energy Management)

A smart grid comprises many interconnected heterogeneous communicating systems that aim to provide and make use of information to save energy and reduce costs in electricity networks of terabytes in th. The smart grid networks continuously expand including more metering devices and it is estimated that the amount of exchanged data will be several thousands in the near future [GRID]. Current wireless networking technologies are not able to support this amount of data efficiently due to the limited spectrum bandwidth. With RDs as the main nodes of the future smart grid the spectrum utilization will be optimized enabling efficient large scale data transmissions, with concurrent lower energy consumption.

2.2 CR-inspired RDs

A basic introduction of the proposed CR-inspired RDs was given in deliverable D2.3 where the RERUM architecture was presented. For the sake of completeness, here we will present an updated version of the CR-inspired agent, before focusing on the specific Dynamic Spectrum Assignment (DSA) mechanisms that are developed within the project.

The CR-agent is the part of the interface selection module that implements the Cognitive Radio mechanisms, allowing the RDs to utilize the radio spectrum in a flexible manner. This is achieved by reconfiguring in real-time the transmission parameters of the interfaces of the devices by software. The internal architecture of the CR-inspired agent is depicted in the Figure 6 below. As it is shown in the figure, the CR-agent comprises four basic modules and can work either in a standalone agent residing in a single RD or can work in a cooperative way, together with other CR-agents of neighbour RDs. For the latter option, the existence of a Trust agent that filters the spectrum occupancy measurements of the neighbour RDs is mandatory in order to ensure that only trusted RDs can impact the decisions of the system regarding spectrum usage. In this respect, attacks like Spectrum Sensing Data Falsification (SSDF) [FTA13] can be easily avoided.

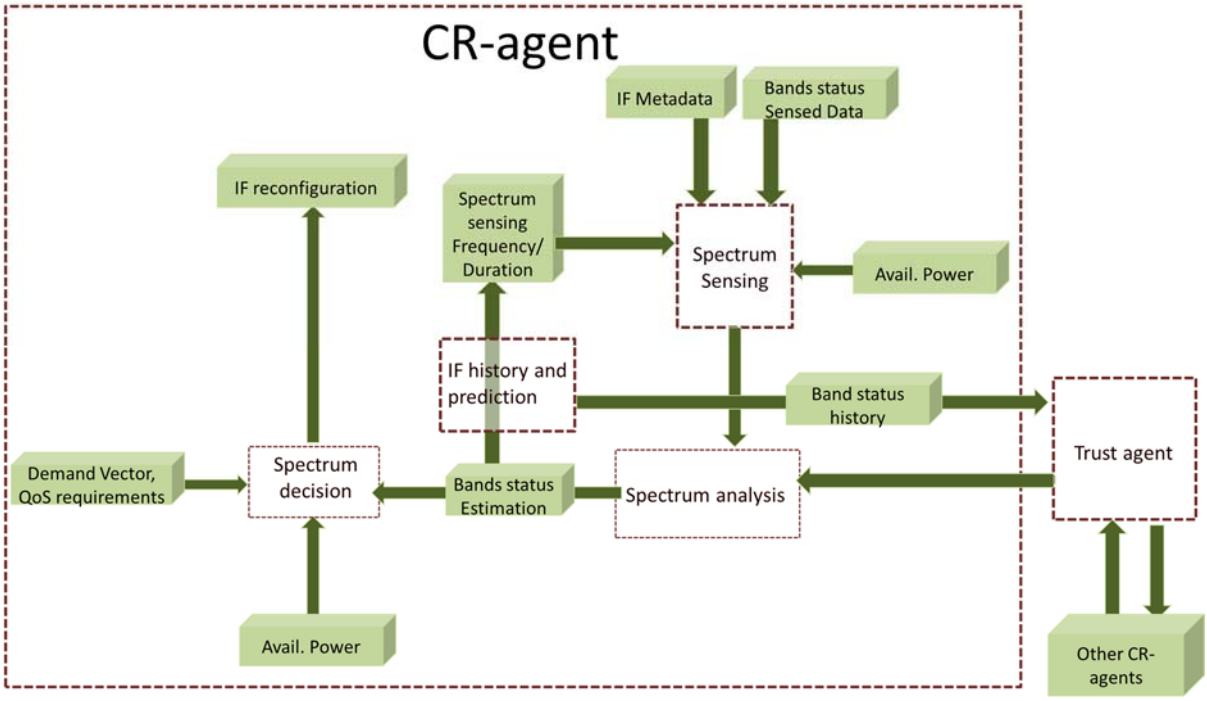

Figure 6. CR-inspired agent internal components

The internal modules of the CR-inspired Agent as described in deliverable D2.3 and in [TA13] are:

- **Spectrum Sensing Module (SSM):** The SSM senses the spectrum in order to identify if a specific spectrum band is being utilized or not, mainly by Primary Users (PUs). Spectrum sensing can utilize several techniques, e.g. energy detection, cyclostationary feature detection, matched filter detection, etc. [A++09]. The most commonly used technique in the literature is by far the energy detection, which measures the energy at a specific spectrum band and comparing the result with pre-defined thresholds it can identify if a transmission exists or not. However, this technique cannot identify what type of transmission it is, namely if it originates from a PU or other cognitive users. On the positive side, this technique is the simplest one that consumes

the lowest amount of resources at the receiver. The SSM takes as direct input from the radio interface the Band Status and some Sensed Data together with some Interface Metadata used for specifying the characteristics of the interface and the spectrum band under consideration (i.e. threshold). Thus, this module is used to establish the wireless channel availability at the time of transmission, but it can't really determine if the transmission comes from a licensed user or not. Spectrum sensing can also be done periodically in order to keep statistics about the utilization of each spectrum band and these statistics are sent to the IF History and Prediction Module. Finally, the SSM is the module that controls the sensing parameters of the interface (spectrum frequencies to examine, sensing duration, sensing period, etc.) according to the input it receives from the IF History and Prediction Module.

- **Spectrum Analysis Module (SAM)**: This module is responsible for analyzing the sensing results that it receives from the SSM in order to identify the characteristics of the examined spectrum frequencies in terms of capacity, condition, interference, occupancy, PU presence and other link-layer related parameters. Since in CR there is no standard definition for a "channel", the RD can select its own central operating frequency and bandwidth, according to its transmission requirements, without being limited by the channel width of i.e. 2MHz for IEEE 802.15.4. In this respect, the RDs can use as much bandwidth as it is available for high resource-demanding applications. The only limitation comes from the hardware capabilities for the band width they can access. The output of the SAM goes both to the Spectrum Decision Module for selecting the spectrum band to use and to the IF History and Prediction Module.
- **IF History and Prediction Module (HPM):** this is the module that identifies statistics about the occupancy model of the spectrum bands/channels and keeps either a respective database or a model of the spectrum occupancy that is updated periodically. The goal is to keep a prioritized table of the spectrum bands in order to see which are more frequently utilized (so that the RD won't try to transmit in these very frequently) and which are mostly underutilized. This module helps also the SSM to adjust the sensing periods per spectrum band, in order to avoid sensing occupied bands very frequently and save energy. The HPM sends also statistics of spectrum occupancy to the neighbour CR-agents (after the required validation of the statistics by the Trust Engine) when cooperative spectrum sensing is utilized.
- **Trust Agent**: this is a module for evaluation of trust for cooperative sensing and access decisions. This module can be connected with the Trust Manager that is described in deliverable D2.3 in order to get the reputations of the measurements of the neighbour devices.
- **Spectrum Decision Module (SDM):** the SDM is the module that does the actual selection of the spectrum band that will be accessed. The prerequisites from the original communication demand according to the service class(es) of the RD are compared with the given sensed information which has been handled in the other modules of the CR-Agent. The available spectrum bands are weighted to the expected energy drainage for the upcoming transmission and the QoS requirements of the service and then, the most suitable spectrum frequency and bandwidth are selected. The selection is also based on the expected validity duration on each channel, i.e. switching spectrum bands too frequently might drain the battery more than finding a more stable band. Utilizing the statistical data from the IF History and Prediction module a preferable and stable spectrum band could be selected. When the most suitable spectrum band and interface have been selected, the most suitable Modulation and Coding (M&C) to be used by the chosen interface for the selected spectrum band are also selected, in order to consider any additional energy savings which still meet the specified requirements. The choice of modulation and coding might face constraints and regulations depending on predefined protocols, e.g. LTE and GSM face big differences in their modulation possibilities.

In such a case the protocol is pre-set and there is no room for M&C variations. In other systems, where there is more room for adaptive M&C, the choice can be further refined to achieve lower energy consumption.

When an RD needs to make a transmission related with some service, a "demand vector" characterizing the traffic volume and QoS requirements is passed to the CR Agent. **The Spectrum Decision Module** will determine if the spectrum selection process needs to be initiated for handling the transmission or the current selected spectrum can satisfy the demand. This means that if the current spectrum selection cannot meet the service demands then the CR-agent has to find another spectrum band (with i.e. larger bandwidth) to access. If the transmission request can be served with the current spectrum selection, then the sensing procedure may be skipped and the CR-inspired agent will directly use the last known transmission option. Furthermore, the sensing procedure can be partly skipped if the transmission has fixed setup requirements or only has to sense a small part of fixed channels.

If a new sensing/selection procedure has to be performed then the demand vector will be passed to the **Spectrum Sensing Module** (SSM), which will determine which spectrum bands will be sensed. The sensed spectrum will be based on a prioritization table generated by the **IF History and Prediction Module** and the available energy, e.g. if the device has a fully charged battery the SSM might sense a wider spectrum over longer periods, while if the battery is low the SSM might only sense statistically safer spectrum channels for a shorter duration. The SSM will send interface parameters such as selected frequencies and sensing duration to interfaces that execute the sensing. The sensing results are analysed by the **Spectrum Analysis Module** to identify the characteristics of the spectrum band and then are stored in the **IF History and Prediction Module**, updating the IF prioritization table before passing on to the **Trust Agent**. The sensed data carries information such as IF Metadata and RF band status. This information will be treated by the Trust Agent to evaluate if the spectrum is trustworthy or not (usually comparing the results with similar results of neighbouring RDs). The results of the spectrum analysis are passed to the **Spectrum Decision Module,** that will select the most applicable option based on the given traffic demand vector, available power and other QoS related information. This process is also shown in Figure 7.

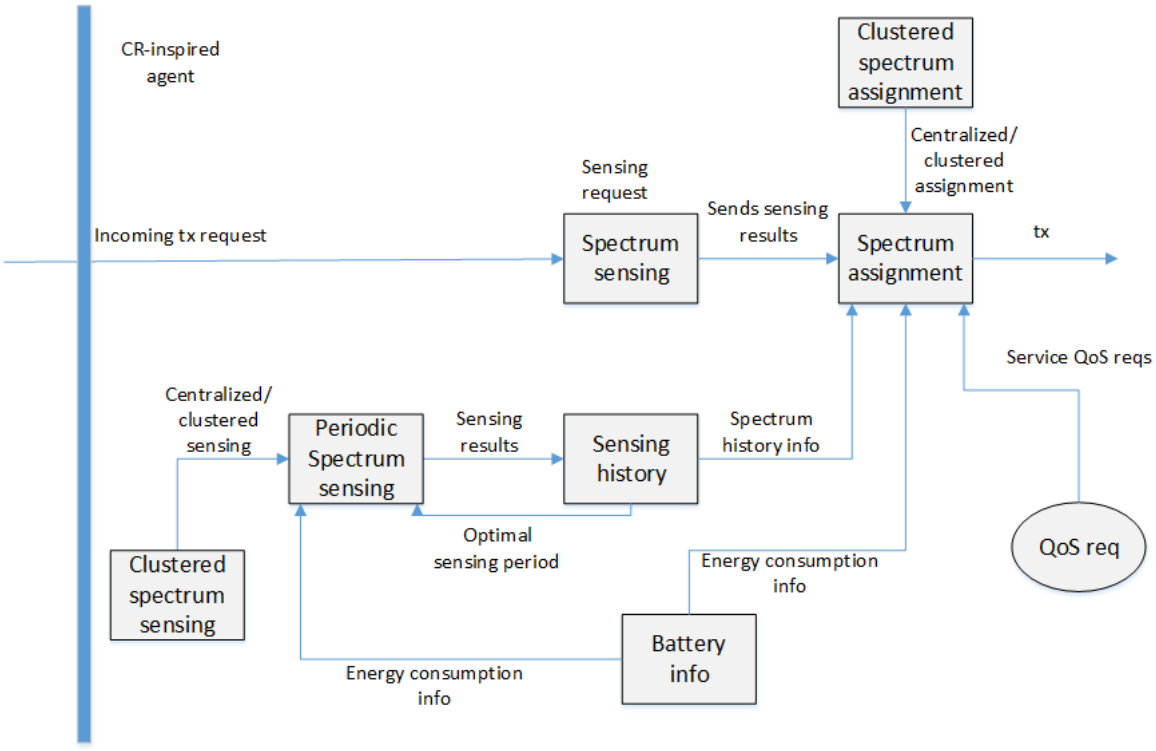

Figure 7. Spectrum management process within the CR-agent

Figure 7 shows the whole process for deciding the spectrum that the device will utilize for the transmission upon an incoming transmission (tx) request. The spectrum sensing module scans the spectrum for **white holes (unused portions)** and sends the results to the spectrum assignment module. The latter analyses the spectrum holes (according to predefined criteria) and selects the most appropriate bandwidth to be used. Battery information, QoS requirements and sensing history about channel conditions are inputs that are taken into account for the decision. Figure 7 also depicts the process for gathering the historic information for periodically sensing the spectrum and minimizing the energy spent in sensing the spectrum. The spectrum sensing and spectrum assignment modules can also utilize information from other devices in a distributed, centralized or hierarchical clustering approach (which will be detailed in the network manager below).

2.3 Lightweight mechanisms for spectrum sensing

2.3.1 Motivation and state of the art

This subsection presents the results of RERUM with regards to the design and development of lightweight spectrum sensing mechanisms. Using the terminology of the CR-agent that was described above, this section presents two methods for the implementation of an integration of the **Spectrum Sensing Module** and the **History and Prediction Module**. Briefly, the idea is to gather statistics about the occupancy of fragments of the radio spectrum in an energy efficient way. The motivation behind this idea stems from the need to minimize the energy consumption of CR-based RDs when they perform spectrum sensing. It is well known that the power consumption of the radio interfaces for "listening" the radio environment is very high and it is even higher than the power consumed for transmitting. Figure 8 depicts the power consumption of the network interface of the Zolertia Z1. It is clear that for listening to the radio spectrum (RX mode) the current consumption is 18.8mA, while for transmitting (TX mode) is slightly lower at 17.4mA.

IC	Operating Range	Current Consumption	Notes
CC2420	2.1V to 3.6V	<1µA	OFF Mode
		20µA	Power Down
		426µA	IDLE Mode
		18.8mA	RX Mode
		17.4mA	TX Mode @ 0dBm

Figure 8. Power consumption for the network interface of Z1

When the advantages of Cognitive Radio techniques are exploited, there is thus a need for an intelligent solution to perform spectrum sensing in a smart way avoiding consuming the whole battery of the devices just for sake of sensing the spectrum. In this respect, the RDs have to minimize the amount of time they spent for sensing and this can be done by keeping track of the spectrum occupancy statistics in order to derive occupancy models. This will help to perform periodic spectrum sensing and avoid sensing the band at each time slot. With an accurate occupancy model, the spectrum sensing period (which is the time duration between two consecutive spectrum sensings at the same band) is being maximized. This means that the RD can have an accurate estimation of the spectrum occupancy with a minimum amount of energy consumed in spectrum sensing, allowing the RD to "sleep" between the consecutive sensings. Thus, the battery life of the RDs is maximized, ensuring a

longer lifetime of the RD, which can be of significant help in smart city deployments to minimize the maintenance costs.

The idea of using channel measurements for predicting the channel usage in the near future is a concept that has emerged the last few years and there are many works in the literature that have researched this issue from various perspectives. The most common approach is to use Markov models in almost all of their variations. In [AT07] this problem is solved using a Hidden Markov model, which is trained by the channel measurements in order to develop an accurate model that can give good channel estimation predictions. In [L++10] a two-state high-order Markov chain is used for predicting spectrum usage, assuming that the next state of the channel depends on a number of previous states and not only on the previous state. In [LGC12] the problem is solved using a semi-Markov process, but assuming perfect knowledge of the traffic parameters. In [ZZ11] a prior probability-based scheme that enables sequence spectrum sensing with utilization of the statistics of licensed band occupancy is proposed. An architecture for proactive spectrum access is proposed in [YCZ08] and introduces a module for channel history and modelling that is used for predicting the channel usage and take decisions regarding switching channels.

In this section, we present two solutions for a lightweight framework to gather spectrum occupancy measurements. The proposed solutions are both lightweight because the amount of resources (in terms of CPU, RAM and battery) required for them to run are extremely low, so they can be applied even in very resource constrained devices.

2.3.2 Lightweight spectrum occupancy measurement framework using multi-armed bandit[2]

The first solution solves the problem of intelligent spectrum sensing based on historical data in the context of CR-based RDs by modelling the problem as a pure-exploration multi-armed bandit problem with dependent arms [BMS09]. We use two well-known arm pulling strategies (uniform sampling and upper confidence bounce exploration) to tackle it and then we identify the existing correlation structure among the candidate solutions and utilize it to significantly speed up the learning process. Numerical results shown below verify the effectiveness of the proposed strategies.

2.3.2.1 System Model

A. Preliminaries

We consider a single RD that performs spectrum sensing in a specific band (channel) licensed to a Primary User (PU) with the purpose of learning how to optimally collect future channel occupancy data for the specific channel. The PU's channel usage can be modelled as an ON/OFF source alternating between active (ON) and idle periods (OFF). For the target channel we model the sojourn time of an ON (OFF) period as a random variable T_{on} (T_{off}) that has a probability density function (p.d.f.) f_{Ton} (t_{on}), $t_{on} \geq 0$ (f_{Toff} (t_{off}), $t_{off} \geq 0$).

We assume that the probability density functions of ON and OFF periods are stationary, that the duration of ON periods is independent of the duration of OFF periods and, finally, that the durations of both ON and OFF periods are independent identically distributed (i.i.d.) random variables. Let Z(t) denote the state of the PU at time t. Then, Z(t), $t \geq 0$ becomes a semi-Markov process: whenever the process enters either the ON or the OFF state, the time until the next state transition is governed by f_{Ton} (t_{on}) and f_{Toff} (t_{off}). Figure 9 shows the state transition model for this semi-Markov process.

[2] We have to note here that the work presented here has been published in Wireless Vitae 2014 and has received the best paper award in this conference [STT14].

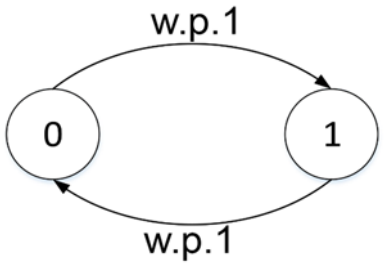

Figure 9. State transition diagram for the PU semi-Markov process.

Channel sensing is a sampling procedure $\{S(kT), k \in \mathbb{N}\}$ of the given channel process Z(t); t≥0 to discover the state of the PU at sensing instant kT, where T (time-slot) is the shortest time interval between two successive samples taken from the channel by the RD. We assume that T is larger than the sensing time, i.e. the predefined amount of time necessary to assess the state of the PU accurately. We also assume that the probability of a PU's state change within T is negligible. Finally we assume that whenever the RD senses the channel it always detects the correct state of the PU. If we represent with 1 the PU's ON state and with 0 the PU's OFF state, then channel sensing produces a binary random sequence. Figure 10 presents the sensing process of an ON-OFF alternating PU.

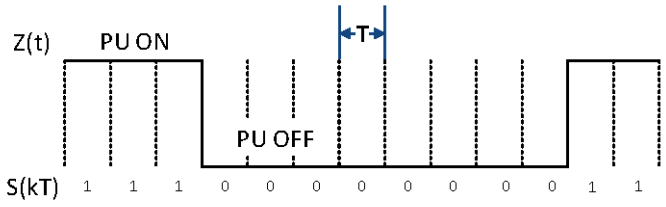

Figure 10. Periodic sampling of the PU's ON/OFF process Z(t), t ≥ 0 with period T produces a binary random sequence

B. Problem Formulation

In the classic multi-armed bandit problem, originally proposed by Robbins [R52], a player must decide which one of K machines to activate at each discrete step in a sequence of trials so as to maximize his long term reward. Every time he plays an arm, he receives a reward. The structure of the reward for each arm is unknown to the player a priori, but in most prior work the reward has been assumed to be independently drawn from a fixed, yet unknown, distribution. The reward distribution is in general different for each arm, therefore the player must use all his past actions and observations to learn the quality of these arms in terms of their expected reward so that he can keep playing the best arm.

Here, we are interested in a special instance of the multi-armed bandit problem called the *pure exploration* problem [BMS09,ABM10], whereby the player is asked to recommend an arm after a given number of pulls and he is evaluated only by the average reward of his recommended arm. We are particularly interested in efficient arm pulling strategies that can identify the optimal arm within a small number of pulls since the PU will be performing energy consuming channel sensing at every time-slot during the learning phase.

We model the problem of energy-efficient collection of spectrum occupancy data as a pure exploration, multi-armed bandit problem with the objective of finding **two periods, one for each state of the PU**, that have **maximal expected reward**. These periods will be used to collect future spectrum occupancy data for the PU. In our problem's setting each arm of the bandit corresponds to a candidate sensing period and can be found in one of two states; **state s = 0** corresponds to PU's OFF state and

state s = 1 corresponds to PU's ON state. In a typical multi-armed bandit problem the duration of all arm pulls is considered to be equal to 1, while in our model each candidate sensing period is expected to have a different duration. In order to align our model with its typical representation we incorporate all time related rewards (and costs) directly into the reward function and introduce the notion of a discrete sequence of rounds.

At the n-th round the RD selects a single arm (period) a_s^n from a finite set A_s of available periods at PU's current state s, according to a sampling strategy, and performs channel sensing at every time-slot of a_s^n. At the end of the round, period a_s^n is granted with a reward. The current state of the PU, '***s'***, is defined to be the state of the PU at the last time-slot of the previous round or the time-slot before the first round begun. The reward granted to period as at the end of round n is given by Equation (1)

$$r_{a_s}^n = E_c + \left(E_g - E_c\right) \cdot X_{a_s}^n \quad (1)$$

where $X_{a_s}^n$ is a random variable that expresses the accuracy of the chosen period at round n and is given by Equation (2):

$$X_{a_s}^n = \begin{cases} 1 & if\ W_{a_s < \gamma}^n \\ 0 & if\ W_{a_s \geq \gamma}^n \end{cases} \quad (2)$$

$W_{a_s}^n$ is a random variable whose value represents the number of time-slots that the channel was found to be in a different state than the current state s of the PU and γ is a pre-defined threshold. We account for all time-slots of period a_s^n except for the last one, at which the channel is sensed and the current state of the PU is updated. Period a_s^n will be characterized as accurate if $W_{a_s}^n$ is less than a predefined threshold g and as inaccurate otherwise. For the single time-slot period, $W_{a_s}^n$ will always be zero and the period will always be accurate. E_g in Equation (1) is the reward granted to an accurate period and corresponds to the energy gain associated with choosing a long, yet accurate, sensing period, compared to sensing the channel at every time-slot. E_g is calculated as the product of the duration $d(a_s)$ of the chosen period measured in time-slots and the energy e_c spent during a single channel sensing time, i.e.,

$$E_g = e_c \cdot d\left(a_s^n\right)$$

or in normalized form

$$E_g = \frac{e_c \cdot d(a_s^n),}{max_{a_s \in A_s} \cdot d(a_s)}$$

Similarly, E_c is a negative reward or penalty associated with an inaccurate period and corresponds to the energy cost that could be induced to the RD due to an inaccurate period. Such an energy cost may result for several reasons: (i) the RD must migrate from a channel that was supposed to be free for the next as timeslots, yet the PU changed its state form OFF to ON, or (ii) because an optimal channel, in terms of its transmission characteristics, was assumed occupied while now is free. Here we assume a constant energy cost for either case.

Whether period a_s^n will be accurate or not depends on its duration and the distribution of the residual time that the PU will remain in its current state. The probability density function of the residual time t_r in the current state 's' for an alternating renewal process, such as the one presented for the PU in Paragraph A above, assuming that it started a long time ago, is given by

$$f(t_r^s) = \frac{1 - F_{T_s}(t_r^s)}{E[T_s]}$$

where, T_s corresponds to T_{off} when s = 0 and to T_{on} when s = 1, $F_{T_s}()$ is the cumulative distribution function (c.d.f.) of the corresponding T_s duration calculated at t_r^s and $E[T_s]$ is the expected duration of T_s. For a detailed treatise on the subject please see [G13] and the references therein. As already stated before we assume that p.d.f's $f_{T_{off}}$, $f_{T_{on}}$ and the corresponding c.d.f.'s $F_{T_{off}}$ and $F_{T_{on}}$ are unknown to the RD. As a consequence we will not be able to speed up the learning process by calculating the probability of accuracy of future periods, by utilizing the p.d.f. $f(t_r^s)$ of PU's residual time in its current state. Therefore, we assume that for each period a_s, random variables $W_{a_s}^n$ are independent of previous rounds and identically distributed (i.i.d.) random variables. This assumption is accurate if the durations of PU's ON and OFF states are exponentially distributed due to the memoryless property of the exponential distribution.

On the other hand, we will be able to utilize the correlation structure between the accuracy of period a_s^n, chosen in round n, and all the periods that were not chosen in the same round. This makes our problem a multi-armed bandit problem with statistically dependent arms. Previous work on multi-armed bandit problems with statistically dependent arms includes [ABM10, MRT09]. Specifically in [MRT09] Tsitsiklis et al. demonstrate that known statistical structure among arms can be exploited for higher rewards and faster convergence. Although the authors use a fairly specific model, it exemplifies a broader class of bandit problems where there is a known prior functional relationship among the arms' rewards. In our problem setting this functional relationship among arms' rewards can be expressed as follows:

If $a_s^n \in A_s$ is the period chosen at round n and $X_{a_s}^n = 1$, then for all candidate periods $c_s^n \in A_s$ with $d(c_s^n) < d(a_s^n)$ it holds that $X_{c_s}^n$ would have also been 1 if period c_s^n had been chosen at round n. Similarly, if $X_{a_s}^n = 0$, then for all candidate periods $g_s^n \in A_s$ with $d(g_s^n) > d(a_s^n)$ it holds that $X_{g_s}^n$ would have also been 0 if g_s^n had been chosen at round n. The above propositions are true because $W_{a_s}^n$ is a counting process and its value can never decrease when periods with increasing duration are considered. Thus, given threshold γ, if period a_s^n is accurate then so would have been all other periods with a smaller duration, had they been chosen for the same round n. On the other hand, if period a_s^n is inaccurate then so would have been all periods with a larger duration, had they been chosen for the same round n. What is more, during the learning phase, the RD performs spectrum sensing at every time-slot in order to find the value of $W_{a_s}^n$, thus the necessary information to assess the accuracy of all periods with smaller duration than a_s^n, even if the latter was found to be inaccurate, i.e., $X_{a_s}^n$ =0, is readily available at the RD. At the end of round n, all periods for which their accuracy status can be assessed by utilizing the existing correlation structure will be granted a reward according to Equation (1).

C. Period sampling strategies

In this section we describe two strategies with low memory, storage and computational cost that sample one period at each round n and update its expected reward estimate. Both strategies have been adapted to the problem of pure exploration in multi-armed bandits and will recommend for each state of PU, after a predefined number of rounds, the period with the largest estimates of expected reward. See [BMS09, ABM10, ACF02] for a detailed analysis of the strategies and their achievable bounds on error probability.

1) Uniform sampling strategy: Let N be the total number of rounds, $\hat{r}_{a_s}$ the empirical mean reward of period a_s and k_{a_s} the number of times period as has been sampled by round n. The uniform sampling strategy can be described as follows:

1: **for all** $n \in [1, N]$ **do**

2: Identify current state of PU s

3: Select period a_s with uniform probability from set A_s

4: Get $W_{a_s}^n$ by spectrum sensing at every time-slot

5: Get reward $r_{a_s}^n$ using Equation (1) and Equation (2)

6: $k_{a_s} \leftarrow k_{a_s} + 1$

7: Update estimate $\hat{r}_{a_s} \leftarrow \hat{r}_{a_s} + \frac{1}{k_{a_s}}[r_{a_s}^n - \hat{r}_{a_s}]$

8: **end for**

9: **Return** $a_s^* \leftarrow \underset{a_s \in A_s}{\operatorname{argmax}} \hat{r}_{a_s}$

where, a_s^* is the recommended period for collecting future spectrum occupancy data when PU's current state is s. Care must be taken when choosing N so that all periods are sampled enough times. This is especially true in our case because each state 's' of the PU is expected to have a different average duration, thus periods corresponding to the state with the largest average duration will be sampled more frequently as the PU will be found in that state more often.

2) Highly exploring strategy based on upper confidence bounds: A disadvantage of the uniform sampling strategy is that it samples all periods (including the best and the worst) equally often resulting in a waste of RD's scarce resources. The Upper Confidence Bound Exploration (UCB-E) [AMS09] strategy is similar to the uniform sampling strategy with the exception of the mechanism being used for sampling periods. UCB-E strategy utilizes the estimates of the expected rewards $\hat{r}_{a_s}$ for each period a_s and samples, more frequently, periods with higher expected rewards. More specifically, period sampling in UCB-E is based on updating quantity B_{a_s} for the selected period a_s at the end of each round by using the following equation

$$B_{a_s}(k_{a_s}) = \hat{r}_{a_s} + \sqrt{\frac{b_s}{k_{a_s}}}$$

All B_{a_s} values have been initialized to ∞. Parameter b_s of UCB-E controls how often the UCB-E strategy will sample sub-optimal periods. A small value of parameter b_s results in low exploring, i.e., rarely selecting sub-optimal periods, while a large value results in selecting frequently sub-optimal periods. In the first case, UCB-E may not recommend the optimal period if the latter gets an abnormally large number of low rewards during the first rounds. In the second case, UCB-E behaves similarly to the uniform sampling strategy and thus wastes resources. In [ABM10] the authors prove certain bounds on the probability of UCB-E recommending a sub-optimal period when

$$0 < b_s \leq \frac{25}{37} \frac{N_s - |A_s|}{H_1}$$

N_s is the total number of times periods from set A_s will be sampled ($N_0 + N_1 = N$), $|A_s|$ is the number of periods in set A_s and

$$H_1 = \sum_{a_s=1}^{|A_s|} \frac{1}{\bar{r}^*{}_{a_s} - \bar{r}_{a_s}}$$

where, $\bar{r}_{a_s}$ is the expected reward of period as and $\bar{r}^*{}_{a_s} = max_{a_s \in A_s} \bar{r}_{a_s}$. Although, expected rewards $\bar{r}_{a_s}$ are unknown to the RD beforehand, one can empirically find values for the parameter b_s that successfully balance exploration and exploitation by utilizing the fact that both H_1 and N_s should be of

the same order of magnitude and b_s should increase linearly with N_s. The UCB-E strategy is derived from the uniform sampling strategy by changing step 3 as follows:

3: For current state s select period $a_s \in \underset{a_s \in A_s}{\operatorname{argmax}} \in B_{a_s}$ with ties broken arbitrarily.

After N rounds UCB-E strategy will recommend the periods with the largest empirical mean reward, one period for each state of the PU, similarly to the uniform sampling strategy.

3) Uniform sampling and UCB-E strategies enhanced with statistical correlation data: Both uniform sampling and UCB-E strategies can be further enhanced utilizing the correlation structure between the accuracy of period a_s^n, selected at round n, and all the other periods that were not selected in the same round. As described in Paragraph B above, we can utilize the existing correlation structure to update k_{h_s} and the empirical mean $\hat{r}_{h_s}$, not only for the selected period as in round n but for all periods h_s for which $X_{h_s}^n$ can be decided by the propositions stated in Paragraph B above. Thus, in each round we have many updates that speed up the learning process.

2.3.2.2 Evaluation

In this part we present simulation results for the two sampling strategies and their enhanced counterparts. Our simulation model is comprised of a PU, whose OFF and ON state durations are exponentially distributed with mean μ_0 =30 and μ_1 = 20 time-slots respectively and a RD that utilizes the sampling strategies described in the previous section in order to decide which periods, among the ones included in the following sets A_s = {1,2, … ,1000}, $s \in \{0,1\}$, have the highest estimates of expected reward. Both sets A_0 and A_1 include a large number of candidate periods as we do not expect the sensor to have any prior knowledge of the PU's model. We assume a unit energy cost e_c = -1, for sensing the channel at a single time-slot and normalized energy gains, i.e., E_g : [1,1000]→[0,1], as described in Paragraph B above. Figure 11 presents estimates of the expected reward for the first 200 candidate periods of A_0 over a range of values of the normalized energy cost E_c associated with a period being inaccurate. E_c is assumed to be a multiple l of the energy spent for sensing a channel at a single timeslot e_c, i.e.,

$$E_c = \frac{l \cdot e_c}{max_{a_s \in A_s} d(a_s)} = \frac{-l}{1000}$$

Figure 11 shows that as the cost of being inaccurate increases, the duration of the optimal period decreases. Eventually, if E_c becomes too high, the optimal solution is to sense the channel at every single time-slot to avoid the possibility of being inaccurate. What is more, the expected reward of periods that have a much longer duration than the optimal periods, will be equal to E_c as they will be inaccurate with very high probability.

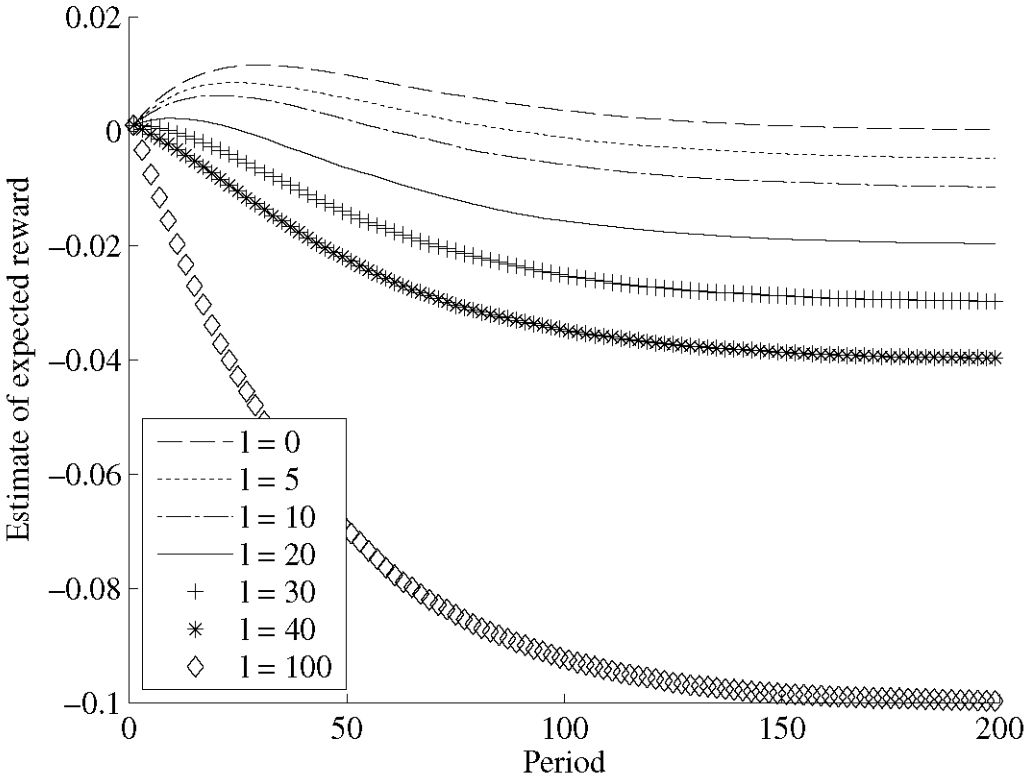

Figure 11. Estimates of the expected rewards of each candidate period for varying energy costs when the PU is in OFF state. The same behaviour is observed when the PU is in an ON state.

Evaluation of the period sampling strategies presented in paragraph C above will be done in terms of absolute regret '**R**'$_s$ [ABM10], i.e. the difference between the expected reward of the optimal period and the empirical mean reward of the recommended period at round n for each PU state,

$$R_s(n) = \left|\bar{r}^*{}_{a_s} - \hat{r}^*{}_{a_s}(n)\right| \quad (3)$$

where, $\bar{r}^*{}_{a_s} = max_{a_s \in A_s} \bar{r}_{a_s}$ and $\hat{r}^*{}_{a_s}(n) = max_{a_s \in A_s} \hat{r}_{a_s}(n)$.

Mean rewards at round n are computed with data from 200 repetitions of each experiment. We can calculate analytically the optimal expected reward for our simulation scenario by approximating the exponentially distributed ON and OFF periods with their discrete time counterparts, i.e., geometric distributions.

We know that if duration T_s is an exponentially distributed random variable then $T_s^g = \lfloor T_s \rfloor$ will be a geometrically distributed random variable with parameter $p = 1 - \exp(-\frac{1}{\mu_s})$, where p is the probability that the PU will change its state at a specific time-slot within period as, i.e., we assume γ = 1 for our evaluation scenario. Utilizing the memoryless property of the geometric distribution we can now calculate the probability that we will have at least γ state changes in $d(a_s)$-1 timeslots (as explained in paragraph B above) and thus an inaccurate period using Equation (4).

$$p_{inac}^{a_s} = \sum_{w=\gamma}^{d(a_s)-1} \binom{d(a_s)-1}{w} p^w (1-p)^{(d(a_s)-1-w)} \quad (4)$$

The expected rewards for each period are given by Equation (5).

$$\bar{r}_{a_s} = p_{inac}^{a_s} \cdot E_c + (1 - p_{inac}^{a_s}) \cdot E_g^{a_s} \quad (5)$$

Finally, for both UCB-E and enhanced UCB-E with correlations the parameter b_s was set to 10 in order to have a medium exploration strategy (after performing extensive simulations with different values).

In Figure 12 and Figure 13 we compare the four period sampling strategies in terms of their speed of convergence to zero regret for a scenario with $E_c = 0.01$($l = 10$) and γ = 1. The inset figures that appear in these two figures offer a closer look at the transient phase of Uniform and UCB-E sampling strategies with correlations during the first rounds. It is clear from the figures that utilizing the existing correlation structure concerning the accuracy of periods impressively reduces the number of rounds required to learn the optimal periods and minimize regret. Furthermore, UCB-E sampling processes perform better than uniform sampling strategies as they do not sample often periods with low expected rewards. Finally, we must note that in these two figures the absolute regret is initially equal to $|\bar{r}_{a_s}|$ for all strategies and increases with the number of rounds until it reaches a peak. This is due to the fact that the empirical mean rewards of all periods were initialized to zero, and thus from Equation (3) we had $R_s(n) = |\bar{r}^*{}_{a_s}|$.

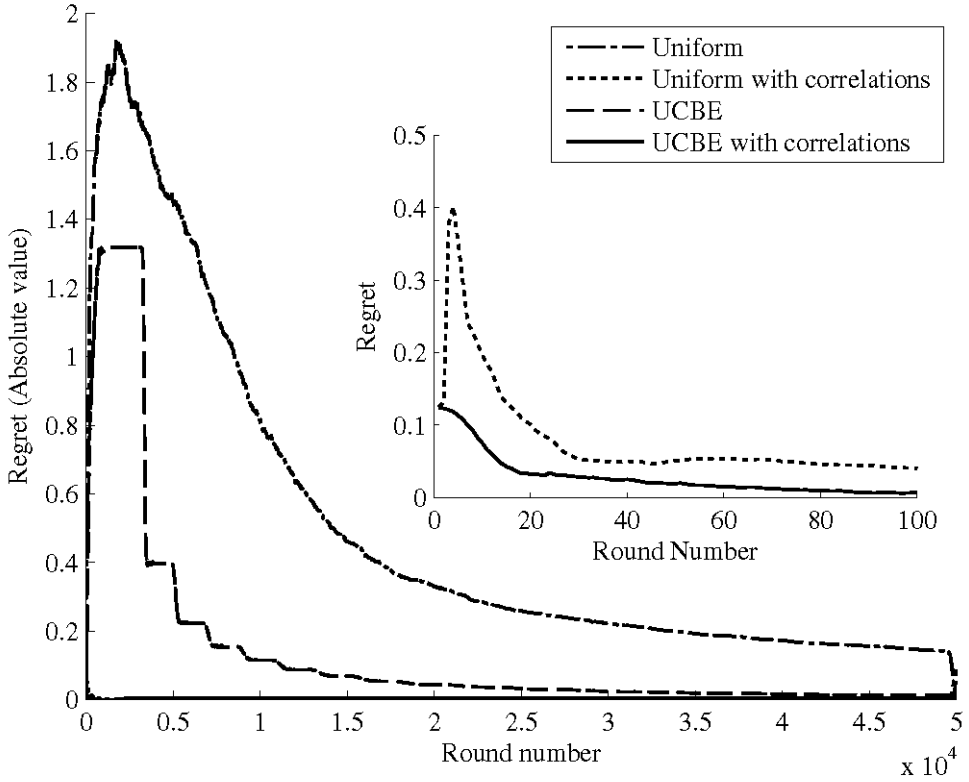

Figure 12. Regret vs. number of rounds when the PU is in OFF state.

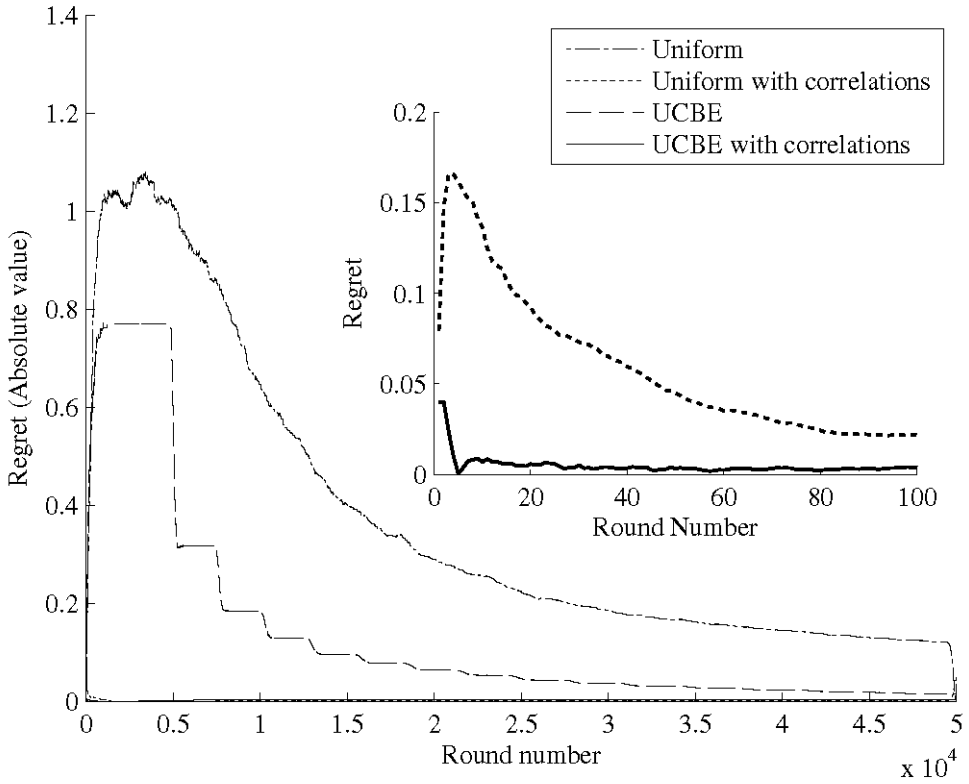

Figure 13. Regret vs. number of rounds when the PU is in ON state.

From the above presented results it is evident that our proposed solution speeds up significantly the learning process of the RDs. The idea of exploiting the correlations between the periods minimized dramatically the amount of rounds required for the algorithms to extract a very accurate model of the spectrum occupancy. This indeed is of high importance, because as mentioned before in the learning phase, the RD senses the spectrum at each time slot, which consumes a large amount of energy. Thus, by doing this only for a few rounds in comparison with the thousands of rounds required by the standard version of the algorithms (without correlation) is of paramount importance for minimizing the energy consumption of the RDs, prolonging their lifetime.

2.3.3 Lightweight CR mechanism for spectrum sensing using reinforcement learning [3]

In the multi-armed bandit formulation of the intelligent spectrum sensing problem, there is no association between the two states of the PU and the reward granted to each selected period depends solely on the accuracy of the period itself. However, we do expect that the probability to find the PU in a specific state at the end of the currently selected period depends on both the length of the selected period and the current state of the PU. In order to consider this association between the states of the PU and consider its implications in the optimal period selecting scheme, we utilize the framework of Reinforcement Learning and more specifically that of semi-Markov Decision Processes [SB98]. Thus, the model presented in this section is more elaborate albeit more computationally demanding compared to the multi-armed bandit model of the problem.

Within this framework the goal is to enable the RD to learn an optimal spectrum sensing "policy" by interacting with the spectrum. To this purpose we utilize two variants of the Q-learning algorithm

[3] This work has been submitted to the IEEE IWCMC 2015 conference [STT15].

[SB98] that are known to converge to optimal policies and we propose enhancements that significantly speed up the learning process. Numerical results verify the effectiveness of the proposed enhanced learning strategies.

2.3.3.1 System Model

A. Preliminaries

The basic system model is similar than the one presented in the previous section for the multi armed bandit problem. Briefly, a CR-enabled RD that performs sensing in a specific spectrum band licensed to a PU, aiming to learn how to optimally collect future spectrum occupancy data for the specific spectrum band.

However, unlike section 2.3.2.1, the state of the PU is directly represented as a *discrete time* semi-Markov process $Z(t)$, $t \geq 0$ that denotes the state of the PU at the k-th time-slot. Whenever the process enters either the ON or the OFF state, the time until the next state transition is governed by probability density functions $f_X(x)$ and $f_Y(y)$ respectively. Figure 14 presents a sample path of $Z(t)$ where $(X_i), (Y_i), i \in \mathbb{N}$ are the sequences of sojourn times in the corresponding states.

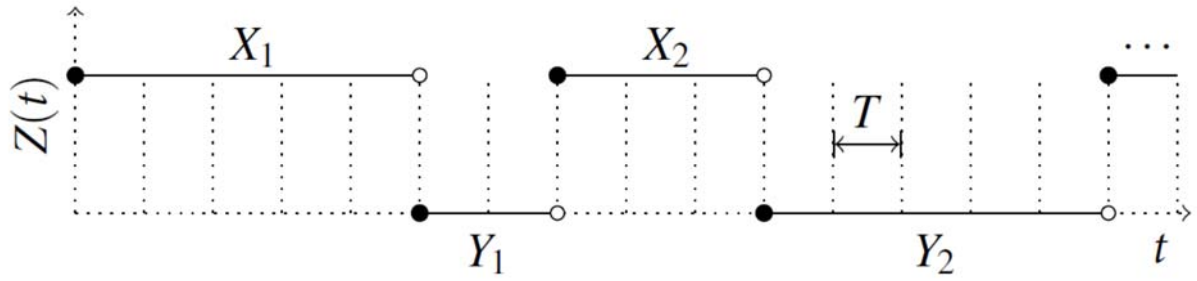

Figure 14. Sample path of PU's semi-Markov discrete time stochastic process.

B. Problem Formulation

We model the problem of energy-efficient collection of spectrum occupancy data as a *Reinforcement Learning (RL)* [SB98, BT96] problem, whereby the RD learns an optimal spectrum sensing policy i.e., a policy that dictates the optimal spectrum sensing period given the current state of the PU. Furthermore, the RD also implements the architecture of a Dyna-Q agent [SB98] that integrates (i) learning through real experience, i.e., experience gained by sensing the spectrum, and (ii) simulated experience, i.e., experience generated by a model. The RD learns the model over the course of time by sensing the spectrum and uses it to produce simulated experience in order to speed up the learning process. To put it simply, the RD exchanges energy intensive spectrum sensing for CPU processing time [SHC++04, C++05] in order to reduce the total energy consumed in the learning process.

The RD, as a typical RL agent, gets real experience by interacting with the spectrum at each of a sequence of discrete rounds, $n = 1, 2, 3, \ldots$; with each round lasting one or more time-slots. At the first time-slot of the n-th round the RD receives some representation of the spectrum's *state*, $s_n \in S$ where S is the set of PU's states and on that basis selects a *period* $a_n \in A(s_n)$, where $A(s_n)$ is the set of periods available in state s_n. Let $m \in [1, M]$ be the duration of the selected period in time-slots, then, m time-slots later the RD receives a reward $r_{n+1} \in \mathbb{R}$ and finds itself in a new state s_{n+1}, where the next round begins. The reward function for r_{n+1} will be defined below in Equation (16). We just note here that r_{n+1} depends on a measure of energy efficiency, which is the length of the period

between two consecutive times of spectrum sensing and on a measure of how accurately the PU's state is represented by its initial state s_n during period a_n, which is the number of times the PU changed its state during period a_n times the number of time-slots it spent in that changed state each time. Finally, the value of r_{n+1} depends on a discount rate γ which expresses the present value of future energy savings. The note above implies that during the learning phase the RD must perform spectrum sensing at every time-slot of period a_n in order to get the information needed to calculate r_{n+1}. We further note here that, in RL terminology, the set of periods $A(s)$ defined above is typically called the set of *actions* available in state S.

What is more, the RD maintains a mapping from each state $s \in S$ to probabilities of selecting each possible period a available at state S. This mapping can be updated at the end of each round and is called the RD's **policy**. We denote the RD's policy by μ, where $\mu(a \mid s)$ is the probability that the RD selects period $a \in A(s)$, given that the spectrum is in state $s \in S$. The objective of the RD is to find an optimal spectrum sensing policy μ^* such that the expected cumulative reward it receives over an infinite time horizon is maximized, i.e.,

$$\begin{aligned} q_{\mu^*}(s,a) &= \max_{\mu} q_{\mu}(s,a) \\ &= \max_{\mu} \mathbb{E}_{\mu}[\sum_{n=0}^{\infty} r_{n+1} \mid s_n = s, a_n = a], \end{aligned} \tag{6}$$

where $q_{\mu}(s,a)$ represents the *value* of choosing period a when the PU is in state S, i.e. the expected cumulative reward that the RD will receive, over an infinite time horizon, in case it senses the PU to be in state S, selects period a, and thereafter follows policy μ.

In case the one step dynamics of the environment depend only on the current state s_n and period a_n i.e.

$$\mathbf{Pr}\{r_{n+1}=r, s_{n+1}=s' \mid s_n=s, a_n=a\} = \mathbf{Pr}\{r_{n+1}=r, s_{n+1}=s' \mid s_0, a_0, \ldots, s_{n-1}, a_{n-1}, s_n, a_n\}$$

the RL problem presented above constitutes a semi-Markov Decision Process (SMDP). In [SBPW04, SPS99, BD95] the authors prove that under certain conditions related to the rewarding scheme of the finite SMDP there exists at least one such optimal policy μ^*. We analyse the conditions under which the RL problem at hand satisfies the Markov property at the end of this section, while the rewarding scheme that satisfies the necessary conditions for convergence is presented in subsection "D. Markov Options" below. Having $q_{\mu^*}(s,a)$ the RD can easily find the optimal period a^* to select at each state S by simply finding the period(s) for which $q_{\mu^*}(s,a)$ is maximized [SB98], i.e., for state S, $a^* = \arg\max_{a \in A(s)} q_{\mu^*}(s,a)$. Finally, any policy μ that assigns non-zero probability only to optimal period(s) a^* available at each state S is an optimal policy μ^*.

The RD utilizes the SMDP version of the well known Q-learning algorithm [SB98, SPS99] in order to find $q_{\mu^*}(s,a)$ and thus the optimal policy μ^* as described above. The SMDP Q-learning algorithm

maintains a table $Q(s,a)$ of period-value estimates for all periods available at state $s \in S$, where $Q(s,a)$ has been initialized arbitrarily. The RD will update the $Q(s,a)$ estimates, based on the real experience of selecting period $a \in A(s)$ in state $s \in S$ at the n -th round and observing the resulting reward r_{n+1} and next state s', as follows:

$$Q_{n+1}(s,a) = Q_n(s,a) + \alpha_n[r_{n+1} + \gamma^m \max_{a' \in A(s')} Q_n(s',a') - Q_n(s,a)] \quad (7)$$

Where m is the duration of period a and step-size parameter $\alpha_n = 1/[1+D_n(s,a)]$, where $D_n(s,a)$ is the total number of times state-period pair (s,a) has been visited up to and including the n-th round. The estimate $Q(s,a)$ converges to $q_{\mu^*}(s,a)$ with probability 1 for all $s \in S$ and $a \in A(s)$ under the given form of α_n and the additional conditions that each state-period pair is visited an infinite number of times and that the discount rate satisfies $0 < \gamma < 1$ [SBPW04,BD95]. The part of the RL agent that utilizes real experience in order to learn $q_{\mu^*}(s,a)$ by use of SMDP Q-learning is presented in Algorithm 1, lines 1 to 6 where the index n of the current round is omitted to keep the presentation simple.

Algorithm 1 Dyna-Q algorithm with SMDP Q-learning

1: Initialize $Q(s,a)$ and $\mathrm{Model}(s,a)$ for all $s \in S$ and $a \in A(s)$

2: **Do forever**:

3: $s \leftarrow$ sense current state of PU

4: $a \leftarrow$ select randomly from $A(s)$

5: Execute action a; Get the resulting reward r and state s'

6: $Q(s,a) \leftarrow Q(s,a) + \alpha \cdot [r + \gamma^m \cdot \max_a Q(s',a) - Q(s,a)]$

7: Update model $M(s,a)$ using the resulting reward r and next state s'

8: **For** $\{i = 1$ to $l\}$ **do**

9: $s \leftarrow$ random state from S

10: $a \leftarrow$ random action from $A(s)$

11: $r, s' \leftarrow M(s,a)$

12: $Q(s,a) \leftarrow Q(s,a) + \alpha \cdot [r + \gamma^m \cdot \max_a Q(s',a) - Q(s,a)]$

13: **End for**

The model learning function of the RL agent appears in line 7 of Algorithm 1. Learning model $M(s,a)$ involves utilizing the real experience produced at line 5 of Algorithm 1, i.e. the resulting next state s' and reward r. These outcomes will be averaged to approximate the expected values for $r(s,a)$ and $p(s' \mid s,a)$ by

$$\hat{r}_{n+1}(s,a) = \hat{r}_n(s,a) + \alpha_n[r - \hat{r}_n(s,a)], \quad (8)$$

and

$$\hat{p}_{n+1}(x|s,a) = \hat{p}_n(x|s,a) + \alpha_n[\delta(s',x) - \hat{p}_n(x|s,a)], \quad (9)$$

for all $x \in S$, where $\delta(s',x) = 1$ if $s' = x$ and 0 otherwise, and where the step-size parameter α is defined the same way as in Equation (7).

Learning through simulated experience (Dyna-Q) is described in lines 8 to 13 of Algorithm 1, where the RD selects randomly a starting state S and a period $a \in A(s)$ and the model outputs the resulting reward r and the next state s' according to $\hat{r}_n(s,a)$ and $\hat{p}_n(x|s,a)$. The output of the model is used to update $Q(s,a)$ as if the transition was produced by real experience. The process described above will be repeated l times, where l is the number of planning steps of the Dyna-Q algorithm. We must note here that in order to use the model we must first allow for a warm up period during which the RD performs only model learning so that estimates for all state-period pairs exist. Furthermore, the warm up period permits for real experience to be accumulated and alleviates the initial bias in the model that affects the quality of the produced policy.

Finally, we note that the RL problem at hand will always satisfy the Markov property in the special case of geometrically distributed sojourn times for the ON and OFF periods of the PU due to the memoryless property of the geometric distribution. However, we can still calculate $\mathbf{Pr}\{s_{n+1} = s' | s_n = s, a_n = a\}$, under the more general assumption of stationarity for the distributions of the sojourn times for the ON and OFF periods, even though the age, i.e. the time since the first time-slot the PU entered state s_n, X_i^a or Y_i^a (see Figure 15), is not represented in state s_n. We can do so by utilizing the cumulative distribution function (c.d.f.) $F_{X^e}(x^e)$ or $F_{Y^e}(y^e)$ of the residual or excess time, X^e or Y^e, in the current state s_n along with the point availability function [VL08] of the PU for delayed alternating renewal chains, i.e. the probability that the PU will be in the ON state at any future time-slot. The time average c.d.f. of the residual time in state s_n is derived in [G13] for the continuous time case and the same method can been applied to find the corresponding c.d.f in the discrete time case.

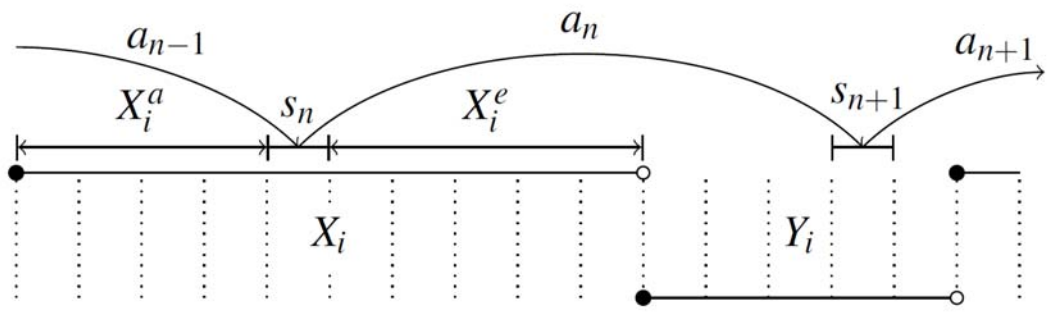

Figure 15. Example of age and excess times in PU's ON state..

The point availability function for a discrete time alternating renewal process is derived in [BL08]. We have extended the point availability function to the case of "delayed" alternating renewal processes under the assumption of stationary distributions of the ON and OFF periods' sojourn times, however, we do not present here the corresponding extension due to scarcity of space. What is more, the

solution to our SMDP problem does not involve direct calculations of $\mathbf{Pr}\{s_{n+1}=s' \mid s_n=s, a_n=a\}$, we rather consider the one-step dynamics of the environment as unknown quantities that the RD learns through interaction with the spectrum. The Markov property in the case of reward r_{n+1} is a direct consequence of the Markov property of state transitions and its definition that will be given in subsection "C. Base MDP" below. Figure 16 illustrates the SMDP for the problem of energy-efficient collection of spectrum occupancy data.

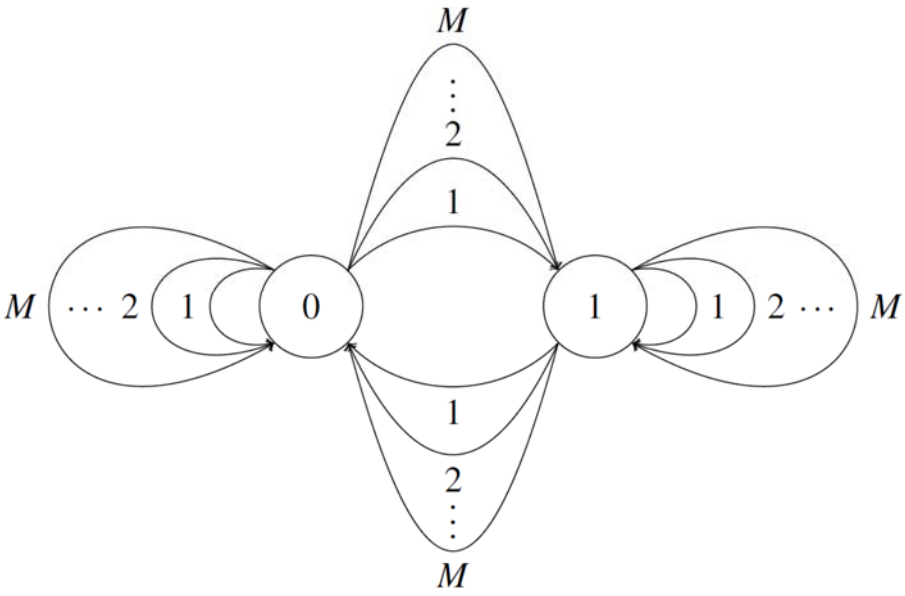

Figure 16. A SMDP for the problem of energy efficient collection of spectrum occupancy data.

C. Base MDP

Classical SMDP theory treats each candidate period of the SMDP model of Figure 16 as an indivisible and unknown unit. Thus, for the RD to learn about each period it must select it multiple times in each state and perform spectrum sensing at each time-slot during the period. This presents a serious shortcoming, especially in the case of a large number of candidate periods due to the excessive energy consumption involved in spectrum sensing.

However, we expect that from a fragment of experience, related to a particular period, we could learn simultaneously about all periods consistent with it, thus significantly boosting the learning process. A framework for SMDPs, that provides the means to look inside each period, was introduced by Sutton et al in [SPS99]. Within that framework the authors extended the usual notion of actions in MDPs to include "options", which are closed-loop policies taken over a period of time that are comprised of primitive actions, each lasting a single time-step. A fixed set of options can be used to define a discrete-time SMDP, such as the one presented in Figure 16, embedded within a base Markov Decision Process (MDP) in a way that periods are not black boxes anymore, but polices in the base MDP. Next, we present the base MDP for the problem of the energy efficient collection of spectrum occupancy data and define all periods in the SMDP of Figure 16 using the notion of options. Finally, we use options to speed up the learning process through intra-option model learning.

The base MDP for the optimal spectrum sensing problem is defined by its state and action sets, its transition rewards and by the one-step dynamics of the environment as follows:

- *States*

The set of states of the base MDP is defined as:

$V = \{(b_i, b_p, b_c, n_a, n_c, n_w, n_\phi) : b_i, b_c, b_p \in \{0,1\} \text{ and } n_a, n_c, n_w, n_\phi \in [1, M]\}$, where b_i, b_c, b_p represent the state of the PU during certain time-slots, n_a, n_c, n_w, n_ϕ are counters and M is the duration of the largest period available.

More specifically, b_i, b_c and b_p represent the state of the PU during the first, the current and the previous than the current time-slot of a spectrum sensing period (see Figure 17). Counter n_a represents the age of the PU's process in the current state, i.e. the number of time-slots the PU has remained in the same state since the first time-slot of the period or since the last state change of the PU. A state change of the PU is represented in the base MDP by $b_c \neq b_p$, i.e. when the current state of the PU is different than its previous state. Counter n_c counts the total number of time-slots since the period begun and n_w is the total number of time-slots the PU was found to be in a different state than the state it had during the first time-slot of the period, b_i. Finally, n_ϕ equals the value of n_c at the first time-slot $n_w \geq \phi$ becomes true. Parameter ϕ is a predefined threshold that indicates the tolerance of our system in having the PU in a different state than the one it had when the period begun, i.e, b_i for more than ϕ time-slots within the same period. Threshold ϕ is measured in time-slots and will be used to assess the accuracy of a period.

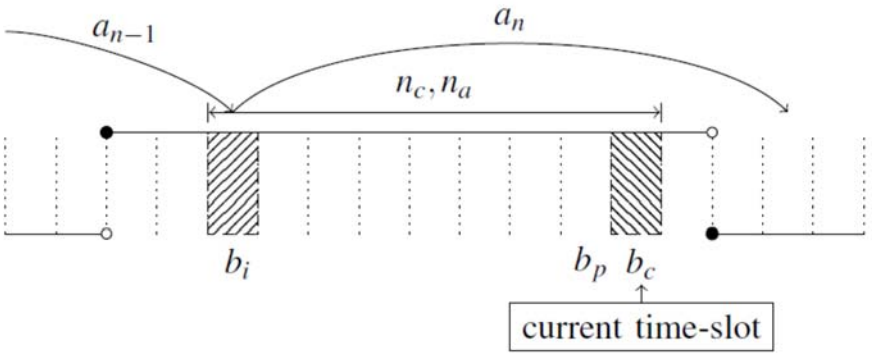

(a) Base MDP state $(b_i, b_p, b_c, n_a, n_c, n_w, n_\phi) = (1,1,1,9,9,0,0)$

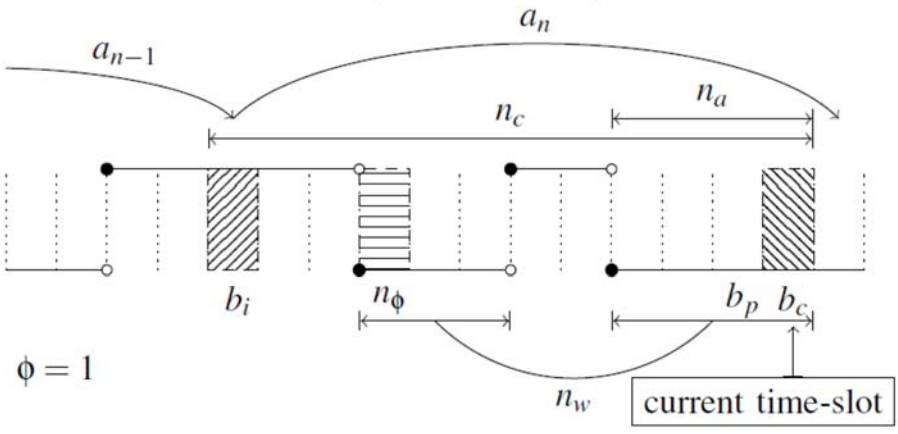

(b) Base MDP state $(b_i, b_p, b_c, n_a, n_c, n_w, n_\phi) = (1,0,0,4,12,7,4)$

Figure 17. Example of base MDP states representation.

We note that during the first time-slot of a spectrum sensing period we assume that $b_i = b_c = b_p$ and $n_w = n_\phi = 0$. Furthermore, we define the initiation set of states for the base MDP to be $V_I = \{v_{I_0}, v_{I_1}\}$ where $v_{I_0} = (0,0,0,1,1,0,0)$ and $v_{I_1} = (1,1,1,1,1,0,0)$. The states of the initiation set represent the two possible states of the base MDP at the commencement of a new period, where no information concerning the state of the PU during past time-slots exists. There is a one-to-one correspondence between states $v \in V_I$ of the base MDP and states $s \in S$ of the SMDP, more specifically state v_{I_0} (v_{I_1}) of base MDP corresponds to state $s = 0$ ($s = 1$) of the SMDP.

- *Actions*

For all states of the base MDP with $n_c < M$, i.e. the current number of time-slots since the initiation state has not exceeded the largest supported period M, the set of available actions is $U = \{u^e, u^r\}$, where action u^e corresponds to extending the current period by one time-slot, while action u^r corresponds to terminating the current period and returning to one of the initiation set states. Action u^e involves spectrum sensing during the next time-slot, storing the value of b_c to b_p, updating the value of b_c to the current state of the PU and properly incrementing the value of all state counters, i.e., incrementing n_a by one if $b_c = b_p$ or setting it to 1 otherwise, incrementing n_c by one, incrementing n_w by one if $b_c \neq b_i$ and finally setting $n_\phi = n_c$ if $n_w \geq \phi$. Action u^r involves spectrum sensing during the next time-slot and, based on the assessed state of the PU, a transition to the appropriate initiation set state. For all states of the base MDP with $n_c = M$, i.e. for all states that may result from a period of maximum duration the only available action is u^r, i.e., the current period will terminate necessarily and a new period will commence at the next time-slot.

- *Transition Rewards*

We define the following rewarding scheme for transitions from any state $v \in V$ to any state $v' \in V$ using action u^e:

$$g(v, u^e, v') = \begin{cases} e_c, & \text{if } n_w < \phi, \\ E_c - \dfrac{1-\gamma^{n_c}}{1-\gamma} \cdot e_c, & \text{if } n_w = \phi, \\ 0, & \text{otherwise,} \end{cases} \quad (10)$$

where e_c represents the energy spent for sensing the spectrum during a single time-slot and E_c is a negative reward or penalty, i.e. $E_c < 0$, associated with an inaccurate period, i.e. a period that is not accurately represented by b_i according to the accuracy criterion $n_w < \phi$, and corresponds to the energy cost that could be induced to the RD due to an inaccurate period. Such an energy cost may result for several reasons: (i) the RD must migrate from a spectrum band that was supposed to be free for the next m timeslots, yet the PU changed its state from OFF to ON, or (ii) because an optimal spectrum band, in terms of its transmission characteristics, was assumed to be occupied while it was free.

According to the presented rewarding scheme, for each time-slot we extend the current period we are rewarded with e_c. However, if the accuracy criterion is violated, i.e. $n_w = \phi$, then the aggregate reward gained up to that point, is subtracted and an additional penalty of E_c is imposed on the current transition. A transition from any state $v \in V$ to any state $v' \in V_I$ using action u^r always results in zero reward:

$$g(v,u^r,v') = 0, \text{for all } v, v' \in V, \tag{11}$$

- *One step dynamics of the environment*

The one step dynamics of the environment can be completely determined by the one-step transition probabilities,

$$p(v' \mid v,u) = \mathbf{Pr}\{v_{k+1} = v' | v_k = v, u_k = u\}, \tag{12}$$

where v_k, $v_{k+1} \in V$ represent the states of the base MDP at the k-th and $(k+1)$ time-slots respectively and $u \in U$ represents the action taken at the k-th time-slot, and the expected value of the next reward:

$$g(v,u,v') = \mathbb{E}[r_{k+1} | v_k = v, u_k = u, v_{k+1} = v'], \tag{13}$$

where, r_{k+1} is the reward granted to the agent for the one time-slot transition from current state v to state v' by taking action $u \in U$ at the k-th time-slot. As was the case with the semi-MDP we assume for the base MDP that the one-step dynamics of the environment are unknown to the RD and the agent will learn them through interaction with the environment. Establishment of the Markov property for the base MDP can be done under the same assumptions presented above in the Problem Formulation section for the semi-MDP. Under these assumptions $p(v' \mid v,u)$ and $g(v,u,v')$ can be completely determined for any pair of states $v, v' \in V$ and action $u \in U$ by utilizing the time average distribution of the excess or residual time in the appropriate state of the initiation set V_I together with the age information stored in counter n_a.

D. Markov Options

Options [SPS99] constitute sequences of primitive actions that produce temporally extended courses of action. In our problem setting, for example, period $a \in A$ of the SMDP with a duration of m time-slots can be constructed as a sequence of $(m-1)$ u^e actions of the base MDP followed by a u^r action.

Let V and U be, respectively, the sets of available states and actions of a MDP, then a Markov option is made of the following three components: an initiation set $I \subseteq V$, a policy $\pi : V \times U \to [0,1]$ and a termination condition $\beta : V \to [0,1]$. An option defined by $\langle I, \pi, \beta \rangle$ is available at state $v_k \in V$, at the k-th time-slot, if and only if $v_k \in I$. An option, selected at initiation state $v_k \in I$, executes iteratively as follows:

- First, an action $u_k \in U$ is selected with probability $\pi(v_k, \cdot)$ and then the environment is sensed to identify it's new state v_{k+1} and the resulting reward g_{k+1}.

- Next, the option either terminates with probability $\beta(v_{k+1})$ or continues by determining the next action u_{k+1} according to $\pi(v_{k+1},\cdot)$.
- Eventually, the option will terminate at some state and the agent will have the chance to select another option available at that state.

We will define the set of available periods A for each state $s \in S$ of the PU using the notion of options over the base MDP. In the Problem Formulation section above, we defined the set of states of the PU to be $S = \{0,1\}$. Now let $V_s = \{v \mid v \in V, s \in S, b_i = s\}$ be the subsets of states of the base MDP, where the initial state of the PU is OFF and ON respectively. Then a period $a \in A(s)$ commencing at PU's state $s \in S$ with a duration of m time-slots can be defined as an option $\langle I, \pi, \beta \rangle$ over the base MDP with,

1) Initiation set: $I = \{v_{I_s}\}$.
2) Policy: For all states $v \in V$ and actions $u \in U$:

$$\pi(v,u) = \begin{cases} 1, & v \in V_s \text{ and } u = u^e \text{ and } n_c < m, \\ 1, & v \in V_s \text{ and } u = u^r \text{ and } n_c = m, \\ 0, & \text{otherwise}, \end{cases} \quad (14)$$

3) Termination condition:

$$\beta(v) = \begin{cases} 1, & \text{if } v \in V_I, \\ 0, & otherwise, \end{cases} \quad (15)$$

Thus, period $a \in A(s)$ can be defined as option $\langle I, \pi, \beta \rangle$ that commences in initiation state v_{I_s}, $s \in S$, performs action u^e until the total number of time-slots counted by n_c becomes equal to period a's duration $m-1$, and then action u^r will be performed in order to a return to an initiation state. At that point, the option will terminate, as dictated by its termination condition, and a new option will be selected by the RD. From this point on we will use a_o to denote option $\langle \mathcal{I}, \pi, \beta \rangle$, defined over the base MDP, that corresponds to period a and A_o to denote the set of all available options.

- *Expected reward of a period*

In [SPS99] the authors prove that for any MDP and any set of options defined on that MDP, the decision process that selects only among those options, executing each to termination, is a SMDP. Having established the equivalence between the SMDP process and base MDP combined with the set of options A_o we can now define the expected reward for a period a with duration m time-slots selected at the first time-slot of round n when the state of the PU is $s \in S$, in terms of its corresponding option defined over the base MDP as follows:

$$r(s,a) = r(v_I, a_o) = \mathbb{E}\{g_{k+1} + \gamma \cdot g_{k+2} + \cdots + \gamma^{m-1} \cdot g_{k+m} \mid \varepsilon(a_o, v_I, k)\} \quad (16)$$

where, $\varepsilon(a_o, v_I, k)$ denotes the event of option a_o being initiated in state $v \in V_I$ at the k-th time-slot, $k+m$ is the random time at which option a_o terminates and $\gamma \in [0,1)$ is the discount rate of the SMDP. Finally, the discounted rewarding scheme presented in Equation (16) satisfies [JBPW04, SPS99, BD95] the necessary conditions for the SMDP Q-learning algorithm to converge to q_{μ^*} with probability 1.

- *Intra-option model and Q-value learning*

For Markov options special temporal-difference techniques [SB98, SBPW04, SPS99] can be used to learn their model, i.e., $r(v_I, a_o)$ or equivalently $r(s,a)$ and $p(v_{I'} \mid v_I, a_o)$ or equivalently $(p(s' \mid s,a))$, while executing a different option, given that some of the selections of the latter option are consistent with the selections of the first one. These techniques are called intra-option learning techniques and are examples of off-policy learning techniques whereby we learn about the consequences of one option's policy while actually behaving according to another. For example, in our problem setting, consider an option $a_{o'}$ with duration m' time-slots that would have terminated a certain number of time-slots earlier than the currently selected option a_o with duration m time-slots, i.e., $m' < m$. Formally it is a different option than a_o and formally it was not executed, yet all the experience gained by executing a_o can be used to improve the model of $a_{o'}$ and thus significantly boost the learning rate of the SMDP Q-learning algorithm.

For the problem at hand we utilize two such intra-option learning techniques, called intra-option model learning with eligibility traces and intra-option Q-value learning with eligibility traces, respectively, that are known [SPS99, SB12] to converge to the optimal solution with probability 1 under the same condition presented in the Problem Formulation section for the SMDP Q-learning algorithm. Both techniques involve storing the rewards g_{k+i} and states s_{k+i} of the PU, for each time-slot $i \in [1, m]$ of the currently executing option a_o. The stored trace renders eligible to update all state-option pairs (s_k, $a_{o'}$), where option $a_{o'}$ has a shorter duration than the currently executing period. In the case of intra-option model learning we update the model of each eligible to update state-option pair (s_k, $a_{o'}$) by applying Equation (8) and Equation (9), while in the case of intra-option Q-value learning we update directly the Q-value of the corresponding state-option pairs by applying Equation (7). More specifically, for each option $a_{o'}$ with duration m' we use $r = g_{k+1} + \cdots + \gamma^{m'-1} \cdot g_{k+m'}$ and $\delta(s', x) = \delta(s_{m'}, x)$ to perform the necessary model or Q-value updates.

In terms of Algorithm 1 and in the case of intra-option model learning we apply step 7 of Algorithm 1 for all eligible to update state-option pairs while in the case of intra-option Q-value learning we apply step 6 of the algorithm for the same set of state-option pairs. We can further accelerate the learning rate of the RD by slightly increasing the probability to select options with a large duration since both intra-option learning methods will always perform updates for the periods with a shorter duration. This selection bias will not affect the convergence property of the SMDP Q-learning or that of the Dyna-Q algorithm as long as all state-option pairs are visited infinitely many times in the limit. However, it will significantly affect the convergence speed of the algorithm as it will increase the learning rate of the whole set of options.

2.3.3.2 Evaluation

In this section we present simulation results for the SMDP Q-learning algorithm and its enhanced versions with either intra-option Q-value learning or intra-option model-learning (Dyna-Q). Our simulation model is comprised of a PU, whose OFF and ON state durations are geometrically distributed with mean $\mu_0 = 70$ and $\mu_1 = 10$ time-slots respectively. For each state S of the PU the set of candidate periods is $A(s) = \{1, 2, \ldots, 100\}$, $s \in \{0,1\}$, tables $Q(s,a)$ and $M(s,a)$ are initialized to zero and the accuracy threshold ϕ is set to 1 for all simulations. With $\phi = 1$ any state change of the PU within a selected period renders the period inaccurate and incurs a cost of E_c. Furthermore, we assume a unit energy cost $e_c = 1$, for sensing the channel at a single time-slot and a discount factor $\gamma = 1 - \frac{1}{10^6}$.

Figure 18 presents the optimal period for both states of the PU versus the energy cost E_c.

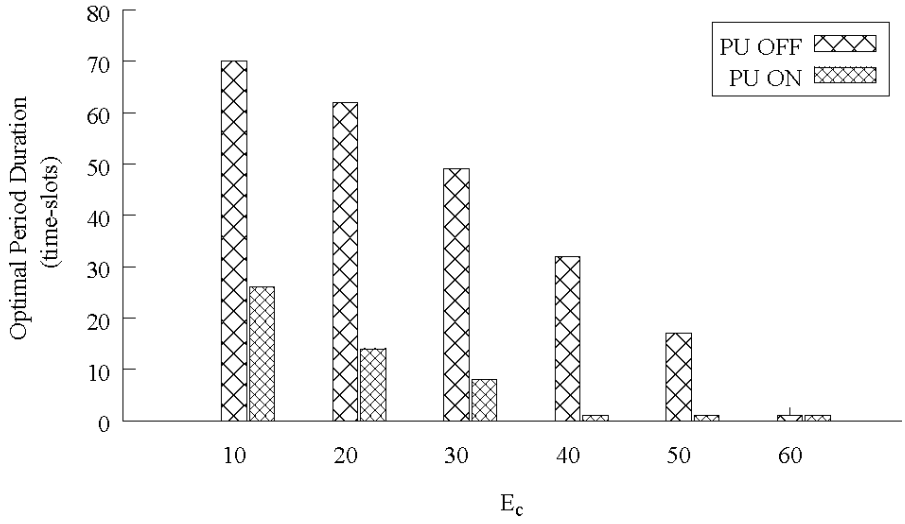

Figure 18. The duration of the optimal period for each state of the PU will diminish as the cost of an inaccurate period Ec increases.

We considered E_c values that are integer multiples of the energy, e_c, spent for sensing the channel at a single time-slot.

Figure 18 shows that for larger values of E_c the duration of the optimal period decreases and eventually the optimal solution is to sense the channel at every single time-slot in order to avoid the high cost of being inaccurate. However, sensing a channel frequently may, by itself, be prohibitive for the energy constrained wireless sensor due to the associated spectrum sensing cost e_c. Thus, the RD may decide to completely disregard specific regions of the spectrum by considering the length of their recommended sensing periods. This concept can be generalized so as to construct a partial ordering of all channels under consideration based on their recommended periods' length.

SMDP Q-learning and Dyna-Q algorithms are known to converge with probability 1 in an infinite number of rounds under the conditions described in the Problem Formulation section. Actual simulation results indicate that the SMDP Q-learning algorithm requires a number of rounds in the

order of millions to practically converge to the optimal pair of spectrum sensing periods. However, we expect that the energy required for such a feat is prohibitive for a wireless sensor and it will have to stop the learning process within a much smaller number of rounds and accept the recommended periods even if they are suboptimal. The proposed intra-option model learning and Q-value learning techniques aim to increase the convergence rate of the Dyna-Q and SMDP Q-learning algorithm to the optimal solution and thus reduce the error in the recommended periods.

Figure 19 and Figure 20 illustrate the effect of the number of planning steps on the convergence rate of the Dyna-Q algorithm with intra-option model learning in our simulation scenario when we set E_c to 20, i.e, the optimal pair of periods is $(62,14)$ and use a warmup period of 1000 rounds.

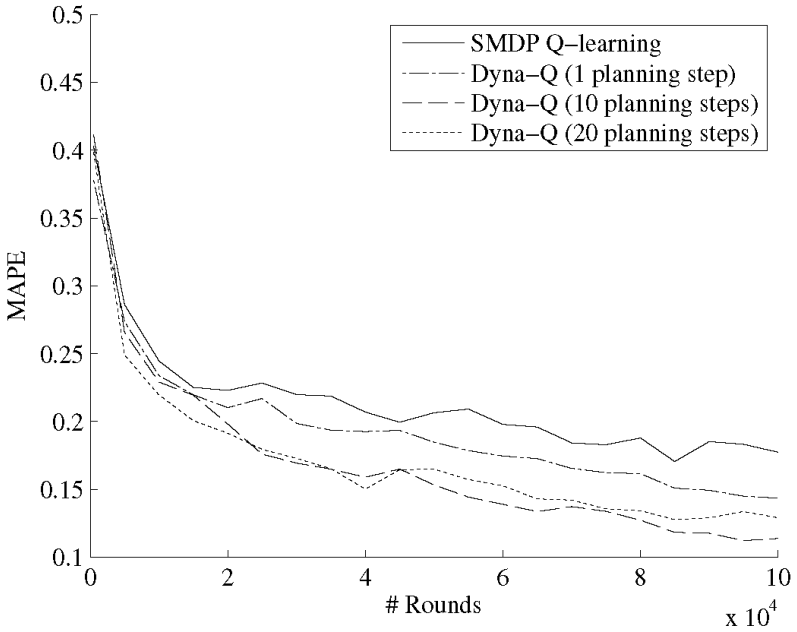

Figure 19. The Dyna-Q algorithm will converge faster to the optimal period for the PU OFF state due to the use of simulated experience.

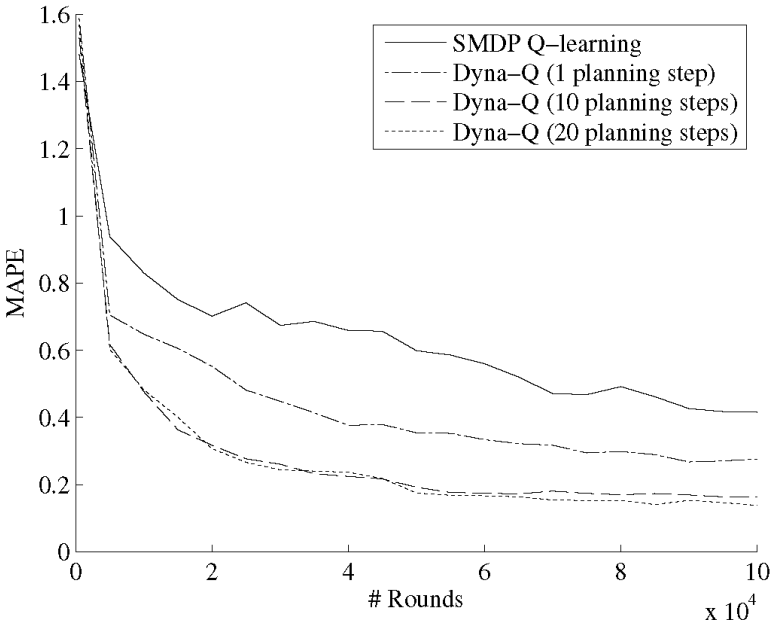

Figure 20. The Dyna-Q algorithm will converge faster to the optimal period for the PU ON state due to the use of simulated experience.

More specifically in Figure 19 and Figure 20 we plot the mean absolute percentage error (MAPE), calculated for each round n over $K = 200$ repetitions of the same scenario, as follows:

$\mathrm{MAPE}_s(n) = \frac{1}{K}\sum_{j=1}^{K}\frac{|\hat{m}^*_{j,s}(n) - m^*_s|}{m^*_s}$ where, $\hat{m}^*_{j,s}(n)$ is the duration of the best period estimate of the j-th scenario repetition at the n-th round for state $s \in S$ of the PU and m^*_s is the duration of the optimal period for the same state.

The results in Figure 19 and Figure 20 indicate that the Dyna-Q algorithm with intra-option model learning will converge to the optimal period faster than the SMDP Q-learning algorithm which does not use simulated experience. This is due to the fact that the proposed intra-option model learning technique permits the RD to create an accurate enough model, within a small number of rounds, and use it to speed up learning via simulated experience. However, increasing the number of planning steps beyond 10 will not increase any further the convergence rate of the Dyna-Q algorithm due to the unavoidable inaccuracies in the model and thus the excess energy spent on the CPU for simulation is wasted. We must further note that the number of planning steps may also be constrained by the processing power of the sensor and the interval between two successive time-slots. Finally, we note that if the sensor has access to an accurate model $M(s,a)$ of the PU, e.g. a model provided by another RD, then it could decide the optimal pair of periods based solely on simulated experience.

In Figure 21 and Figure 22 we present comparative results for the SMDP Q-learning algorithm with and without the intra-option Q-value learning technique. We observe in Figure 21 and Figure 22 a significant reduction in MAPE, within only a few hundred rounds, when the RD enhances SMDP Q-learning with the intra-option Q-value learning technique. What is more, the SMDP with intra-option Q-value learning converges faster to the optimal solution than the Dyna-Q algorithm with intra-option model learning presented in Figure 21 and Figure 22 and it is more memory and energy conserving as well, since it does not require the storage of the model table $M(s,a)$ and the generation of simulated experience. On the other hand model $M(s,a)$, once it becomes accurate enough, can be useful to the wireless sensor in the special case parameter E_c changes and a new optimal pair of spectrum sensing periods must be derived without resorting to spectrum sensing again.

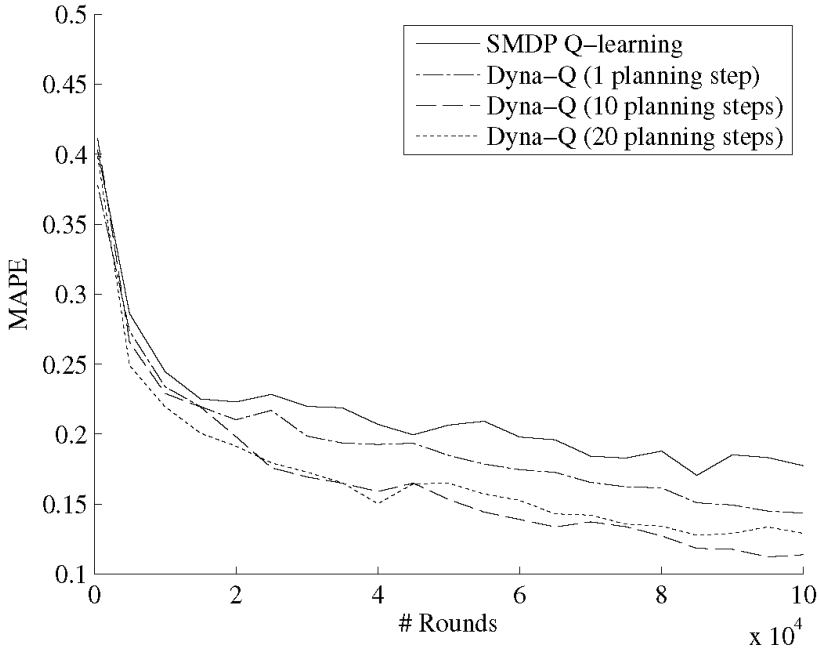

Figure 21. SMDP with intra-option Q-value learning will converge faster to the optimal period for the PU OFF state due to the increased learning rate.

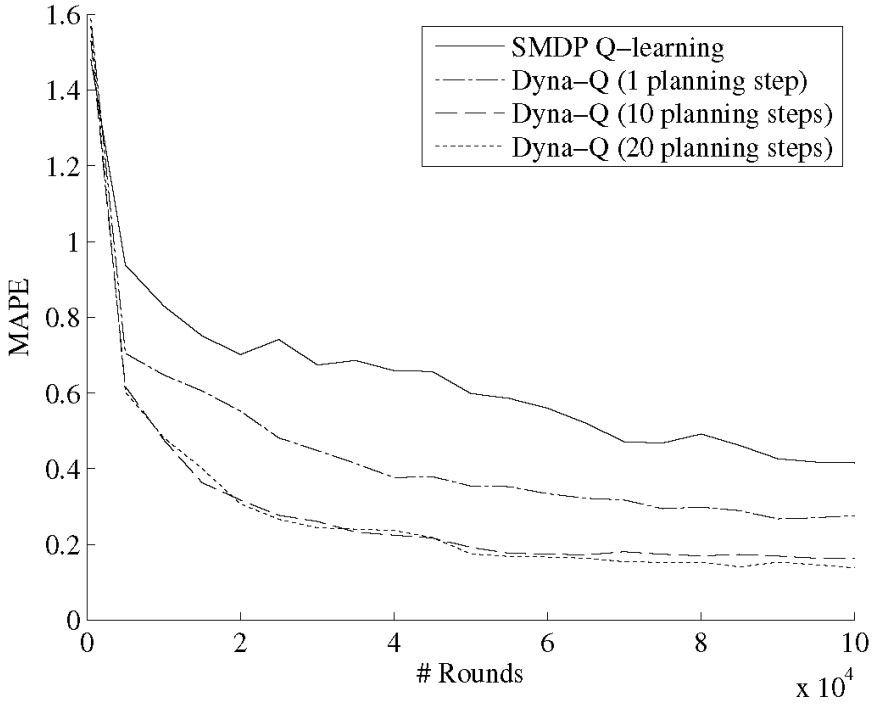

Figure 22. SMDP with intra-option Q-value learning will converge faster to the optimal period for the PU ON state due to the increased learning rate.

As it can be easily acknowledged from the above results, by taking advantage of intra-option learning techniques we can achieve significant gains in learning speed, which results in important energy savings in the constrained IoT devices, prolonging their lifetime. This can have great benefits for network operators decreasing their deployment and operational costs.

2.4 Lightweight mechanisms for spectrum assignment[4]

2.4.1 Introduction and state of the art

As described in Figure 6, the module that takes the decision about the spectrum that the RD will use for a specific transmission is the SDM. Selecting the optimal spectrum is of outmost importance to ensure that the QoS of the services can be provided efficiently and that no other transmissions will be affected. In the case of CR-inspired RDs there is another constraint that has to be taken into account and this is the fact that licensed users (or Primary Users) have specific frequencies that they operate and by law their transmissions should not face interference by RDs. This is indeed the case in the RERUM outdoor use cases, where the RDs will be deployed in city areas, where multiple PUs are operating (e.g. TV and Radio stations). So, the RDs will have to select frequencies that are not assigned to the PUs or transmit only in the timeslots that are not utilized by the PUs to avoid creating interference to them.

Since their conception [MM99] cognitive radio networks have been the topic of intense research. A key challenge in the design of such networks is that of *dynamic spectrum allocation*, i.e., the opportunistic assignment of spectrum bands to wireless nodes so that they are able to communicate

[4] Part of this work has been submitted to the IEEE QoMex 2015 conference [STT15b].

without incurring harmful interference to the incumbent users of these spectrum bands. Existing works regarding dynamic spectrum allocation can be classified by their optimization objectives into performance-centric and energy-centric.

In performance-centric approaches, the objective of the dynamic spectrum assignment scheme is usually expressed in terms of throughput maximization or delay minimization and in some cases in terms of fair resource allocation. The authors typically attain their objectives through interference minimization and efficient spectrum utilization by means of dynamic spectrum assignment. Examples of performance-centric dynamic spectrum assignment schemes include [XJL08, XJL12, LA11] and many of the references in survey [TZFS13]. By design, performance centric approaches do not consider power consumption and thus the proposed dynamic spectrum assignment schemes are prohibitive for the power-constrained devices considered in this work.

In contrast, energy-centric approaches aim to minimize the power consumption of the devices in order to maximize their operational lifetime. In [GQV08], the authors propose a distributed algorithm to identify the optimal number of channels that a cognitive secondary user must access so as to transmit with the least power consumption while satisfying a given data rate requirement. In [YLH10] the authors aim to maximize total capacity while minimizing power consumption of their cognitive network by means of transmission power minimization on the channels selected by their proposed dynamic spectrum assignment scheme while not causing interference to PUs, not selecting a channel used by another SU and not exceeding the power constraint of the device. In [LWM11] the authors' objective is to prolong the operational lifetime of a cognitive sensor network by predicting the residual energy of its nodes and minimizing their total transmission power by means of a model for the energy spent during transmissions as a secondary user.

Although the proposed schemes reduce energy consumption in cognitive networks the aforementioned approaches do not consider the power consumption involved in the cognitive functions that support dynamic spectrum assignment itself such as spectrum sensing, spectrum switching and the coordination with other devices or the base station at the PHY and MAC layers. Different than previous work, however, we constraint the energy consumption involved in dynamic spectrum assignment itself, by assuming a two-stage spectrum assignment process, which we further constraint in a small region of the whole opportunistic access band while still considering the power consumption required for the transmission over different regions of bandwidth.

2.4.2 Preliminaries

In this section we consider the problem of dynamically assigning spectrum bands to the RD in a way that satisfies certain QoS criteria. To facilitate our presentation we begin with the **simple case of a single RD** that has to transmit L bits of data within T_d time units **with the least possible power consumption**. Subsequently, we generalize our approach to the case of multiple RDs that share the same spectrum region for communication.

In order to achieve the desired objective, in the simple case of the single RD, we must assign it an appropriate number of sub-channels within a spectrum region B that is available for opportunistic access. We assume that spectrum region B is divided into N channels. Each of these channels is occupied by a Primary User (PU) and is subdivided into M sub-channels (Figure 23).

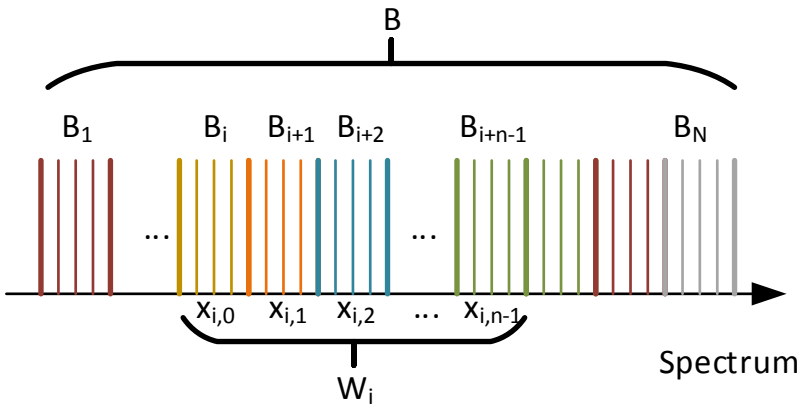

Figure 23. Model of the spectrum available for opportunistic access and the accessible spectrum by the RD.

Without loss of generality we assume that all channels have the same bandwidth $B_i = \frac{B}{N}$ and consequently the same number of sub-channels M. This is a common system description in many licensed spectrum bands [YB++07]. Each sub-channel supports a fixed transmission rate r when using a fixed transmission power c. Transmission power c may differ between channels due to the varying levels of noise and interference and of course the different path loss exponents.

Each RD can access a spectrum region of bandwidth W, where W is assumed to be smaller than or equal to B. Again, to keep presentation simple we assume that W consists of n channels, i.e., it is an integral multiple of channel's bandwidth. We denote with W_i the i-th region of bandwidth W within B (see Figure 23). In this work we consider only regions W_i whose first sub-channel coincides with the first sub-channel of a channel, i.e., we do not consider spectrum regions of bandwidth W that would include partially a channel. Hence, for example we would not consider a spectrum region W_i that would begin at the middle of channel B_i as a candidate region. There exist $(N+1)-n$ valid regions to consider and thus $i \in \{1, 2, \dots, (N+1)-n\}$.

The proposed spectrum assignment scheme will assign zero or more sub-channels, as indicated by the value of variable x_i^j from each channel $B_{i+j}, j = 0, 1, \dots, n-1$ that lies within W_i. **Aggregated**, these sub-channels will form a link with adequate transmission rate. The objective of our proposed spectrum assignment scheme is to find both the optimal region $W^*$and the optimal spectrum assignment within it so as to satisfy the QoS constraint while minimizing the transmission power consumption.

2.4.3 Primary User Model

The PU model considered in this section is identical to the discrete model presented in section 2.3.3.1, i.e., that of a discrete time semi-Markov process [YB++07]. For the sake of clarity and economy in notation in this section we denote with $\mathbb{S} = \{0,1\}$ the set of PU states and with $S_k \in \mathbb{S}$ the random variable that indicates the state of a PU at the k-th time-slot, where $S_k = 0$ corresponds to the OFF state and $S_k = 1$ corresponds to the ON state.

The model assumed for the PU activity is generic enough to support the majority of probability distributions used in the literature (see [BL08, LA11], the survey paper [J++12] and the references therein) to characterize PU's behaviour as well as those suggested by the large measurement campaigns presented in [SR14, WRM09].

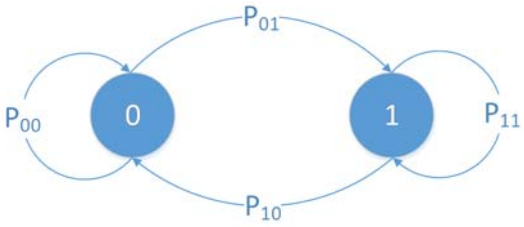

Figure 24. A two state Markov chain model for the PU activity on its channel.

The same is also true for the model of a two-state Markov chain (Figure 24) that is a special case of the semi-Markov model presented in section 2.3.3.1. More specifically, the two state semi-Markov model reduces to the two state Markov chain model when the sojourn time in each state is equal to 1. For the two state Markov chain model the durations of the ON/OFF periods are geometrically distributed. Obviously, not all PUs activity and inactivity sojourn times can be modelled accurately by a geometric distribution. However, extensive measurement studies conducted by the authors in [SR14] indicate that geometric distribution can model effectively traffic distributions for PUs of multiple communication systems. Furthermore, we make a special mention to this model because it is used extensively in the literature in the context of dynamic spectrum assignment (see [BL08, LA11], the survey paper [J++12] and the references therein).

For both models mentioned above we can derive a steady-state probability distribution $\pi_s = P(S_n = s)$ when n is large. Probability distribution π_s indicates the probability that the PU will be in state S given that the process started a long time ago and is independent of the initial state of the PU. Actually, a convergence to a steady state probability distribution may only occur if the Markov chain is comprised of a single, aperiodic recurrent class with finite mean sojourn times for all its states. To put it simply, for a RD that transmits data infrequently, the probability that the PU of a channel it accesses will be in state S can be approximated with π_s independently of its last known state as long as the PU is not always ON, always OFF or alternates between the ON/OFF states periodically. We assume that the aforementioned conditions are satisfied and that the steady-state probability distribution is known for all PUs.

2.4.4 Dynamic spectrum assignment as a stochastic two-stage integer program

In this section we consider the problem of selecting **both an appropriate region W and the number of sub-channels from each channel within W** that should be assigned to the RD so that it can transmit L bits of data within T_d time units with the least the least possible power consumption. A static spectrum assignment scheme would not suffice to this objective as the assigned sub-channels may be occupied by their respective PUs when the data to be transmitted become available. Consequently, a dynamic spectrum assignment scheme is imperative in order to satisfy the QoS requirements posed by the problem.

Within RERUM we propose a dynamic spectrum assignment scheme that is based on a **two-stage stochastic integer program with non-fixed recourse** [WM07, BL11, KW94, KW12]. According to this scheme spectrum assignment will be conducted in two stages:

a) In the first stage a number of sub-channels will be assigned to the RD. This assignment will be conducted before any information concerning the true state of their respective PUs is known. The assignment will be based solely on the probabilistic models of the PUs occupying spectrum region B.
b) Subsequently, when the data to be transmitted become available, the RD will assess the state of the assigned sub-channels and, if necessary, appropriate recourse actions will be taken so that the RD may transmit its data on time. This is called the second-stage spectrum assignment and the corresponding period is called the second stage.

We must emphasize here that what separates the first from the second stage is the assessment of the assigned spectrum's state and that each stage may in fact contain a whole sequence of decisions and actions instead of just a single decision. After the transmission of the data the RD will revert back to the first-stage spectrum assignment, i.e., its default configuration, and will wait for new data to transmit.

To anticipate the case of one or more sub-channels, from the set assigned to the RD in the first stage, being occupied by their respective PUs we provide two types of recourse actions. First, let us denote with W^* the spectrum region where the sub-channels, assigned to the RD at the first stage, reside. The first type of recourse action is to search within W^* for potentially available sub-channels to muster the necessary aggregate transmission rate. This recourse action is efficient in terms of power consumption since it does not involve spectrum switching and sensing over other regions of B for potentially available sub-channels. All the sensor has to do is perform spectrum sensing over the channels within W^* whose state is yet unknown.

We expect that in certain cases the number of available channels within W^* may not be adequate to satisfy the required demand in transmission rate. In such a case the second type of recourse action will be taken, i.e., the cognitive sensor will migrate to a free ISM band. What is more, since we cannot expect the RD to be able to utilize at the same time both the available parts of the opportunistic spectrum access band and the ISM band, we assume that all necessary capacity must be allocated within the ISM band.

We note here that another potential recourse action could be the migration to a dedicated band where the RD will assume the role of a PU. However, for such a type of recourse action we should consider monetary costs besides transmission power consumption and thus it is beyond the scope of RERUM.

The **motivation** to complete the spectrum assignment in two stages stems from the fact that both processes involved in dynamic spectrum assignment, i.e., spectrum sensing and switching are costly in terms of power consumption and delay and thus incongruous with the **power constrained nature** of RDs and the QoS requirements of the problem. By restricting the dynamic spectrum assignment to take place in only two stages we significantly reduce power consumption due spectrum sensing for spectrum holes discovery and incorporate a single spectrum switching action in the whole process.

Finally, before we proceed with the details of the proposed dynamic spectrum assignment scheme we note that the entity responsible for the spectrum assignment can be either the gateway, the cluster head or the RD itself, given that it possesses the necessary computational power.

2.4.4.1 QoS constraint in the case of a single RD

In this section we express the QoS constraint of a single RD that has to transmit L bits of data within T_d time units in terms of the aggregate transmission rate of the virtual link set up by the two stage process described in the previous section. More specifically the QoS constraint in its initial form can be written as:

$$\frac{L}{R_x + R_y} + \tilde{\tau}_o \leq T_d \Leftrightarrow \tag{17}$$

$$R_x + R_y \geq \frac{L}{T_d - \tilde{\tau}_o} \tag{18}$$

where R_x, R_y are the aggregate transmission rates of the sub-channels assigned to the RD in the first and second stage respectively. Random variable $\tilde{\tau}_o$ represents the overhead time spent in the MAC and PHY layers and includes the overhead from all cognitive functions involved in the second stage spectrum assignment such as spectrum sensing, spectrum switching and coordination with the receiver. Now, R_x is defined as follows:

$$R_x = \sum_{j=0}^{n-1} r \cdot \xi_i^j \cdot x_i^j \tag{19}$$

where x_i^j is the number of sub-channels assigned to the RD in the first stage from the $(i+j)$-th channel $B_{i+j}, j = 0, 1, \dots, n-1$ within W_i, r is the fixed transmission rate supported by each sub-channel and ξ_i^j is a binary random variable indicating the availability of channel B_{i+j}. The value of ξ_i^j depends on the state S_i^j of the PU of the $(i+j)$-th channel B_{i+j} as follows:

$$\xi_i^j = \begin{cases} 0, & if\ S_i^j = 1, \\ 1, & if\ S_i^j = 0, \end{cases} \tag{20}$$

The aggregate transmission rate from the sub-channels assigned at the second stage is similarly defined as:

$$R_y = \sum_{j=0}^{n-1} r \cdot \xi_i^j \cdot y_i^j + r \cdot h_d \cdot y_f, \tag{21}$$

where, y_i^j is the number of sub-channels assigned in the second stage from the $(i+j)$-th channel $B_{i+j}, j = 0, 1, \dots, n-1$ within W_i, y_f is a binary variable indicating whether the free ISM band will be utilized or not and the product $r \cdot h_d$ equals the total transmission rate required for the timely transmission of the data entirely through the ISM band, i.e.,

$$r \cdot h_d = \frac{L}{T_d - \tilde{\tau}_o}, \tag{22}$$

Equation (22) can be written as:

$$h_d = \left\lceil \frac{L}{(T_d - \tilde{\tau}_o) \cdot r} \right\rceil \tag{23}$$

Indicating that h_d is the number of sub-channels of transmission rate r necessary to transmit the L bit of data within T_d time units. From equation (23) we have that T_d cannot be equal or less than the overhead time $\tilde{\tau}_o$, i.e., the guaranteed delay must be larger than the overhead time of the process. We further note that $h_d = \left\lceil \frac{L}{(T_d - \tilde{\tau}_o) \cdot r} \right\rceil = \left\lceil \frac{L/T_d}{(1 - \tilde{\tau}_o/T_d) \cdot r} \right\rceil = \left\lceil \frac{R}{(1 - \tilde{\tau}_o/T_d) \cdot r} \right\rceil$ where R denotes the transmission rate of a dedicated link, i.e. with zero overhead, that can transfer L bits within T_d time units. By Equations (18), (19), (21), (23) we get the following expression for the QoS constraint:

$$\sum_{j=0}^{n-1} r \cdot \xi_i^j \cdot x_i^j + \sum_{j=0}^{n-1} r \cdot \xi_i^j \cdot y_i^j + r \cdot h_d \cdot y_f \geq \frac{L}{T_d - \tilde{\tau}_o} \tag{24}$$

$$\sum_{j=0}^{n-1} \xi_i^j \cdot x_i^j + \sum_{j=0}^{n-1} \xi_i^j \cdot y_i^j + h_d \cdot y_f \geq \frac{L}{(T_d - \tilde{\tau}_o) \cdot r} \tag{25}$$

$$\sum_{j=0}^{n-1} \xi_i^j \cdot x_i^j + \sum_{j=0}^{n-1} \xi_i^j \cdot y_i^j + h_d \cdot y_f \geq h_d \tag{26}$$

Equation (26) indicates that, given a fixed transmission rate for each sub-channel, the number of sub-carriers assigned in the first stage plus the number of sub-carriers assigned in the second stage should be greater than or equal to the demand in sub-channels. In Equation (26) we use an inequality sign to anticipate for the case that we resort to the ISM band in the second stage. In such a case, we always assume that we bind the total number of required sub-channels h_d from the ISM band. This number of sub-channels, bound from the ISM band, plus any available sub-channels assigned in the first stage, i.e., sub-channels with PU in OFF state, as indicated by $\sum_{j=0}^{n-1} \xi_i^j \cdot x_i^j$ will certainly surpass the demand h_d. However, as we have already mentioned, the sensor will use exclusively the ISM band in such a case and therefore we have properly formulated the cost function to account only for the transmission power cost spent in the ISM band.

2.4.4.2 Formulation of a two-stage stochastic integer program for the dynamic spectrum assignment problem

For each candidate spectrum region W_i, we formulate a two-stage, stochastic integer program (IP) with non-fixed recourse. The solution to each of these programs will provide us with the **optimal spectrum assignment within the corresponding region W_i**. Subsequently, we will choose the optimal region W^* among the ones that provide the spectrum assignment with the least power consumption. The two-stage, stochastic IP is stated as follows:

$$\min\ z_i = \boldsymbol{c}_i^T \cdot \boldsymbol{x}_i + \mathbb{E}_{\boldsymbol{\xi_i}}\{\min[\![(\boldsymbol{q}_i^T \cdot \boldsymbol{y}_i(\boldsymbol{\xi_i})]\!] + q_f \cdot y_f(\boldsymbol{\xi_i}) - \boldsymbol{b}_i^T(\boldsymbol{\xi_i}) \cdot \boldsymbol{w_i})\}, \tag{27}$$

s.t.

$$\sum_{j=0}^{n-1} \xi_i^j \cdot x_i^j + \sum_{j=0}^{n-1} \xi_i^j \cdot y_i^j + h_d \cdot y_f \geq h_d, \tag{28}$$

$$x_i^j = w_i^j \text{ for all } j \in \{0, 1, \dots, n-1\}, \tag{29}$$

$$x_i^j + y_i^j \leq M, \text{ for all } j \in \{0, 1, \dots, n-1\}, \tag{30}$$

$$x_i^j \leq M \text{ for all } j \in \{0, 1, \dots, n-1\}, \tag{31}$$

$$y_i^j \leq M \text{ for all } j \in \{0, 1, \dots, n-1\}, \tag{32}$$

$$y_f \in \{0, 1\} \text{ and } x_i^j, y_i^j, w_i^j \in \mathbb{N}, \tag{33}$$

The objective function of the two-stage stochastic IP is presented by Equation (27). It expresses our objective to minimize the total power consumed for the transmission of data, z_i, by properly assigning sub-channels within region W_i in the first and second stage. More specifically, $\mathbf{c}$ is a cost vector whose elements represent the transmission power consumption for the transmission of data over a single sub-channel denoted by $P_{tx}^{i,j}$. Indices i, j on $P_{tx}^{i,j}$ indicate that transmission power can be different for sub-channels that belong to different channels. The inner product $\boldsymbol{c}_i^T \cdot \boldsymbol{x}_i$ represents the total power cost of the first stage when assigning and using sub-channels $\boldsymbol{x_i} = (x_i^0, x_i^1, \dots, x_i^{n-1})$ from the i-th region W_i to the RD. As mentioned in the previous section x_i^j indicates the number of sub-channels assigned to the RD in the first stage from the $(i+j)$-th channel $B_{i+j}, j = 0, 1, \dots, n-1$ within W_i.

The **expected power cost** of the second stage spectrum assignment is determined by the state of the PUs within spectrum region W_i, which is represented by random vector $\boldsymbol{\xi_i} = (\xi_i^0, \xi_i^1, \dots, \xi_i^{n-1})$. For each possible value of random vector $\boldsymbol{\xi_i} = (\xi_i^0, \xi_i^1, \dots, \xi_i^{n-1})$ we have a potentially different second stage assignment indicated by vector $\boldsymbol{y}_i(\boldsymbol{\xi_i})$ and thus a different cost $\boldsymbol{q}_i^T \cdot \boldsymbol{y}_i(\boldsymbol{\xi_i})$, where $\boldsymbol{q}$ is a vector whose elements represent the power consumed by each sub-channel in the second stage. We assume that the power consumed by a sub-channel in the second stage is larger than the power consumed by

the same sub-channel in the first stage due to the extra power consumption involved in the process of setting up this new channel for operation, i.e., spectrum sensing and switching, processing power for CPU operation and coordination with the receiving end, etc. Furthermore, the scalar product $q_f \cdot y_f(\boldsymbol{\xi_i})$ represents the power consumed for the transmission of data through the ISM band, i.e., when all $\mathrm{h_d}$ sub-channels are allocated in the ISM band. In this cost factor $y_f$is a binary variable indicating whether the ISM band will be used or not and $q_f$is the corresponding total power consumption for transmission in the ISM band. We assume, that the use of an ISM band is costs power-wise more than the opportunistic spectrum band either due to the increased noise and interference expected in the ISM band as well as the necessary coordination process for spectrum migration.

Finally, we consider the fact that the number of sub-channels $\boldsymbol{x_i}$ that will be actually used for transmission is unknown before the second stage, whereby we assess the state of each PU within W_i. For z_i to account for the total power consumption in both stages, we must subtract any power cost that was accounted for in the first stage, through inner product $\boldsymbol{c}_i^T \cdot \boldsymbol{x}_i$, involving sub-channels that were not eventually unavailable for transmission. The quantity to be subtracted is represented by the inner product $\boldsymbol{b}_i^T(\boldsymbol{\xi_i}) \cdot \boldsymbol{w_i}$, where vector $\boldsymbol{w_i}$ is set to be identical $\boldsymbol{x}_i$ as indicated by Equation (29) and the elements of vector $\boldsymbol{b}_i^T(\boldsymbol{\xi_i})$ are defined according to:

$$b\left(\xi_i^j\right) = \begin{cases} 1, & if\ \sum_{j=0}^{n-1} \xi_i^j \cdot M < h_d \\ \left(1 - \xi_i^j\right), & \text{otherwise} \end{cases} \quad (34)$$

The first branch of Equation (34) indicates that if the total number of available sub-channels within spectrum region W_i is less than the demand in sub-channels, then the ISM band will be used and thus the power consumed by all sub-channels in the first stage must be subtracted from z_i as none of them was actually used for transmission. The second branch of Equation (34) indicates that when there exists an adequate number of sub-channels within W_i to accommodate the demand in sub-carriers, these channels will be used from the opportunistic access band since it is cheaper power-wise. Thus we will subtract the amount of power accounted in the first stage only for those sub-channels that were eventually unavailable.

Equation (30) states that from a specific channel the total number of sub-channels assigned in the first and second stage **cannot exceed its total amount of sub-channels $\boldsymbol{M}$**. Similarly, Equations (31) and (32) state that the amount of sub-channel assigned either in the first stage or in the second stage from a specific channel cannot exceed the number of its available sub-channels. Equation (33) states that y_f is a binary variable.

We note that although the presented two-stage stochastic IP has non-fixed recourse as indicated by the fact that the coefficients of vector $\boldsymbol{y}$ in the constraint presented in the Equation (28) depend on random vector $\boldsymbol{\xi_i}$ the convex hull of the feasible set of the problem is a convex polyhedron due to the discrete nature of random vector $\boldsymbol{\xi_i}$ (see [BL11] Chapter 3, Section 3.1b).

Finally, as mentioned in the beginning of this section, the solution to each of the $i \in I$ stochastic programs, where $I = \{1, 2, \dots, (N+1) - n\}$ will provide us with the optimal spectrum assignment within the corresponding region $\mathrm{W_i}$. Subsequently, we will choose the optimal region W^* among the ones that provide the spectrum assignment with the least power consumption, i.e, $z^* = \min_{i \in I}\{z_i\}$ and adopt their corresponding optimal spectrum assignments $\boldsymbol{x}^* = \boldsymbol{x}_m, \boldsymbol{y}^*(\boldsymbol{\xi}^*) = \boldsymbol{y}_m(\boldsymbol{\xi}_m)$ and $\boldsymbol{y_f}^*(\boldsymbol{\xi}^*) = \boldsymbol{y}_{\boldsymbol{f},m}(\boldsymbol{\xi}_m)$, where $m = \arg\min_{i \in I} z_i$.

2.4.5 Deterministic Equivalent of the Stochastic IP

To solve the stochastic IP problem of the previous section we must convert it to its deterministic equivalent form. Let ω be an index variable for the set of all possible states of spectrum band $\mathrm{W_i}$ as determined by the binary vector $\boldsymbol{\xi_i} = (\xi_i^0, \xi_i^1, \dots, \xi_i^{n-1})$. Consequently ω will take values in the set $\Omega =$

$\{1, \dots, 2^n\}$. Furthermore, for each possible scenario ω we create a new set of variables $x_i^{\omega j}, y_i^{\omega j}, w_i^{\omega j}$, $\omega \in \Omega, j \in \{0, 1, \dots, n-1\}$ and the deterministic form of the stochastic IP will be:

$$\min\ z_i = \sum_{j=0}^{n-1} c_i^j \cdot x_i^j + \sum_{\omega=1}^{2^n} P_\omega \left(\sum_{\omega=1}^{2^n} q_i^j \cdot y_i^{\omega j} + q_f \cdot y_f^\omega - \sum_{\omega=1}^{2^n} b_i^j \cdot w_i^{\omega j}\right), \tag{35}$$

s.t.
$$\sum_{j=0}^{n-1} \xi_i^{\omega j} \cdot x_i^{\omega j} + \sum_{j=0}^{n-1} \xi_i^{\omega j} \cdot y_i^{\omega j} + h_d \cdot y_f^\omega \geq h_d, \quad \text{for all } \omega \in \Omega \tag{36}$$

$$x_i^{\omega j} = w_i^{\omega j}, \quad \text{for all } \omega \in \Omega \text{ and } j \in \{0, 1, \dots, n-1\}, \tag{37}$$

$$x_i^{\omega j} + y_i^{\omega j} \leq M, \quad \text{for all } \omega \in \Omega \text{ and } j \in \{0, 1, \dots, n-1\}, \tag{38}$$

$$x_i^{\omega j} \leq M, \quad \text{for all } \omega \in \Omega \text{ and } j \in \{0, 1, \dots, n-1\}, \tag{39}$$

$$y_i^{\omega j} \leq M, \quad \text{for all } \omega \in \Omega \text{ and } j \in \{0, 1, \dots, n-1\}, \tag{40}$$

$$y_f^\omega \in \{0, 1\} \text{ and } x_i^{\omega j}, y_i^{\omega j}, w_i^{\omega j} \in \mathbb{N}, \tag{41}$$

To tackle the problem at hand we relax the integer constraints on $x_i^{\omega j}, y_i^{\omega j}$ and $w_i^{\omega j}$ so that a stochastic two-stage Mixed Integer Linear Program (MILP) is formulated from the initial integer program. The resulting problem is given as input to a MILP solver that implements various MILP solving techniques such as linear programming preprocessing, relaxation of integer constraints, MILP preprocessing and generation of Gomory cuts [GC07]. Equations (35) to (41) indicate that the number of constraints of the MILP is exponentially increasing with the number of channels, or equivalently, the number of PUs that comprise each spectrum region W_i. The number of sub-channels on the other hand do not have such an effect on the complexity of the MILP. On the contrary, the larger the number of sub-channels per channels the smaller the error introduced by the relaxation of integer constraints for $x_i^{\omega j}, y_i^{\omega j}$ and $w_i^{\omega j}$. For the type of applications we consider in RERUM and given the power constrained nature of RD we expect the number of channels comprising candidate spectrum regions W_i to be rather limited.

2.4.6 A greedy approach to the two-stage dynamic spectrum assignment

A spectrum assignment strategy proposed by many authors in the literature [SDR14,JPCV12,BF14, YCZ08, HPM08] is to always assign to SUs those channels for which PU has the highest probability to be idle. Obviously, the rationale behind this proposition is to assign SUs the channels that are idle most of the time and thus result in fewer spectrum switching actions. In this section, we formulate a stochastic two-stage, dynamic spectrum assignment algorithm that greedily assigns channels based solely on the probability to find the PU in the idle state. Next, we will compare the efficiency of the greedy and MILP spectrum assignment programs in terms of power consumption.

The greedy algorithm will visit sequentially each candidate region W_i, $i \in \{1, 2, \dots, (N+1) - n\}$ within B and will assign an appropriate number of sub-channels for the first $(\boldsymbol{x_i})$ and second stage $(\boldsymbol{y_i}, y_f)$ so as to minimize the total power consumption for data transmission, z_i, while satisfying the QoS requirements. Finally, the algorithm will choose as optimal the region W_g, where $g = \arg\min_i(z_i)$,i.e., the region whose spectrum assignment program results in the lowest power consumption, denoted by z_g.

In the first stage, for any region W_i the algorithm will sort all available channels in descending order of PU OFF state probability. Subsequently, it will sequentially assign sub-channels from the sorted list of channels until the demand is satisfied. Vector $\boldsymbol{x_i}$ will be properly updated to represent the first stage spectrum assignment for the initial, unsorted list of channels.

In the second stage, the state of the spectrum is assessed by sensing in order to determine the elements of vector $\boldsymbol{\xi_i} = (\xi_i^0, \xi_i^1, \dots, \xi_i^{n-1})$. As a first step, applied in each possible realization of vector $\boldsymbol{\xi_i}$, the greedy algorithm will identify the channels assigned to the RD in the first stage that have now become unavailable and will properly increment unsatisfied demand. Furthermore, it will decrement power consumption, accounted for in the first stage that corresponds to the now unavailable channels by properly updating vector $\boldsymbol{w_i}$. We note here that vector $\boldsymbol{w_i}$ is not used in exactly the same way as in the two stage stochastic IP. Vector $\boldsymbol{w_i}$ now explicitly accounts for the channels assigned in the first stage that have now become unavailable or $\boldsymbol{w_i}$ will be equal to $\boldsymbol{x_i}$ in case the ISM band is eventually selected for transmission, indicating that none of the channels in $\boldsymbol{x_i}$ will be used. Subsequently, in case an adequate number of sub-channels is available in the opportunistic access band, the greedy algorithm will assign the SU enough sub-channels to satisfy the unsatisfied demand. The sub-channels will be selected from the channel list that is sorted in descending order of PU OFF state probability and the corresponding elements of vector $\boldsymbol{y_i}$ will be updated. If, on the other hand, the number of available sub-channels is not adequate to cover demand, all unsatisfied demand in sub-channels will be satisfied from the ISM band and $\boldsymbol{w_i}$ will be updated to reflect that no channel from the first stage assignment will be used.

Finally, for all $i \in \{1, 2, \dots, (N+1) - n\}$ transmission power cost for the stochastic, two-stage greedy algorithm will be calculated using:

$$z_i = \sum_{j=0}^{n-1} c_i^j \cdot x_i^j + \sum_{\omega=1}^{2^n} P_\omega \cdot \left(\sum_{\omega=1}^{2^n} q_i^j \cdot y_i^{\omega j} + q_f \cdot y_f^\omega - \sum_{\omega=1}^{2^n} b_i^j \cdot w_i^{\omega j} \right) \quad (42)$$

and region W_m, where $m = \arg\min_i(z_i)$, along with its spectrum assignment program will be selected as optimal.

2.4.7 Dynamic spectrum assignment for average throughput guarantees in a CSMA/CA RD network

In this section we consider a single hop network of **multiple RDs sharing the same spectrum bands**. The shared nature of the wireless medium prohibits the provision of strict delay guarantees as described in the previous sections. However, we can still guarantee the average throughput of each flow in the network by properly setting the transmission rate of the wireless link. The transmission rate of the link can be controlled by the number of sub-channels assigned to the RDs in the two-stage spectrum assignment process. To this end we utilize the analytical framework presented recently by Kleinrock et al. in [LK13] for the throughput of wireless CSMA/CA networks. The authors in [LK13] present a framework that models throughput accurately based on a small set of assumptions and, furthermore, surpasses typical modelling inefficiencies in CSMA/CA wireless networks due to the strong interdependence among the state of transmitters across the network, the distributed nature of the CSMA/CA protocol, the dynamics of unsaturated traffic sources and the dependence of downstream traffic on upstream traffic.

More specifically, the authors assume that nodes have a unique transmission queue for each flow which acts as an individual collocated transmitter k with its own backoff counter, that no hidden terminals exist, that carrier sensing is instantaneous and that the transmission rate r_k of the k-th transmitter is fixed. No restrictions are posed on the distribution of packet lengths L_k and their generation process, the distribution of backoff times B_k or the saturation levels of the flows. Especially the latter assumption is of great significance to a RD network since typically most sensors transmit

data rather infrequently. Furthermore, the fact that traffic in RD networks is typically unsaturated greatly reduces the collision probability in many scenarios and thus the inaccuracy introduced in the model by the assumptions of no hidden terminals and instantaneous spectrum sensing is diminished. The interested reader may refer to [LK13] for a detailed presentation of the model.

According to [LK13] the average throughput G_k of the k-th transmitter, i.e. of the k-th flow, is given by:

$$G_k = \left(\frac{\gamma_k}{1+\gamma_1+\gamma_2+\cdots+\gamma_n}\right) r_k p_k \tag{43}$$

Where, r_k is the fixed transmission rate of the wireless link used by the k-th transmitter, p_k is the probability of correct transmission of a packet by the k-th transmitter and $\gamma_k = \rho_k \theta_{k,}$ where $\rho_k \in [0,1]$ is a parameter of the system indicating the saturation level of flow k, i.e., when $\rho_k = 1$ the flow has always packets to transmit, while when $\rho_k = 0$ the flow has never packets to transmit, finally $\theta_k = \frac{\mathbb{E}[T_k]}{\mathbb{E}[B_k]}$ is the ratio of the expected packet transmission time to the expected backoff duration. We note that packets have random length L and are transmitted with fixed rate r_k, thus $\mathbb{E}[T_k] = \frac{\mathbb{E}[L_k]}{r_k}$.

With simple algebraic manipulations Equation (43) can be written as:

$$(1+\gamma_1+\gamma_2+\cdots+\gamma_n) = \frac{\gamma_k r_k p_k}{G_k}, \forall k \in \{1,2,\dots,n\}$$

and by defining $\phi_k \triangleq \frac{\rho_k \mathbb{E}[L_k]}{\mathbb{E}[B_k]}$ we get $\gamma_k = \rho_k \frac{\mathbb{E}[L_k]}{r_k \mathbb{E}[B_k]} = \frac{\phi_k}{r_k}$ and thus the equation above can be written as:

$$\frac{\phi_1}{r_1} + \frac{\phi_2}{r_2} + \cdots + \frac{\phi_n}{r_n} = \frac{\phi_k p_k}{G_k} - 1, \forall k \in \{1,2,\dots,n\}$$

Finally, by letting $u_k = \frac{1}{r_k}$ we form the following linear system of n equations with n unknowns:

$$\phi_1 u_1 + \phi_2 u_2 + \cdots + \phi_n u_n = \frac{\phi_k p_k}{G_k} - 1, \forall k \in \{1,2,\dots,n\} \tag{44}$$

Given the average throughput demand for each flow, G_k, the indicator parameter of its saturation level ρ_k and the expected backoff and packet length values $\mathbb{E}[B_k]$ and $\mathbb{E}[L_k]$, respectively, we can calculate the demand in transmission rate r_k for all flows $k \in \{1,2,\dots,n\}$. Finally, we may assign each RD an adequate number of sub-channels so as to satisfy the demand of its flows in an optimal way by properly modifying the stochastic two-stage MILP formulation for the dynamic spectrum assignment problem so that it may handle random demand.

Next, we focus on the simple scenario whereby all RDs serve a single flow of the same type. Thus, by setting $\phi_1 = \phi_2 = \cdots = \phi_n$, $p_1 = p_2 = \cdots = p_n, G_1 = G_2 = \cdots = G_n$ and $u_1 = u_2 = \cdots = u_n$ in Equation (44) we get the following demand in transmission rate for all RDs:

$$r = \frac{n\rho\mathbb{E}[L]G}{\rho\mathbb{E}[L]p - G\mathbb{E}[B]} \tag{45}$$

The demand in transmission rate r can be expressed in terms of a discrete number of sub-channels by letting $h_d = \lceil r \rceil$.

2.4.8 Evaluation

We begin the evaluation of the proposed dynamic spectrum assignment scheme by means of an example that will clarify the application of the proposed scheme as well as exhibit its lightweight character. We set the number N of channels of the opportunistic access band B equal to 10, the number (M) of sub-channels per channel equal to 10, the number of channels accessible by the RD (n) equal to 3 and the demand in sub-channels (h_d) equal to 14. These numbers are selected arbitrarily for minimizing the computational cost; the results are not depending on these numbers. The first stage power costs for each channel, i.e., the elements of vector $\boldsymbol{c}$, were selected randomly from the set of available transmission powers for the **mica2 mote** as described in [SHC++04] and presented in Table 2 and are equal to

$$\boldsymbol{c} = (6.310, 0.032, 6.310, 0.032, 10.000, 0.158, 0.013, 0.032, 2.511, 0.316).$$

The second stage power costs for the opportunistic spectrum access bands are increased, relatively to the first stage costs, by a factor of 0.1 mW, i.e., $\boldsymbol{q} = \boldsymbol{c} + 0.1\text{mW}$. We assume that the extra power consumption of 0.1mW times the expected transmission duration over the sub-channel is equal to the energy consumed on cognitive functions such as spectrum sensing, switching and coordination with the receiving end for the setup of a single sub-channel in the second stage. Finally, the second stage power cost for each sub-channel in the ISM band is equal to 12mW, i.e., sub-channels in the ISM band are the most expensive in terms of power consumption. The vectors of PU ON and OFF probabilities for each channel were selected randomly to be

$\boldsymbol{\pi_1}$ = (0.7513, 0.2551, 0.5060, 0.6991, 0.8909, 0.9593, 0.5472, 0.1386, 0.1493, 0.2575)

and

$\boldsymbol{\pi_0}$ = (0.2487, 0.7449, 0.4940, 0.3009, 0.1091, 0.0407, 0.4528, 0.8614, 0.8507, 0.7425) respectively.

Table 1: The two-stage spectrum assignment program produced by the two-stage stochastic IP for an example scenario. First stage spectrum assignment is presented above the double horizontal line, while second stage spectrum assignments are presented below the double horizontal line.

ω	Channel availability			Sub-channels assigned from channel				Total
	ξ_8^0 (channel 8)	ξ_8^1 (channel 9)	ξ_8^2 (channel 10)	Channel 8	Channel 9	Channel 10	ISM band	
-	Unknown	Unknown	Unknown	10	0	4	0	14
1	0	0	0	0	0	0	1	14
2	0	0	1	0	0	0	1	14
3	0	1	0	0	0	0	1	14
4	0	1	1	0	4	6	0	14
5	1	0	0	0	0	0	1	14
6	1	0	1	0	0	0	0	14
7	1	1	0	0	4	0	0	14
8	1	1	1	0	0	0	0	14

The output of the two-stage stochastic IP scheme for the example scenario described above is presented in Table 1. The optimal first stage spectrum assignment $\boldsymbol{x}^*$ is presented above the double horizontal line while the optimal second stage spectrum assignment $(\boldsymbol{y}^*(\boldsymbol{\xi}^*), y_f^*(\boldsymbol{\xi}^*))$ below it. The optimal triad of channels W^*, proposed by the two stage stochastic IP, is comprised of channels 8, 9 and 10 and corresponds to the 8-th possible triad of channels, i.e., $m = \arg\min_{i \in I} z_i = 8$. Within W^* sub-channels are assigned only from channels 8 and 10. All sub-channels of channel 8 were assigned for transmission in the first stage since it exhibits the highest probability for the PU to be in the OFF state among all channels and at the same time requires a very low transmission power. On the other hand, although channel 9 has a higher probability to be found in the OFF state ($\pi_0^9 = 0.8507$) compared to that of channel 10 ($\pi_0^{10} = 0.7425$) 4 sub-channels are bound from the latter one because it requires a lot less transmission power.

During the first stage the states of the PUs are unknown and thus the availability of their channels, denoted by the values of vector $\boldsymbol{\xi}_8$, is marked as unknown. However, once the RD enters the second stage, the state of the PUs is assessed and one of the 8 possible outcomes presented in Table 1 will be realized. For each such outcome the two stage program assigns the necessary number of sub-channels so that z is minimized and the QoS criterion is satisfied. In Table 1 we get the total demand of 14 sub-channels either by adding the number of sub-channels assigned in the first stage, given that their corresponding channel is available, i.e., $\xi_i^j = 1$, with the number of sub-channels assigned in the second stage or utilize the ISM band, i.e., set $y_f(\omega) = 1$. So, for example, when $\omega = 1$ none of the sub-channels assigned in the first stage is available and thus we get the necessary capacity from the ISM band, i.e., $y_f(1) = 1$. The same is true for $\omega \in \{2, 3, 5\}$ whereby at most one channel is available. In all these cases the available number of sub-channels within the opportunistic access band is not adequate to cover demand ($h_d = 20$). Another interesting case appears when $\omega = 4$. In that case channels 9 and 10 are available but channel 8 that has the best operational characteristics is not, thus the optimal program will assign 6 more sub-channels from the more favourable channel 10, additionally to the 4 sub-channels that were assigned during the first stage and then proceed to assign the remaining channels from the less favourable channel 9.

Next, we evaluate the efficiency of the two-stage stochastic IP dynamic spectrum assignment scheme over the same opportunistic access spectrum band B that is comprised of 10 channels. This time however, we vary the number of sub-channels (M) per channel, the number of channels accessible by the RD (n) and the demand for sub-channels (h_d) over a predetermined range so as to assess the efficiency of the proposed scheme in a wide range of scenarios. For each combination of M, n and h_d we run 1000 experiments and in each one of them we select uniformly the first stage power costs for each channel from the set of available transmission powers presented in Table 2.

Table 2: Power model for Mica2 mote.

Transmission Power	**dBm**	-20	-19	-15	-8	-5	0	4	6	8	10
	mW	0.01	0.013	0.032	0.158	0.316	1	2.511	3.98	6.31	10

Similarly, for each experiment the probability for each PU to be in the ON state (π_1) is selected uniformly in the range $[0,1]$ and the probability of the PU's OFF state is derived by $\pi_0 = 1 - \pi_1$. The second stage power costs for the opportunistic spectrum access bands are again increased, relatively to the first stage costs, by a factor of 0.1 mW, i.e., $\boldsymbol{q} = \boldsymbol{c} + 0.1\text{mW}$. Finally, the second stage power cost for each sub-channel in the ISM band is again fixed and equal to 12mW.

Figure 25, Figure 26, Figure 27 and Figure 28 present the average of the optimal power consumption z^* for the 1000 experiment results when the two-stage stochastic IP is used and M equals 10, 20, 30 and 40 sub-channels per channel respectively. In Figure 25 we see that as the number n of accessible channels within the opportunistic access band increases, the expected power consumption decreases for all levels of demand h_d. Of special interest are the cases where n is inadequate to cover demand. For example, in Figure 25, when demand h_d is set to 40 sub-channels then it cannot be satisfied when

n is equal to 1, 2 or 3 channels, i.e., 10, 20 and 30 sub-channels respectively. In such a case the RD can only get the necessary number of sub-channels within the ISM band. The power consumption for all these scenarios is deterministic and it is analogous to the product of power consumption per sub-channel within the ISM band and demand itself, i.e., occupying 40 sub-channels within the ISM band requires more power than occupying 30 or 20 sub-channels. At some point, n becomes large enough to satisfy demand within the opportunistic access band and the expected power consumption exhibits a significant decrease. From that point on, further increasing n results in an ever increasing number of combinations of the states of PUs of the n channels, as described by the random vector $\boldsymbol{\xi}$, where demand can be satisfied within the opportunistic access band either in the first or in the second stage. The effect of this is that expected power consumption further decreases since we avoid the ISM band.

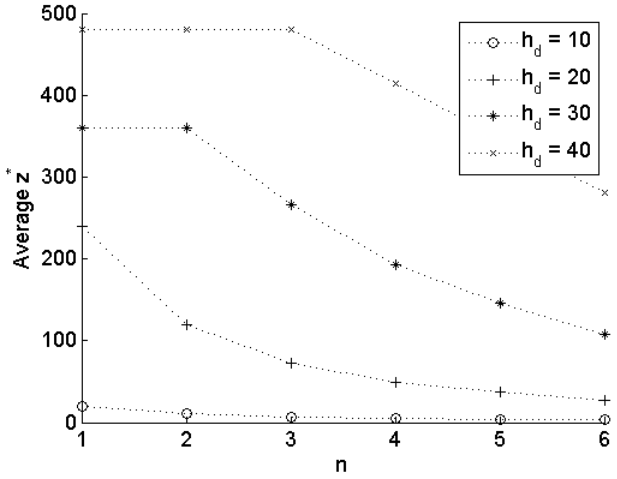

Figure 25. Average of expected power consumption z^* for the two-stage stochastic IP as a function of n when $M = 10$.

Furthermore, Figure 26, Figure 27 and Figure 28 depict the effect of increasing the number of sub-channels M per channel on expected power consumption. We assume that the transmission rate r remains the same for all sub-channels as M increases, i.e., increasing M is equivalent to increasing the bandwidth of each PU channel and thus its number of available sub-channels. As can be seen from the Figure 26, Figure 27 and Figure 28, power consumption drastically reduces with n as M increases. This is due to the fact that now, even for small values of n, the demand in sub-channels can be satisfied with spectrum assignments within the opportunistic access band either in the first or the second stage without resorting to the ISM band. Thus, for example, in Figure 26, for $n = 1$ demands h_d of 30 and 40 sub-channels cannot be satisfied within the opportunistic access band resulting in an excessively large power consumption, in Figure 27 where M increases to 30, for $n = 1$ only demand h_d of 40 sub-channels cannot be satisfied within the opportunistic access band, and finally in Figure 28, where M increases to 40, all classes of demand are satisfied from the opportunistic access band even when $n = 1$, resulting in a significant reduction in expected power consumption.

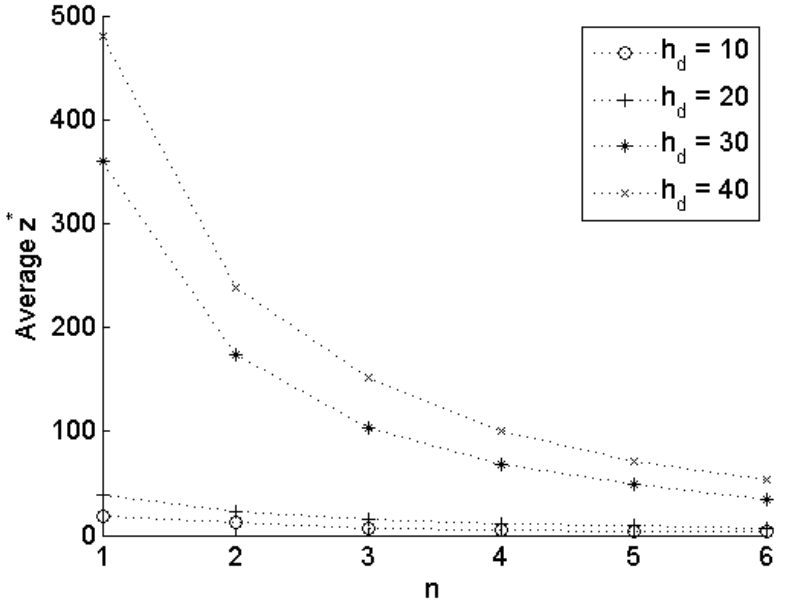

Figure 26. Average of expected power consumption z^* for the two-stage stochastic IP as a function of n when $M = 20$.

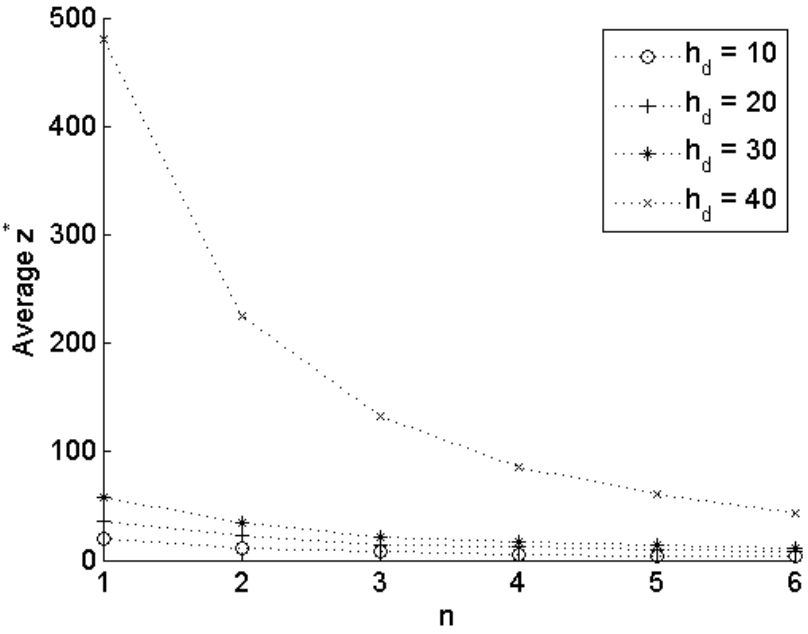

Figure 27. Average of expected power consumption z^* for the two-stage stochastic IP as a function of n when $M = 30$.

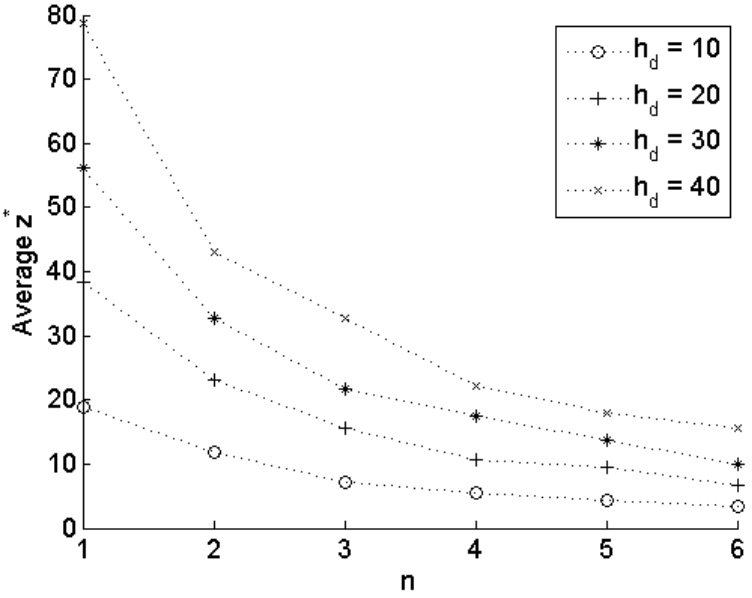

Figure 28. Average of expected power consumption z^* for the two-stage stochastic IP as a function of n when $M = 40$.

Next, we compare the efficiency, in terms of power consumption, of the two-stage stochastic IP to that of the two-stage greedy algorithm using numerical simulations. Figure 29, Figure 30, Figure 31 and Figure 32 present the percentage increase in expected power consumption when we use the two-stage greedy spectrum assignment algorithm instead of the two-stage stochastic IP for spectrum assignment. As before Figure 29, Figure 30, Figure 31 and Figure 32 correspond to M being equal to 10, 20, 30 and 40 sub-channels per channel respectively.

Starting with Figure 29, we note that when demand is low ($h_d = 10$) use of the two-stage greedy algorithm results in a significant increase of power consumption especially for large values of n. For larger values of demand however ($h_d \in \{20, 30, 40\}$), the phenomenon becomes much less intense. To explain this result we remind that when M equals 10 there are not enough sub-channels to cover demand within the opportunistic access band for low values of n and thus, regardless of the spectrum assignment algorithm used, the RD will have to resort to the ISM band to get the required number of sub-channels. As was explained previously for Figure 25 the cost in such cases is deterministic, independent of the spectrum assignment algorithm used and completely determined by the demand and the power cost per sub-channel in the ISM band. Figure 29, Figure 30, Figure 31 and Figure 32 depict that as the number M of sub-channels per channel increases demand can be satisfied within the opportunistic access band and the greedy algorithm will result in a significant increase of power consumption for all levels of demand since it disregards power consumption when selecting channels within the opportunistic access band.

Finally, a special case that is of particular interest and appears in all figures and for all levels of demand is when $n = 1$. In that case both spectrum assignment algorithms perform the same even when demand can be satisfied in the first stage by a single channel. This is so because the RD will have to bind all its sub-channels from a single channel. There are two possible events for such a case, either the PU will be in the OFF state and a channel within the opportunistic access band B will be used or the PU will be in the ON state and the ISM band will be used. We remind that the ISM band costs more than any channel within the opportunistic access band in terms of power consumption. Thus, the optimal strategy for both SA algorithms is to select that single channel within opportunistic access band B that has the largest probability to be found in the OFF state since the main objective is to avoid the expensive ISM band. However, if there are more than one channels, with the same probability for

the PU to be found in the OFF state but different power consumption per channel then the two-state stochastic IP will find the optimal solution unlike the two-stage greedy algorithm.

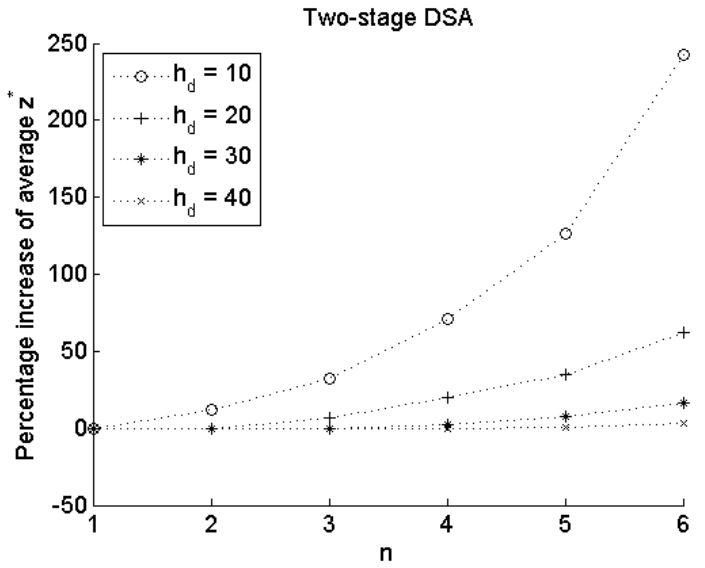

Figure 29. Percentage increase on the average of expected power consumption z^* when the greedy algorithm is used instead of the two-stage stochastic IP for M = 10.

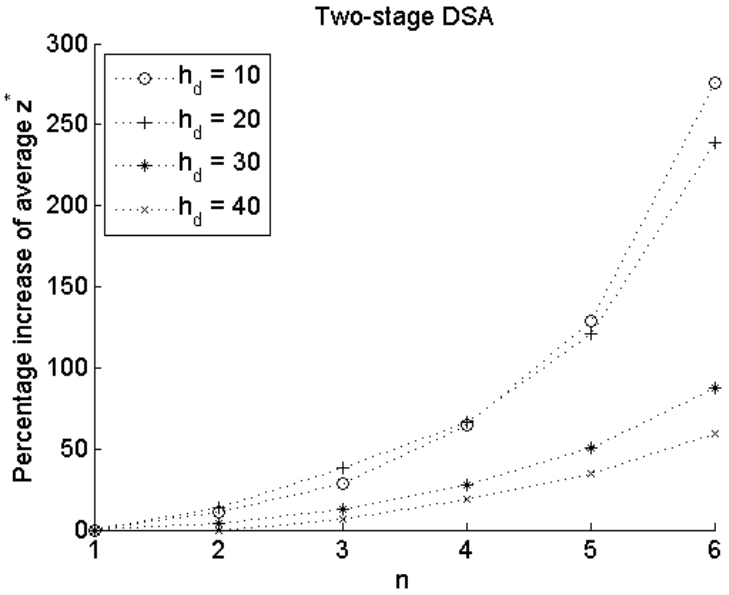

Figure 30. Percentage increase on the average of expected power consumption z^* when the greedy algorithm is used instead of the two-stage stochastic IP for M = 20.

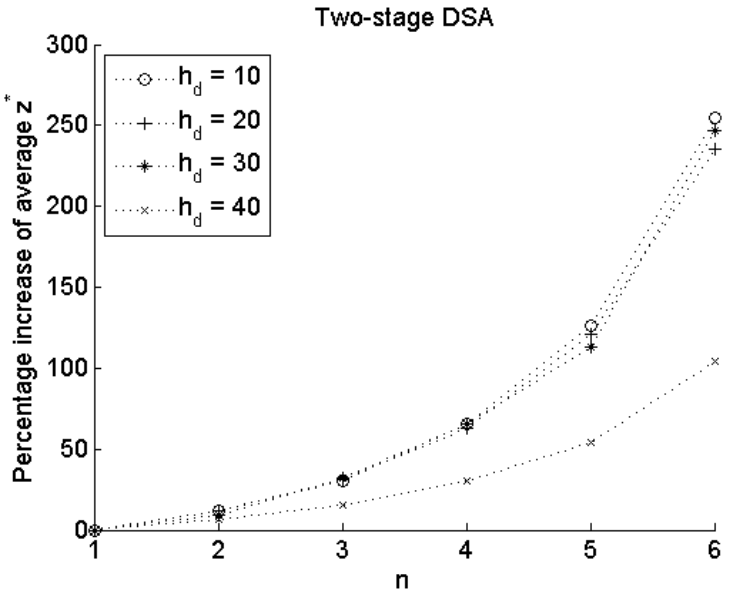

Figure 31. Percentage increase on the average of expected power consumption z^* when the greedy algorithm is used instead of the two-stage stochastic IP for M = 30.

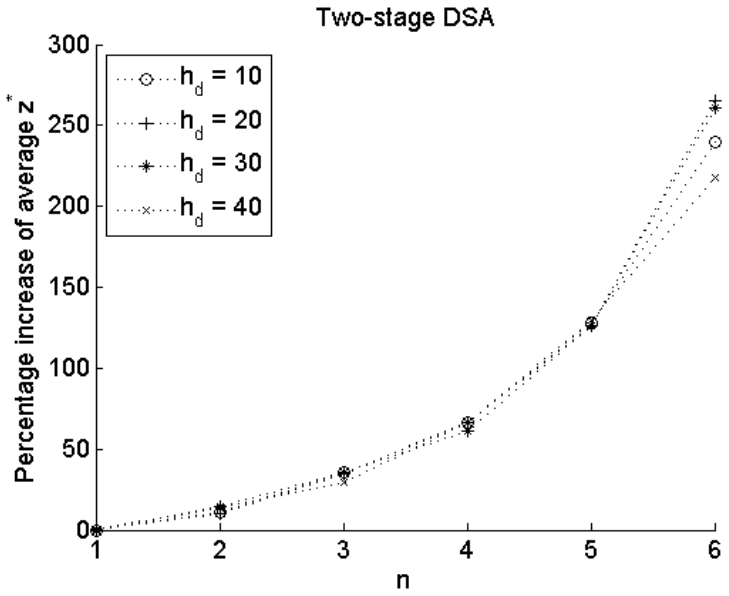

Figure 32. Percentage increase on the average of expected power consumption z^* when the greedy algorithm is used instead of the two-stage stochastic IP for M = 40.

2.5 A prototype implementation of an SDR-based Gateway [5]

This section presents an implementation of an SDR-based RERUM Gateway that can be utilized for the implementation of the Cogitive Radio techniques presented above. The goal of this prototype is to show that using simple hardware it is possible to create an advanced gateway, able to be fully reconfigurable by software and to support the interconnectivity of various technologies without the need of multiple interfaces. One of the key objectives of RERUM is the efficient interconnection of RDs of heterogeneous networking technologies The majority of the current IoT deployments include devices that use either the IEEE 802.15.4 or the IEEE 802.11 standard, both sharing the same 2.4 GHz unlicensed band. Aiming to advance the existing approaches where gateways provide connectivity to devices of different standards using discrete radio modules for each technology, we describe here the adaptation of SDR technology for minimization of the number of radio components (i.e. transceivers, antennas) needed to support a heterogeneous environment, and for easy adaptation to any future standard or modification.

2.5.1 Related Work

GNU Radio [GNU] is the most popular open-source software development and simulation platform for SDR. The implementation of the IEEE 802.15.4 on the GNU Radio platform is presented in [S06]. It provides the Physical Layer (PHY) layer of the standard and a basic, network socket-based interface in order to send/receive data. Further than the standard PHY layer implementation no connectivity with higher layers, like the 6LoWPAN, is provided. Comparably, there are some works that partially implement the IEEE 802.11 standard on the GNU Radio [WW11,B++13], but the only fully functional transceiver has been presented by K. Tan et al. in [T++11] on the Sora SDR platform. The main disadvantages of the Sora-based implementation are the restricted compatibility to their own SDR hardware and the Windows Operating System (OS). On the other hand, the Sora 802.11 implementation appears as a typical network interface in the Windows OS, but there is no SDR implementation in the Linux OS of any networking standard that is available as normal network interface. Finally, to the best of our knowledge, there is no single-device concurrent multi-standard SDR platform presented to date. In this section the following contributions are presented: We introduce a fully software implementation of the PHY of IEEE 802.11a/g transceiver assisted by a fully compliant SoftMAC layer in the GNU Radio platform in the Linux OS. The in-house IEEE 802.11 implementation, along with the existing implementation of the IEEE 802.15.4, which is adopted and adjusted to use the available higher layer 6LoWPAN module, are used by the Linux OS as normal network interfaces, maximizing the usability and minimizing the adaptation effort. An additional technique is proposed for the concurrent transmission of multiple standards and channels on adjacent frequencies on SDR devices. A complete platform which is based on off-the-shelf products and open source software is demonstrated; a single SDR device plus a host computer solution that is able to encode/decode in real time and concurrently transmit/receive multiple channels and standards, in a large bandwidth span

2.5.2 Physical layer implementations

2.5.2.1 IEEE 802.11a/g PHY

We implemented completely from scratch our own IEEE 802.11a/g PHY software implementation based on the GNU Radio software stack. It supports all the data rates of the standard, together with the half- and quarter-clocked operation modes utilizing 10 and 5 MHz of bandwidth respectively. In order to fulfil the demanding computational and timing requirements of the standard we applied several software optimization techniques. One of those techniques was the parallelism exploitation. Due to the overhead for synchronizing threads, which in several tasks was heavier than the actual

[5] The work presented in this section was accepted to be published in the VTC 2015 Spring conference [SMP15].

computation, some processing tasks were grouped together into a single thread. In addition, our implementation heavily uses lookup tables when applicable. The majority of these tables have small memory footprint, easily fitting inside the caches of modern CPUs, providing a significant performance boost. Significant performance gains achieved using the Single Instruction, Multiple Data (SIMD) instructions. These instructions can provide not only increased processing throughput, but also help to eliminate conditional branches, which is quite essential for a real-time implementation of the Viterbi decoding algorithm [VIT67]. With these techniques applied our software can operate real-time, leaving to the host enough CPU resources for other tasks as well.

2.5.2.2 IEEE 802.15.4 PHY

For the IEEE 802.15.4 PHY we used the publicly available code of [S06] for GNU Radio. While this implementation provides some basic communication with off-the-shelf IEEE 802.15.4 devices, we extended its capabilities in order to support a basic Carrier Sense Multiple Access (CSMA) mechanism, Acknowledgments (ACKs) and re-transmissions. Furthermore, we applied some software optimization techniques due to the unexpected, based on the low complexity of the standard, high CPU resources utilization.

2.5.3 MAC layer integration

The Media Access Control (MAC) layer of both IEEE 802.11 and IEEE 802.15.4 is responsible for managing a great part of the corresponding standard. It implements among others mechanisms for medium access, ACKs and re-transmission, Station (STA) management, and security. Over the last years, even more vendors have decided to follow a SoftMAC approach for their wireless devices allowing the MAC Sublayer Management Entity (MLME) to be managed by software. In this approach, hardware becomes simpler and with lower cost. In addition, SoftMAC allows a unified way of configuration, providing the ability to the OS to handle wireless devices **robustly and efficiently**. In this section we propose an architecture that allows the SDR PHY implementations to co-operate with the IEEE 802.11 and IEEE 802.15.4 compliant SoftMAC layer implementations.

2.5.3.1 IEEE 802.11 MAC layer

The Linux SoftMAC solution for the IEEE 802.11 is the mac80211 [LWi] kernel module, providing an API that can be used from the wireless hardware driver in order to properly transmit and receive data. For dealing with the user defined settings a second helper kernel module is used, the cfg80211. This module acts as a middleware between the userspace and the mac80211 driver, allowing the user to configure several parameters of the IEEE 802.11a/g standard. Such parameters may be the supported data rates, transmission power, operating mode (managed, ad-hoc, mesh, monitor) and other supplemental ones. Each SoftMAC device derives data and configuration using the API of mac80211. In addition, hardware devices report back data to the mac80211 driver, which is responsible to propagate them to the upper layers. The architecture is summarized in Figure 33a. Our approach uses the mac80211 driver to provide a fully standard compliant MAC layer to the PHY software implementation, by exploiting the SoftMAC capabilities. A great feature of this design is that the OS is unaware of fact that the PHY actually runs on software, hence it identifies the wireless interface as a normal one, with all the capabilities of parameterization enabled.

2.5.3.2 IEEE 802.15.4 MAC layer

Similarly, Linux systems can take advantage of the corresponding SoftMAC solution for the IEEE 802.15.4 based devices. We used the 6lowpan [6LOW] kernel module, in order to create an arbitrary number of 6LoWPAN compliant wireless interfaces.

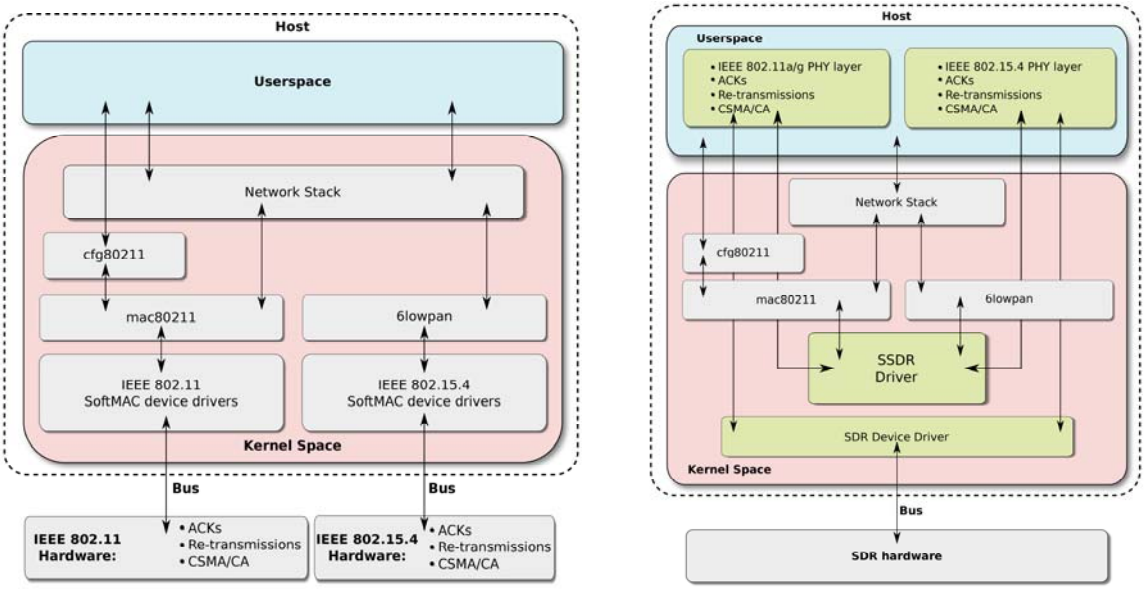

(a) *(b)*

Figure 33. (a) Conventional SoftMAC devices (b) Proposed approach for SDRs.

2.5.3.3 The SSDR Driver

The software implementations of both the IEEE 80211 and 802.15.4 PHY encode a Physical Layer Convergence Procedure (PLCP) Service Data Unit (PSDU) derived from the MAC layer and transmits it over the air. During reception the PHY propagates the decoded PSDU towards the MAC in the reverse manner. Running on userspace, the PHY software should be able to retrieve and send the PSDU data to the mac80211 or 6lowpan driver. For this reason we developed the Super Software Define Radio (SSDR) driver, a special kernel module that its main purpose is to allow the data exchange between the userspace PHY implementations and the corresponding kernel modules. The proposed architecture is illustrated in Figure 33b.

An arbitrary number of virtual IEEE 802.11 or 802.15.4 wireless network interfaces may be created by loading the SSDR driver, along with a character device for each one. Writing to the character device allows data to be delivered to the mac80211 or 6lowpan driver for the corresponding network interface. Respectively, PSDUs are retrieved by reading from the appropriate character device. This is a very convenient and efficient solution, because the PHY reads and writes PSDU data as it would read them from a regular file. To simplify even further the data exchange between the SSDR and the PHY, the following conventions were made. Data exchanges between the SSDR and the PHY software are equal in size. The size of the fixed size data chunk can be derived based on the Maximum Transmission Unit (MTU) of the interface that is always less than the maximum allowed PSDU. While reading or writing a fixed size data chunk is very convenient, optimally the actual length of the PSDU should be known. This information among with additional information required by the PHY implementations of IEEE 802.11a/g & IEEE 802.15.4 and the corresponding kernel modules, is encapsulated in the header of each data chunk.

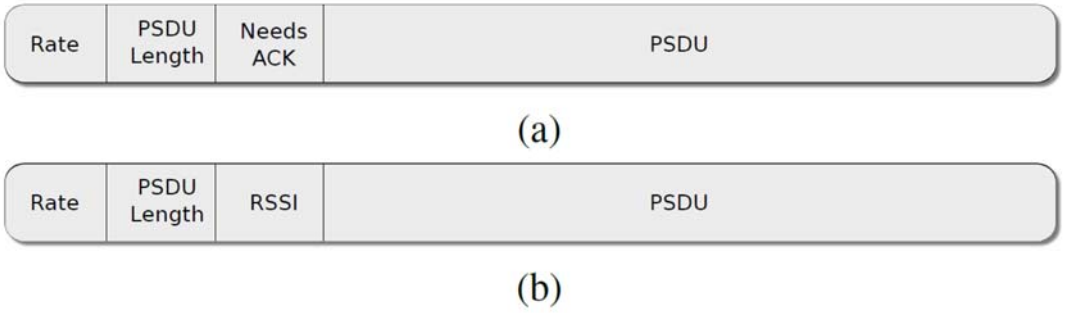

Figure 34. Structure of the PSDU data: (a) as retrieved by the SSDR driver and (b) form the PHY to the SSDR driver.

The formats of each fixed size data chunk during transmit and receive operations are presented at Figure 34a and Figure 34b respectively. The data rate of the frame and the length of the PSDU should be known to both sides. The data rate is used by the PHY in order to use the appropriate modulation scheme, whereas the mac80211 uses this information to feed its rate control algorithm. The 'Needs ACK' field is indicating that the current PSDU should be acknowledged by the other side. When the PHY transmitter handles a frame that has to be acknowledged, it should enable the re-transmission mechanism for this specific frame. In addition, during a PSDU reception, the PHY calculates and reports back to the MAC the Received Signal Strength Indication (RSSI) of the frame placing its value to the corresponding field of the SSDR header. The mac80211 makes use of this in order to report back to userspace applications the signal strength. However, as noted previously, even in the case of SoftMAC, the time critical operations are still implemented in hardware. These operations are completely ignored by the mac80211 and 6lowpan drivers, because the network devices are expected to implement their functionality. Thus in our case these operations should be implemented in the PHY software in order to achieve minimum response times. These operations fall under three categories: (1) CSMA/CA, (2) ACKs and (3) re-transmissions.

2.5.3.4 Time critical MAC operations

Software MAC implementations face large latency values due to the processing, buffering and data transfer overheads between software and hardware. This latency can reach the order of microseconds [SSS07]. This magnitude of latency greatly degrades the performance of the MAC protocols, as they require to be able to identify changes on the wireless medium or transmit data as soon as possible. The existing software MAC approaches can be classified into two main categories depending on their design space: the fully software [T++11,APG10] and the hybrid software-hardware approaches [N++09, F++14,B++14]. While the hybrid solutions are preferred due to the predictable and small delays of the Field Programmable Gate Array (FPGA), which is used by the SDR hardware, these approaches are device specific and reduce the flexibility and portability. The great throughput and sub-microsecond latency of the PCIe bus satisfy the requirements for SDR implementations of various wireless telecommunication standards. Therefore, for a competitive system performance where the SDR device should provide the minimum possible latency, PCIe interconnect is the best solution.

1) CSMA/CA - Backoff mechanism: The system ensures that the channel is idle, before a frame is transmitted. In the case of IEEE 802.15.4, carrier sensing is performed with a cross-correlation processing block that calculates the correlation of the received signal and the known preamble of every IEEE 802.15.4 frame. Upon a frame transmission, if the cross-correlation resulting values are above a threshold, the medium is marked as busy and a random backoff counter is set. The counter is decreased in a per sample basis until zero is reached. Then the cross-correlation output is checked again.

If the values are below a threshold, the channel is marked as idle and the transmission of the frame proceeds, otherwise the procedure is repeated. On the other hand the CSMA mechanism of the IEEE 802.11 standard is more complicated. *Similarly in this case, the cross-correlation of the received signal and the known preamble of each frame is computed. The IEEE 802.11 standard specifies that the random backoff counter should be decrease*d while the medium is idle. For this case, when the frame length information becomes available, the remaining samples for both the end of frame and the backoff time can be obtained. The channel is set to idle when this number of samples is processed.

The use of software carrier sensing provides unparalleled flexibility, allowing Cognitive Radio (CR) techniques to be applied, better multi-standard coexistence and efficient use of hardware.

2) Acknowledgments: Some frame types should be acknowledged by the receiver. When a receiver successfully decodes a frame and according its type and the destination MAC address for the case of IEEE 802.11, decides if the frame should be acknowledged. In this case, an ACK frame transmission is immediately issued utilizing the transmit chain of the PHY software. To minimize response times, especially for the case of IEEE 802.11a/g, the construction of the ACK frame is initiated before the received frame is fully decoded. Going even further, ACK caching techniques are applied, due to the high probability that another ACK with the same destination station (STA) will be required in the near future.

3) Re-transmissions: The re-transmission mechanism is responsible for transmitting again those frames whose ACKs were not received successfully within a specified period of time. This functionality is implemented in the last processing block of the PHY transmit chain. When this block transmits a frame that should be acknowledged, it blocks waiting on an asynchronous blocking message queue. The receive chain is responsible to enqueue a message at this queue whenever an ACK is received, unblocking in this way the receiver. The transmitter should also be able to perform a re-transmission if time elapsed while waiting for the ACK. For this reason, a time limited version of the blocking dequeue operation was added in the asynchronous message queue functionality of the GNU Radio core system, taking advantage of the conditional variables of Boost [BC++] library as described in the following algorithm.

```
1:  function DEQUEUE TIMED(time usec)
2:          scoped lock(mutex)
3:          while queue.is empty() do
4:                  not empty is a condition variable
5:                  if !not empty.timed wait(mutex, time usec) then
6:                          return NULL
7:                  end if
8:          end while
9:          return queue.dequeue()
10: end function
```

If the dequeue returns NULL, the transmitter should initiate a re-transmission. The procedure is repeated until an ACK is received or the limit of re-transmissions is reached, dropping the frame. In order to minimize the re-transmission delay and the CPU utilization, each frame is stored in a pre-allocated buffer, based on the fact that its maximum possible size is known a priori. Thus, when a re-transmission occurs, the buffered samples are sent directly to the SDR device.

2.5.4 SDR support

Many SDR devices transmit the I/Q samples in a stream oriented way. Unfortunately, frame transmissions of the IEEE 802.11a/g and IEEE 802.15.4 are bursty in nature. They send a number of samples and stop, until transmission of the next frame. This can cause several problems in the stream oriented SDR devices. To deal with this problem, USRP devices introduced the 'Burst tag', which is the information of the number of samples that comprise the frame and the device should wait for. But this solution requires additional changes to the PHY software implementation, limiting its portability. Based on the aforementioned limitation, we opted to alter/improve the driver and the firmware of the FPGA of the SDR card, in order all the produced samples from the userspace software to be immediately available to the FPGA. When the Digital to Analog Converter (DAC) consumes all the available samples, the FPGA automatically continues to feed the DAC with zero samples. Therefore, there is no need to explicitly inform the device about the frame size and is guaranteed that the whole frame will be transmitted.

Supporting multiple channels or standards that are stream oriented is quite straight forward. The time domain signals from all the streams can be added to form a single stream, before the samples are propagated at the SDR device. However, in bursty protocols this is quite difficult due to the fact that each different channel or protocol can initiate a frame transmission at arbitrary time instants. To efficiently deal with this problem, we create a different flow of samples for each network interface. Each channel/protocol writes the samples of a burst to the corresponding flow independently. The FPGA is responsible to check the available samples of each flow, sum them together on a single stream and propagate the resulting samples to the DAC. The architecture of this approach, which is relatively trivial to implement and requires minimum amount of the FPGA resources, is illustrated in Figure 35. However, each I/Q sample is represented by a 16-bit integer, so there is always the case of an arithmetic overflow after the summation of each flow. We do not handle this case in the FPGA, assuming that PHY software takes care of it.

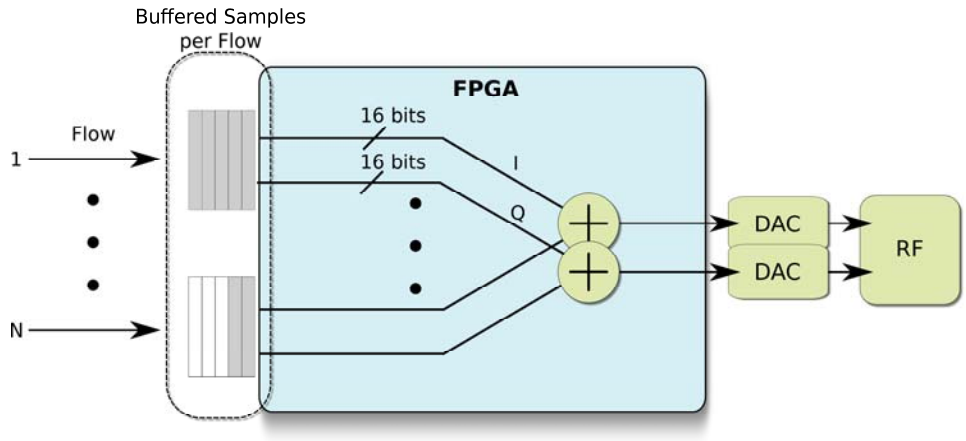

Figure 35. Multiple flows architecture for multiple bursty channels/protocols support.

Furthermore, the SDR device can tune and transmit the signal at a specific centre frequency; each flow should have its own centre frequency. While this is complex to achieve via shared RF hardware, the time domain signal of each flow can be up/down-converted easily by multiplying it with a software Numerically Controlled Oscillator (NCO) and an appropriate frequency offset, as long as the offset lies within the bandwidth span that the SDR device can provide. Finally, it is essential that all flows operate

at the same sampling rate. For example, IEEE 802.11a/g requires a sample rate of 20 Mega Samples per Second (MSPS), whereas IEEE 802.15.4 requirement is only 4 MSPS. Therefore appropriate Sample Rate Conversion (SRC) techniques should be applied at each flow to meet the aforementioned restriction.

2.5.5 Evaluation

Our platform is based on a Dell T7610 workstation PC running Linux (kernel 3.7.10) and GNU Radio (ver. 3.7.6), equipped with a Per Vices Noctar SDR card [NSDR]. The T7610 workstation has two Intel Xeon E5-2630 CPUs at 2.6 GHz, assisted by 16 GB of 1866 MHz DDR3 RAM. The Noctar SDR is a PCIe v1.1 x4 card, that utilizes an Altera Cyclone IV EP4CGX22C FPGA, a dual channel 250 MSPS 16bit DAC and two 125 MSPS 12bit ADCs. The card can be tuned into any frequency between 100kHz and 4.4GHz. It provides four common transmission and reception sampling rates through four steps of interpolation/decimation; 125 MHz, 62.5 MHz, 31.25 MHz and 15.625 MHz, thus it can cover the whole 2.4 GHz ISM band. The Noctar SDR has the advantage of wide bandwidth reproduction based on the high speed PCIe interface that allows the use of high sampling rates.

The platform was tested in our laboratory under real, congested spectrum conditions. The IEEE 802.15.4 transmission and reception performance was tested by exchanging packets with the off-the-shelf sensors of the Zolertia Z1 platform [Z1] equipped with the CC2420 radio module. The IEEE 802.11g was tested by transmitting and capturing packets of the existing wireless networks within our laboratory. The transmitted packets were successfully decoded within an acceptable packet error rate comparable to the commercial devices.

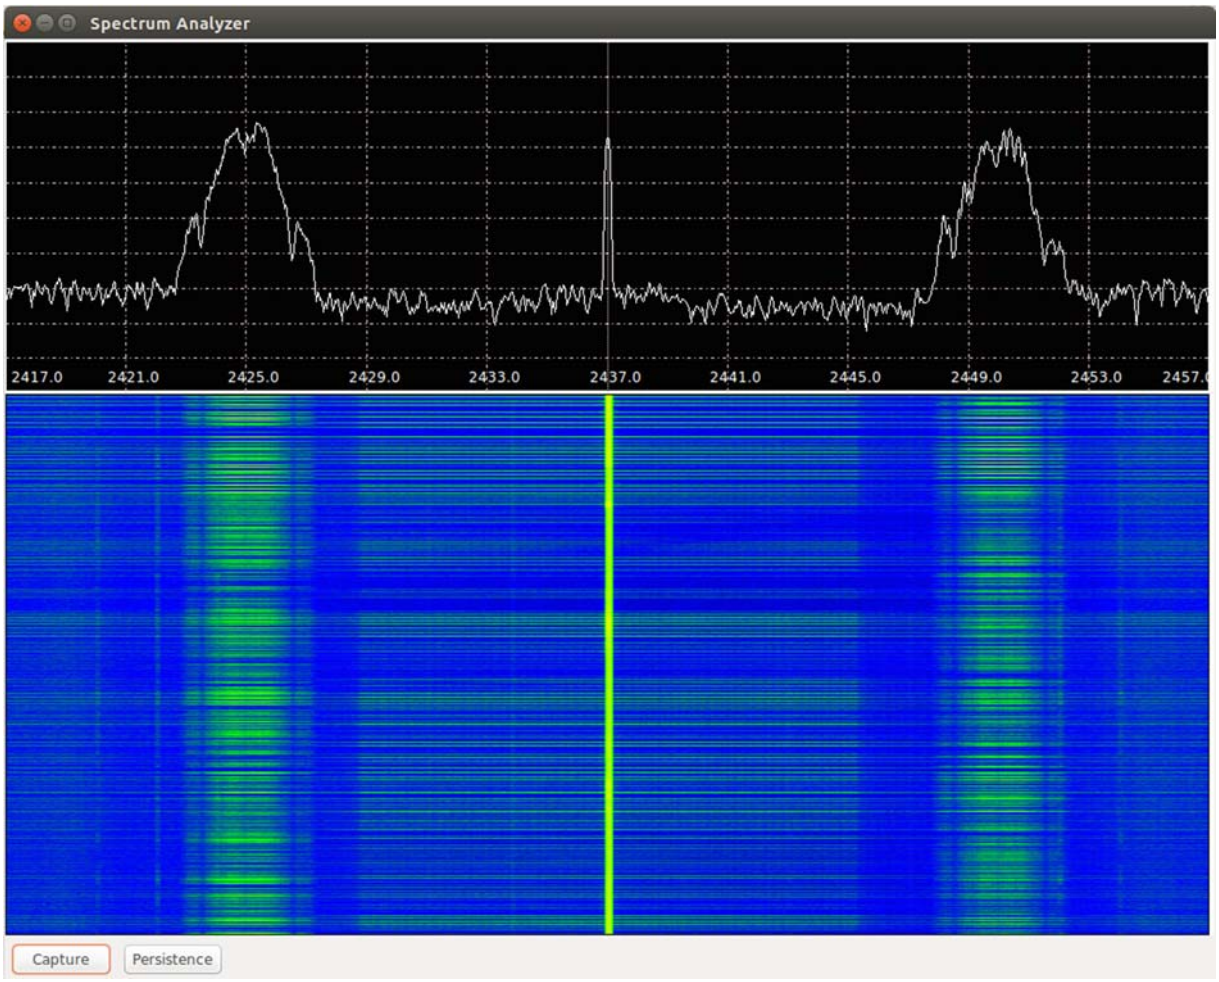

Figure 36. RTSA screenshot of a concurrent transmission of one 802.11 channel and two adjacent 802.15.4 channels.

An example of the concurrent transmission of one 802.11 and two 802.15.4 channels is presented in Figure 36, where the captured signal using our SDR-based Real-Time Spectrum Analyzer (RTSA) implementation is depicted. The captured bandwidth span is 40 MHz, the centre frequency is 2437 MHz and the time duration is one minute. Although there are various minor signal impairments, due to hardware imperfections, which are clear in the waterfall trace, the result shows the high potential of the platform. As far the CPU utilization concerns, the IEEE 802.11a/g implementation which is the most demanding, utilizes 29% of the available CPU resources while being idle, whereas under heavy bidirectional UDP traffic 36%. The unexpected high CPU utilization when the system is idle has to do with the continuous computations required for frame detection and the CSMA algorithm. To accomplish the tasks required by the IEEE 802.11a/g transceiver, we exploit parallelism using 22 threads. Each of these threads executes a GNU Radio processing block or specific tasks of the standard. The tasks are organized in a pipelined way in order to reduce the latency.

As it is shown, the flexibility and future proof design of the proposed prototype can be a key enabler for RERUM for achieving an efficient interconnectivity of heterogeneous RDs. It goes a step ahead in the usability, by exposing the available standards and channels as separate typical interfaces of the Linux operating system. Therefore any Linux based tool, component, and protocol may be used. Moreover, as a software based solution it is fully customizable and cooperative cognitive techniques may be employed incorporating the finest spectrum detail processing available by the SDR device. It is a robust, unparalleled solution capable of serving any present and future requirements of the IoT wireless networking.

3 Optimization of wireless resources of networks of RDs for interference avoidance

This section deals with interference avoidance techniques in a dense network of RDs, when served by RERUM gateways or cellular infrastructures. In the former case, we consider that multiple gateways can be utilized in a small area thus giving rise to severe interference which can limit the connectivity performance (in terms of both downlink and uplink throughput of the RDs, as well as packet collisions and errors). Considering the RERUM use cases, this work can resolve the interference and congestion challenges that exist in the outdoor deployments (in the environmental monitoring use case) when many RDs are deployed in dense city areas. However, within buildings, the deployments tend to be more dense with multiple gateways at narrow areas to mitigate signal propagation losses because of the existence of walls (no-line of sight environments). Thus, the wireless networks are becoming more dense and the interference in the free ISM bands can be a significant challenge. As a result, there is an increasing need for a technique to avoid neighbour cells or clusters of IoT devices interfering with each other.

The work herein assumes that in the case of gateway deployments, efficient clustering has been performed by the Clustering module of the Network Manager (see Figure 4 in the introduction). Similarly, in cellular-served RDs the clustering is mapped to the cell assignment. In what follows the terms cluster and cell are used interchangeably and denote the set of RDs served by a single backhaul-linked device (Gateway or Base Station - BS).

3.1 Frequency reuse in dense RD radio environments

Orthogonal Frequency Division Multiple Access (OFDMA) is in use in the dominant IoT radio standards (including the IEEE 802.11 suite and 802.15.4 PHY layer). OFDMA divides the available system spectrum in a number of orthogonal[6] subcarriers (see Figure 37), enabling flexible bandwidth usage and increased spectrum efficiency. In multi-cell wireless networking with OFDMA, while intra-cell interference is not a problem for either the uplink or downlink –due to the subcarriers' orthogonality, inter-cell interference, control, and mitigation become critical.

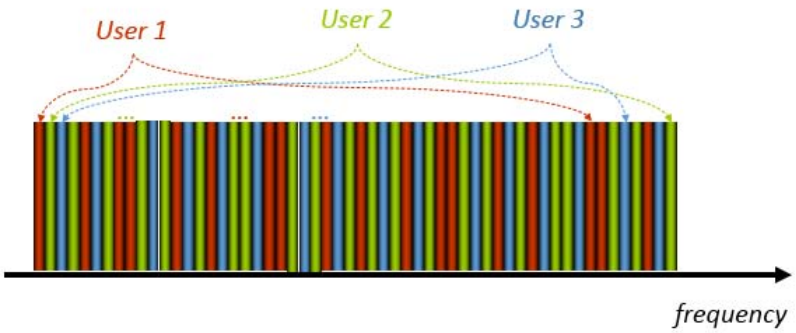

Figure 37. The basic OFDMA principle – subcarriers allocated without grouping to different users.

Reusing all subcarriers in each cluster, i.e. employing the “Reuse 1” scheme, gives the best performance in terms of the system aggregate throughput. In this case though, users near the cluster edges suffer severely, due to poor signal-of-interest strength (due to the distance from the transmitter) and the high interference (due to the proximity to potential other transmitters). Reuse schemes of higher order can be employed to limit interference, but this causes significant bandwidth under-utilization and brings down the system total throughput.

[6] i.e. their power profiles over the occupied bandwidth are insignificant.

To address these issues the Fractional Frequency Reuse (FFR) technique was introduced originally in the cellular 4G context of WiMax. FFR, supported in both the LTE and WiMAX specs, enables interference management through a re-definition of the notion of frequency reuse. Each serving cluster is partitioned into two zones: the centre and the edge (Figure 38). The available frequency band is also partitioned into two respective parts: a "centre band" – utilized with Reuse 1 i.e. in all clusters of our infrastructure, and an "edge band", which is further partitioned to be utilized with a higher reuse factor. The definition of what constitutes an edge, the frequency reuse factor and the band partitioning are parameters of the scheme that, despite research efforts, are implementation- and scenario-specific.

In any case, enhancing the edge throughput with the FFR scheme comes at cost: centre RDs need to "sacrifice" a portion of their available bandwidth. Those having favourable channel conditions are not interference-limited, but rather bandwidth limited. Therefore their throughput decreases. Since they constitute the majority of the cell population the overall cell throughput also decreases.

Initially, in this section we assume that power is evenly distributed over the subcarriers in every cluster and we consider the power assignment scheme where each cell can select, among a given set of power levels, the power to be used per OFDMA subcarrier. The total power used for the cell-edge zone of a cell is thus the product between the chosen power level and the number of allocated sub-bands.

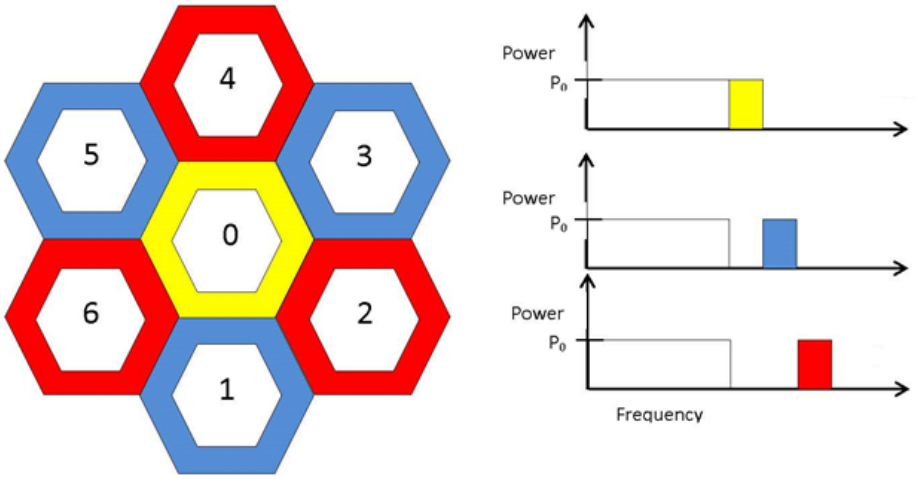

Figure 38. The FFR Reuse Scheme in an idealized hexagonal plane tessellation: white is the centre area served under the same band, while the edges are served by the edge band comprising the yellow, blue and red sub-bands

3.1.1 System model

We denote $\mathcal{L}$ the set of clusters served by OFDMA network of radio gateways or cellular towers. The service area is considered to be discretized by a grid of a large number of square tiles. Each tile $k \in \mathcal{K}$ is assumed to have uniform radio conditions for a served RD. We use g_{ck} to denote the total gain between the serving transmitter $l \in \mathcal{L}$ served tile $k \in \mathcal{K}$. The gain value can be assumed known either by measurements (performed by the sensing engine of the CR agent of an RD in our case) or assuming textbook radio signal propagation models. The service area of a cluster is divided into a centre zone and an edge zone. The latter represents locations that get high interference by surrounding gateway/cellular transmitters. We use $\mathcal{K}_{Ek}$ to identify the tiles forming the edge area of cluster l. Note that, for real networks, the shape of both the centre and edge will be quite irregular, unlike Figure 38 above and more like Figure 39 below. Moreover, it may even happen that a cell does not have an edge zone. This occurs if the signal of the cell antenna is strong over the entire cell area. Simplifying the notation, we take all cells to have an edge zone in our system modelling.

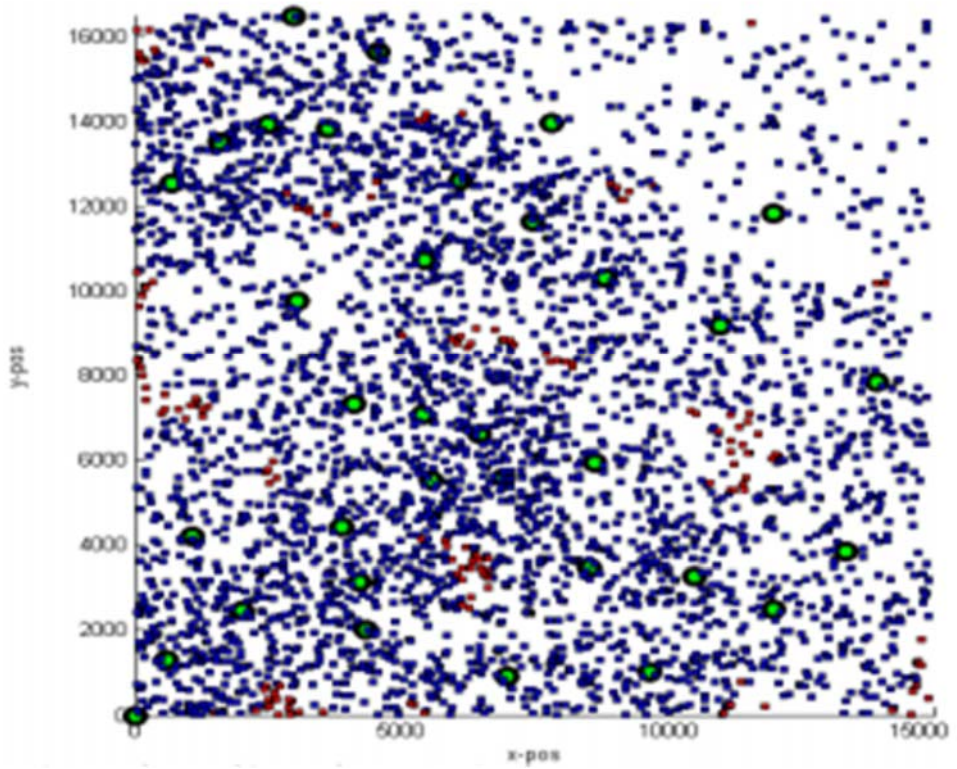

Figure 39. A realistic scenario with green nodes being the serving transmitters, tiles in blue & red are active user tiles, and red denotes edge tiles with a served RD.

The total downlink transmit power of any cluster serving antenna is denoted by P. We assume that P does not vary by cell/cluster, again towards simplifying notation. The available spectrum is split in two bands to be used by the centre and edge, respectively (white and coloured parts in Figure 38). We denote the system bandwidth B, while that allocated to center and edge B_C, and B_E, respectively.

The centre bandwidth B_C is allocated with Reuse 1 to all centre zones. The B_E band is split into N equal-sized sub-bands. In what follows η is the noise power, which we assume is the same on all edges.

Our centralized solution model focuses on optimizing the edge performance via mitigating interference. In our model, we allow multiple edge bands of the B_E split to be allocated to an edge. For each band $n = 1, \dots, N$ we book-keep the clusters that our solution will have it in a vector $\boldsymbol{E_n}$. Taking each vector $\boldsymbol{E_n}$ we construct the allocation of bands to clusters edges which we shall denote by $\boldsymbol{e}$. By $\boldsymbol{e}|^{\boldsymbol{i}}$ we denote the number of sub-bands allocated to the edge zone of cell i; this equals the number of set elements in containing cell i.

3.1.1.1 Power allocation

We consider a uniform power over the edge sub-bands in each cluster. Thus, the power used per sub-band in by the cluster head of cluster i is denoted by p_i. This is our optimization variable, and its value is to be selected from a given set of available power levels $\{P_1, P_2, \dots, P_L\}$. The final power allocation solution is given by the vector $\boldsymbol{p} = (p_1, .. p_{|\mathcal{L}|})$.

The number of bands of an edge zone be up to M, thus, the maximum value that p_i can take depends on the number of bands utilized in i. We assume that the largest available power P_L can be the total power available to any cell for its edge zone, and set its value to be equal to the total cell power scaled by the edge bandwidth. Clearly the power allocation is under the constraint that the product of p_i times the number of bands allocated in cluster i does not exceed P_L.

3.1.1.2 Edge Throughput

Throughput is typically determined via a function of the Signal to Interference-plus-Noise Ratio (SINR). These are monotonous and increasing in SINR, in many cases the Shannon formula serves for purpose of evaluation. Thus in a tile t served by antenna i for a sub-band allocation $\boldsymbol{E_n}$ and power assignment $\boldsymbol{p}$, the SINR for band n becomes:

$$SINR_n^{it}(\boldsymbol{c},\boldsymbol{p}) = \frac{p_i g_{it}}{\sum_{h \neq i} p_h g_{ht} + \eta}$$

Here, the numerator is the tile received power from the serving antenna, while the denominator sums all interference from clusters using sub-band n, and the noise.

3.1.1.3 The optimization model

Given the monotonicity of the throughput function discussed above the model proposed attempts to find a band allocation $\boldsymbol{c}$ and a power allocation $\boldsymbol{p}$, such that the overall SINR of the edges areas is maximized: $\max_{\boldsymbol{c},\boldsymbol{p}} SINR_n^{it}(\boldsymbol{c},\boldsymbol{p})$

$$s.t.\ \boldsymbol{e}|^i \geq 1, \forall i$$

$$\boldsymbol{e}|^i p_i \geq P_L, \forall i$$

The optimization is subject to the conditions that each cell-edge zone is allocated at least one sub-band, and that the power used by the sub-bands together does not exceed the power limit in any cell. The model defines a combinatorial optimization problem, which in its general form, is NP-hard.

3.2 A negotiation protocol for wireless resource bargaining

Given the hardness of the problem presented above, here, we introduce a scheme for cooperatively adapting the FFR subcarrier allocation based on the load in the served clusters' centres. We enable overloaded clusters to acquire additional sub-bands by borrowing their neighbours' edge bands. Edge band borrowing requests are triggered by the clusters' load and are thus generated in an asynchronous fashion. A serving antenna (cluster head - which may be a RERUM Gateway or mobile broadband enabled base station – i.e. an ENodeB), upon receiving such a request, will estimate the expected added interference its edge users would suffer, for a set of initiator-proposed power levels.

The decision of whether or not to grant a permission to borrow the edge band at one of the proposed power levels takes into account the *local* knowledge of

(i) neighbouring clusters' requests and

(ii) existing subcarrier allocation

If the expected interference, at one of the proposed power levels, result in a tolerable throughput loss for the lending cluster edge, then its serving cluster head will grant a permission to the requesting cluster head to use the band at that power level. A cluster that issued a request may clearly receive, from its neighbouring cluster heads, several edge band usage options.

In its turn, it also estimates how many of these bands to use and at which power. The flexibility in sub-band usage is accomplished by modelling the edge band borrowing grant decision as a small set of small-sized 0/1 knapsack problems, to be solved at each antenna when a request arrives. The goal is to maximize the requesting cluster centre throughput, while maintaining acceptable degradation on the granting sector's edge with respect to the allocation at the time the request is received.

3.2.1 Protocol description preliminaries

We now assume an IoT- serving cellular-based system of N clusters utilizing FFR. The available system bandwidth is divided into a center band and a set of disjoint edge sub-bands. At each cluster the cluster head (eNodeB) utilizes the centre band for its centre RDs, and one of the edge bands for its edge RDs. Therefore the centre band is used with the Reuse 1 scheme, while the edge bands have a higher reuse factor (here *for presentation simplicity* we employ a reuse factor of N= 3). The original assignment of the edge bands to clusters is assumed to be given, and falls out of the scope of this protocol. We finally assume that each eNodeB has a fixed transmit power which is initially distributed equally over all resource blocks. Each cluster is assumed to initially be the "owner" of one edge band. This edge band ownership notion conveys that (i) a cluster head can issue grants only for the edge band it owns, and that (ii) once an edge band owner revokes a grant, the borrower is expected to release the band.

Since our goal is to address overloaded centres, we consider that a borrowed edge band is used by the overloaded cluster to extend its own centre band. The power that will be used in the borrowed bands is reduced from that originally used in borrowing cluster centre. Note that our scheme does not make any assumptions, nor impose any restriction on the location and the number of cluster heads, the shape of the clusters, or the edge frequency, sub-bands and pre-allocation method.

The bargaining protocol consists of 3 stages: (i) loan request for neighbours' edge band, (ii) lending decision, and (iii) borrowing decision.

3.2.1.1 Requesting

When a cluster centre becomes overloaded, the cluster head (Gateway/eNodeB) node issues a request to all neighbouring clusters. This is done through a message that contains up to $M = 16$ pairs of: 1) a power level $p_k,\ k = 1, \dots, M$ to utilize an edge band, 2) the mean expected center RD throughput g_{ki} , that the issuing cluster i expects from utilising one extra sub-band at the above power level. For this value the cluster head issuing the request assuming that (i) the borrowed band will be utilized for its centre RDs, (ii) the borrowed edge band will experience interference coming from the band's original owners. All neighbouring owners of an edge band take part in the lending process. This is because each cell may well have more than one neighbours that own a given edge band (e.g. in Figure 38 clusters 1, 3 and 5 own the same edge band). If the neighbours are themselves overloaded, or under high interference they will produce replies with zero power of activation. If that is not the case, then power level p_1 is such that all subcarriers utilized for the centre of the requesting cluster will have equal power.

As a bargaining protocol, the rationale for providing directly more power-gain pairs in the request message is that this "equal power" assumption may result in high interference to the edge lender and thus prevent them from giving a positive lending decision. Having more pairs directly in the message we limit any overhead of an actual back-and-forth communication process in case of negative lending decisions.

3.2.1.2 Lending Decision Solving small knapsack

Upon arrival of any new request, a cluster head examines $k = 1, \dots, M$ to evaluate whether it is beneficial to lease its edge band to the requester. To evaluate this, it must examine its edge band's lease condition and the potential gains from altering it. We have modelled the lending decision at cluster $n, n = 1, \dots, N$ as a 0/1 knapsack problem using the following notation and assumptions: K_n is the set of neighbor clusters of n. Variable x_i is the set of decision variables indicating whether the BS should lend its edge band to neighbour i when $x_i = 1$ it will borrow n's edge band otherwise it is 0. Parameter g_i is the expected average center RD throughput gain for cluster i. This parameter is provided in the request message issued by the cluster. The c_i parameter captures the average edge user throughput loss that will be observed at cell n due to increased interference if i borrows the edge band. Note that neighbours that have already borrowed the edge band from the current cluster must

be also included in the knapsack formulation, assuming the same gain and loss values they had before. Finally the T_n parameter is the acceptable average edge user throughput loss for this cluster. We can now model the lending decision at n as a 0/1 knapsack problem

$$\max \sum_{i=1}^{|K_n|} g_i x_i$$

$$\text{s.t.} \sum_{i=1}^{|K_n|} c_i x_i \leq T_n$$

$$x_i \in \{0,1\}$$

Note that c_i calculated locally at the Sensing Engine of the CR-agent of a RERUM Gateway n can be a *pessimistic* estimation of interference. This overestimation will occur in the cases that (i) the borrower will end up acquiring more than one bands, and hence will use n's band at a power level lower than the one it is assuming; (ii) in the case that not all owners of a band unanimously decide to lease their edge band to the cluster that issued the request. As we shall see in the next section, this leads to a "soft" decision: The requesting node will actually go ahead and borrow the band using a power level far lower than the one originally assumed.

3.2.1.3 Borrowing band selection

With the solution of a knapsack problem instance, cluster head n generates a new boolean 0/1 lending decision for each neighbour cluster in K_n and respectively informs the requesting clusters. For an illustrative example assume the topology of Figure 38, and consider the central cluster being overloaded issued a request to borrow an edge band. All six neighbours would solve the respective knapsack problems and provide their responses back to the central cluster. Consider the simplest case where all decisions for a single band come in cell 0 unanimously. Then clearly the borrowing decision is fairly simple: the central cluster will indeed borrow the band and set the power level to the one described earlier.

3.2.1.4 Simulation Results

The simulation results in Figure 40 below provide the initial proof-of-concept for the protocol depicting in one plot the performance on the users of an overloaded cluster utilizing Reuse 1, the classic FFR3 and the proposed protocol. It can be observed that in Reuse 1 RDs in the centre benefit (have higher throughput) while those on the edge suffer. In Reuse 1 a continuous throughput decrease is observed as the distance of the RDs from the serving transmitter increases, which is due to higher path loss and more interference from the neighbouring clusters transmitters. FFR3 improves the throughput of the edge users but at the cost throughput to central RDs. Furthermore in FFR3 the few edge RDs enjoy the increase in the throughput while a significant number of centre RDs experience throughput decrease. Our protocol does better than FFR3 for the edges and delivers results between the Reuse 1 scheme and FFR3 scheme for the centre. It improves the centre users' throughput significantly as compared to FFR3 centre users' throughput which is represented by blue lines while at the same time it is not performing so bad in edge. It shows marginal decrease in the throughput of edge users as compared to FFR3 edge users' throughput and still shows significant improvement in the throughput as compared to Reuse 1 users' throughput in the edge area which is represent by the green lines.

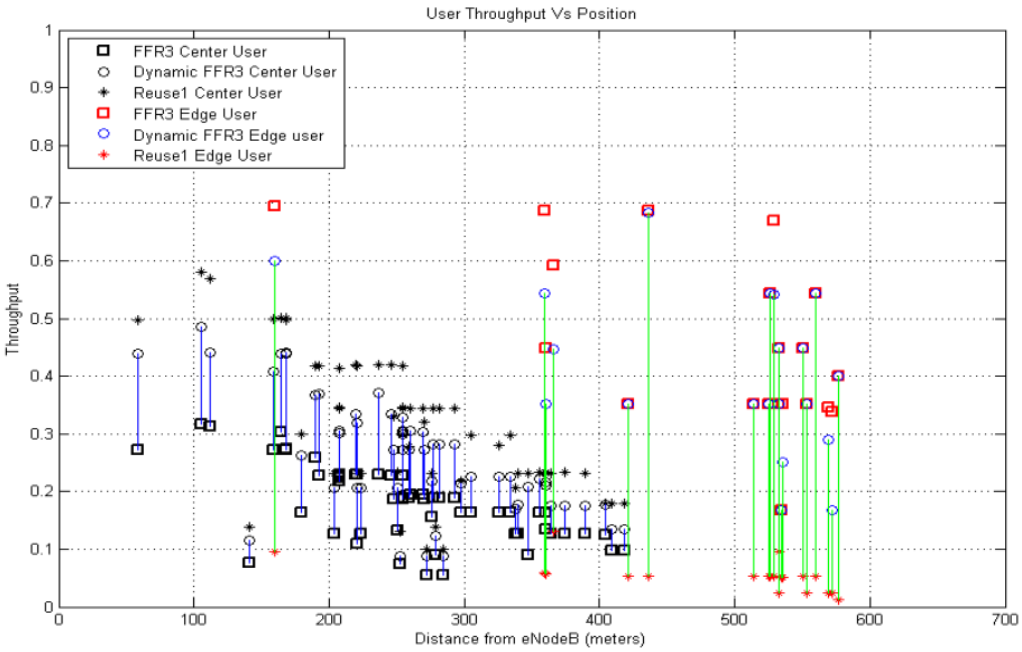

Figure 40. For a single, overloaded cluster (in a topology similar to that of Figure 38 – with fading) the effect of running the reallocation schemes on the centre and edge RDs.

Figure 41 shows the centre throughput gain for overloaded cells, aggregate system throughput gain and the edge loss with against the FFR3 scheme. On the top we see the result when only 2 of 21 clusters are overloaded, while on the bottom figure all 21 clusters are overloaded. In case of a lightly overloaded network, our protocol achieves up to an over 45% throughput gain in the centres with a cost of below 10% edge throughput loss when compared to the FFR3 scheme. For a heavily overloaded network the percent gain in the centre throughput still is a considerable 36%, but to an edge loss of 34%. This increase in the edge loss and decrease in the gain of centre RDs throughput is due to the utilization of the borrowed edge bands in the centre of greater number of overloaded cells. Although the centre throughput gain decreases in this case a significant increase in the overall system throughput of 26% is observed. This is due to the fact that a greater number of sectors are utilizing the borrowed edge bands in the centres of overloaded clusters.

The run time of the knapsack solutions was quite efficient, depending directly on the number of overloaded sectors and the U parameter. This is clearly because of increased number of knapsack objects when the number of overloaded sectors increases, and increased number of knapsack solutions when the U parameter decreases, where due to our implementation, the runtime also includes recalculation of estimated path-losses..

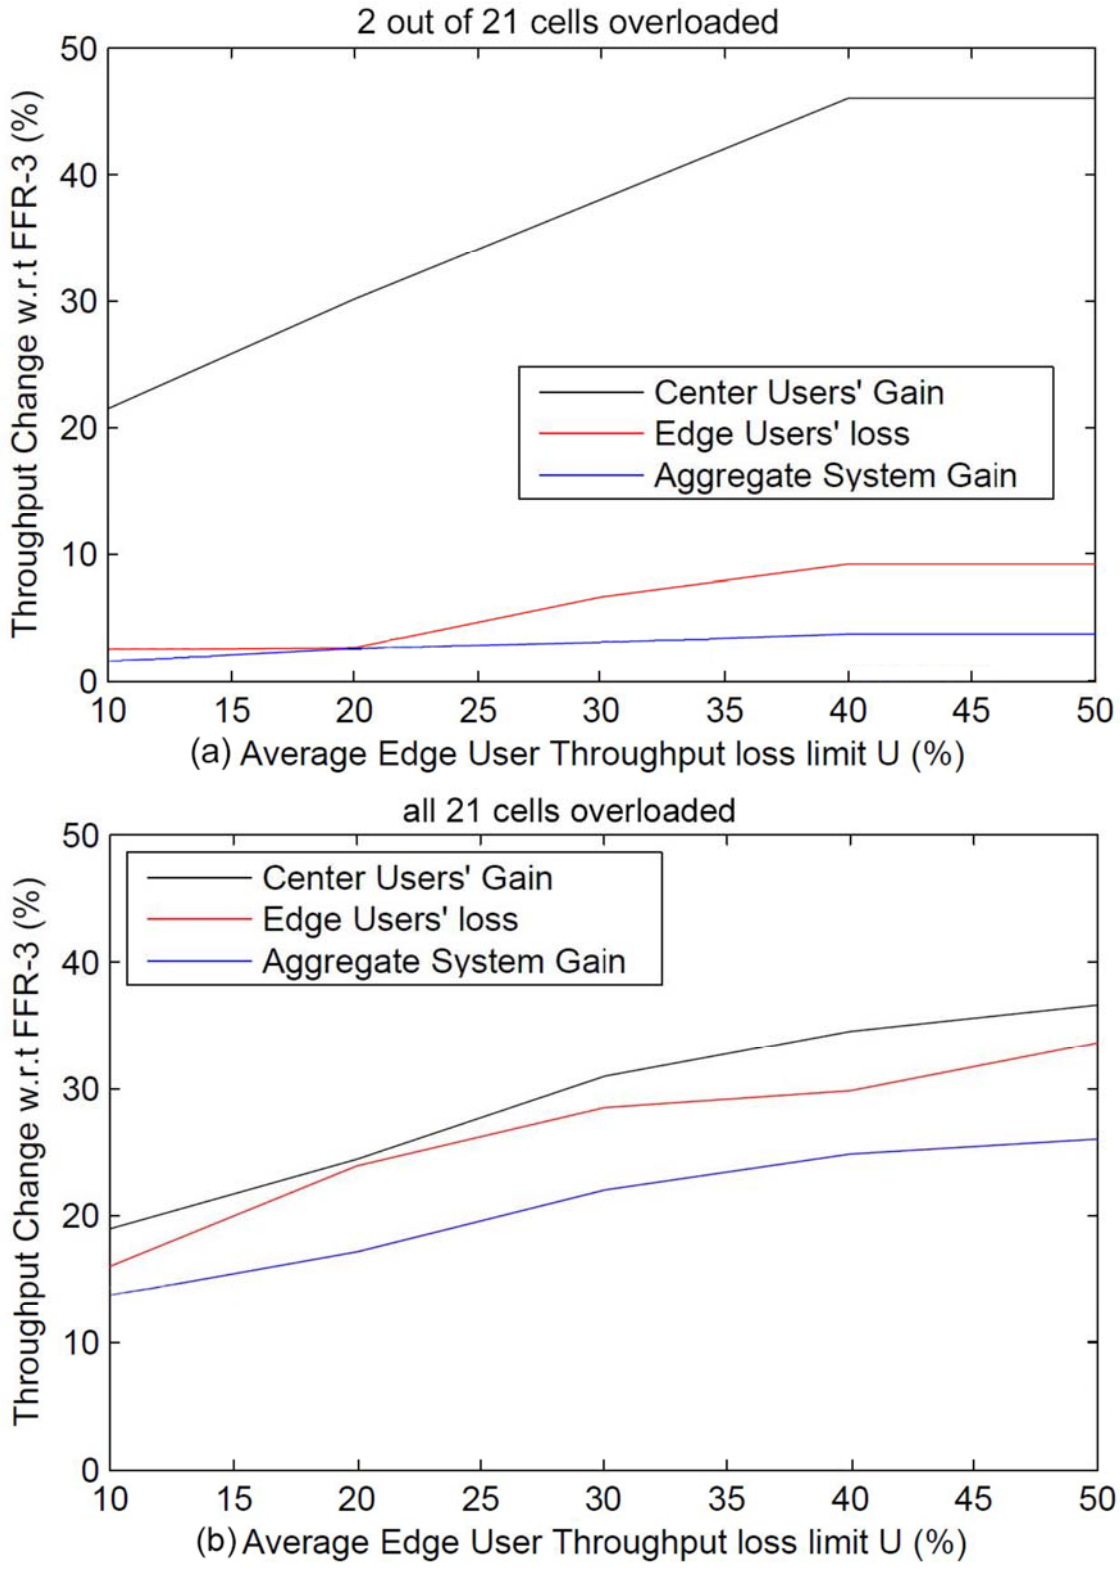

Figure 41. Performance analysis for FFR3 in overloaded cells

4 Modelling and assigning resources in virtualised networking environments

IoT applications are expected to introduce content-aware communication flows over the top of the IP level, which creates the need for coupling the network layer with the upper layers through Network Aware Applications (NAA). This type of content-aware communications are based on intelligent devices (e.g., routers), which process and forward data with respect to the content type or the name of the data objects, in contrast to conventional routers that consider mainly the IP address for identifying the destination device [CH++11, JST09].

IoT mobilizes various communication technologies and sensor platforms, which raises the need for the seamless interoperability from different vendors. To this end, the virtualisation concept has been proposed for enabling a uniform management of devices and sensors through the interposition of an abstraction layer between the application logic and the sensor driver [ASSC02, LIN09]. Virtualisation can also include the formation of Virtual Sensor Networks, which enable multi-purpose, collaborative, and resource-efficient exploitation of the physical infrastructure. Furthermore, software abstraction layers are used as well for facilitating the interoperability, the uniform management of different platforms [YK10], the dynamic reconfiguration of sensor nodes [JHI07], and the combination of sensor data through federations of sensors [AHS07].

This section presents the techniques developed within RERUM that enable the network virtualisation of the RDs and the efficient discovery and assignment of their Resources. These Resources can be either network Resources (providing measurements with regards to the status of the network the the RD's network interfaces) or IoT Resources (that describe or act on a Physical Entity providing i.e. the sensing measurements that the RD gathers from its on-board sensors; see also the Definitions section at the beginning of the document). There are many IoT-related works that focus on the discovery and assignment of IoT Resources, so within RERUM we will utilize existing solutions to avoid re-inventing the wheel. Thus, the focus of this section (and of RERUM) is on the network Resources of the RDs that represent the physical and networking capabilities (and measurements) of the devices. Considering this, throughout this chapter (unless stated otherwise) the term "Resource" will represent the physical and networking Resources of the RDs. In an abstracted way, if we want to use the IoT-A terminology, we can also assume that the RDs are the Physical Entities (PEs) that are of interest for this type of Resources. The RDs include have embedded the "network interfaces", which are also PEs using an inheritance like the box that contains another box in Figure 9.7 in the IoT-A book [IOTA]. The network interfaces have various attributes (channel, bandwidth, buffer queue, link quality, link throughput, etc.) that are of interest for monitoring and acting upon from the network mechanisms' point of view for optimizing the interconnectivity of the RDs. The networking Resources can indeed be abstracted, virtualised and exposed by Services in the same way as the standard IoT Resources that describe/act on Physical Entities like the house, the room, a road segment, etc., as described in IoT-A [IOTA].

The development of new approaches for semantic representation of resources and services for their virtualisation is out of the scope of RERUM. However, to ensure that the reader gets a better understanding on existing work in this area that will be re-used within RERUM for the representation of the network Resources and Services, we present at the beginning of this chapter (Section 4.1) an overview of existing literature for semantic resource representation methods and their applicability in RERUM. Then, in Section 4.2 we present a novel method for virtualising the network interfaces of the RDs to allow a more flexible and efficient way of allocating heterogeneous Services to interfaces for ensuring their QoS. This method builds upon the efficient description of the Resources that are available on each of the network interfaces in order to make the appropriate assignment of Services to interfaces, so a detailed semantic description of these Resources is of paramount importance. Similarly, the Services should also be semantically represented in a rich and efficient way so that the mechanism will know specifically their requirements in order to accurately map them to the appropriate virtual interfaces. In Section 4.3 a method for discovering and assigning network Resources

to RDs is described. The target is to meet the requirements of the Service requests and to achieve an overall minimization of the network interference when the RDs have both 3G and WiFi interfaces. The interference can easily be assumed as one attribute that describes the "wireless spectrum" that can be considered as the Physical Entity of interest for the RD. So there is a need for rich representation of the interference in order to ensure that the gateway takes the correct decisions about assigning the network Resources to the RDs.

4.1 Semantic representation, naming and addressing of Resources and devices

4.1.1 Semantic representation of Resources

A. State of the art

From the early beginning of Machine to Machine (M2M) and IoT platforms "semantics" played a crucial role. IoT platforms identified from the beginning the need to provide uniform exposure of resources via virtualisation and to overcome heterogeneity and interoperability issues with commonly designed semantics/metadata for data types, structures and taxonomies. These aspects proved to be useful to perform search and reasoning in a uniform way. However, semantic is not only useful to describe static capabilities, addressing structure and manipulation of resources, but also data and data streams produced and consumed in IoT ecosystems. Linked Data streams [LDAT] face an increase interest due to the richness added to classical streaming solutions.

Despite recent progress, current implementations of IoT architectures are confined to particular application areas and tailored to meet only the limited requirements of their specific applications. Current standardization activities as a result of research projects or industrial efforts work to unify this approach at the level of processes and metadata. Hence the provisioning of IoT enabled business services (in short IoT Services) is a time- and cost- extensive process. The complexity involved in data acquisition, quality control, context interpretation, decision support, and action control hinders the uptake and penetration of specific tailored services for "Internet of Things" applications.

As per any systemic work on information systems we can identify an impact of semantic technologies at both: (i) design time, when resource provisioning is happening in targeted virtualisation platform and (ii) run time, when virtualised resources should be discovered, aggregated and used by domain oriented services and applications. For both aspects, common and specialized types of semantics should be considered. In RERUM, context semantics aspects are considered as enablers for key project objectives, such as security and privacy (as described in [VF14]) or networking reliability, as it is described in the following parts of this chapter. As mentioned above, RERUM will investigate, adapt and reuse the optimal existing approaches, as it is not within the project goal to extend the state of the art in this area.

Scientific literature in this area documents a significant volume of results at practical level that are also reflected in IoT standards, for example, the SSN ontology [W3C11] for annotating sensors and sensor networks; Linked Data [BTT09] for sensor data publishing [BPM10] and discovery [PHPS10], and semantic sensor observation services (SemSoS) [HPS09]. Some recent results such as ones of IoT-A [IOTA] and IOT.est [IOTEST] projects propose a modelling approach in which resources on the IoT are able to expose standard service interfaces. This coincides with the principle of the Service Oriented Computing [OASIS06] that can provide scalable, distributed and service-oriented means to access and process IoT information. More importantly, existing methods for service discovery and composition can easily access IoT based services to create context-aware and personalised services and applications.

Semantic modelling for the IoT aims to provide a common ground for interoperating among different systems and applications by providing machine understandable and processable annotations; however, current work has mostly focused on IoT resources management, but not on how to access and utilise information generated in IoT. The ontologies help exploit the synergy of the existing efforts and provide support for crucial tasks in such as IoT resource and service discovery, composition, adaptation and etc.

As previously noted, there are also past and present efforts that aim to address the IoT data management issue. However, further research on the semantic service oriented approach for data access, abstraction and processing is still needed. Furthermore, as many IoT resources operate in the highly dynamic and volatile physical environments, the services exposed by those resources become less reliable compared to most of the well designed and maintained Web and Internet services.

An early work on defining common interfaces and descriptions for IoT related resources and data are provided by the Sensor Web Enablement (SWE) group at OGC [OGC07]. The main specifications defined by OGC are: (i) Observations & Measurements (O&M), which defines a standard model and XML Schema for encoding real-time and archived observations and measurements of sensor data; (ii) Sensor Model Language (SensorML) [OGC07], which is a standard model to describe sensor systems and processes associated with sensor observations in an XML-based schema; (iii) Sensor Observations Service (SOS), which is a standard Web service interface for requesting, filtering, and retrieving observations and sensor system information; (iv) Sensor Planning Service (SPS), which is a standard Web service interface and acts as an intermediary between a client and a sensor collection management environment; (v) PUCK Protocol, which defines how to retrieve a SensorML description and other information and can enable automatic installation, configuration and operation of sensor devices; (vi) SWE Common Data Model, which is used in nodes to exchange sensor related data; (vii) SWE service model, which defines data types used across SWE services. The PubSub Standards Working Group is implementing the SWE standards to enable publish/subscribe functionality for OGC Web Services and define the methods to realise the core publish/subscribe functionality for a specific service binding (e.g. using SOAP, RESTful).

The W3C Semantic Sensor Networks Incubator Group has developed an ontology for describing sensors and sensor network resources, called the SSN ontology [CBB+12]. The ontology provides a high-level schema to describe sensor devices, their operation and management, observation and measurement data, and process related attributes of sensors. It has received consensus of the community and has been adopted in several projects. To model the observation and measurement data produced by the sensors, the SSN ontology can be used along with other ontologies such as the Quantity Kinds and Units ontology [LL05] and the SWEET ontology [RP03]. The SSN has also been used with domain ontologies to develop various smart Things ontologies, such as the smart product ontology.

B. Applicability on RERUM

The IoT-A project [IOTA] has identified entities, resources and IoT Services as key concepts within the IoT domain [W++11]. The Physical Entity is the main focus of interactions between humans and/or software agents. RERUM adopts the same terminology and assumes that the digital representation of the physical entity is termed as a Virtual Entity (VE), which is accessed by a Virtual RERUM Device (VRD) that is the digital representation of a RERUM Device. Due to the variable possible types of connected devices in the IoT domain, a Resource abstraction allows for both hardware (e.g. sensor, actuator) and software specification (e.g. in the case of storage device) of the connected device. The Resource is the actual software component that provides information on the entity or enables controlling of the device. Thus, a Resource is not only related with the measurements that are gathered by the RD, but can also be related with some capabilities or attributes of the RD, i.e. for characterizing the network interfaces, the buffer, the CPU, the memory, etc. Thus, all these Resources can be represented with semantics in order to provide more meta-information that will be used by networking/communication

techniques to take their decisions. In this respect, the RD's network Resources (that are described in the following two subsections in this section) can be semantically represented and accessed via a respective Service from the RERUM gateway or the MW (by invoking the Service of the respective VRD). The IoT Service has been modelled to provide a uniform abstraction for exposing the functionalities provided by such resources. The Entity, Resource and Service models are detailed in [D++12]. Detailed explanation of the important features and terms is presented in D2.3.

A VRD is modelled to have attributes that tie it to the domain (i.e. observable or actionable features), as well as type and identifier specifications. Also captured are optional temporal features and links to known vocabularies for specifying ownership. The VRD location could be defined in terms of a modelled WGS-84 Location concept (hasLatitude, hasLongitude, has Altitude) and could have properties that link to global (hasLocationURI) and indoor location (hasLocation) models. To specify the global location, an instantiation of the Entity Model could specify a URI from existing standards such as Geonames [GEON] that model well-known location aspects such as cities, districts, countries and universities. The indoor location model would provide detailed location description at a higher degree of granularity. This type of representation of the location of a VRD is useful in clustering mechanisms or in the technique for discovering and assigning resources presented in section 4.3. The VRD location represented in such way is basically used in the Smart transportation use case, but can also be considered on the environmental monitoring use case (when we install mobile sensors on buses), as well as on the indoor use cases for identifying in which room each RD is located.

A Resource model captures (i) different RERUM resource types (e.g. sensor, actuator, RFID tag), (ii) the geographic location of the hosting device and (iii) a link to the service model that exposes the resource capabilities.

The Service Model for a RERUM Service reflects the Service related aspects of IoT-A's Domain Model [W++11] as also described in D2.3. The Service Model contains information needed for discovering and looking up the Service as well as information on how to invoke the Service. The Service model exposes Resource functionalities in terms of the IOPE (input, output, precondition, effect) aspects. The type of the Service should specify the actual technology used to invoke the Service (e.g. OWL-S, REST etc.). Two important aspects to be revealed by Service Model are the service area and the service schedule. For sensing services, the service area would be the observed area, while actuating Services would specify the area of operation. The possibility of specifying time constraints on Service availability is captured through the service schedule feature.

A generic ontology to model and describe RERUM world may contain some relevant area of relevance as follows:

- **IoT Resources**: as the ones listed in SSN ontology [W3C11] and adding other relevant ones such as **Actuator**, **IoT Gateway** and **Server**.
- **IoT Services**: the scale and distributed nature of the IoT requires scalable and interoperable means for managing and accessing information pertaining to the physical world. With the service interfaces exposed by the IoT Resources, existing business applications and services need intelligence and context awareness to be easily integrated with the low level IoT services. An important consideration is to model IoT Services, or more generally, Services, independent of any particular service technology (e.g., SOAP and REST based services). The service modelling in RERUM is based on the **Profile-Model-Grounding** pattern, which is a variant of the Profile-Process-Grounding pattern. The idea is to reduce the complexity related to process modelling which results in heavyweight service description model. The process modelling is realised in the service composition phase.
- **Deployment, Systems and Platforms**: This module provides descriptions on how the IoT resources are organised and deployed as well as the system they form. Modelling and linking together these concepts enable a high-level view on relationships among the IoT resources and the systems and platforms that support them.

- **Observation and Measurement**: concepts in this module represent the information collected from the physical world by the IoT resources. Concepts related to Observation and Measurement from the SSN ontology [W3C11] and those from the Quantity Kind and Unit ontology are reused [QU-Ontology].
- **Physical Entity and Physical Locations**: Physical Entity (or Entity of Interest) represents an object in the physical world that is of interest to a user or application. Physical locations are associated with entity of interest and essential for IoT resource and service discovery.

Modelling methods on the IoT Resources, Physical Entity and Physical Locations, Deployment, Systems and Platforms, and Observation and Measurement have been extensively discussed in existing works such as [W3C11] and [D++11].

Semantic service modelling provides a machine interpretable framework for representing many aspects (e.g., functional, non-functional and transactional attributes) of services. The semantic Web service community has developed several models for semantically describing general Web services such as the Ontology Web Language for Services (OWL-S) [OWL] and Web Service Modelling Ontology [RKL++05]. The work in [D++11] proposed an "Entity-Device-Resource" model for representing IoT resources and services based on the SSN ontology [W3C11] and OWL-S ontology [OWL]. However, these ontologies are heavyweight and complex (e.g., the OWL-S ontology includes the complex process modelling) and not suitable for describing IoT services.

IoT Services exposed by IoT Resources have mostly limited computation capabilities and often operate in dynamic and constrained physical environments (e.g., network, platform and system); compared to the Web services, they are less reliable and stable; their logic is also much simpler and their output usually represents observation and measurement of features of interest of physical entities (therefore, service models have to be associated with IoT resources). Despite these characteristics, in a service oriented IoT, they also need to participate in service composition and the issues on effective service adaptation and compensation mechanisms become prominent. However, these issues are out of the scope of RERUM. It has to be noted here, that the RERUM ontology model, the Services and Resource model and the semantic service modelling will be included in the final version of the RERUM System Architecture deliverable (D2.5) that will be delivered end of August 2015.

4.1.2 Naming and addressing of devices

Naming and addressing of IoT devices are two challenging issues, mainly because of the vast number of devices that are expected to be globally interconnected and the fact that different platforms and networks shall cooperate in order to provide a seamless experience to the end-user. The lack of a standardized architecture for a naming, addressing and profile entity is a limiting factor for efficiently deploying global IoT services and applications. The design of a new addressing and naming scheme is out of the focus of RERUM. However, for completeness, here we present existing solutions that can be applicable for the techniques presented in sections 4.2 and 4.3.

Regarding the naming of IoT devices, key information such as the device meta-data, the device type and the domain are considered to be essential. On the other hand, the addressing format should consider including the necessary granularity of efficient accessibility and addressability to the physical infrastructure, while the profile services shall enable an easy application query and discovery as well as system configurations (e.g., device status) [LYL14].

The Internet Protocol (IP), either IPv4 or IPv6, is considered as possible solution for addressing of IoT devices. Nevertheless, the fact that most of the IoT devices are expected to be constrained devices impedes the application of the IP, pushing for other solutions. 6LoWPAN is the IETF working group for enabling IPv6 communication for low power devices with 802.15.4 networking capabilities. On the other hand, ETSI has proposed an overall service architecture for internetworking with 3GPP machine-

type communications, where each device has a unique International Mobile Subscriber Identity (IMSI) for signalling and charging and is IP addressable [3GPP11].

EPCglobal [EPCG] focuses on using RFID in information-rich networks (e.g., logistics) and employs an Object Name Service (ONS) that is similar to DNS and translates a 96-bit binary sequence to a Uniform Request Identifier (URI). The latter directs the name query to a set of databases in order to retrieve information about the devices. OLE for process control – Unified Architecture [SKS01] focuses on device addressing, which are interconnected in mesh topology.

Besides the industrial standards, naming and addressing has been widely investigated for wireless sensors networks, following either the approach of efficiently allocating addresses among nodes, or the approach to reuse spatially the assigned addresses and represented by codewords of variable length [LYL14], [FR++05, HS++01]. For example, the device addresses may be encoded using the Huffman encoding relating the address length with the energy consumption. In another approach [SL++10], the service client may request a service using a specific name without an explicit relation to the end-device that will serve the request, in order to allow the applications to seamlessly communicate with the end-device even when the mapping between naming and end-device changes during the service session. Distributed hash tables have been used [X++11] to improve the performance of identifier management with unique identifier generation.

In [LYL14], the addressing and naming problem is investigated as a service layer which is part of a middleware at the back-end data centre, trying to provide a homogeneous and unique naming and addressing convention, taking into account the heterogeneity of different platforms for IoT services. The overall architecture is depicted in Figure 42. It becomes obvious that the path from the physical world to the service layer is characterized by a high level of heterogeneity in terms of communication, transport and routing technologies. In order to deal with this issue, an IoT-application infrastructure is proposed in [LYL14], where new entities are introduced (application gateway, naming/addressing/profile server (NAPS), service registration portal and real-time operational database). Besides the design of the NAPS middleware, a unique device naming and addressing is given, which works with any existing system/platform to assist the upstream data collection and identification, content-based data filtering and matching, downstream control message delivery, as well application query. Indicative results have shown that a single node commodity server can achieve average throughput around 500–600 tps, with system response time 61 ms to search devices within five million devices.

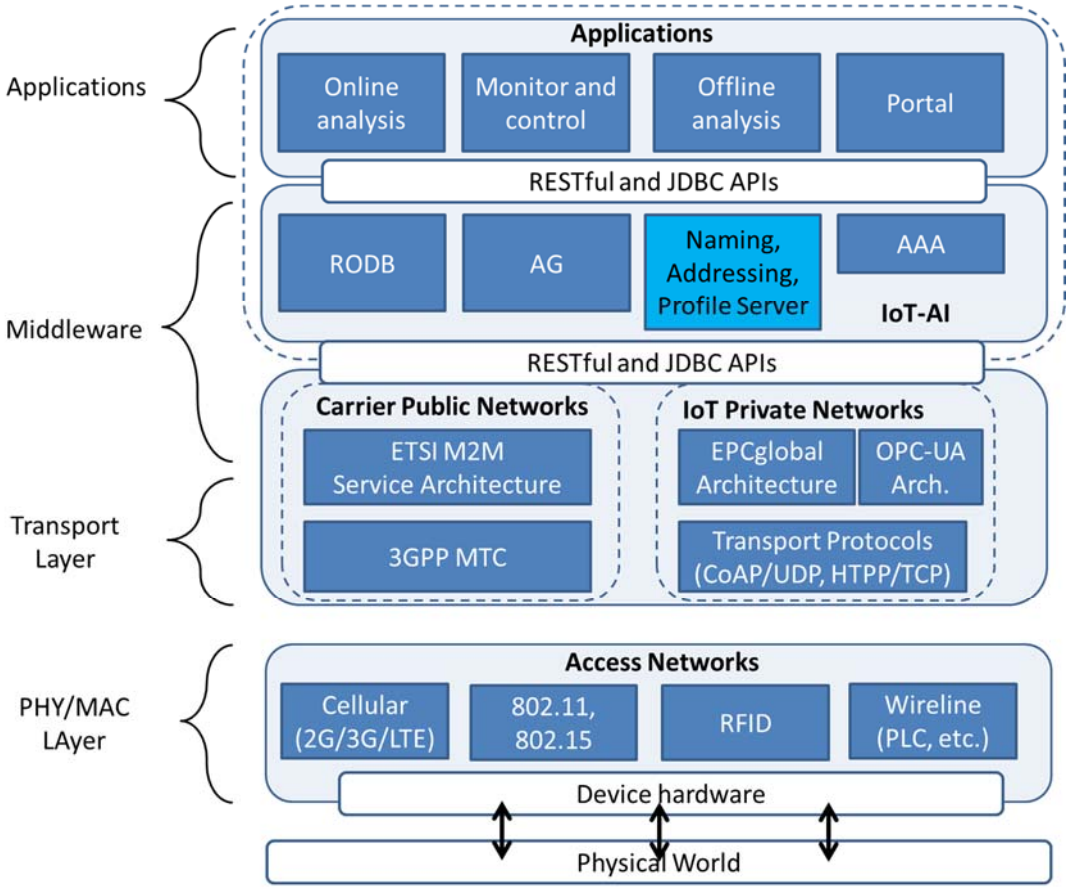

Figure 42: Naming/addressing architecture

4.2 Virtualisation modelling of interfaces and links[7]

4.2.1 Motivation and relation with the architecture and the use cases

In this section we address the problem of virtualisation of the RD networking interfaces and computational resources. We approach the problem from the perspective of a multiple resource allocation problem. Specifically, we consider that the set of resources utilized by any service become bound, regardless of the physical entity serving them, and thus create a virtual interface mapped to that service. Thus the work contained herein can be considered to be part of the "interface selection" of the communication manager (see Fig. 38), and more specifically the module of "IF Selection parameters" of the CR-Inspired agent (Fig. 45, D3.2). Furthermore, when the decisions are taken centrally (if the RDs are very resource constrained and can't run the optimization algorithm), this technique can also be assumed as part of the Spectrum Management module of the Networking Manager.

Consider for example an RD (as that of Figure 1 in D.2.2). We model its resources as a vector of independent quantities (e.g. uplink/downlink capacity per interface, available CPU cycles per CPU, memory pages and storage space). A service requesting to use the RD may end up utilizing for example the wi-fi interface for its uplink traffic, and split its downlink traffic over the wired IPv4 and 6LoWPANinterfaces, provided that the request is feasible (i.e. there is a sum or resources available). In our model we consider that **this resulting allocation of service to interfaces is the virtual interface**

[7] Part of this work has been accepted to be published at IEEE INFOCOM Student workshop 2015 [AAPY15]

utilized by the service. This process in the RD takes place in the Commmunication Interfaces Management (see Figure 3 D2.2), and the result is the generation of virtual of devices (interfaces) as per Figure 6. in D2.2.

The classical problem of bandwidth allocation in a congested network can be viewed as a special case of multi-resource allocation. Given a network of physical devices and its topology, we can view each link as a separate resource with a distinct capacity. Each service is represented by a network flow, which uses a predefined subset of links. In this special case, a service demands the same amount of bandwidth on each link. In general, multi-resource allocation cannot be turned into single-resource allocation by interchanging different resources. For example, if an RD needs both transmit buffer and bandwidth for each service, adding more buffer will not reduce the amount of required uplink throughput. Even similar resources such as bandwidth on different links are non exchangeable, as services may require bandwidth on each utilized link.

In general, multi-resource allocation problems, such as scheduling jobs, are often treated as a single-resource problem (e.g., see the Hadoop and Dryad schedulers [I++07]). A recent work [W++12] extends the max-min fairness measure to multiple-resource settings. Our work within RERUM develops a multi-objective mixed integer program aiming to minimize the number of physical interfaces contained in a virtual interface. The rationale is that involving multiple interfaces by splitting services across them will incur significant overheads in terms of coordination between the communication and network manager (see Figure 2).

4.2.2 Problem Setup and model

Our problem amounts to optimally allocating the demands of a feasible set of services to the RD network interfaces. We consider that we have a set $\mathcal{I}$ of I interfaces. The interfaces are characterized by a set $\mathcal{K}$ of K resources associated with them (for example {*CPU_cycles, Downlink capacity, Buffer_size*}). We assume that each service $j \in \mathcal{J}$ is associated a K-dimensioned demand integer vector $\boldsymbol{d_j}$, and that each interface has a likewise K-dimensioned resources availability integer vector $\boldsymbol{b_i}$. We consider the case in which services are flexible and can utilize resources of different interfaces, with appropriate costs to model the job management overheads that will be imposed upon the operating system of the device carrying the interfaces. We finally make the assumption that the given assignment is a feasible one, i.e. $\sum_{j\in\mathcal{J}} d_{jk} \leq \sum_{i\in\mathcal{I}} b_{ik}\ \forall k \in \mathcal{K}$. Our goal is to assign the services to the physical interfaces, such that the activation cost of the interfaces is minimized and all services are assigned. We call this the Service to Interface Assignment (SIA) problem.

In the model that follows we use variable x_{ijk} for the value of the k-th resource of the i-th interface utilized by job $j \in \mathcal{J}$. We Consider these values to be integer like the ones in the per-unit cost to activate resource $k \in \mathcal{K}$ on interface $i \in \mathcal{I}$ we denote c_{ik}, while with f_{ij} we denote the activation cost to assign service $j \in \mathcal{J}$ on interface $i \in \mathcal{I}$. We finally assume that each job $j \in \mathcal{J}$ utilizing the k-th resource of the i-th interface incurs an overhead denoted α_{ijk}. Thus, our model amounts to:

$$min \sum_{k\in\mathcal{K}} \sum_{i\in\mathcal{I}} c_{ik} \sum_{j\in\mathcal{J}} x_{ijk} + \sum_{i\in\mathcal{I}} \sum_{j\in\mathcal{J}} f_{ij} ACT_{ij} \tag{46}$$

$$s.t: \sum_{j\in\mathcal{J}} x_{ijk} = d_{jk}, \forall j \in \mathcal{J}, \forall k \in \mathcal{K} \tag{47}$$

$$\sum_{i\in\mathcal{I}} (1 + \alpha_{ijk}) x_{ijk} \leq b_{ik}, \forall i \in \mathcal{I}, \forall k \in \mathcal{K} \tag{48}$$

$$x_{ijk} \geq 0, \forall i \in \mathcal{I}, \forall j \in \mathcal{J}, \forall k \in \mathcal{K} \quad (49)$$

$$ACT_{ij} = \mathbb{I}\left(\sum_{k \in \mathcal{K}} \mathbb{I}(x_{ijk} > 0),\right) \forall i \in \mathcal{I}, \forall, j \in \mathcal{J}, \quad (50)$$

Where the objective of Equation (46) is to minimize the total cost of two terms: the first aims to capture the total cost (say power) incurred by the utilization of the resources over heterogeneous interfaces, the second term captures the cost introduced by splitting the service over multiple interfaces, since with each additional interface utilized the overall cost is encumbered by another f-term. The set of constraints in Equation (47) ensures that all services demands are met, while the constraints of Equation (48) ensures that the service allocation will be performed on interfaces with available resources. In Equation (50) the $\mathbb{I}$ symbol denotes the indication function becoming 1 if the argument is positive, 0 otherwise, thus, ACT_{ij} is one if and only if there is at least one resource utilizing at interface i for service j.

4.2.3 Complexity Proof

In this section we derive the proof on the NP-completeness of the problem discussed above. We base the proof on the construction of an instance of the problem from any instance of the Partitioning Problem. The Partitioning Problem (PP) amounts to determining if a set of $\mathcal{J}$ integers, of sum S can be partitioned into two subsets, each having a sum of $\frac{S}{2}$. The PP is a well-known NP Complete problem.

Assume that we have only one resource type on the interfaces available ($K = 1$). Take each element in the set of the PP to be a service of the SIA problem instance and let the value of each element be the resource demand d_j of the corresponding j-th service. Now let there be just two interfaces ($I = 2$), each with resource availability $b_i = \frac{S}{2}, \forall i \in \{1, 2\}$. We set the overhead coefficients to zero $\alpha_{ij} = 0\ \forall i \in \{1, 2\}, \forall j \in \{1, \dots, J\}$ and the resource activation cost to zero likewise $c_{ij} = 0$, while we fix the interface activation cost to one $f_{ij} = 1$.

This SIA instance constructed is feasible, because (i) by construction the total resource availability on the two interfaces suffices to serve the aggregate demand and (ii) splitting service demand on more than one interfaces is allowed.

Consider a solution of the constructed SIA instance where no splitting occurs. If such a solution exists, then each service is assigned to one interface and by Equation (46) the cost will be equal to $\mathcal{J}$. Furthermore, since any service split in two interfaces gives a cost of two, if splits exist in a solution the cost will be at least $\mathcal{J} + 1$. Hence, by the construction of the instance, it becomes obvious that no value lower than $\mathcal{J}$ can be achieved. The recognition version of the SIA instance is to answer whether or not there is a solution for which the cost is at most some value, in our case $\mathcal{J}$.

In any solution of SIA the service demand assigned on each interface will be $\frac{S}{2}$. If there is no split of services in a solution then their assignment to the two interfaces is a partitioning of the integers in the PP with equal sum. So, if the answer to the original instance of the PP is yes, then by assigning to the two interfaces the services mapping to the elements of the solution subsets, no split will exist and the cost will be $\mathcal{J}$. Thus the answer to the recognition version of the SIA instance is yes. Conversely, if the answer to the SIA is yes, then there cannot be a split service so the assignment of services to the two interfaces is a valid PP solution. Therefore, solving the constructed SIA instance is equivalent to solving an the arbitrary PP instance.

4.2.4 A greedy fast algorithm

Since the SIA problem has been shown to be computationally complex for large input sizes (with respect to the dimensions of the vectors of interfaces, services, and resources) we propose a modular greedy heuristic to enable fast solutions in such cases. The proposed algorithm comprises two parts. By its construction outlined below it becomes evident that the algorithm **(a)** does indeed provide a solution (i.e. it terminates), **(b)** does indeed take into account the problems' costs, striving to **(b.1)** keep the assignment cost low **(b2)** not split the services - or at least the demand of a service's single resource over more than one interfaces.

- Step 1: Transform demands into an ordered set

Since the demands as presented in the $\boldsymbol{d}$ matrix they may have wildly arbitrary values from a numerical point of view[8]one needs to map these values into something that can be directly ordered so that one can have a reasonable sequence for visiting the elements of $\boldsymbol{d_j}$.

Example

The simplest ordering that can be performed is to normalize the column vectors of $\boldsymbol{d}$. Thus, the matrix gets new $\boldsymbol{d}'$ real values in [0, 1]. In this case a slight problem is that each column will always have at least one '1' value. A simple fix is to further divide the new $\boldsymbol{d}'$ values by the respective lowest cost $c_k = c_{\{\mathrm{argmin_i}(c_{ik})\mathrm{k}\}}, \forall k$. This will further boost the potential for the most "bulky" demand to get a higher ranking value by considering of the "cheapest" possible interface it might be allocated. Alternatively one can simply throw a dice to order the equally ordered d'_{jk} values

Both methods discussed in the example will conclude transforming the matrix of demands into a (0, 1] elements matrix that is an ordered set.

- Step 2: allocate demands sequentially to the cheapest resource available

do

for the ranking above find the highest unserved demand for service j

find the resource with the lowest c_{ik}.

if $d_{jk} \leq b_{ik}$

allocate j at interface i

update $b_{ik} \coloneqq b_{ik} - d_{jk}$

else

sum the cost for allocating *fully* $\mathrm{d_{jk}}$ to a split interface of ordered higher costs $b_{ik} + b'_{ik} + b''_{ik}$(the b's follows the ordering of the respective c's here) and calculate into this cost also the appropriate f_{ij};

compare this allocation with the allocation of d_{jk} to b_{ik}, where b_{ik} is used to denote the cheapest possible resource availability that can accommodate d_{jk} without splitting of the two solutions (split/non-split) keep the cheaper one;

update the $\boldsymbol{b}$ vector accordingly;

remove service **;**

While there are demands still unserved.

[8] e.g. consider the resources modelled to be **{**downlink capacity, uplink capacity, buffer size**}** an instance for an LTE interface could then well be {100x10^6, 50x10^6, 1.3x10^6}, while for a 6lowpan these numbers are at least three orders of magnitude less for the two first, and close to four in the last one.

A slight twist is that one can instead of following the order created in step 1 by d_{jk}, to follow per d_j attempting to allocate directly the entire demand vector onto a single interface: this needs some thought for how to order the interfaces.

4.2.5 Performance & Results

In what follows we provide results for the cost of activation of 3 to 10 services with different demand vectors on an RD of having three interfaces. The interfaces resource vectors have been taken to be mapping to a 6LoWPAN, a Wi-Fi, and an LTE interface. The services come in 3three classes of demands that map to a low rate/low buffer (e.g. mapping to the traffic use case data collection), a low rate high buffer (e.g. a highly compact streaming service (e.g. one enabled by CS), and a high rate low buffer (e.g. video streaming).

The first figure (Figure 43) displays the cost behaviour vs the number of services running on the RD. We present cost behaviour for different activation cost modes, which can be applied in dynamic cases where the interfaces need to be reserved or when power constraints of a battery-operated RD impose preference on the activation of specific interfaces.

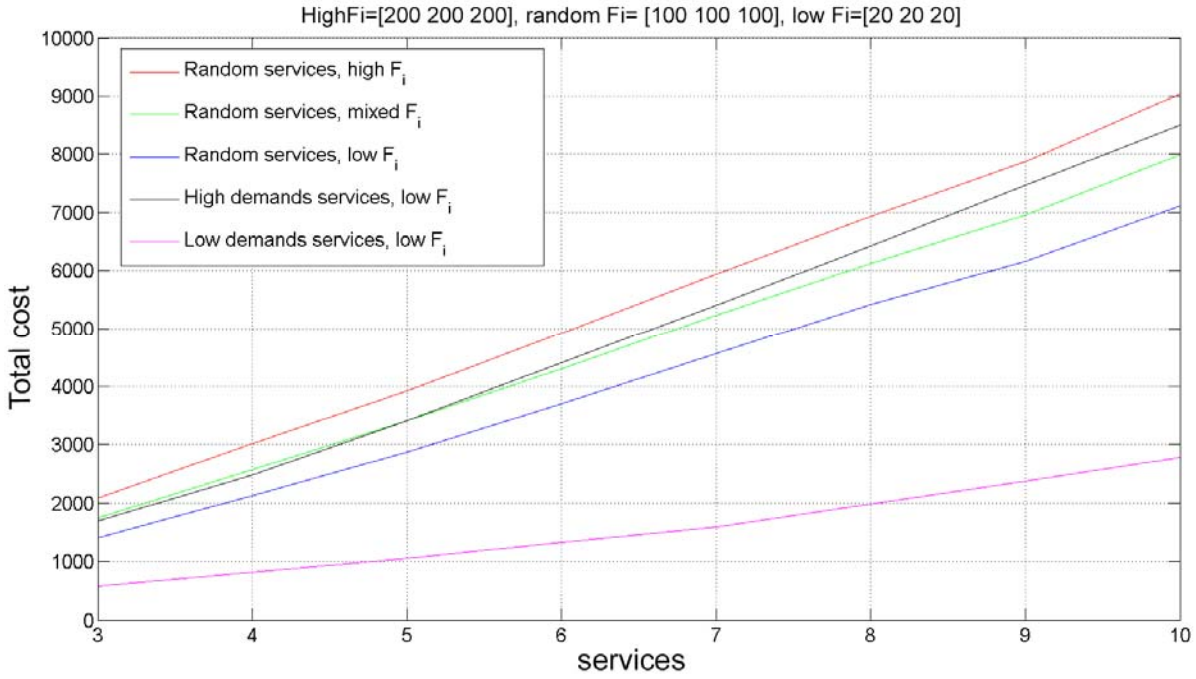

Figure 43. Cost behaviour vs number of services

In Figure 44 we display the number of total splits across physical interfaces incurred in the solution. This mean that any bar presented here implies the number of physical interfaces contained in the virtual interfaces arising from the solution. Notice that for low activation cost, when the demand vector remains feasible, the virtual interfaces are not comprising of more than one physical interfaces. In fact each virtual interface is merely a part of a physical RD interface.

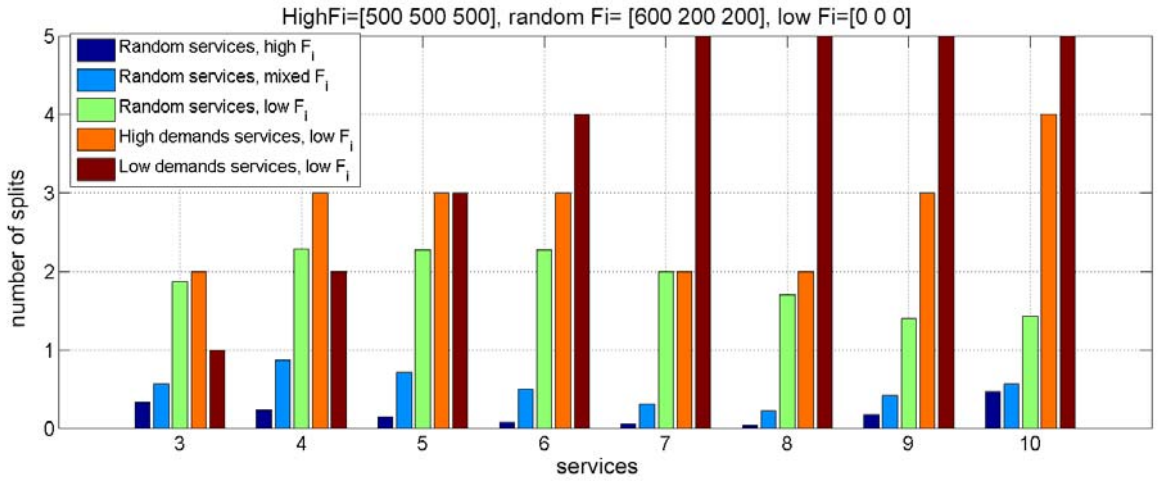

Figure 44. Total number of service splits between multiple interfaces

The plots presented in the above figures reflect the fact that the higher the activation cost (Fi) is, the more expensive it is to split a service to multiple interfaces than to keep it in one interface (as expected). The optimal costs for different set of services are depicted in Figure 43. If the activation cost is much higher than the power cost of the interfaces, then the optimal total cost is higher in the case of services of random requirements than in the case of high demands services (with low activation cost). Using a mixed activation cost yields an optimal cost close to that of high demands services, because the latter causes more splits (see Figure 44) while the former has at most one job split on average. When the activation cost (Fi) is high in comparison to the power cost of the interfaces, less splits of services among interfaces happen. For instance, in Figure 44 services of random requirements with high activation cost split at most once on average. On the other hand, when the activation cost is low (compared to the power cost) more splits occur. Low requirements services split more than any other case to exploit the inexpensive (with regard to activation cost) interfaces. More splits happen as the number of such services increase. However, this is not the case regarding high demands services. More than six of such services causes a split less. The increase of splits for the case of nine and ten services can be attributed to the chosen interfaces' capacities - a split more was necessary to accommodate the demands of the extra served jobs.

4.3 Algorithms for network Resource discovery and allocation[9]

4.3.1 Motivation

The seamless and efficient interconnectivity of a large number of connected devices is a fundamental requirement for IoT applications. The access networks, either wired or wireless, have to be adapted to these new circumstances and conditions in order to be able to support those devices under specific QoS requirements. In this context, spectral efficiency is a critical research topic especially for wireless access network technologies. In this section, we face this challenge and introduce a novel network Resource allocation scheme, which aims at increasing the spectral efficiency of hybrid cellular/WLAN

[9] Parts of this work were presented in the IEEE CAMAD 2015 conference held in Athens, Greece, December 2015 [BLKS14].

access networks and hence the number of connected devices within a specific area that is served by a cellular network and various WLAN access points.

The proposed scheme introduces many benefits for the smart cities applications, since it is an enabler for increasing the spectral efficiency and hence the number of the end-users that can simultaneously be connected receiving an adequate QoS. The benefit of utilizing such a technique is obvious for the network operators that can increase their revenue because of the increased number of served subscribers. On the other hand, the citizens will experience lower costs due to the exploitation of the free unlicensed bands instead of the costly cellular networks, as well as improved network connectivity in both terms of attainable bit rates and reliability, ensuring that whenever a device has to send data (for addressing a Service request) it will indeed have adequate connectivity and will be able to send them on time.

Considering the RERUM smart transportation use case, the RDs are mostly mobile devices (e.g., devices on buses or subscribers' smartphones) participating in a crowdsourcing application. These devices have both wireless and cellular radio interfaces and will be able to utilize either their cellular radio resources or WiFi radio resources, whenever available, in an optimum way. The goal is to minimize interference, which will also result to maximizing the number of users that have access to the same radio resources. Furthermore, even in indoor use cases, the devices may utilize cellular radio resources in the case that the WiFi is not available. In this way the service reliability is increased, with the minimum cost on the performance of every user or device that is served in the same area. Throughout this subsection, the term "resources" refers to "radio or network resources" that characterize the wireless spectrum. Each RD gathers its own statistics for the radio resources and transmits them to a central entity that takes the decisions about the optimal allocation of resources to each RD.

4.3.2 Resource discovery and allocation

The European FP7 project ALICANTE [X++11] developed a system architecture for content-aware networks, which operates over multi-domain IP by creating unicast or multicast Virtualised Content-Aware Networks (VCAN)s. The network Resources are provided by the network providers and are managed quasi-statically or dynamically based on service layer agreements and specifications. Metadata describing the service/content are allowed to be inserted to the data plane by the content servers, which then can be identified and processed by the intelligent devices (e.g., routers) of the VCANs [PBZ13].

A Topology and Network Resource Discovery Protocol (TNDRP) is proposed in [PBZ13], which runs between the CAN Managers, to collect inter-domain and abstracted intra-domain information, in order to enable the VCAN mapping algorithms onto network resources. As depicted in Figure 45, [PBZ13], the architecture includes three layers: (i) the core and access network layer, (ii) the CAN management layer and (iii) the service layer. It supports the co-existence of several different network providers with their core and access networks domains, which can be seen as autonomous systems as well. The CANs are managed by the CAN managers (CAN providers), which exist for each IP domain in order to ensure the consistency of VCAN planning, provisioning and advertisement. Each domain has an Intra-domain Network Resource Manager (Intra-NRM), as the ultimate authority configuring the network nodes. This layer cooperates with some local entities called Home Boxes (HB) and SE by offering them CAN services [BI10, MB13].

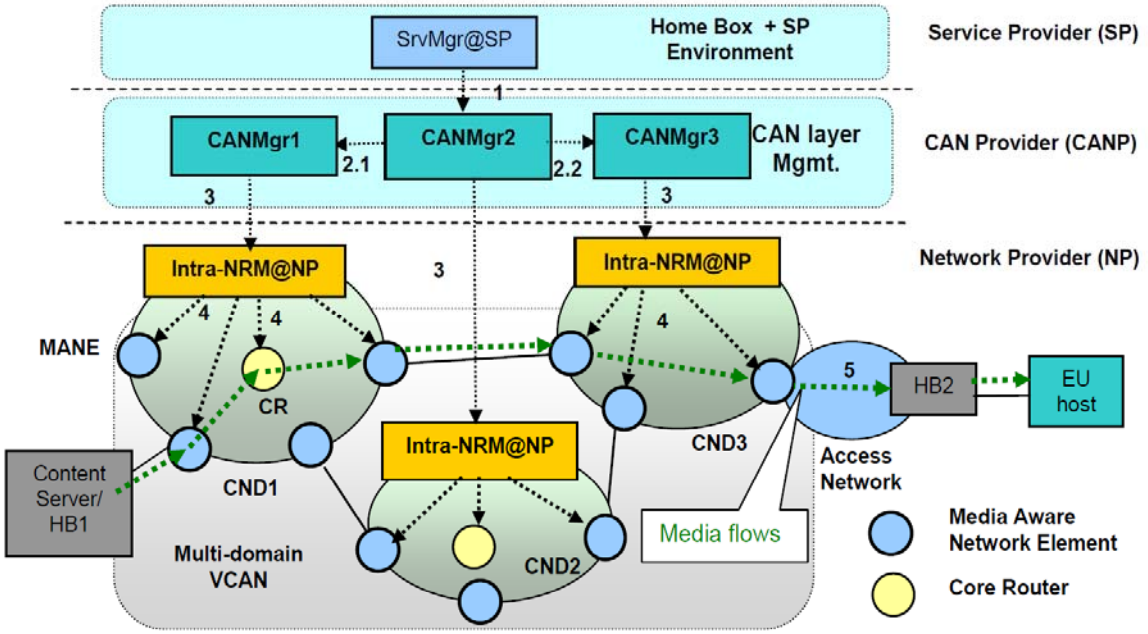

Figure 45. Example of network infrastructure (three domains) and management actions.

Due to the increase of the number of connected devices (e.g., IoT and M2M applications), the demand for wireless data traffic is growing faster than the capacity provided by the network. In addition, the link efficiency has reached its performance limits, while the increase of spectrum seems to be impractical. In this context, the adoption of suitable network architectures and the efficient allocation of resources to devices is of critical importance for the sustainability of future networks [BLKS14].

It is widely accepted that further spectral efficiency improvement is more likely possible by increasing the node deployment density [MW++11]. Towards denser network deployments, two major approaches exist: a) network splitting into smaller macro-cells, and b) adoption of Heterogeneous Networks (HetNet)s concept. Cell splitting may not always be an efficient solution, especially in already dense deployments, where the additional inter-cell interference is prohibitive [MW++11]. The concept of HetNets, relies on the deployment of heterogeneous low power nodes (LPN)s within the macro-cell [DBZ09, SMY12]. The HetNet deployments provide a wide area coverage through the macro cell and a more targeted coverage of special zones through the LPNs [LF12, SA12]. The HetNet concept is also part of the 3rd Generation Partnership Project (3GPP) Long Term Evolution (LTE) network architecture, where the LPNs include femtocells, home eNodeBs (eNB)s and relay nodes [3GPP10].

HetNets provide many advantages, including the increased attainable downlink transmission rates, mainly due to the decreased distance between the served nodes and the serving stations and the diversity that is gained due to the utilization of more than one different access technologies. Nevertheless, the deployment of HetNets may not always be advantageous, unless a specific strategy for sharing the resources between the served nodes and the serving stations is followed. Interference may become extremely severe and the results may be the opposite form the desired ones. Because of the utilization of different access technologies, conventional interference mitigation techniques cannot be employed. On the contrary, new mechanisms must be developed (the proposed strategy will be presented in subsection 4.3.3.), which should take into consideration the cost (in terms of interference) and the gain (in terms of bit rate) of the joint HetNet that is formed.

In what follows, a **hybrid cellular/WLAN communication is proposed** [BLKS14], where the mobile users can be served by either the eNB or a WLAN AP, depending on the selection strategy (e.g., the signal to interference plus noise ratio (SINR)). In contrast to the conventional offloading approach,

where the WLAN APs have a wired backhaul, a tethering approach is adopted. According to this scheme, the WLAN APs are wirelessly connected to the eNB and share this broadband connection with specific users over WLAN frequencies. The users select their serving node, i.e., the macro-cell eNB or a WLAN AP, based on a performance criterion. The aim of this architecture is to reduce the transmission power from the eNB to users with bad channel conditions (e.g., cell edge users) and thus the interference at both the cellular and WLAN part of the hybrid network, while avoiding modifications to the existing cellular network.

The downlink of a multi-cell wireless network is considered [3GPP06], where the macro-cell coexists with a number of WLAN access points (see Figure 46 and Table 3 for the notations). These access points are connected to the core network of the network operator, according to the 3GPP TS23.234 WLAN Internetworking specifications (I-WLAN), which defines two new procedures (the overall architecture is illustrated in Figure 47):

- **WLAN 3GPP IP Access**, which allows WLAN UEs to establish connectivity with External IP networks, such as 3G operator networks, corporate Intranets or the Internet via the 3GPP system.
- **WLAN Access, Authentication and Authorisation**, which provides for access to the WLAN and the locally connected IP network (e.g. Internet) to be authenticated and authorised through the 3GPP System. Access to a locally connected IP network from the WLAN, is referred to as WLAN Direct IP Access.

I-WLAN enables the seamless interconnectivity between 3GPP-based cellular networks (e.g., LTE or GPRS) and the WLAN and the ability to have a joint resource management for both networks at the core layer through signalling between the BRAS/AC and the RNC.

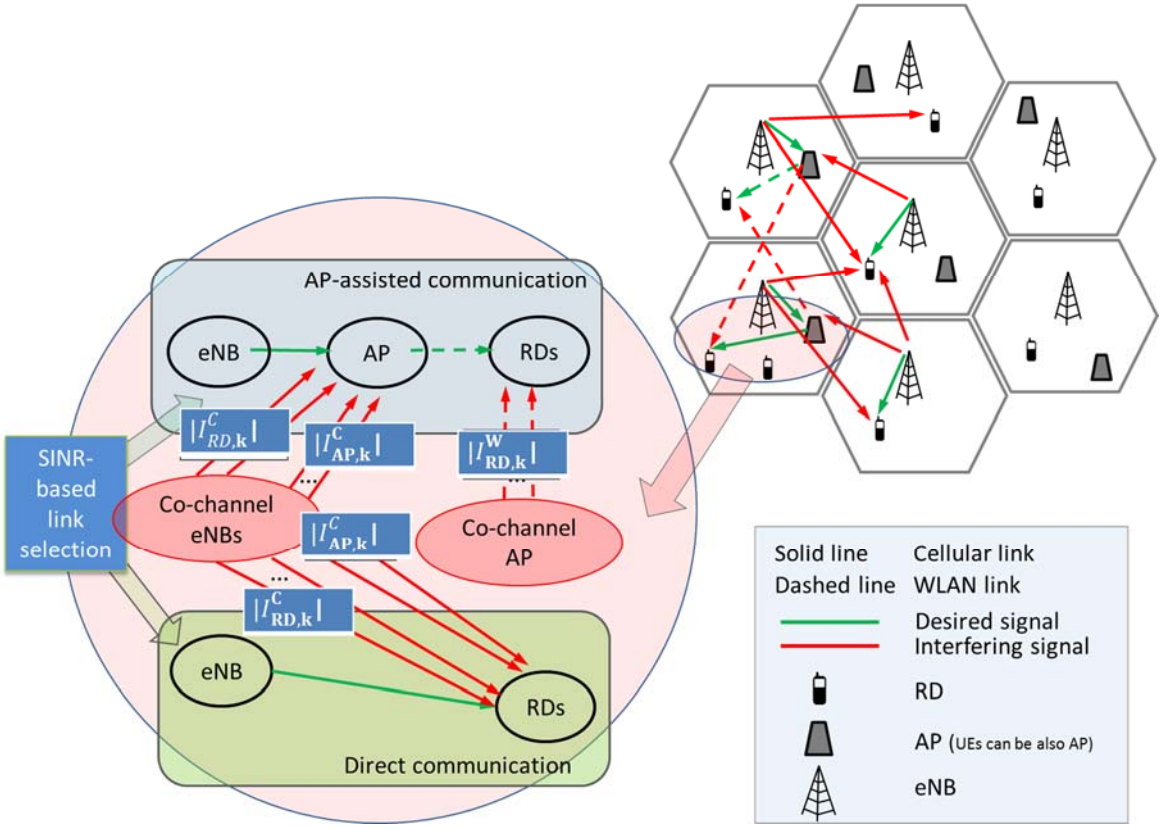

Figure 46. An overview of the proposed multi radio access scheme.

Due to the frequency reuse pattern among the macrocells, the User Equipments (UE)s and the APs in each macrocell are subject to co-channel interference due to the number of nodes present. Depending on the communication strategy (e.g., based on the received SINR), each RD might be served either through:

- a direct communication link to the eNB (one phase communication, eNB →RD) utilizing cellular frequencies, or
- an indirect link to the eNB via a WLAN AP, which decodes and forwards the received signal to the target UE (two-phase communication, eNB→AP→RD), utilizing cellular frequencies for the eNB→AP link and WLAN frequencies for the AP→RD link.

Table 3: Notations for the multi-radio access scheme.

$z_{rd_{j,i}}$	Complex channel gain between the j[th] eNB and the i[th] RD
$z_{ea_{j,i}}$	Complex channel gain between the j[th] eNB and the i[th] AP
$z_{au_{j,i}}$	Complex channel gain between the j[th] AP and the i[th] UE
$s_{rd_{j,i}}$	The complex symbol transmitted by the j[th] eNB, targeting the i[th] RD
$s_{ea_{j,i}}$	The complex symbol transmitted by the j[th] eNB, targeting the i[th] AP
$s_{au_{j,i}}$	Complex channel gain between the j[th] eNB and the i[th] AP
$I^{C}_{RD,i}$	The set of interfering RDs using the same cellular frequency as the i[th] UE or AP
$I^{C}_{AP,i}$	The set of interfering APs using the same cellular frequency as the i[th] UE or AP
$I^{W}_{RD,i}$	The set of interfering RDs using the same WLAN frequency as the i[th] RD
$Z_{Ird_{j,i}}$	Complex channel gain of the j[t]h RD ϵ $I^{C}_{RD,i}$
$Z_{Iau_{j,i}}$	Complex channel gain of the j[th] AP ϵ $I^{C}_{AP,i}$
$Z_{Iea_{j,i}}$	Complex channel gain of the j[th] RD ϵ $I^{W}_{UE,i}$
S_{Ird_j}	The complex symbol transmitted by an eNB, targeting the j[th] RD ϵ $I^{C}_{UE,i}$
S_{Iea_j}	The complex symbol transmitted by an eNB, targeting the j[th] AP ϵ $I^{C}_{AP,i}$
S_{Iau_j}	The complex symbol transmitted by an AP, targeting the j[th] RD ϵ $I^{W}_{UE,i}$
$P_{rd_{j,i}}$	The power of the complex symbol transmitted by the j[th] eNB, targeting the i[th] RD $P_{rd,j,i} = E < \lvert s_{rdj,i}\rvert^2 >$
$P_{ea_{j,i}}$	The power of the complex symbol transmitted by the j[th] eNB, targeting the i[th] AP $P_{ea,j,i} = E < \lvert s_{eaj,i}\rvert^2 >$

$P_{au_{j,i}}$	The power of the complex symbol transmitted by the j[th] AP, targeting the I[th] RD $P_{au,j,i} = E < \lvert s_{auj,i} \rvert^2 >$
$P_{rd_{j,i}}$	The power of the complex symbol transmitted by an eNB, targeting the j[th] RD $\epsilon\ I_{RD,i}^{C} P_{rd,j,i} = E < \lvert s_{rdj,i} \rvert^2 >$

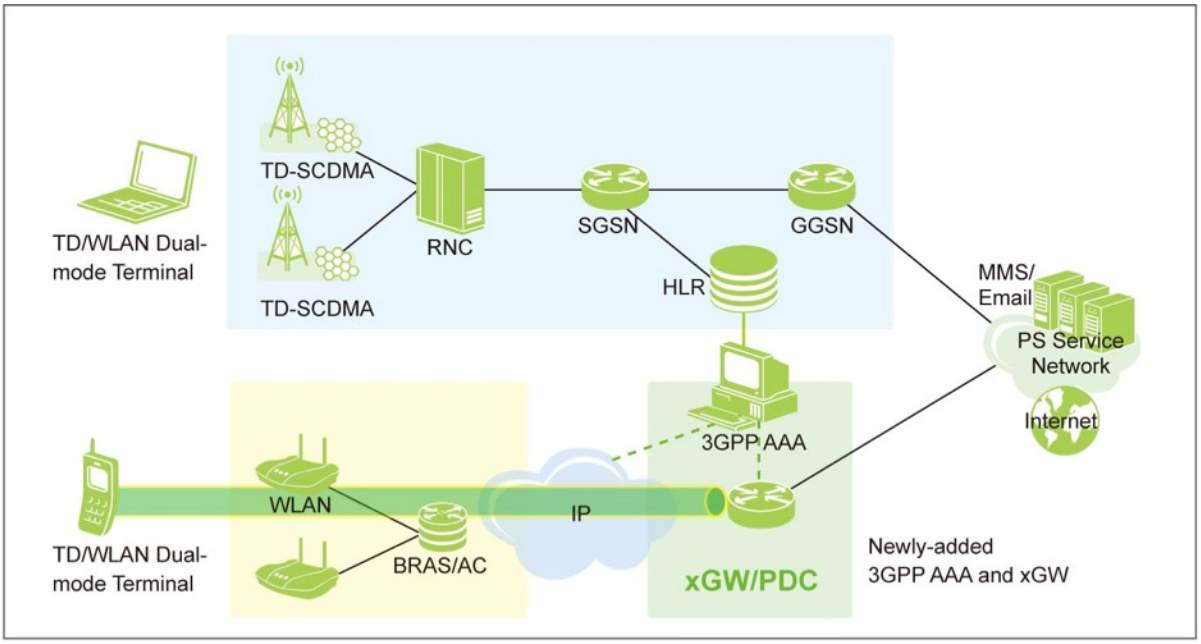

Figure 47. 3GPP I-WLAN Architecture.

The RDs that are directly connected to the cellular network will interfere with other RDs that are also connected to the same cellular network, since adjacent cells reuse frequencies. Similarly, in the indirect scenario where the RDs are connected to a WLAN AP, the APs will experience interference from adjacent APserving other RDs and APs at the same cellular frequencies, while the RDs will interfere with other APs that serve RDs at the same WLAN frequencies. Hence, the complex baseband signals transmitted by the nth eNB and received by the k[th] RD will be

$$y_{rd_{n,k}} = z_{rd_{n,k}} s_{rd_{n,k}} + \sum_{i=1}^{|I_{RD,k}^{C}|} z_{I_{rd,i}} s_{I_{rd,i}} + \sum_{i=1}^{|I_{AP,k}^{C}|} z_{I_{ea,i}} s_{I_{ea,i}} + w_{rd,k} \quad (51)$$

The complex baseband signal transmitted by the m[th] AP and received by the k[th] RD is given by

$$y_{au_{m,k}} = z_{au_{m,k}} s_{au_{m,k}} + \sum_{i=1}^{|I_{RD,k}^{W}|} z_{I_{au,i}} s_{I_{au,i}} + w_{au,k} \quad (52)$$

and that transmitted by the n[th] eNB and received by the k[th] AP

$$y_{ea_{n,k}}=z_{ea_{n,k}}s_{ea_{n,k}}+\sum_{i=1}^{|I_{RD,k}^{C}|}z_{I_{rd,i}}s_{I_{rd,i}}+\sum_{i=1}^{|I_{AP,k}^{C}|}z_{I_{ea,i}}s_{I_{ea,i}}+w_{ea,k} \quad (53)$$

where $w_{rd,k}$, $w_{au,k}$ and $w_{ea,k}$ denote the complex Additive White Gaussian noise (AWGN) with zero mean and variance N_0 at the corresponding RDs and APs. It is noted that the channel gains of both desired and interfering signals are related to the distance as well as the propagation path loss between the transmitter and the receiver. It is assumed that the envelopes of channel complex gains follow the Rayleigh distribution and hence their squares are exponentially distributed.

Compared to a conventional cellular network where the instantaneous signal-to-interference-and-noise-ratio (SINR) of the k[th] RD served by the n[th] eNB equals to

$$\gamma_{c,k}=\frac{\gamma_{rd_{n,k}}}{\underbrace{1+\sum_{i=1,i\neq k}^{|I_{RD,k}^{C}|}\gamma_{I_{rd,i}}}_{eNB\rightarrow RDs}} \quad (54)$$

where $\gamma_{rd,n,k}$ is the instantaneous received signal-to-noise ratio (SNR) of the kth RD and $\gamma_{Ird,i}$ is the interference-to-noise ratio (INR) due to the cellular co-channel interferers, which follows the exponential distribution, the SINR in the proposed hybrid resource allocation scheme is given by

$$\gamma_{rd,k}=\frac{\gamma_{rd_{n,k}}}{1+\underbrace{\sum_{i=1}^{|I_{AP,k}^{C}|}\gamma_{I_{ea,i}}}_{eNB\rightarrow AP}+\underbrace{\sum_{i=1}^{|I_{RD,k}^{C}|}\gamma_{I_{rd,i}}}_{eNB\rightarrow RD}} \quad (55)$$

for the direct connection scenario and

$$\gamma_{ea,k}=\frac{\gamma_{ea_{n,m}}}{1+\underbrace{\sum_{i=1}^{|I_{AP,k}^{C}|}\gamma_{I_{ea,i}}}_{eNB\rightarrow AP}+\underbrace{\sum_{i=1}^{|I_{RD,k}^{C}|}\gamma_{I_{rd,i}}}_{eNB\rightarrow RD}} \quad (56)$$

for the indirect scenario, where $\gamma_{Iea,i}$ is the INR due to the cellular co-channel interferers, which follows the exponential distribution and $\gamma_{ea,n,m}$ is the instantaneous SNR at the m[th] AP, following the exponential distribution. For the indirect scenario, the desired signal from one AP and interfering signals coming from the mth APs, the instantaneous SINR at the k[th] RD can be expressed as

$$\gamma_{au,k}=\frac{\gamma_{au_{n,m}}}{1+\underbrace{\sum_{i=1}^{|I_{UE,k}^{W}|}\gamma_{I_{au,i}}}_{AP\rightarrow UE}} \quad (57)$$

4.3.3 Resource allocation strategy

The communication mode of each RD (i.e., direct communication with the eNB or via an WLAN AP) is determined by the eNB, which has a full knowledge of the channel state information for all the links within the cell. In this work, we assume that this decision is based on the SINR, namely the UE connects to that node (eNB or WLAN AP) that results in the maximum individual instantaneous SINR. Signalling details for realizing this scenario is beyond the scope of this proposal.

The algorithm that is employed takes into consideration the following input:

- The number of UEs, N, the number of WLAN APs, M.
- The number of UEs, Q, for the optimal topology.
- The channel gains between the eNB and the UEs.
- The channel gains between the eNB and the APs.
- The channel gains between the APs and the UEs.

At the initialization stage, the algorithm finds the optimal topology that maximizes the overall SINR. Then continuing to the Q + 1 RD, the algorithm looks for the AP that this UE will be connected to, searching among the M APs, in order to maximize the overall SINR given that the previous Q RDs cannot change their AP. Similarly, considering that the previous Q+1 UEs cannot change their APs, the algorithm continues to the Q + 2 RD and searches for that AP that the UE will be connected to, in order to maximize the overall SINR. The high-level overview of the resource allocation scheme is illustrated in Figure 48.

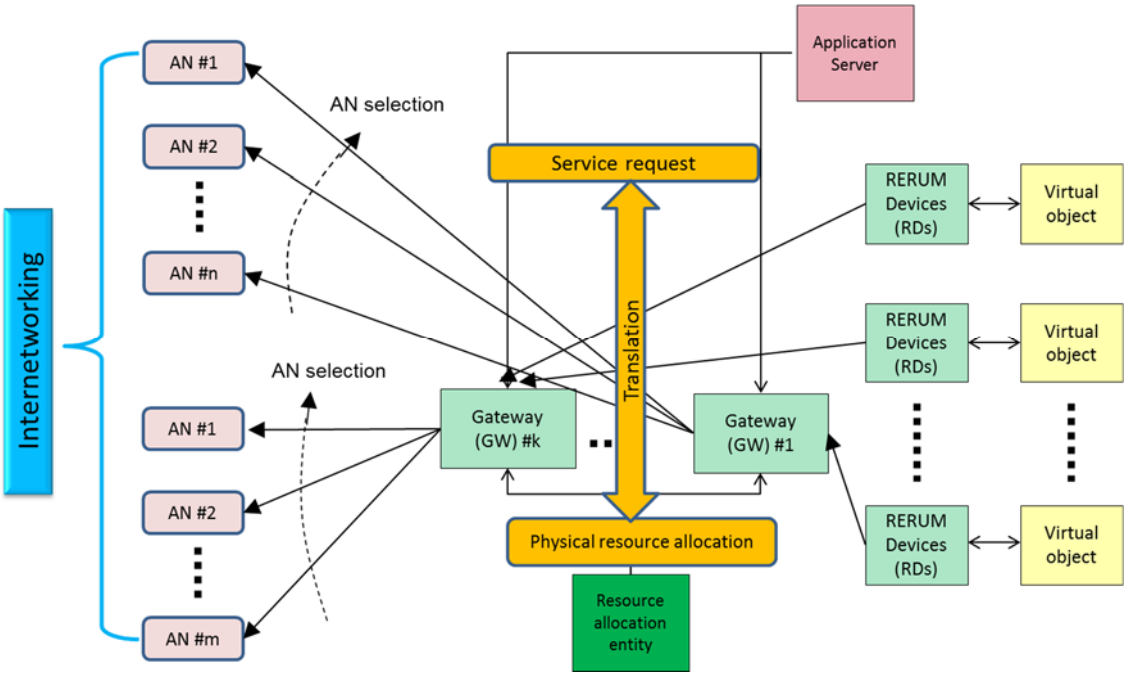

Figure 48. Mode of operation of the resource allocation algorithm and mapping to RERUM deployment.

4.3.4 Evaluation of the proposed scheme

In this subsection, the evaluation of the proposed algorithm is provided in order to illustrate how it can benefit a RERUM deployment in two ways:

- By enabling the RDs to utilize simultaneously a WLAN and a cellular access interface
- By minimizing the total interference of this hybrid network deployment and hence increase the number of RDs that can be served simultaneously without decreasing the performance of the overall network.

The evaluation of the proposed scheme is done both theoretically (i.e., through the derivation of closed form expressions for the total interference) and via computer simulations. The metrics that are used

include the SINR gains compared to a single access network scheme. The number of APs and the number of devices, serve as the parameters against which the evaluation is done.

The total output end-to-end received SINR can be expressed as

$$\gamma_{tot,k}(\gamma) = \max(\gamma_{rd,k}, \gamma_{eau,k}) \tag{58}$$

Hence, the CDF and PDF of $\gamma_{tot,k}(\gamma)$ will be given by

$$F\gamma_{tot,k}(\gamma) = F\gamma_{eu,k}(\gamma) F\gamma_{eau,k}(\gamma) \tag{59}$$

and

$$f\gamma_{tot,k}(\gamma) = f\gamma_{eu,k}(\gamma) F\gamma_{eau,k}(\gamma) + F\gamma_{eu,k}(\gamma) f\gamma_{eau,k}(\gamma) \tag{60}$$

where the form of corresponding CDF and PDF of the individual SINRs are given in the next equations:

$$f_C(\gamma) = \frac{\frac{1}{Y}}{Z_{2,k}^{X_2}\Gamma(X_2)} \frac{\exp(-\gamma/Y)}{Z_{1,k}^{X_1}\Gamma(X_1)} \sum_{j=0}^{X_2-1} \frac{\binom{X_2-1}{j}(-1)^j \Gamma(X_1+j)}{\left(\frac{1}{Z_{1,k}} - \frac{1}{Z_{1,k}}\right)^{j+X_1}} \left\{ \frac{\Gamma(X_2-j)}{\left(\frac{\gamma}{Y} + \frac{1}{Z_{1,k}}\right)^{X_2-j}} \times \left(1 + \frac{X_2-j}{\frac{\gamma}{Y} + \frac{1}{Z_{1,k}}}\right) - \sum_{k=0}^{X_1+j-1} \frac{\Gamma(X_2+k-j)}{k!} \frac{\left(\frac{1}{Z_{1,k}} - \frac{1}{Z_{1,k}}\right)^k}{\left(\frac{\gamma}{Y} + \frac{1}{Z_{1,k}}\right)^{X_2+k-j}} \left(1 + \frac{X_2+k-j}{\frac{\gamma}{Y} + \frac{1}{Z_{1,k}}}\right) \right\}$$

$$F_C(\gamma) = \frac{1}{Z_{1,k}^{X_2}\Gamma(X_2)} \frac{1}{Z_{1,k}^{X_1}\Gamma(X_1)} \sum_{j=0}^{X_2-1} \frac{(-1)^j \binom{X_2-1}{j} \Gamma(X_1+j)}{(1/Z_{1,k} - 1/Z_{2,k})^{X_1+j}} \left\{ \left[\frac{\Gamma(X_2-j)}{Z_{2,k}^{j-X_2}} - \sum_{n=0}^{X_1+j-1} \frac{(1/Z_{1,k} - 1/Z_{2,k})^n}{n!} \frac{\Gamma(X_2+n-j)}{Z_{1,k}^{j-X_2-n}} \right] - \exp\left(-\frac{\gamma}{Y}\right) \left[\frac{\Gamma(X_2-j)}{(\gamma/Y + 1/Z_{2,k})^{X_2-j}} - \sum_{n=0}^{X_1+j-1} \frac{(1/Z_{1,k} - 1/Z_{1,k})^n}{(\gamma/Z_{1,k} + 1/Z_{1,k})^{X_2+n-j}} \frac{\Gamma(X_2-j+n)}{n!} \right] \right\}$$

Considering the multi-cell MU case, the SINR gain of the hybrid Cellular/WLAN scheme over the baseline conventional cellular is given in Figure 49 as a function of the number of co-channel interference. In the conventional cellular scheme, the UEs connect to the eNB that maximizes the overall SINR. We observe the gain increases as the number of interferers increases until it reaches a ceiling. The greedy solution is close to the optimum one obtained via exhaustive search, especially for a small number of APs.

In Figure 50, the SINR gain of the proposed architecture over the cellular scheme is illustrated, considering both the optimal solution (i.e., through the exhaustive search) and the sub-optimal one provided by the greedy algorithm (with Q=2). As observed the overall cell's SINR may quadruple compared to the cellular scheme, while the gain diminishes as the number of interferers increases.

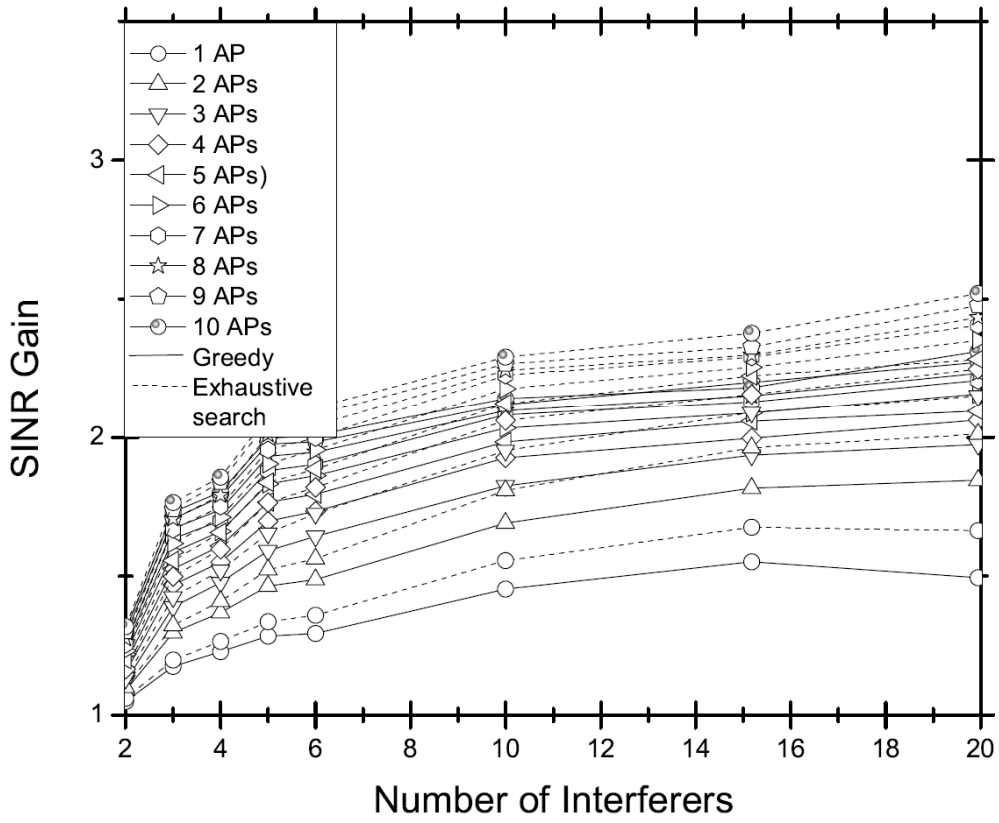

Figure 49. The SINR gain of the proposed Cellular/WLAN scheme over the cellular (MU multi-cell scenario).

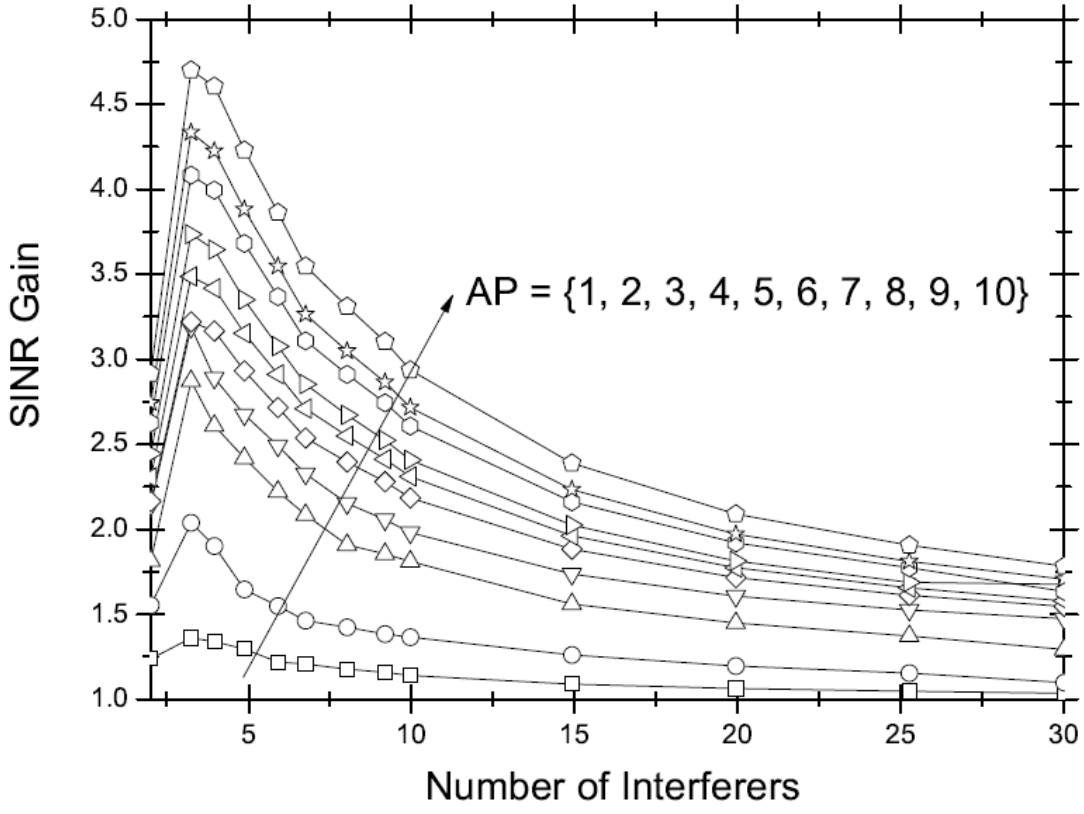

Figure 50. The SINR gain of the proposed scheme over the cellular (single cell)

<this page is intentionally left blank>

5 Mechanisms for optimizing the use of virtual resources

Section 4 of this deliverable presented a host of techniques developed by RERUM that enable the virtualisation of RDs in order to hide heterogeneity and allow discovery and assignment of their networking and physical capability in an abstract fashion. Those virtualisation techniques rely on communication over a network in the underlying layers and one of the key networking components is routing. In order to achieve RERUM's goals of security, reliability and low energy consumption we have developed novel techniques for unicast as well as multicast routing. The motivation for developing those techniques, their design and relationship with RERUM's use-cases is discussed in this section.

More specifically, Section 5.1 focuses on unicast routing and more specifically on trusted routing overlays, a set of algorithms which assures that data are sent from RDs to their respective gateway through a route of trusted intermediate nodes. This ensures that data will be neither overheard nor altered by malicious nodes. Our trusted routing overlays are based on a scheme that considers the following: (i) behavioural statistics of sensors reported by their one-hop neighbours, (ii) link quality statistics, and (iii) a fusion scheme based on the Dempster Shafer theory of evidence. These are described in detailed in sections 5.1.4, 5.1.5 and 5.1.7 respectively.

Subsequently, in Section 5.2, we shift our focus to multicast routing, by presenting a RERUM-developed algorithm that allows point-to-multipoint communication (from one sender to multiple RDs) in an energy-efficient fashion. This algorithm has been designed and implemented in a way that makes it suitable for severely constrained devices in terms of processing power and networking capability. Section 5.2.3 presents the design of our multicast forwarding algorithm and Section 5.2.4 discusses the details of its implementation for the Contiki Operating System.

5.1 Trusted routing overlays

5.1.1 Introduction and relationship with the RERUM use cases

The Internet of Things aims to interconnect numerous heterogeneous devices through virtualisation for concealing the underlay networking technology that is used for providing connectivity to the devices. In the same way, for concurrently providing many applications on top of a shared physical environment, virtualisation can create different virtual topologies for each one of the supported applications. As presented in [IHLH12] that focuses on virtualisation of wireless sensor networks (which are the main subset of IoT) there are differences between the terms Virtual Sensor Network (VSN) and Overlay Sensor Network.

According to [IHLH12] a Virtual Sensor Network is formed by an appropriate subset of sensor nodes of a standard Wireless Sensor Network (WSN). This subset collaborates in order to provide a specific service as requested by an application. A virtual sensor network is basically formed by interconnecting logically (and not always physically) collaborative sensor nodes. An Overlay Sensor Network is a network that is built on top of a standard Wireless Sensor Network in a similar way that the Internet was built on top of a telephone network. The overlay network creates a virtual topology comprising nodes that are connected with virtual links that correspond to paths of the underlying network. Overlay networks are usually formed at the application layer and are related to a specific application. As it can be easily seen, overlay networks are based on virtualisation of sensor networks. Due to the fact that the differences between these two terms are not easily distinguishable, throughout this document (and within RERUM) these two terms are considered similar and are used interchangeable.

RERUM assumes that for ensuring the efficient provision of applications with a specific Quality of Service there is a need for creating overlay virtual sensor networks. However, in a more generalized approach than in standard virtualised sensor networks, the virtual paths that are formed don't have

to include only sensors that provide the same type of service when collaborating for a specific application. The virtual paths may also include required nodes (RDs) that only have the responsibility to forward the data from the leaf device to the gateway when there is a multihop network.

Wireless Sensor Networks and IoT deployments in general are highly susceptible to attacks, due to both the wireless open nature of the network and the usually limited resources of the nodes [FAT14, GSRV08]. Examples of such attacks can be for the disruption of the routing procedure [KW03, KN++07, SHL08] in a cooperative multihop network. Thus, for ensuring the efficient routing operation in multihop deployments, there is a specific need to select an appropriate and trusted route for sending the packets from the RD to the gateway and the middleware, because intermediate malicious nodes may steal the data or alter them for their benefit. For this reason, within RERUM a key solution is the design and development of a trusted routing mechanism for overlay sensor networks, aiming to ensure that sensitive data will only be sent from the device to the gateway (and vice versa) through trusted intermediate nodes (when a multihop deployment exists). Trust-aware routing is very important for the design of a secure RERUM network. A basic requirement is to define appropriate trust (or reputation) metrics to evaluate the trustworthiness or the reputation of neighbouring RDs.

Within RERUM, the proposed trusted routing mechanism is centralized because we want to reduce the complexity of the modules that will run on the RDs because they are mostly constrained devices. We assume that we form an overlay network, because the RDs can utilize different networking technologies for their interconnectivity, so we don't want to be limited to only one underlying network technology. The proposed mechanism is highly related to the RERUM use cases and most Smart City applications in general. Considering the RERUM use cases, except from the smart transportation, all the others include multihop deployments to minimize the costs of installing multiple gateways at each specific location of interest (i.e. squares, parks, or building floors). In the indoor use cases, the data should be exchanged only through trusted RDs because all of them carry personal information of the inhabitant of the house that could be exploited from an attacker to gain knowledge about the preferences and the habits of the resident. Even information about the temperature of a room or of the energy consumption of a TV can be linked together with other information and could result in the identification of when a user is at home or not. Thus, the information should only be transferred through trusted routes to avoid being stolen by malicious users. Similarly, in the outdoor use case for environmental monitoring, several measurements can be important and their integrity should be ensured, e.g. alerts when the air quality is lower than a threshold or when there is a fire. Malicious users that are playing the role of forwarders in insecure routes can intercept and change this information, affecting the decision making processes of the system. In the smart transportation use case, the multihop phenomenon is not present (in the scenario as it is considered within RERUM). However, in a more generalized use cases for Intelligent Transportation Systems (ITS) with users in vehicles exchanging information for deciding, i.e. which lane to take, if there is an accident ahead, etc. intermediate malicious nodes may both intercept private user data (e.g. identify who is driving ahead) or alter the exchanged messages for their benefit, i.e. to clear the lane for driving much faster. Thus, in this scenario, a trusted routing overlay is of high importance for ensuring that only trusted nodes are included in the network.

Our goal within RERUM is to create trusted routing overlays based on a scheme that considers the following: (i) behavioural statistics of sensors reported by their one-hop neighbours, (ii) link quality statistics, and (iii) a fusion scheme based on the Dempster Shafer theory of evidence. The respective mechanism is described below.

5.1.2 State-of-the-Art

Node reputation has been introduced as a trust mechanism in several publications [JIB07, ZYI04, SS05]. These provide classification schemes and information about how to evaluate reputation and trust, but they are quite focused only on users' ratings. They also include analysis of online traditional models

for reputation (like eBay and Amazon) and other based on these ones with concrete improvements, such as NICE [LSB03] or Sporas and Histos [Z99]. There are many ways to calculate the reputation associated to an agent or to a service, based on the ratings received from users. Probabilistic theory [LWL09], Hidden Markov Model [MAB09], ontologies [SS10] semantics [HF02] and fuzzy models have been employed in reputation calculations and modelling. Robustness is delegated in the way to calculate trust and concrete filters, instead of controlling the data involved in a meaningful way, amending it when inconsistencies arise. Gligor and Wing [GW11] present a theory of trust in networks of humans and computers that consists of elements of computational trust and behavioural trust. Two types of trust are usually considered [CSC09]: (1) social trust (extracted from personal relationships or history data), and (2) QoS trust (related to performance issues) to define different metrics to quantify the trust level in each CR, based on hosts behaviours and their impacts on the performances of the CR network. Social trust is a subjective metric based on friendship, honesty, social interactions, social reputation and history, while QoS trust is a more objective metric based on QoS measurements.

Trust management has been used as an easy and efficient approach to mitigate a wide number of attacks in routing of wireless sensor networks. Nodes create trust relationships based on the expectation that their neighbours will cooperate on particular tasks (e.g. packet forwarding) and decide on routing paths based on them. In [LAK07] a trust management protocol is applied in geographical routing for WSNs that is able to avoid attackers by routing through benevolent neighbours. Yao et al. [YKD06] propose a parameterized and localized trust management scheme for WSN security, where each node maintains highly abstracted parameters to evaluate its neighbours. In [ZSD12] the authors design a trust-aware routing framework that can cope with various harmful attacks with no time synchronization or geographic information, while [RE07] describes a routing scheme that expresses routing reputation by combining link quality and node behaviour. Finally, in [BCCC12] a hierarchical (social and QoS) trust management scheme for heterogeneous sensor networks is introduced. Application to trust-based geographical routing shows performance comparable to the ideal case of flooding-based routing.

5.1.3 Network model

We model the overlay network as a bi-directed graph G(*V*,*E*) where *V* is the set of nodes (smart objects), including a single sink node, and *E* is the set of edges representing wireless links between nodes, as depicted in Figure 51. The sink node, which also plays the role of the gateway, is a more powerful RD with increased capabilities with regards to processing, memory and energy. We assume that each RD has a limited communication range, thus can directly communicate only with the nodes that reside inside this range, namely its *neighbours*. This is of course something quite reasonable in IoT deployments due to the fact that devices want to utilize low transmission power in order to consume less energy and prolong their lifetime. Communication with the sink/gateway node is possible for each RD in a multi-hop fashion with each one following a path calculated through a path calculation algorithm at the sink/GW.

Figure 51. Network topology example

5.1.4 Behavioural statistics

The trust model that we use within RERUM is built on the evaluation of trust degree values based on the observed behaviour of an RD by its neighbour RDs. Any RD observes and records a neighbour's abnormal behaviour each time there is an interaction between them, while it also keeps track of the total amount of incoming packets. For the sake of simplicity and in order to keep the overhead low, two statistics that capture the forwarding behaviour of an RD are used: (i) *packet drop rate* (PDR) and (ii) *packet modification rate* (PMR) defined as follows:

$$PDR = \frac{\#\text{ of dropped packets}}{\text{total }\#\text{ of received packets}}$$

$$PMR = \frac{\#\text{ of modified packets}}{\text{total }\#\text{ of forwarded packets}}$$

Assume a sender RD 'i' and a receiver neighbour RD 'j'. Each time i sends a packet to j, it enters promiscuous mode to overhear the forwarding behaviour of j and it updates accordingly the values of $PDR_{i,j}$ and $PMR_{i,j}$. Then, the aggregate *misbehaviour rate* (MBR) of RD *j* as perceived by RD *i* is calculated as:

$$MBR_{i,j} = w \times PDR_{i,j} + (1-w) \times PMR_{i,j} \tag{61}$$

where $w \in [0,1]$ is a user-defined weight controlling the balance between the behavioural statistics.

5.1.5 Link quality statistics

Although the previously described trust framework is able to isolate malicious RDs in the network, it is necessary to take into consideration a link quality metric that will be used in conjunction with the trust metric for the optimal path calculation of the routing algorithm. In the following, we use the *estimated transmission count* (ETX) metric [DC++05, BM05] to quantify the reliability of the link between any two RDs.

ETX is one of the most widely used routing metrics due to its simplicity and computational efficiency that has been proven in a large number of applications [DC++05, BM05]. Concretely, ETX calculated for RD *i* by RD *j* is defined as:

$$ETX_{i,j} = \frac{1}{f_{i,j} \cdot r_{i,j}} \tag{62}$$

where $f_{i,j}$ is the forward delivery ratio, i.e. the measured probability that a packet sent from RD *i* is received by RD *j*, and $r_{i,j}$ is the reverse delivery ratio, i.e. the measured probability that the acknowledgement packet from RD *j* will be received by RD *i*. Thus, ETX expresses the average number of transmissions needed for a packet to successfully reach its destination in cases when there are transmission failures due to degradation of link quality (e.g. interference, collisions, etc.).

5.1.6 Behavioural trust fusion based on Dempster-Shafer theory of evidence

In the proposed routing scheme a crucial step is the combination of MBRs of each RD, as they are observed and recorded by its interacting neighbours. Since there is a possibility that some of the reported values are not reliable, we need an appropriate combination technique to fuse the neighbour views.

In our framework, each node builds a local view of the neighbours' behaviour from its own observations. As a result, the reported views may be biased or intentionally wrong, leading to the existence of uncertainty. Therefore, we propose the use of Dempster-Shafer (DS) theory of evidence [S76, D08, SF02] as a fusion technique. DS is a mathematical framework able to handle uncertainty of the complete probabilistic model describing the system under consideration.

Essential terminology related to DS and used throughout this section is presented below to ease the reader to understand our contribution:

1. **Frame of discernment *Θ*.** This is the set of all possible mutually exclusive and complete states of a system. For our trust management system the frame of discernment is defined as *Θ* = {trust, no-trust}.
2. **Mass function or probability assignment function *m*.** This function is a primitive of theory of evidence. If *Θ* is the frame of discernment then $m: 2^{\theta} \rightarrow [0,1]$ is a mass function if $m(\varnothing) = 0$ and $\sum_{A \subset \Theta} m(\mathrm{A}) = 1$, where A is any proposition that is a subset of *Θ*.
3. **Belief function**. This function measures the belief of a proposition A, and it computes the sum of all the nonempty subsets of A. It is given by the formula $Bel(A) = \sum_{B \subseteq A} m(B)$, where B are the nonempty subsets that compose A. In the current work, since the belief on an RD's trust depends only on RD's MBR value, the belief function and the mass function are equivalent.
4. **Focal elements.** The focal elements of a frame of discernment *Θ* are comprised of all hypotheses for which the observers can provide evidence. In our case, where *Θ* = {trust, no-trust}, the corresponding focal elements form the set [trust, no-trust, {trust, no-trust}], where the last element shows the uncertainty for the trustworthiness of one RD.

Individual beliefs on RD's trust as reported by each of its interacting neighbours can be combined into a single belief by applying Dempster's rule of combination:

$$m_1(\mathrm{A}) \oplus m_2(\mathrm{A}) = \frac{\sum_{B \cap C = A} m_1(\mathrm{B}) m_2(\mathrm{C})}{1 - \sum_{B \cap C = \varnothing} m_1(\mathrm{B}) m_2(\mathrm{C})} \quad (63)$$

In cases where there are in total N neighbours reporting individual beliefs for a specific node, the belief functions can be combined sequentially in any order (due to the associative property of DS), as expressed by the following formula:

$$m_{1 \ldots K}(A) = \bigoplus_{k=1}^{K} m_k(A) \quad (64)$$

Each neighbour produces a single belief for each focal element (b^T : belief that a node is trusted, b^U : belief that a node is untrusted, b^{TU} : belief expressing ambiguity) according to Table 4:

Table 4: Belief computation based on Dempster-Shafer algorithm.

Belief \ MBR	$MBR < h(1-\mathrm{r})$	$MBR > h(1+\mathrm{r})$	$h(1-\mathrm{r}) \le MBR \le h(1+\mathrm{r})$
b^T	$c+\lvert h-MBR\rvert / h$	$(1-b_U)/2$	1/3
b^U	$(1-b_T)/2$	$c+\lvert h-MBR\rvert / h$	1/3
b^{TU}	$(1-b_T)/2$	$(1-b_U)/2$	1/3

The uncertainty area lies between $h(1-\mathrm{r})$ and $h(1+\mathrm{r})$, where h is a user-defined threshold, $0 < r < 1$ is a user-defined scalar controlling the width of the uncertainty area and c is a scalar controlling how fast trust/distrust belief increases.

By combining the believes of trust reported by each RD for their neighbours we calculate the trust degree of a node *j*, denoted as b_j^{CT} , by using Equation (64). It is reminded that the belief function and the mass function are equivalent, since node trust is characterized only by a single basic event (MBR).

5.1.7 Trust-based routing with link quality

Considering both the trust degree b_j^{CT} and the link quality as quantized by $ETX_{i,j}$ a combined routing metric on a route r from an RD to the sink can be defined as

$$RM_r = \sum_{i,j \in r} ETX_{i,j}\left(1-b_j^{CT}\right) \tag{65}$$

We find the shortest path from each node to the sink using the classic algorithm of Dijkstra [SS90, NS00]. The "shortest path" in this context shows a path that has both a low number of retransmissions at the links and comprises of RDs that have a high trust degree. In the following, the steps of the proposed trust-based routing scheme are described.

Initialization

Each RD measures the ETX metric by broadcasting a probe packet periodically with a fixed time interval. Additionally, it sends a packet containing the count of probe packets from its neighbours, so that all RDs can calculate the loss of probe packets on both directions of a link. Subsequently, the calculated ETX values are flooded to the sink/GW. Since at this time no behavioural statistics are available, all RDs are assigned a trust degree of 0.5 that shows that they are neither trusted nor untrusted.

Finally, optimal paths are computed using Dijkstra's algorithm and are broadcasted to the network.

Trust belief update

The behaviour of malicious (or malfunctioning) RDs is possible to change over time; as a result, it is necessary to update periodically the trust values assigned to each RD. While a node interacts with its

neighbours, it observes and records their behaviour and recalculates accordingly the behavioural statistics, the total misbehaviour rate and, finally, the trust belief. After a predefined trust update interval Δt we have:

$$b_{i,j}^{'\mathrm{T}}(t+\Delta t)=\theta\times b_{i,j}^{\mathrm{T}}(t)+(1-\theta)\times b_{i,j}^{\mathrm{T}}(t+\Delta t),$$

where $\theta\in[0,1]$ is a weighting factor that balances current and previous estimation of misbehaviour rate.

We choose to use a different weighting factor depending on whether the MBR has increased or decreased, as follows:

$$\theta=\begin{cases}\theta_1, & \text{if } b_{i,j}^{T}(t+\Delta t)>b_{i,j}^{T}(t)\\ \theta_2, & \text{if } b_{i,j}^{T}(t+\Delta t)\le b_{i,j}^{T}(t)\end{cases}, \text{ with } \theta_1<\theta_2.$$

As a result, trust degree decreases quickly in case of misbehaviour but increases slowly in case of improved behaviour.

Routes update

After updating the trust degree, each RD calculates the beliefs for the neighbours based on Table 4. New beliefs are sent to the sink where they are fused based on Dempster-Shafer rule of combination, so that a combined trust belief b_j^{CT} and distrust belief b_j^{CU} is assigned to each RD *j*. We use a threshold B_{mal} to detect malicious or misbehaving RDs, in other words if $b_j^{CU}>B_{mal}$, then RD *j* is considered untrusted and it is omitted from the routing paths. Afterwards, the optimal routing paths are found through Dijkstra's algorithm, by considering the combined metric defined in Equation (65).

5.1.8 Performance evaluation

In order to evaluate the proposed scheme we assume the indoor sensor network of Figure 52, which corresponds to a deployment of 54 Mica2Dot sensors at Intel Berkeley lab employed with weather boards to monitor low frequency phenomena (humidity, temperature, light and voltage). Both topology information (nodes' coordinates) and bi-directional link quality (probability of successful transmission that can be used to easily compute ETX values) are available online[10].

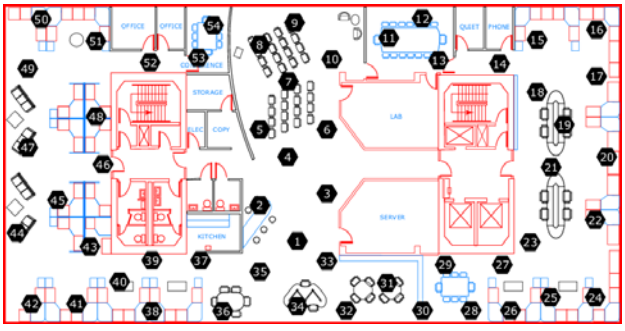

Figure 52. Intel Berkeley lab sensor network topology.

[10] http://db.csail.mit.edu/labdata/labdata.html

We choose randomly a node as the sink/gateway and we conduct simulations for data aggregation rounds. Trust belief update interval is set to 20 rounds, while we choose $c = 0.1$, $h = 0.5$, $r = 0.05$ and $w = 0.5$. In order to account for some variance of wireless channel characteristics, in each round we vary randomly the ETX of all links in [90%, 110%] of their nominal values.

Figure 53 and Figure 54 show the combined trust belief for a normal and a malicious node, respectively, and for various values of weights θ_1 and θ_2 . Obviously, an increase of θ_1 reduces the rate of trust belief increase, applying more significance on history data. Similarly, an increase of θ_2 causes trust belief to decrease slower. Furthermore, we can see that the disbelief for a node (who has low trustworthiness) is calculated very quickly, which means that in only a few rounds we can identify which nodes are untrusted and avoid them in future routing decisions. Generally, by choosing the appropriate values for θ_1 and θ_2 we can assure that our scheme becomes more sensitive to malicious actions, enabling the quick omission of the malicious nodes from the routing paths.

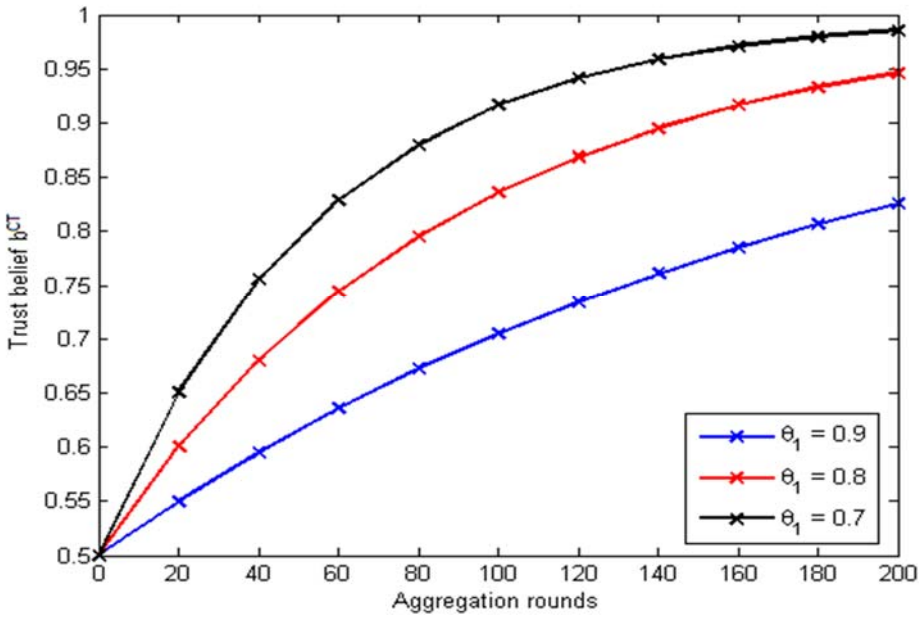

Figure 53. Trust belief of normal node

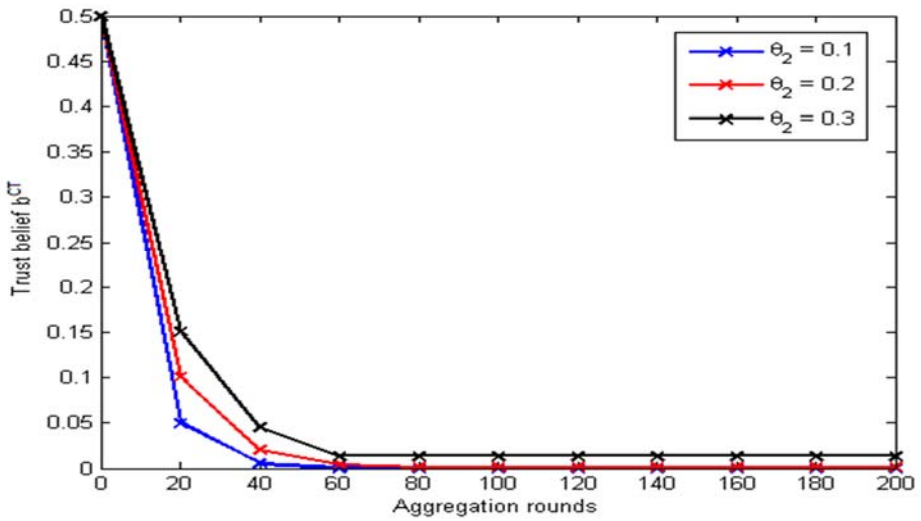

Figure 54. Trust belief of malicious node

Next, we proceed by evaluating our scheme by means of packet delivery ratio at the sink. We assume two scenarios, namely *highly malicious behaviour* and *low malicious behaviour*. In first case, malicious nodes exhibit PDR and PMR with values in [0.8,1] while in second case the values vary in [0.3,0.5]. For each scenario we execute 20 Monte Carlo runs each consisting of 2000 data aggregation rounds. The proposed trust-based routing scheme is compared against a simple routing scheme with no trust management mechanism.

Mean packet delivery ratio (averaged over all runs) vs percentage of malicious nodes for highly malicious behaviour and low malicious behaviour is depicted in Figure 55 and Figure 56, respectively. Observe that in both cases the trust-based scheme outperforms the simple scheme, with the difference being more profound for small number of malicious nodes and highly malicious behaviour. For example, in case of 5-20% of highly malicious nodes there is an increase of more than 30% in packet delivery ratio. Normally, as the number of malicious nodes increases the packet delivery ratio for both routing schemes is almost equal.

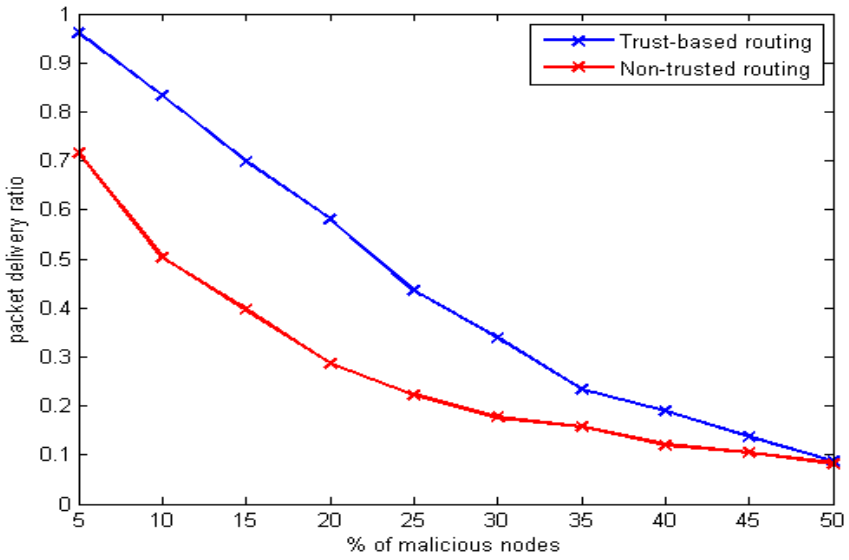

Figure 55. Packet delivery ratio (highly malicious behaviour)

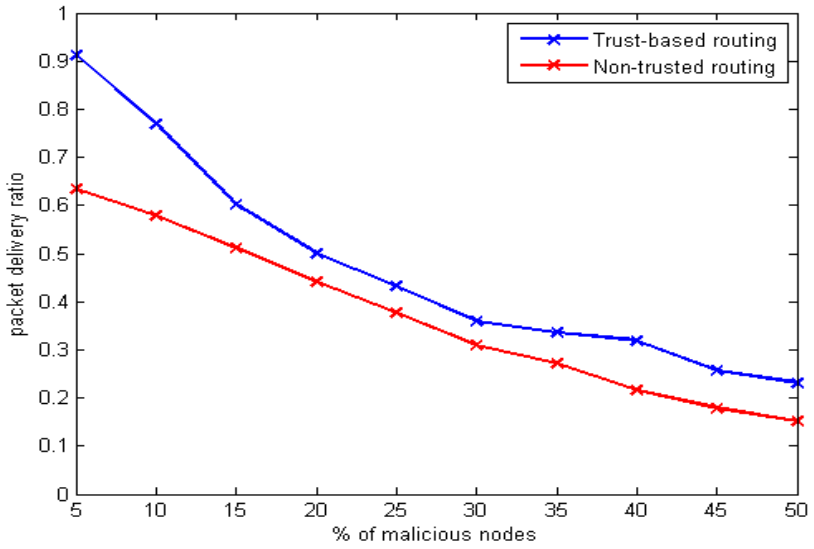

Figure 56. Packet delivery ratio (low malicious behaviour)

5.2 Bi-Directional Multicast Forwarding Algorithm (BMFA) for RPL-based 6LoWPANs

In scenarios involving point-to-multipoint network traffic, transmitting to each destination individually with unicast leads to (i) poor utilization of network bandwidth, (ii) excessive energy consumption caused by the high number of packets and (iii) suffers from low scalability as the number of destinations increases.

This section documents the Bi-Directional Multicast Forwarding Algorithm (BMFA) for 6LoWPANs, which is an algorithm designed and implemented by the RERUM project. It also documents performance evaluation results. Experiments evaluating BMFA's energy efficiency are conducted as part of Task 4.3 and will be documented in D4.2.

The section is structured as follows: Section 5.2.1 highlights this innovation's relationship with RERUM's Use-Cases. In Section 5.2.2 we present an overview of current state-of-the-art. We proceed with a detailed description of BMFA's design in Section 5.2.3, followed by implementation details in section 5.2.4. The sub-section concludes with a discussion of current open issues in Section 5.2.5.

5.2.1 Relationship with smart city applications and RERUM use cases

For UC-O2 - Environmental monitoring in particular, it is expected that networks will be formed by a potentially very high number of RDs and therefore scalability is a requirement. In cases when RDs are powered by batteries, it is impractical or outright untenable to replace batteries very frequently due to high management cost and possibly hard-to-reach installation locations. Thus, long battery life is important. For devices powered from mains, low energy consumption is also important in order to reduce financial cost, but also in order to comply with national and international regulations where applicable.

For those reasons, RERUM has designed and implemented a new multicast forwarding algorithm for 6LoWPANs as presented below. This algorithm addresses the needs of RERUM use-cases by achieving very low energy consumption, as will be demonstrated in deliverable D4.2.

5.2.2 Current State-of-the-Art

This section outlines existing state-of-the-art in the area of multicast forwarding for traditional WSNs with bespoke, application-driven network stack designs (sec. 5.2.2.1), ZigBee networks (Section 5.2.2.2) and 6LoWPANs (sec. 5.2.2.3). BMFA's design is heavily inspired by its predecessor, the "Stateless Multicast RPL Forwarding (SMRF)" algorithm [OP12, OPT13]. For this reason, an entire subsection is dedicated to SMRF (sec. 5.2.2.4), with an additional sub-section (sec. 5.2.2.5) dedicated to IETF's recommendation, called "Multicast Protocol for Low power and Lossy Networks" (MPL) [HK14].

The aim of this section is to bring out the limitations of existing approaches and to therefore back the decisions made during BMFA's design.

5.2.2.1 Multicast in Traditional Wireless Sensor Networks

Previous research efforts in the area of multicast for WSNs have primarily been focusing on traditional, application-centric network designs and as such do not address IPv6-specific challenges. Multicast forwarding algorithms based on geographic routing are dominant in existing bibliography and can be broadly classified as either purely geographic [SRLS07, SMPR+12, CKK12, FZDH12] or hybrid [KDHS10, SCK10], whereby the geographic component is complemented by features of other approaches, such as hierarchical routing.

The geographic multicast routing (GMR) algorithm builds on existing unicast geographic routing approaches. By adapting them, it aims to achieve multicast message delivery to all intended destinations while maintaining minimum bandwidth consumption. Nodes exchange position information with their neighbours through periodic beacons. GMR is characterised by low computational complexity and a small memory footprint [SRLS07]. According to subsequent works, GMR scales better than some of its predecessors [SMPR+12] but still suffers from scalability issues when dealing with large deployments [KDHS10].

Using periodic beacons can have negative side-effects such as collisions and increased energy consumption [SMPR+12]. Based on this observation, BRUMA attempts geographic multi-casting without beacons, whereby neighbour positions are discovered reactively. Next hop selection happens opportunistically through a mechanism which only requires a low number of control messages. The authors demonstrate that BRUMA is more efficient than GMR [SMPR+12].

Carzaniga et al. [CKK12] propose a compact and completely decentralised multicast scheme with asymptotically optimal network congestion properties. It operates by building a multicast forwarding tree over an underlying geographic unicast routing service, which allows nodes to send messages to a destination defined by a coordinate pair (x, y).

Receiver-based multicast (RBMulticast) [FZDH12] is another purely geographic approach. Its principal novelty lies in the next-hop determination phase: potential next hops contend for the channel based on their contribution towards delivering a packet to its destination. Nodes offering the highest forward progress have higher probability of getting selected as next hop. By adopting this approach, RBMulticast can operate without routing tables and without maintaining a forwarding tree. To achieve this, RBMulticast embeds the geographic location of all destinations in the packet header. This raises questions regarding its scalability in large deployments.

The hierarchical geographic multicast routing (HGMR) [KDHS10] is a hybrid algorithm combining the key design concepts of GMR [SRLS07] and the hierarchical rendezvous point multicast (HRPM) protocol [DPH08]. The resulting HGMR algorithm is further optimised to be more energy efficient and scalable. HGMR divides multicast groups into subgroups by using HRPM's geolocation hashing. It takes advantage of layer 2 reliability mechanisms by using HRPM's unicast forwarding approach for long, sparse paths and reverts to layer 2 broadcasting in areas of high density in order to reduce the number of transmissions. When unicast forwarding is in use, HRPM (and therefore HGMR) uses source routing along the branches of an overlay tree generated by the traffic source. In this work the authors conducted a performance evaluation of the three algorithms in a simulated IEEE 802.11 network. It is therefore difficult to understand how the algorithms would behave under the frame size and bandwidth limitations or the loss characteristics related to IEEE 802.15.4 networks.

The multicast routing with branch information nodes (MR.BIN) [SCK10] protocol is a hybrid approach, combining geographic unicast routing with state-based multicast. It maintains multicast states only on branch nodes of the forwarding tree. Communication between non- branching nodes takes place with geographic unicast. In order to perform datagram forwarding, each node maintains a potentially long list of next hops. Messages are tagged with a 2-byte multicast group identifier but the management scheme for those ids is not discussed; it is unclear how a node can choose which group id to join and how two different groups are prevented from having the same id.

Adaptive geo-source multicast routing (AGSMR) [SKC09] is a geographic unicast and source multicast hybrid. It relies on generating a forwarding tree at the traffic source. Path information is embedded in packet headers in a compressed format which uses 2 bytes per hop, an additional 2 bytes per branch and reduced by $2(n - 2)$ bytes for each n-node long, non-branching path. The authors demonstrate that, with this compression scheme in place, AGSMR's packets are smaller than GMR's. However, reported packet sizes have an order of magnitude of KBytes (e.g. about 9 KB for a tree with about 50 subscribers in a network of 1024 nodes). Packet size increases with the number of subscribers as well as the total number of network nodes.

Branch aggregation multicast (BAM) [OIM05] differs from the aforementioned efforts in that it operates without knowledge of node geolocation. BAM's design has two components: S-BAM achieves single-hop aggregation at branching nodes and M-BAM aims to reduce the number of branches. The former tries to combine multiple layer 2 unicasts within a single radio frame, by extending headers to list the addresses of all intended destinations. Recipients reply with an acknowledgement frame, using a random back off to avoid ACK collisions. M-BAM relies on existing forwarding tables, which can be populated by any routing protocol. The authors make the assumption that the routing table contains multiple candidate next hops for the same destination and they propose an algorithm for path aggregation. However, routing table size has an impact on scalability with increasing network size [OP11] even when the table only lists a single next hop per destination. Storing multiple candidate next hops per destination would impose further memory overheads and should not be considered common practise. Furthermore, BAM does not manage multicast groups internally. Instead, it assumes that the network adopts a data centric routing model, whereby nodes broadcast requests for data and their neighbours remember querying node addresses in order to forward relevant data accordingly [OIM05].

All aforementioned algorithms assume that the message source is aware of and maintains a list of all destinations, uniquely identified by an attribute such as a node id, address or name. This is untrue in the case of IP and IPv6 multicast, whereby datagram destination is expressed as the multicast group's IP or IPv6 address and the traffic source is oblivious with respect to the unicast address of each individual group member.

The GMR protocols discussed above also require knowledge of the geolocation of all destinations. When this information is not available, a geographic routing service can be used in conjunction with a location service, such as MLS [FW06] or GLS [LJDKM00]. Querying the location service incurs time overheads which have an impact on delivery delay, compared to non-geographic approaches which do not rely on this information. However, this is not considered at all in relevant work.

Many of the above multicast forwarding algorithms embed a list of all destinations in the packet header [SRLS07, SMPR+12, FZDH12, SCK10, SKC09, OIM05]. This increases the byte overhead associated with each data transmission and has an adverse impact on the protocol's scalability [KDHS10].

Lastly, with all those approaches traffic is confined within the deployment's boundaries. In order to be able to communicate beyond the WSN's borders, one would need a dedicated gateway. This does not apply to 6LoWPANs and IPv6 multicast forwarding protocols, such as those discussed later on in this deliverable.

5.2.2.2 Multicast Support in ZigBee

ZigBee is based on cluster-tree WSNs. After the release of Z-cast, the multicast mechanism proposed by the ZigBee alliance, some gaps related to memory consumption and communication overhead have been noted. A suggestion for reducing the memory costs and improving its efficiency was proposed in [BG13].

5.2.2.3 Multicast in 6LoWPANs

Previous research with focus on 6LoWPANs is very limited. Sá Silva et al. [SCPR+08] investigated the applicability and usefulness of traditional multicast paradigms in WSNs. In this work, the authors evaluate Multicast Ad hoc On-demand Distance Vector (MAODV) [RP99] and Source-Specific Multicast (SSM) [B03] in a simulated sensor network. They demonstrate that, in their simulated scenarios, multicast can offer significant advantages over traditional network-wide broadcast flooding and unicast. Results also suggest that MAODV and SSM are comparable in terms of bandwidth consumption and energy efficiency.

Clausen and Herberg [CH10] conducted relevant research with focus on RPL networks. Despite the fact that this work only discusses network-wide broadcast, it offers some important insight on related issues and techniques. In their contribution, the authors conduct experiments by simulating an IEEE

802.11b [IEEE09] network, which cannot capture the duty cycling aspects of modern WSNs, nor the low-power, lossy nature of IEEE 802.15.4-compliant radio hardware.

A very common problem in wireless networks is the one of the "fixed base rate" transmission. In other words, the throughput a multicast group experiences is confined by the node which demands the lowest rate. For example, if a node that is member of a multicast group moves at the limit of the signal coverage boundaries, the transmission rate for all the nodes must be reduced, as much as required, to retain low packet loss to the distant node. Several techniques have been proposed to address this problem but they are characterized by high complexity and communication and computation costs. In [AFRB+13], the authors propose a method which combines multicast and unicast transmissions in order to ameliorate this problem. More specifically, multicast traffic that arrives to an Access Point (AP) is converted and forwarded as unicast traffic into the sub-network. Comparative results to existing algorithms show that the multicast-unicast algorithm outperforms them.

The RPL RFC (RFC 6550) [WTBH+12] briefly discusses built-in multicast support. In RPL networks, nodes advertise unicast downward paths inside destination advertisement object (DAO) messages. An RPL instance is administratively configured with one out of a possible four modes of operation (MOP). Nodes may only participate in the network as routers if they support the advertised mop, otherwise they may only join as leaf nodes. In the Storing with multicast support mop, DAO messages are also used for group management by advertising multicast prefixes. Unfortunately, the approach leaves many open issues.

5.2.2.4 Stateless Multicast RPL Forwarding (SMRF)

The efforts made by IETF's "IPv6 Low Power Wireless Personal Area Networks" (6LoWPAN) work group for transmitting IPv6 datagrams efficiently in low power radio frames allowed the integration of IPv6 stacks in energy constrained Wireless Sensor Networks. The emerging standard for routing such frames is the "IPv6 Routing Protocol for Low Power and Lossy Networks" (RPL) [WTBH+12] which however offers minimal multicast forwarding support. To fill this gap, "Trickle Multicast" (MPL) [HK14] appears to be one of the most suitable algorithms. The lightweight "Stateless Multicast RPL Forwarding" [OP12, OPT13] (SMRF) algorithm aims to counter some of MPL's deficiencies.

SMRF is a multicast forwarding algorithm based on the topology information provided by the RPL protocol. Nodes participating in an RPL network exchange topology information in order to build a Destination Oriented Directed Acyclic graph and construct their routing tables. When the "Storing with multicast support" Mode of Operation is used, a node can join a multicast group by advertising its address in an outgoing DAO message. Upon reception of such a message from one of its children, the parent node registers the multicast address in its routing table; then the same address is advertised by the parent in its own DAO messages. When a multicast packet destined to that address arrives, it will be forwarded downwards the tree until it reaches the recipient nodes.

While the RPL mechanism solves the problem of propagating group membership information towards the DODAG root, it confronts two major pitfalls: a) A node can receive the same datagram more than once. This can happen in the case where another node that is in the signal coverage also propagates the same multicast message to its children. b) A router, as specified in [WTBH+12], must send multicast datagrams to a subset of its one-hop neighbours; either to its parent or only to the children that are members of that particular group. This can only be done by re-transmitting unicast datagrams to each recipient, resulting in bandwidth, delay and processing overhead.

By using the group membership information provided by RPL, SMRF addresses this drawback as follows:

- A node will only accept an incoming multicast message from its preferred parent.
- If the node is member of the specific multicast group, the message will be delivered to the local stack.

- If the routing table has an entry with the multicast address contained in the packet, then it will be forwarded downwards the tree.

In order to avoid the immediate forwarding of an accepting message, SMRF introduces a delay D between two radio-on cycles. Because the underling duty algorithm might set the Channel Check Interval (CCI) to zero, SMRF defines a delay D as the maximum of (F_{min}, CCI), where F_{min} is a non zero value. In this way the RF hardware can never be always turned on. Furthermore, in order to mitigate the hidden-terminal problem, where an AP receives corrupted data due to the concurrent transmission from two nodes, the forwarding delay is randomized using a Spread value.

- SMRF's advantages: Since each node will receive a datagram only from its parent, this message will be forwarded at most once, allowing every next node to avoid using per-packet information in order to distinguish between duplicate packets. A side effect of this is that datagrams are guaranteed to arrive in the correct order. In addition, multicast datagrams will only reach parts of the network where nodes have expressed interest for the specific multicast group.
- SMRF's limitations: Multicast propagation can only travel downwards the DODAG, meaning that nodes of higher level can not be members of multicast source nodes located lower in the tree.

In their works, the authors argue that SMRF's performance is consistent irrespective of changes in the lower layers of the network stack, demonstrates low end-to-end delay, has a very low code and RAM requirements and is very energy-efficient. As a drawback, it is less reliable than Trickle Multicast. They conclude that ultimately, the choice of multicast forwarding algorithm should be based on the anticipated usage of a sensor deployment.

5.2.2.5 Multicast Protocol for Low power and Lossy Networks (MPL)

The Multicast Protocol for Low power and Lossy Networks is an emerging standard for multicast forwarding and group management in LLNs. Unlike SMRF, MPL is independent of the protocol used for unicast routing. It will thus operate in RPL-based 6LoWPANs, as well as in deployments where RPL is not in use. MPL is specified in an Internet Draft published by the IETF [HK14].

In a nutshell, MPL routers maintain a cache of multicast datagrams they have seen. Neighbours exchange information about the content of their caches by exchanging ICMPv6 messages. If one node detects that one of its neighbours has not received a multicast datagram that the former has in its cache, it will schedule subsequent forwarding of the datagram(s) in question. The exchange of ICMPv6 messages is governed by trickle [LPCS04, LCHG11], thus reducing network control traffic when multicast traffic is not present, but reacting very quickly to the arrival of new multicast datagrams.

By design, MPL has some very significant advantages:

- Generality: MPL will work, without modifications irrespective of the protocol used for unicast routing. Thus, it will work in RPL-routed networks
- Reliability: By caching datagrams and maintaining per-packet state information, MPL increases its reliability (high packet delivery ratio / low loss).
- Guaranteed no duplicates: MPL uses a technique to uniquely identify each datagram. A positive side-effect of this technique is that each network will receive each individual datagram at most once.

However, there are also some concerns:

- Arrival Order: Due to MPL's datagram forwarding logic, it is likely that datagrams of the same traffic stream may arrive to their destination out of order.
- Complexity: MPL introduces a new ICMPv6 message type for the exchange of cache information, it requires caching of multiple datagrams alongside meta-data and it also requires from each router to maintain multiple trickle timers. These requirements increase code size and RAM requirements and also raise concerns in terms of MPL's scalability with increasing network size and traffic volume.

The Contiki Operating System supports an older version of the MPL protocol, when it was still called simply "Trickle Multicast" (TM). This implementation has been used here for comparisons with BMFA and was also used for comparisons with SMRF in [OP12, OPT13].

5.2.3 BMFA's Design

SMRF is very lightweight but this comes at the cost of a severe limitation: it is only capable of forwarding traffic "downwards" in the RPL DODAG tree [OP12, OPT13]. BFMA's primary design goal is to alleviate this limitation while at the same time maintaining SMRF's lightweight, energy-efficient nature.

5.2.3.1 Information Available Locally to All Nodes in an RPL-Routed 6LoWPAN

Each node has a list of network interfaces, with each interface associated with one or more unicast and multicast IPv6 addresses. When a process running on a node wishes to subscribe to a multicast group, the group's IPv6 address is added to the list of addresses of one of the node's network interfaces.

In an RPL-routed 6LoWPAN, each node has the following additional information available:

- The link-layer (MAC) and IPv6 addresses of its preferred parent. The preferred parent is essentially a node's next hop in the path towards the DODAG's tree ("upward" path). In Contiki, the preferred parent's IPv6 address is stored as a default route, whereas its link layer address can be found by looking up the IPv6 address in the ND cache.
- The link-layer and IPv6 addresses of all its children. A child in this context is a node that has selected the node in question as its preferred parent. For example, if node A has selected node B as its preferred parent, then node A is listed as one of B's children. IPv6 addresses will be found in node B's routing table ("downward" routes), whereas link-layer addresses can again be resolved by a lookup in node B's ND cache.

5.2.3.2 Avoidance of Forwarding Loops and of Duplicate Datagram Delivery

Upon reception of a multicast datagram, a node first decides whether this datagram should be processed or discarded. A datagram can be rejected for one of the following reasons:

- The datagram is malformed
- The receiving node does not understand the datagram
- The datagram gets detected as a duplicate
- The datagram is part of a forwarding loop

This decision is taken at layer 3 by the multicast engine's implementation. Assuming the datagram is not malformed and can be understood, the decision of whether it is a duplicate or part of a loop is taken by combining the following pieces of information:

- Information available to the node locally. This includes the information discussed in the previous section
- Information available in the datagram itself, in the shape of layer 3 headers.

In broad terms, there are two approaches that can be used to avoid delivery of multicast datagrams more than once:

- Only accept datagrams from a single neighbour and only forward datagrams at most once. This approach is very straightforward but it offers no reliability-enhancing mechanisms. This is the approach adopted by SMRF.
- Adopt a method to uniquely identify each multicast datagram. This is the approach adopted by e.g. TM and MPL. This approach has the advantage that it is possible to attempt re-transmissions in order to decrease packet loss.

In order to avoid duplicate datagrams, BMFA also adopts the latter approach, by using the following logic:

- A datagram's originator does not add any additional information to the datagram.
- Before forwarding a datagram, a router inserts to it the Link Layer address of the node it received the datagram from, either as an IPv6 HBHO header, or as part of the IPv6 Flow Label.

Furthermore, nodes adopt the following logic to determine whether to accept or discard an incoming datagram. A datagram is accepted if:

- The previous hop's Link-Layer address is either not present or does not match the node's own Link-Layer address and
- The datagram's sender is either the receiving node's preferred parent or one of its children.

This is illustrated in Figure 57. The top of the figure illustrates the unicast routing topology, whereby:

- The path from Node 1 to Nodes 3, 4 and 5 are via Node 2.
- The path from Node 2 to Node 5 is via Node 3.

Thus, we have the following relationships:

- Node 2 is the preferred parent for Nodes 3 and 4. Thus, Nodes 3 and 4 are among Node 2's immediate children (more may exist but are not shown for clarity).
- Node 3 is Node 5's preferred parent and Node 5 is among Node 3's immediate children.
- Node 1 is Node 2's preferred parent and Node 1 is among Node 1's immediate children.

Assume now that Node 1 is the original source of a multicast datagram and that Nodes 2 and 5 are subscribed to the datagram's destination multicast group. The bottom of the figure displays BMFA's duplicate detection and loop avoidance logic:

- Node 2 will receive the datagram from Node 1. It will add Node 1's link layer address to the datagram before forwarding it.
- Node 1 will reject the copy sent by Node 2, because it will see its own link layer address inside the datagram. For the same reason, various copies will be rejected among the remaining nodes. The paths of the copies rejected for this reason are illustrated with dotted lines.

- All datagrams between Nodes and 3 and 4 as well as those between Nodes 4 and 5 will be discarded, since the nodes in question do not have a parent-child relationship. Those discarded messages are shown with dashed lines.
- As a result, all nodes will only process a single datagram each. The path followed by those datagrams are displayed with solid lines.

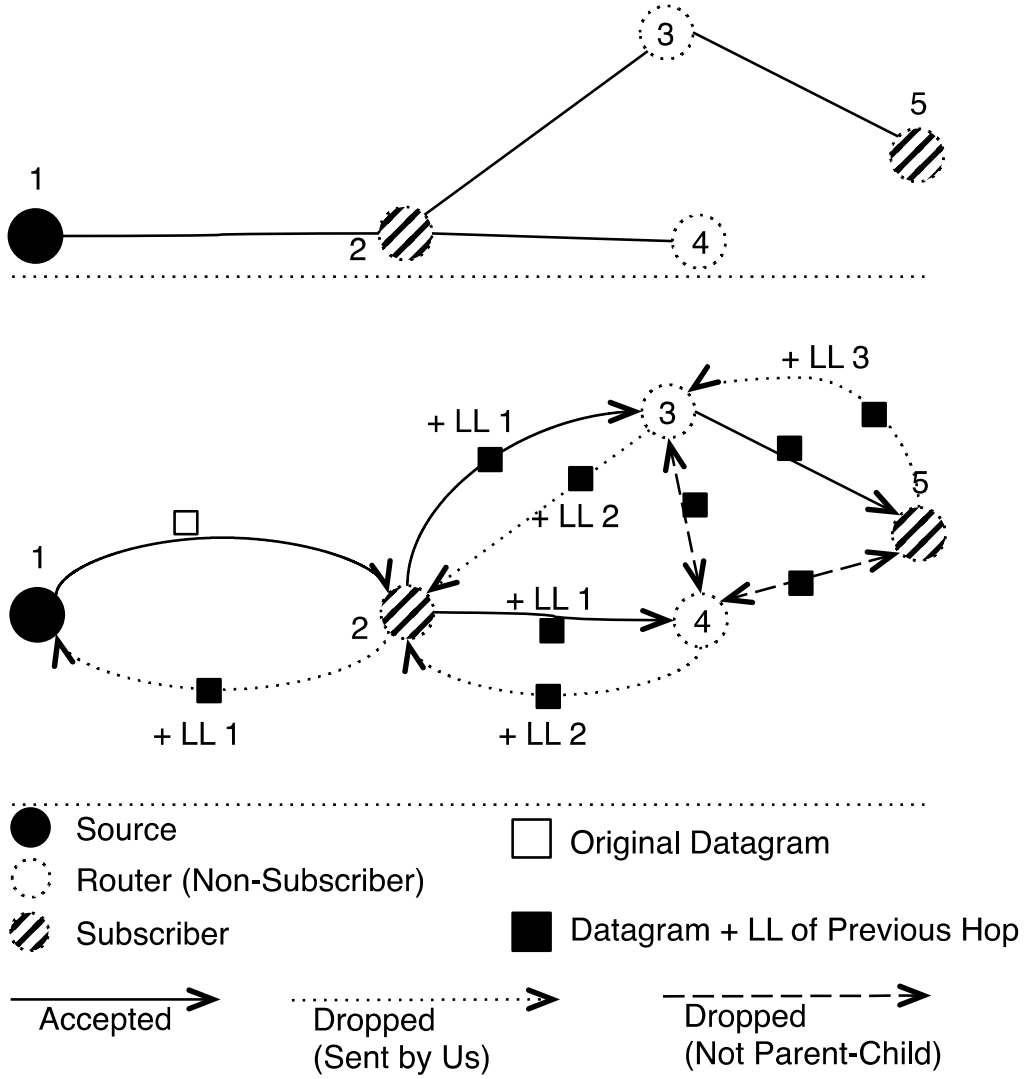

Figure 57. Usage of Sender's LL Address Beyond the 1st Hop.

5.2.3.3 Insertion of Previous Hop's Link-Layer Address

As discussed in the previous sub-section, the decision of whether to accept or discard an incoming datagram is partially based on the Link-Layer address of the previous hop in the datagram's path. This address is added to the datagram itself by each intermediate router. There are a variety of methods to achieve this addition of the source's LL address. The ones we consider here are:

- Utilise an IPv6 Hob-by-Hop Option (HBHO) extension header
- Utilise the IPv6 Flow Label

Using an IPv6 HBHO is a standard approach. In a 6LoWPAN context, this approach is adopted by RPL itself as well as by MPL. However, this approach has some disadvantages when used in 6LoWPANs,

one of the most important ones being the increase to frame size as well as the potential negative consequences it can cause on 6LoWPAN header compression.

Instead of using an IPv6 HBHO, it is possible to convey the same information by using the Flow Label Bits present in the IPv6 header (Figure 58). The flow label is only 20-bits wide, therefore there is not enough space to accommodate the entire Link-Layer address of the previous hop. To counter this hurdle, instead of inserting the entire Link-Layer address, we insert a 2-byte long digest, calculated with CRC16.

A discussion about the advantages and disadvantages of using the IPv6 Flow Label versus using an IPv6 HBHO is out of scope of this deliverable, but it should be mentioned that the IETF is considering a similar approach for RPL-routed 6LoWPANs [T14]. Our current implementation uses the Flow Label, but this is merely an implementation decision. The algorithm will operate with both approaches.

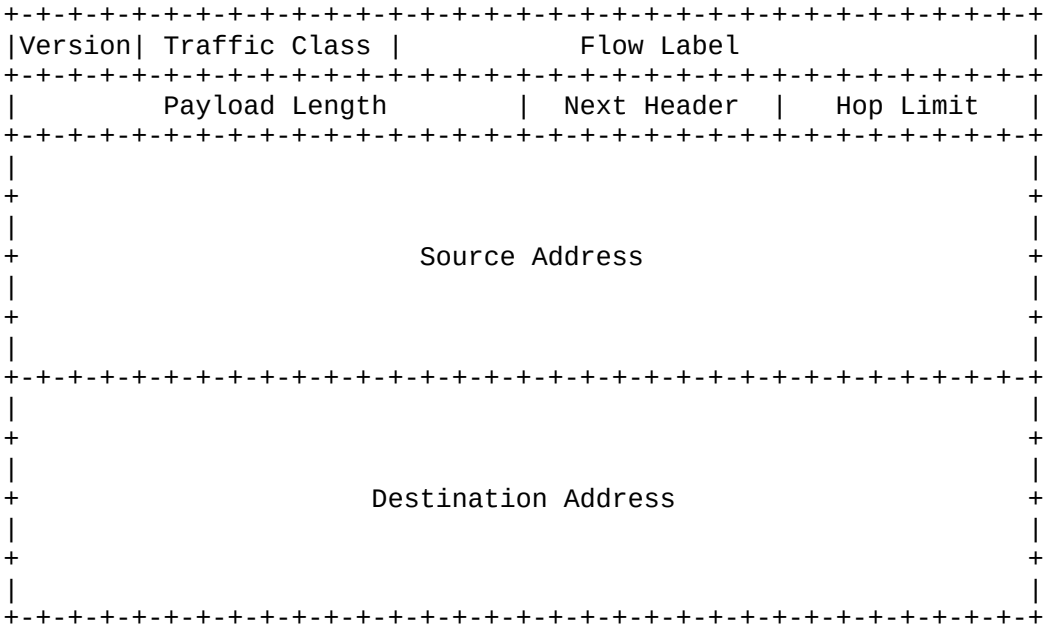

Figure 58. The IPv6 Header Format (Source: RFC2460).

5.2.3.4 Datagram Forwarding and Local Delivery to Upper Layers of the TCP/IP stack

If a datagram is accepted for further processing, the node is further faced with two independent decisions:

- Should this datagram be forwarded?
- Should this datagram be delivered to the upper layers of its own network stack?

The answer to the both questions is very straightforward:

- If accepted, a datagram will also be forwarded
- A datagram is delivered if and only if its destination IPv6 address is listed among the multicast addresses of one of the node's interfaces. In other words, the datagram is delivered up the stack if some process is subscribed to this group.

5.2.4 Implementation for the Contiki OS

This section aims to present BMFA's design details in brief. To do so, the section starts with an overview of Contiki's layered architecture. We then present Contiki's multicast APIs[11] and we conclude this section with a discussion about BMFA's implementation details.

5.2.4.1 Brief Overview of Contiki's Architecture

Conceptually, Contiki's can be broken down into 3 "layers":

- **Hardware layer**: Includes platform-dependent drivers, for example the implementation of drivers for hardware clocks, RF transceivers, UART [O80], SPI [SPI], I^2C [NXP14] etc
- **Core**: Includes platform-independent implementation of Contiki's "kernel". This includes the implementation of software timers, the event infrastructure, process manager and scheduler, and many other libraries. Importantly in the context of this deliverable, this part of Contiki also includes the implementation of all networking.
- **Apps**: Includes platform-independent implementations of various applications, such as a web browser, a CoAP [SHB14] engine, an MQTT [L10] engine and others.

As discussed above, both "Apps" and "Core" components are platform-independent. The main difference is that, while large parts of the "Core" are included in all firmware images, the same is not true for "Apps", which are only included in the developer explicitly requests one of them. This conceptual layered structured is illustrated in Figure 59.

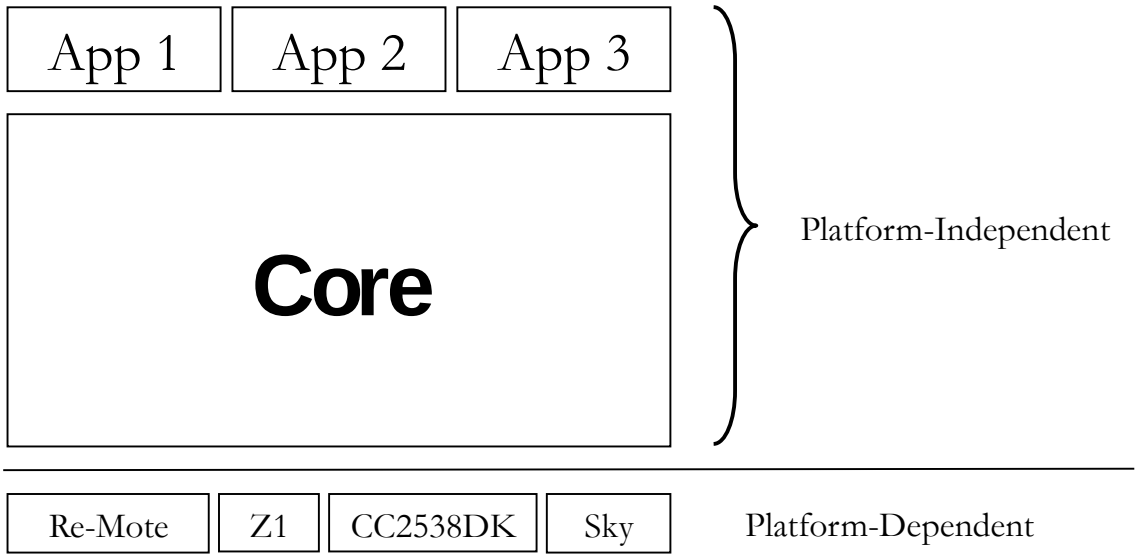

Figure 59. Contiki's Conceptual Layers.

As discussed above, the entire networking functionality is implemented at Contiki's core. This includes implementations of the following modules:

- **Physical Layer**: An API defining the interfaces that platform-dependent RF drivers must expose. This does not include the implementation of RF drivers, which reside in the platform-dependent layer. It simply defines a set of functions that those drivers must implement.
- **Link-Layer (Layer 2)**
    - **Radio Duty Cycling (RDC) Layer**: This includes the implementation of a series of algorithms which handle radio sleep cycles in order to preserve node energy.

[11] https://github.com/contiki-os/contiki/blob/master/core/net/ipv6/multicast/README.md

- **Medium Access Control (MAC) Layer**: An implementation of Carrier-Sense Multiple Access (CSMA).
- **Framer**: Encoding and Parsing of IEEE 802.15.4 headers.
- **Link-Layer Security (LLSEC):** Implementation of AES 128 encryption/decryption as per the IEEE 802.15.4 standard.

- **6LoWPAN**: Implementation of 6LoWPAN header compression, datagram fragmentation and re-assembly as per RFCs 4944 and 6282.
- **Network Layer (Layer 3)**: Implementations of IP, IPv6, ICMP, ICMPv6, IPv6 Neighbor Discovery (ND) and routing protocols, including RPL.
- **Transport Layer (Layer 4)**: Implementations of TCP and UDP, including a TCP sockets API and library.
- **RIME**: Contiki also includes the RIME network stack, which can be used instead of a TCP/IP stack. RIME is out of context of RERUM, but is mentioned here for completeness.

With this layered approach in mind, Figure 60 presents a possible configuration of the network stack that could be used by the Re-Mote platform.

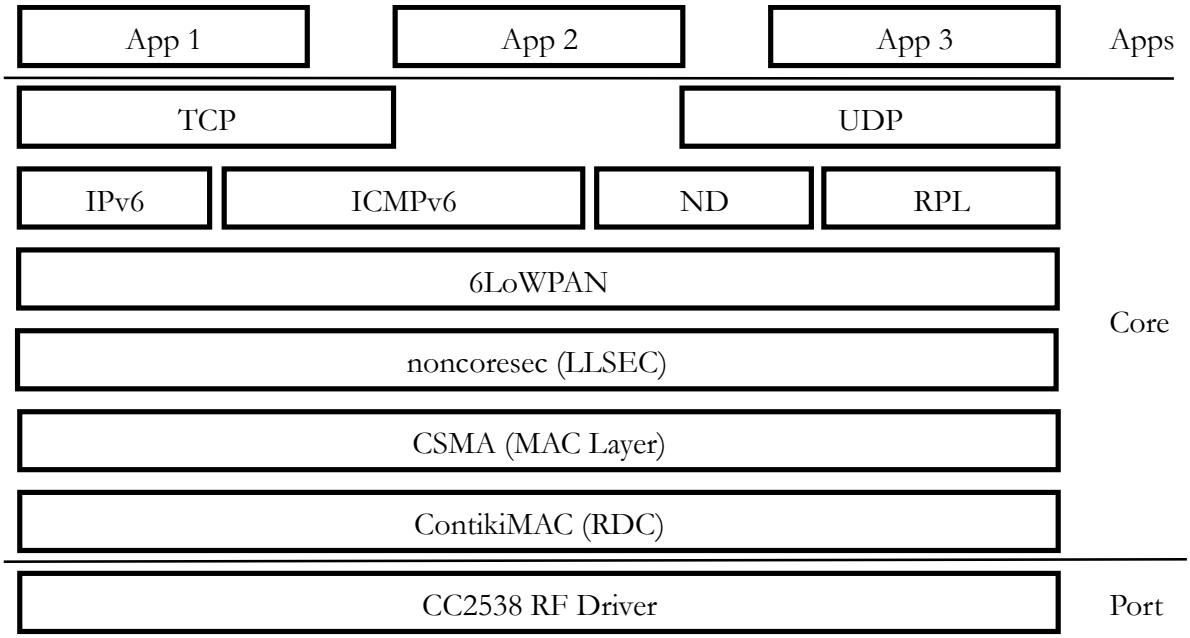

Figure 60. A Possible Configuration of the Contiki Network Stack for the Re-Mote Platform

5.2.4.2 Contiki's Multicast Application Programming Interfaces

Within Contiki's networking core is a generic API for multicast implementations (called "engines" using Contiki's nomenclature). The implementation has adopted this approach in order to allow developers to choose among different engines depending on their needs, or to even implement their own. Currently Contiki officially supports two engines: Trickle Multicast (TM, an older version of MPL) and SMRF.

This API is used to hook engines into Contiki's networking core, in order to achieve the following:

- Initialise the engine,
- Notify a multicast engine about an incoming multicast datagram,
- Allow the multicast engine to transmit an outgoing datagram,
- Maintain statistics,

In line with the above, a multicast engine must expose three functions in the shape of function pointers:

```
void (* init)(void);
void (* out)(void);
uint8_t (* in)(void);
```

Snippet 1: Function Prototypes for the Multicast Engine API.

Each multicast engine is represented by a data structure of type `struct uip_mcast6_driver`. This data type exposes the engine's implementation of the multicast API. The prototype for this data structure is displayed in Snippet 2.

```
struct uip_mcast6_driver {
  char *name;
  void (* init)(void);
  void (* out)(void);
  uint8_t (* in)(void);
};
```

Snippet 2: Multicast Engine Data Type.

Additionally, each engine has a unique integer assigned to it in the shape of a C pre-processor macro:

```
#define UIP_MCAST6_ENGINE_NONE        0 /**< Selecting this
disables mcast */
#define UIP_MCAST6_ENGINE_SMRF        1 /**< The SMRF engine */
#define UIP_MCAST6_ENGINE_ROLL_TM     2 /**< The ROLL TM engine */
```

Snippet 3: Multicast Engine Selection.

At build time, the developer has the choice of selecting a multicast engine or of disabling multicast support altogether. This is achieved by setting a value for the `UIP_MCAST6_CONF_ENGINE` macro. For example, to select Trickle Multicast, the developer will have to configure the build as per Snippet 4.

```
#define UIP_MCAST6_CONF_ENGINE UIP_MCAST6_ENGINE_ROLL_TM
```

Snippet 4: Configuration to Select Trickle Multicast.

To extend Contiki by adding support for a new multicast engine, a developer has to perform the following a series of simple steps, described in Contiki's relevant README[12]. This process was followed in order to implement BMFA, as described in further detail in Section 5.2.4.3.

5.2.4.3 The BMFA Engine's Implementation

In line with the procedure described in Contiki's multicast README, we have implemented BMFA by performing the following steps:

1. Create the BMFA engine's header file and implementation (`bmfa.h` and `bmfa.c` respectively).
2. Assign a new engine number by adding to Contiki's `uip-mcast6-engines.h`, as per Snippet 5.

[12] https://github.com/contiki-os/contiki/blob/master/core/net/ipv6/multicast/README.md

```
#define UIP_MCAST6_ENGINE_BMFA              3
```

Snippet 5: Declaration of BMFA's Engine Number

3. BMFA makes use of RPL's multicast group management, therefore it will only operate in RPL networks using MOP3 ("Storing Mode with Multicast"). This is achieved by extending Contiki's `uip-mcast6.h`, as per Snippet 6.

```
#elif UIP_MCAST6_ENGINE == UIP_MCAST6_ENGINE_BMFA
#define RPL CONF MULTICAST     1
```

Snippet 6: Enabling RPL MOP3 when Using BMFA

4. If the build is configured to maintain multicast statistics, BMFA maintains stats identical to those maintained by SMRF, The BMFA stats data type is shown in Snippet 7.

```
struct bmfa_stats {
  uint16_t mcast_in_unique;
  uint16_t mcast_in_all;
  uint16_t mcast_in_ours;
  uint16_t mcast_fwd;
  uint16_t mcast_out;
  uint16_t mcast_bad;
  uint16_t mcast_dropped;
};
```

Snippet 7: BMFA Statistics Data Type

5. Expose BMFA's implementation with a data structure of type `struct uip_mcast6_driver`, as shown in Snippet 8. `init`, `out`, `in` are pointers to functions of the same name. Those functions are defined in `bmfa.c` and implement the entire engine's logic for processing incoming and outgoing multicast datagrams.

```
const struct uip_mcast6_driver bmfa_driver = {
  "BMFA",
  init,
  out,
  in,
};
```

Snippet 8: BMFA Engine Driver's Definition

5.2.5 Open Issues

BMFA, as well as its predecessor SMRF and MPL, will only handle multicast traffic within the 6LoWPAN's boundaries. In order to traverse 6LoWPAN boundaries, the network's Border Router needs to be able to participate in the formation of a multicast tree by exchanging topology information with its peer routers. In IPv6 networks, this is achieved by exchanging Multicast Listener Discovery (MLD) [DFH99] or MLDv2 [VC04] messages. At the time of writing of this deliverable, Contiki did not support MLD.

6 Conclusions

This document presented the techniques that are developed within RERUM for optimizing the interconnectivity of devices in IoT deployments. These techniques are integral parts of the RERUM System Architecture as presented in deliverable D2.3 [RD23]. More specifically, as also presented in Section 1 in this document, these techniques are parts of the Communication and Networking manager of RERUM, which is the functional entity responsible for enabling the seamless and efficient interconnectivity of RDs in smart city deployments. One of the most important challenges in IoT deployments in smart city environments is the interference that the wireless transmissions face. It is well known that interference can have a severe impact on the network performance and especially for IEEE 802.15.4 devices, since the channel bandwidth is very narrow and any other transmission at the same frequency results in high packet loss.

For addressing the efficient interconnectivity of wireless IoT devices in city environments several techniques have been presented in the document, while all of them were based on the requirement for lightweight and energy efficient operation, since IoT devices are resource constrained. One of the key advances of RERUM with regards to the state of the art is the adoption of the Cognitive Radio technology on the IoT world, designing a CR-agent that can be built on top of RERUM Devices, creating CR-inspired RDs. The key functionalities of the CR-agent for spectrum sensing and for spectrum assignment were presented in Section 2. We described a framework for gathering historical data for spectrum occupancy in an energy efficient way (minimizing the energy consumption spent for sensing) and then we showed how the RD can select the optimal spectrum frequency and band width for meeting the QoS of the services. The results of the framework were very interesting in the way that the spectrum occupancy behaviour is learnt extremely fast and very accurately, thus the RD can decrease significantly the energy it spends for spectrum sensing and then it can use this information for identifying what is the optimal spectrum fragment to select.

Further optimization of the wireless resources within networks of RDs is achieved by grouping the RDs into clusters and executing a novel frequency reuse method among the network of clusters. This method can be applied in cases where the network deployment follows the classic cellular scheme where RDs are connected to the Cluster-Head and all the RDs within the cluster share the same frequency for the wireless connectivity. The proposed technique can propose an optimized way of allocating the frequencies in such network of clusters, in a way that the interference between clusters is minimized, which can result in a maximization of the throughput per cluster and a maximization of the spectral efficiency. Then, in cases when some clusters are overloaded, they can borrow resources from their neighbouring clusters using the proposed negotiation protocol presented in the document.

Another key innovation presented in the deliverable is the virtualisation of the network interfaces of the RDs and the creation of virtual network interfaces for allocating each requested Service in only one virtual interface. The latter can be either a single physical interface or can aggregate resources of multiple physical interfaces for meeting the QoS requirements of the service. Another technique for the optimal allocation of resources to RDs when they have heterogeneous interfaces is also proposed aiming to ensure that the RDs will always select the most appropriate interface/technology in terms of both performance and cost. This technique can be easily used by mobile RDs in city environments where multiple WLAN deployments coexist with mobile networks.

The final goal of this deliverable was to present techniques utilized for optimizing the delivery of data from the devices to either other devices or to the Middleware. In this respect two mechanisms are developed within the project – one for ensuring that only trusted intermediate nodes are selected for routing the data and another one for multicast forwarding. The former technique is more abstracted and can work with overlay networks, since it abstracts the heterogeneity of the networking technology of the RDs, while the latter is specific for 6LoWPAN networks. The latter technique is extremely lightweight and very efficient and a prototype implementation on Contiki OS is also presented.

Parts of the work done in this deliverable will be used for input in the rest of the deliverables of WP4 for either optimizing the mechanisms for their energy efficiency (to be delivered in D4.2 in August 2015) or for assessing their scalability (to be delivered in D4.3 February 2016). Furthermore, as will be described in more detail in D5.1 (due April 2015) most of the techniques presented in this deliverable will be implemented and tested in either experiments or real-world trials, in order to test their performance and their efficiency in real deployments. Early results will be used to define the system interfaces and refine the system architecture to be delivered in D2.5 due August 2015.

This document aims to stimulate more researchers in the IoT world to work towards developing IoT mechanisms for Dynamic Spectrum Access (DSA), because the need for enhanced connectivity of IoT devices nowadays is of outmost importance and existing techniques (without DSA) cannot ensure the high connectivity of the IoT devices. It is assumed that the progress in the IoT hardware in the near future will also assist towards the implementation and the testing of DSA techniques in real-world deployments. A first step towards this has been presented in this document with the SDR-based prototype, which will be enhanced and improved in the next months of the project. However, real-world testing of DSA techniques is not really possible in the TV white spaces due to the lack of regulations to allow transmissions in these spectrum bands. We hope that the EU regulations in the near future will be altered to allow such operation in the TV white spaces, so that IoT can really benefit from the capabilities of Cognitive Radio and ensure the optimum connectivity of the millions of devices in urban environments.

References

[3GPP06] 3GPP TS 23.234 V6.10.0 (2006-09), 3GPP system to Wireless Local Area Network (WLAN) interworking; System description (Release 6)

[3GPP10] 3GPP Evolved Universal Terrestrial Radio Access (EUTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall Description; Stage 2 (Release 10).

[3GPP11] 3GPP TS 22.368 v11.0.0, Service Requirements for Machine-Type Communications.

[6LOW] 6LoWPAN Linux implementation. [Online . Available: https://www.kernel.org/doc/Documentation/networking/ieee802154.txt]

[A++09] I. F. Akyildiz, , et al., ”Spectrum management in cognitive radio ad hoc networks,” IEEE Network, vol.23, no.4, pp.6-12, July 2009.

[A++12] A. Araujo et al., “Security in cognitive wireless sensor networks. challenges and open problems,” EURASIP Journal on Wireless Communications and Networking, vol. 2012, no. 1, pp. 1–8, 2012.

[AAPY15] V. Angelakis, I. Avgouleas, N. Pappas, and Di Yuan, “Flexible Allocation of Heterogeneous Resources to Services on an IoT Device,” IEEE INFOCOM’15 Student Workshop, Hong Kong, Apr. 2015

[ABM10] J. Audibert, S. Bubeck, and R. Munos, “Best arm identification in multiarmed bandits,” Proceedings of 23th Conference on Learning Theory COLT 2010, Haifa, Israel, 2010.

[ACF02] P. Auer, N. Cesa-Bianchi, and P. Fischer, “Finite-time analysis of the multiarmed bandit problem,” in Machine Learning Journal, 2002.

[ACN10] S. Alam, M. Chowdhury, J. Noll, Senaas: an event-driven sensor virtualisation approach for internet of things cloud, in IEEE International Conference on Networked Embedded Systems for Enterprise Applications (NESEA) (Suzhou (China), 2010), pp. 1–6. 2010

[AFRB+13] Abbas M, Fazirulhisyam H, Raja S. Azmir Raja A, Borhanuddin Mohd A, Mohamed O, Sabira K: Multicast-Unicast Data Delivery Method in Wireless IPv6 Networks. Journal of Network and Systems Management, February 2013

[AG99] Abowd, Gregory D., et al. "Towards a better understanding of context and context-awareness." Handheld and ubiquitous computing. Springer Berlin Heidelberg, 1999.

[AHS07] K. Aberer, M. Hauswirth, A. Salehi, Infrastructure for data processing in largescale interconnected sensor networks, in 2007 International Conference on Mobile Data Management (Mannheim (Germany), 2007), pp. 198–205.

[AKE09] O. B. Akan, O. Karli, O. Ergul, “Cognitive radio sensor networks,” IEEE Network, vol.23, no.4, pp.34-40, July- 2009.

[AKJ10] N. Ahmed, S. Kanhere, and S. Jha, “Mitigating the effect of interference in wireless sensor networks,” in Local Computer Networks (LCN), 2010 IEEE 35th Conference on, Oct 2010, pp. 160–167.

[AMS09] Audibert, Jean-Yves, Rémi Munos, and Csaba Szepesvári. "Exploration–exploitation tradeoff using variance estimates in multi-armed bandits." Theoretical Computer Science 410.19 (2009): 1876-1902

[AP++08] V. Angelakis, S. Papadakis, N. Kossifidis, V. A. Siris, A. Traganitis, The Effect of Using Directional Antennas on Adjacent Channel Interference in 802.11a: Modeling and Experience With an Outdoors Testbed, Proc. of WiNMee, Berlin, Germany, March 2008

[APG10] J. R. Gutierrez-Agullo, B. Coll-Perales, and J. Gozalvez, “An IEEE 802.11 MAC Software Defined Radio implementation for experimental wireless communications and networking research,” in Wireless Days (WD), 2010 IFIP. IEEE, 2010, pp. 1–5.

[ASA13] J. A. Abolarinwa, N. Salawu, and A. Achonu, “Cognitive radio-based wireless sensor networks as next generation sensor network: Concept, problems and prospects,” Journal of Emerging Trends in Computing and Information Sciences, vol. 4, no. 8, 2013.

[ASC02] I. Akyildiz, Y. S. W. Su, and E. Cayirci, “Wireless sensor networks: a survey,” in Comput. Networks (Elsevier), vol. 38, no. 4, 2002.

[ASSC02] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. A. Cayirci, survey on sensor networks. IEEE Commun. Mag. 2002, 40, 102-114.

[AT07] I. Akbar and W. Tranter, “Dynamic spectrum allocation in cognitive radio using hidden markov models: Poisson distributed case,” in SoutheastCon, 2007. Proceedings. IEEE, March 2007, pp. 196–201.

[AY07] Abbasi, Ameer Ahmed, and Mohamed Younis. "A survey on clustering algorithms for wireless sensor networks." Computer communications 30.14 (2007): 2826-2841.

[B++13] B. Bloessl et al., “An IEEE 802.11 a/g/p OFDM Receiver for GNU Radio,” in Proceedings of the second workshop on Software radio implementation forum. ACM, 2013.

[B++14] B. Bloessl et al., “Timings matter: standard compliant IEEE 802.11 channel access for a fully software-based SDR architecture,” in Proceedings of the 9th ACM international workshop on Wireless network testbeds, experimental evaluation and characterization. ACM, 2014.

[BCCC12] F. Bao, I. Chen, M. Chang, J. Cho, "Hierarchical Trust Management for Wireless Sensor Networks and its Applications to Trust-Based Routing and Intrusion Detection," Network and Service Management, IEEE Transactions on , vol.9, no.2, pp.169,183, June 2012

[BLKS14] Bithas, P.S.; Lioumpas, A.S.; Karagiannidis, G.K.; Sharif, B.S., "Interference minimization in hybrid WiFi/cellular networks," Computer Aided Modeling and Design of Communication Links and Networks (CAMAD), 2014 IEEE 19th International Workshop on , vol., no., pp.198,202, 1-3 Dec. 2014

[B03] Bhattacharyya, S. (Ed.) (2003). An overview of source-specific multicast (SSM). RFC 3569.

[BT96] D. Bertsekas and J. T. and, Neuro-Dynamic Programming. Athena Scientific, 1996.

[BC++] The Boost C++ Libraries. [Online . Available: http://www.boost.org]

[BD95] S. Bradtke and M. Duff, “Advances in neural information processing systems,” 1995.

[BF14] A Markovian approach for best-fit channel selection in cognitive radio networks. Suzan Bayhan, Alagöz Fatih. 0, s.l. : Ad Hoc Networks, 2014, Vol. 12.

[BG13] Boujelben H, Gaddour O: Enhancement and performance evaluation of a multicast routing mechanism in ZegBee cluster-tree Wireless Sensor networks. In: Systems, Signals & Devices (SSD), 2013 10th International Multi-Conference. 2013;1-8

[BI10] E. Borcoci, and R. Iorga, “A Management Architecture for a Multi-domain Content-Aware Network” TEMU 2010, July 2010, Crete, http://www.temu.gr/2010/.

[BL08] V. Barbu and N. Limnios, “Semi-markov chains and hidden semi-markov models toward applications.” Springer, 2008.

[BL11] John R. Birge, Francois Louveaux. Introduction to stochastic programming, 2ed. s.l. : Springer, 2011. ISBN 978-1-4614-0237-4.

[BLKS14] P. S. Bithas, A. S. Lioumpas, G. K. Karagiannidis and Bayan S. Sharif, "Interference minimization in Hybrid WiFi/Cellular Networks, "IEEE CAMAD, Athens, 2014.

[BM05] Biswas, Sanjit, and Robert Morris. "ExOR: opportunistic multi-hop routing for wireless networks." ACM SIGCOMM Computer Communication Review. Vol. 35. No. 4. ACM, 2005.

[BMS09] S. Bubeck, R. Munos, and G. Stoltz, "Pure exploration in multi-armed bandits problems," Algorithmic Learning Theory, 2009.

[BPM10] Barnaghi, Payam, Mirko Presser, and Klaus Moessner. "Publishing linked sensor data." CEUR Workshop Proceedings: Proceedings of the 3rd International Workshop on Semantic Sensor Networks (SSN), Organised in conjunction with the International Semantic Web Conference. Vol. 668. 2010.

[BPW04] J. Si, A. Barto, W. Powell, and D. Wunsch, Handbook of learning and approximate Dynamic Programming. Wiley-IEEE Press, 2004.

[BTT09] Bizer, Christian, Tom Heath, and Tim Berners-Lee. "Linked data-the story so far." (2009).

[C++05] D. Cabric, et al., "A cognitive radio approach for usage of virtual unlicensed spectrum," in proc. IST Mobile and Wireless Communications Summit, June, 2005.

[C10] P. Cardieri , "Modeling Interference in Wireless Ad Hoc Networks," IEEE Communications Surveys & Tutorials, vol.12, no.4, 2010.

[CB++10] A.P. Castellani, N. Bui, P. Casari, M. Rossi, Z. Shelby, M. Zorzi, Architecture and protocols for the internet of things: a case study, in: IEEE PerCom Workshops'10, 2010, pp. 678–683.

[CBB+12] Compton, Michael, et al. "The SSN ontology of the W3C semantic sensor network incubator group." Web Semantics: Science, Services and Agents on the World Wide Web 17 (2012): 25-32.

[CH++11] Choi, Jaeyoung, et al. "A survey on content-oriented networking for efficient content delivery." Communications Magazine, IEEE 49.3 (2011): 121-127..

[CH10] Clausen, T., & Herberg, U. (2010). Comparative study of RPL-enabled optimized broadcast in wireless sensor networks. In Proceedings of the sixth international conference on intelligent sensors, sensor networks and information processing (ISSNIP 2010). Brisbane, Australia: IEEE.

[CKK12] Carzaniga, A., Khazaei, K., & Kuhn, F. (2012). Oblivious low-congestion multicast routing in wireless networks. In Proceedings of the thirteenth international symposium on mobile ad hoc networking and computing (MobiHoc 2012).

[CSG10] de Castro, Rodrigo Soulé, and Philippe Godlewski. "An overview of DSA via multi-channel MAC protocols." Ad Hoc Networking Workshop (Med-Hoc-Net), 2010 The 9th IFIP Annual Mediterranean. IEEE, 2010.

[CSG10] de Castro, Rodrigo Soulé, and Philippe Godlewski. "An overview of DSA via multi-channel MAC protocols." Ad Hoc Networking Workshop (Med-Hoc-Net), 2010 The 9th IFIP Annual Mediterranean. IEEE, 2010..

[CSC09] Jin-Hee Cho; Swami, A.; Ing-Ray Chen; , "Modeling and Analysis of Trust Management for Cognitive Mission-Driven Group Communication Systems in Mobile Ad Hoc Networks," Computational Science and Engineering, 2009.2

[D08] Dempster, Arthur P. "The Dempster–Shafer calculus for statisticians." International Journal of Approximate Reasoning 48.2 (2008): 365-377.

[D++12] "An Internet of Things Platform for Real-World and Digital Objects", Journal of Scalable Computing: Practice and Experience, vol. 13, no.1, 2012.

[DC++05] De Couto, Douglas SJ, et al. "A high-throughput path metric for multi-hop wireless routing." Wireless Networks 11.4 (2005): 419-434.

[DD10] Dimitrakopoulos, George, and Panagiotis Demestichas. "Intelligent transportation systems." *Vehicular Technology Magazine, IEEE* 5.1 (2010): 77-84

[DBZ09] T. Q. Duong, V. N. Q. Bao, and H.-J. Zepernick, “On the performance of selection decode-and-forward relay networks over Nakagami-m fading channels,” IEEE Comm. Lett., vol. 13, no. 3, pp. 172-174, Mar. 2009.

[DFH99] Deering, S., Fenner, W., Haberman, B., Multicast Listener Discovery (MLD) for IPv6, RFC 2710, Oct. 1999

[DPH08] Das, S. M., Pucha, H., & Hu, Y. C. (2008). Distributed hashing for scalable multicast in wireless ad hoc networks. IEEE Transactions on Parallel and Distributed Systems, 19(3), 347–362.

[EPCG] EPCglobal. www.epcglobalinc.org/

[ETSI11] ETSI Machine-to-Machine Communications (M2M); Functional Architecture, Draft ETSI TS 102 690 v0.10.4 (2011-01).

[F++14] P. Di Francesco et al., “A Split MAC Approach for SDR Platforms,” Computers, IEEE Transactions on, vol. PP, no. 99, pp. 1–1, 2014.

[FAT14] Alexandros Fragkiadakis, Vangelis Angelakis, and Elias Z. Tragos, “Securing Cognitive Wireless Sensor Networks: A Survey,” International Journal of Distributed Sensor Networks, vol. 2014, Article ID 393248, 12 pages, 2014. doi:10.1155/2014/393248

[FR++05] R. Fonseca, S. Ratnasamy, J. Zhao, C.T. Ee, D. Culler, S. Shenker, I. Stoica, Beacon vector routing: scalable point-to-point routing in wireless sensornets, in: USENIX NSDI’05, 2005

[FTA13] Fragkiadakis, Alexandros G., Elias Z. Tragos, and Ioannis G. Askoxylakis. "A survey on security threats and detection techniques in cognitive radio networks." Communications Surveys & Tutorials, IEEE 15.1 (2013): 428-445.

[FW06] Flury, R., & Wattenhofer, R. (2006). MLS: An efficient location service for mobile ad hoc networks. In Proceedings of the 7th ACM international symposium on mobile ad hoc networking and computing (MOBIHOC) (pp. 226–237).

[FZDH12] Feng, C. H., Zhang, Y., Demirkol, I., & Heinzelman, W. B. (2012). Stateless multicast protocol for ad hoc networks. IEEE Transactions on Mobile Computing, 11(2), 240–253.

[G13] R. Gallager, Stochastic processes: Theory for applications. Cambridge University Press, 2013.

[GC07] Cornuéjols, Gérard. "Revival of the Gomory cuts in the 1990’s." Annals of Operations Research 149.1 (2007): 63-66.

[Geonames] GeoNames Ontology, http://www.geonames.org/ontology/documentation.html

[GK00] P. Gupta and P.R Kumar “The Capacity of Wireless Networks”, IEEE Transactions on information theory vol.46, no.2, March 2000.

[GNU] GNU Radio. [Online]. Available: http://gnuradio.org

[GQV08] Distributed energy efficient spectrum access in wireless cognitive radio sensor networks. S. Gao, L. Qian, and D. Vaman. s.l. : Wireless Communications and Networking Conference (WCNC'08), 2008.

[GRID] https://www.sbireports.com/Smart-Grid-Utility-2496610/

[GSRV08] V.C. Giruka, M. Singhal, J. Royalty, and S. Varanasi, “Security in wireless sensor networks”, Wireless Communications Mob. Comput., 2008; 8:1–24.

[GW11] V. Gligor and J. M. Wing, “Towards a theory of trust in networks of humans and computers,” in Proc. 19th International Workshop on Security Protocols, ser. LNCS. Springer Verlag, March 28-30 2011.

[HF02] Huang, J. and Fox, M. S.: An ontology of trust: formal semantics and transitivity. In Proceedings of the 8th international conference on Electronic commerce: The new ecommerce: innovations for conquering current barriers, obstacles and limitations to conducting successful business on the internet (ICEC '06). ACM, New York, NY, USA, pp. 259-270 (2006)

[HK14] Hui J, Kelsey R: Multicast Protocol for Low power and Lossy Networks (MPL), draft-ietf-roll-trickle-mcast-11, Nov. 2014

[HLZ09] Optimal Transmission Strategies for Dynamic Spectrum Access in Cognitive Radio Networks. Senhua Huang, Xin Liu, Ding Zhi. 12, s.l. : IEEE transactions on mobile computing, 2009, Vol. 8.

[HPM08] Performance improvement with predictive channel selection for cognitive radios. Höyhtyä M., Pollin S., Mämmelä A. s.l. : CogART, 2008.

[HPS09] C. Henson, J. K. Pschorr, A. P. Sheth, and K. Thirunarayan, "SemSOS: Semantic Sensor Observation Service," in Proc. of the 2009 International Symposium on Collaborative Technologies and Systems (CTS 2009), Baltimore, MD, 2009.

[HS++01] J. Heidemann, F. Silva, C. Intanagonwiwat, R. Govindan, D. Estrin, D. Ganesan, Building efficient wireless sensor networks with low-level naming, in: ACM Symp. on Oper. Syst. Principles (SOSP'01), 2001, pp. 146–159.

[IEEE06] "Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs), IEEE Standard, 802.15.4-2006", 2006.

[IEEE99] "Higher Speed Physical Layer Extension in the 2.4 GHz band", IEEE standard, 802.11b-1999, Feb, 1999

[IHLH12] Islam, Md Motaharul, et al. "A survey on virtualisation of wireless sensor networks." *Sensors* 12.2 (2012): 2175-2207.

[IOTA] A. Bassi, et. al, Enabling Things to Talk, Springer, ISBN: 978-3-642-40402-3 (Print) 978-3-642-40403-0 (Online)

[IY++11] M. Iqbal, D. Yang, T. Obaid, T.J. Ng, H.B. Lim, Demo abstract: a service oriented application programming interface for sensor network virtualisation, in Information Processing in Sensor Networks (IPSN),10th International Conference (Chicago, Illinois (USA), 2011), pp. 143–144. 2011

[J++12] Joshi, S., et al. s.l. Performance of channel bonding for opportunistic spectrum access networks. GLOBECOM, 2012.

[JHI07] A.P. Jayasumana, Q. Han, T.H. Illangasekare, Virtual sensor networks—a resource efficient approach for concurrent applications, in Fourth International Conference on Information Technology. ITNG '07 (Las Vegas, Nevada (USA), 2007), pp. 111–115. 2007

[JIB07] Jøsang, A., Ismail, R., and Boyd, C.: A survey of trust and reputation systems for online service provision. Decis. Support Syst. 43, 2, pp. 618-644 (Mar. 2007)

[JPCV12] Shaunak Joshi, Przemyslaw Pawelczak, Danijela Cabric, John Villasenor. S.l, .When channel bonding is beneficial for Opportunistic Spectrum Access Networks. : CoRR, 2012.

[JST09] V. Jacobson et al., "Networking Named Content," CoNEXT '09, New York, NY, 2009, pp. 1–12.

[KDHS10] Koutsonikolas, D., Das, S. M., Hu, Y. C., & Stojmenovic, I. (2010). Hierarchical geographic multicast routing for wireless sensor networks. Wireless Networks, 16(2), 449–466.

[KN++07] B. Kannhavong, H. Nakayama, Y. Nemoto, N. Kato, and A. Jamalipour, "A survey of routing attacks in mobile ad hoc networks", IEEE Wireless Communications, October 2007, pp. 85-91.

[KS08] Efficient Discovery of Spectrum Opportunities with MAC-Layer Sensing in Cognitive Radio Networks. Hyoil Kim, K.G. Shin. 5, s.l. : IEEE Transactions on Mobile Computing, 2008, Vol. 7.

[KW03] C. Karlof, and D. Wagner, "Secure Routing in Wireless Sensor Networks: Attacks and Countermeasures", Ad Hoc Networks, Elsevier, ed. Vol. 1, 2003, pp. 293–315.

[KW12] Alan J. King, Stein W. Wallace. Modeling with stochastic programming. S.l. : Springer, 2012. ISBN-13: 978-0387878164.

[KW94] Perter Kall, Stein Wallace. Stochastic Programming, 2ed. Chichester : John Wiley & Sons, 1994.

[L++10] Y. Li et al., "Spectrum usage prediction based on high-order markov model for cognitive radio networks," in Computer and Information Technology, 2010 IEEE 10th International Conference on, June 2010.

[L10] Locke, D., "MQ Telemetry Transport (MQTT) V3.1 Protocol Specification", IBM, Aug. 2010

[LA11] A Spectrum Decision Framework for Coginitive Radio Networks. Won-Yeol Lee, Ian F. Akyildiz. 2, s.l. : IEEE Transactions on mobile computing, 2011, Vol. 10.

[LAK07] K. Liu, N. Abu-ghazaleh, and K. D. Kang, "Location verification and trust management for resilient geographic routing," J. Parallel Distrib. Computing, vol. 67, no. 2, pp. 215–228, Feb. 2007

[LCHG11] Levis, P., Clausen, T.H., Hui, J., Gnawali, O., & Ko, J. (2011). The trickle algorithm. RFC 6206.

[LDP12] Y. Liu, Q. Du, P. Ren, "A Multi-Channel Sensing Order Optimization Algorithm Based on Markov Prediction," in proc. WiCOM, 2012 pp.1-4, Sept. 2012

[LF12] P. Li and Y. Fang "On the throughput capacity of heterogeneous wireless networks," IEEE Trans. Mobile Comput., vol. 11, no. 12, pp. 2073-2086, Dec. 2012.

[LGC12] C.-H. Liu, W. Gabran, and D. Cabric, "Prediction of exponentially distributed primary user traffic for dynamic spectrum access," in GLOBECOM, 2012 IEEE, Dec 2012, pp. 1441–1446.

[LHAY14] M. Lundgren, D. Helgesson, V. Angelakis, and D. Yuan, Energy aware rate selection in cognitive radio inspired wireless smart objects, IEEE Symposium on Computers and Communications (ISCC), 2014, Jun. 2014, Madeira

[LIN09] H.B. Lim, M. Iqbal, T.J. Ng, A virtualisation framework for heterogeneous sensor network platforms, in Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems (SenSys'09) Berkeley (USA, 2009), pp. 319–320. Doi:10.1145/1644038.1644080

[LDAT] Linked Data – Connect Distributed Data across the Web, http://linkeddata.org/

[LJDKM00] Li, J., Jannotti, J., De Couto, D. S. J., Karger, D. R., & Morris, R. (2000). A scalable location service for geographic ad hoc routing. In Proceedings of the 6th annual international conference on mobile computing and networking, MobiCom'00 (pp. 120–130).

[LK13] On the capacity of wireless CSMA/CA Multihop Networks. Rafael Laufer, Leonard Kleinrock. S.l. : IEEE INFOCOM, 2013.

[LL05] Lefort, Laurent. "Ontology for quantity kinds and units: units and quantities definitions." W3 Semantic Sensor Network Incubator Activity (2005).

[LPCS04] Levis, P., Patel, N., Culler, D., & Shenker, S. (2004). Trickle: A self-regulating algorithm for code propagation and maintenance in wireless sensor networks. In Proceedings of the first USENIX/ACM symposium on networked systems design and implementation (NSDI) (pp. 15–28).

[LSB03] S. Lee, R. Sherwood, and B. Bhattacharjee.: Cooperative Peer Groups in NICE. In IEEE Infocom, San Francisco, CA, Apr. (2003)

[Lwi] Linux wireless. [Online . Available: http://wireless.kernel.org]

[LWM11] Residual energy aware channel assignment in cognitive radio sensor networks. X. Li, D. Wang, and J. McNair. S.l. : IEEE Wireless Communications and Networking Conference (WCNC'11), 2011.

[LYL14] C. H.Liu, B. Yang, and T. Liu. Efficient naming, addressing and profile services in Internet-of-Things sensory environments. Ad Hoc Networks, 18, 85-101, 2014.

[M2M] "Machine-to-Machine (M2M) – The Rise of the Machines," Juniper Networks Whitepaper, 2011.

[M95] J. Mitola, “The software radio architecture,” Communications Magazine, IEEE, no. May, pp. 26–38, 1995.

[MAB09] Malik, Z., Akbar, I., and Bouguettaya, A.: Web Services Reputation Assessment Using a Hidden Markov Model. In Proceedings of the 7th international Joint Conference on Service-Oriented Computing (Stockholm, November 24 – 27, 2009). L. Baresi, C. Chi, and J. Suzuki, Eds. Lecture Notes In Computer Science, vol. 5900. Springer-Verlag, Berlin, Heidelberg, pp. 576-591 (2009)

[MAS11] Machine-to-machine device connections: worldwide forecast 2010–2020, Analysys Mason report December 2010.

[MAS14] T. Rebbeck, M. Mackenzie, N. Alfonso, Low-powered wireless solutions have the potential to increase the M2M market by over 3 billion connections, Analysys Mason report September 2014

[MB13] R. Miruta and E. Borcoci (UPB) “Optimization of Overlay QoS Constrained Routing and Mapping Algorithm for Virtual Content Aware networks” ICNS 2013 – Lisbon, Portugal. http://www.thinkmind.org/index.php?view=article&articleid= icns_2013_4_30_10174.

[MM99] Cognitive Radio: Making software radios more personal. Joseph Mitola, Gerarld Q. Maguire Jr. 4, s.l. : IEEE, Personal Communications, 1999, Vol. 6.

[MPPT12] N. Mitton, S. Papavassiliou, A. Puliafito, A and K. S. Trivedi, Combining Cloud and sensors in a smart city environment. EURASIP Journal on Wireless Communications and Networking, 2012(1), 1-10.

[MRT09] A. J. Mersereau, P. Rusmevichientong, and J. N. Tsitsiklis, “A structured multiarmed bandit problem and the greedy policy,” in Proc. IEEE Transactions on Automatic Control, vol. Vol. 54, No. 12, 2009.

[MW++11] Damnjanovic, J. Montojo, Y. Wei, T. Ji, T. Luo, M. Vajapeyam, T. Yoo, O. Song, and D. Malladi, “A survey on 3GPP heterogeneous networks,” IEEE Wireless Commun., vol. 18, no. 3, pp. 10-21, Jun. 2011.

[N++09] G. Nychis et al., “Enabling MAC Protocol Implementations on Software-Defined Radios.” In NSDI, vol. 9, 2009, pp. 91–105.

[NFEV] http://netfutures2015.eu/rerum/

[NS00] Noto, Masato, and Hiroaki Sato. "A method for the shortest path search by extended Dijkstra algorithm." Systems, Man, and Cybernetics, 2000 IEEE International Conference on. Vol. 3. IEEE, 2000.

[NSDR] Noctar PCIe SDR card. [Online . Available : http://www.pervices.com]

[NXP14] NXP Semiconductors: “I2C-bus specification and user manual”, Rev. 6, 2014

[O80] A. Osborne, An Introduction to Microcomputers Volume 1: Basic Concepts, Osborne-McGraw Hill Berkeley California USA, 1980

[OASIS06] MacKenzie, C. Matthew, et al. “Reference model for service oriented architecture 1.0.” OASIS Standard 12 (2006).

[OGC07] OGC, “Open Geospatial Consortium (OGC) Sensor Web Enablement: Overview and High Level Architecture,” OGC white paper2007.

[OIM05] Okura, A., Ihara, T., & Miura, A. (2005). BAM: branch aggregation multicast for wireless sensor networks. In: IEEE international conference on mobile ad hoc and sensor systems conference (pp. 354–363).

[OP12] Oikonomou G, Philips I: Stateless Multicast Forwarding with RPL in 6LoWPAN Sensor Networks. In: Proc.2012 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops). Lugano, Switzerland (2012). 2012; 272-277

[OPCA] OPC Unified Architecture (UA). <www.opcfoundation.org/ua/>.

[OPT13] Oikonomou G, Philips I, Tryfonas T: Ipv6 Multicast Forwarding in RPL-Based Wireless Sensor Networks. Wireless Personal Communications. 2013; 73(3): 1089-1116

[OWL] OWL-based Web service ontology, http://www.ai.sri.com/daml/services/owl-s/1.2/

[PBZ13] V. A. Poenaru, E. Borcoci and M. Zotea. Topology and Network Resources Discovery Protocol for Content-Aware Networks. In INNOV 2013, The Second International Conference on Communications, Computation, Networks and Technologies (pp. 23-29), Nov. 2013.

[PHPS10] Pschorr, Josh, et al. "Sensor discovery on linked data." (2010).

[PHY08] "Proactive channel access in dynamic spectrum networks," Physical Communication, vol. 1, no. 2, pp. 103 – 111, 2008.

[PL05] C. Park and K. Lahiri, "Battery discharge characteristics of wireless sensor nodes: An experimental analysis," in In Proceedings of the IEEE Conf. on Sensor and Ad-hoc Communications and Networks (SECON, 2005.

[PP++09] Performance Analysis of Multichannel Medium Access Control Algorithms for Opportunistic Spectrum Access. Przemyslaw Pawelczak, Sofie Pollin, Hoi-Sheung Wilson So, Ahmad Bahai, Venkatesh Prasad, Ramin Hekmat. S.l. : IEEE transactions on vehicular technology, 2009. Vol. 58.

[PPC11] Peformance of Joint Spectrum Sensing and MAC Algorithms for Multichannel Opportunistic Spectrum Access Ad Hoc Networks. Jihoon Park, Przemyslaw Pawelczak, Danijela Cabric. 7, s.l. : IEEE transactions on mobile computing, 2011, Vol. 10.

[PPGC12] Park, Jihoon, et al. "Analysis framework for opportunistic spectrum OFDMA and its application to the IEEE 802.22 standard." Vehicular Technology, IEEE Transactions on 61.5 (2012): 2271-2293..

[QZ++07] L. Qiu, Y. Zhang, F. Wang, M. K. Han, and R. Mahajan, A General Model of Wireless Interference, Proc. Of MobiCom, September 2007.

[R52] H. Robbins, "Some aspects of the sequential design of experiments." Bulletin of the American Mathematical Society, pp. 58:527–535, 1952.

[RD22] J. Cuellar et. al, System requirements and smart objects model, RERUM Deliverable D2.2, May 2014

[RD23] E. Tragos et. al. System Architecture, RERUM Deliverable D2.3, August 2014

[RD31] D. Ruiz et. al. Enhancing the autonomous smart objects and the overall system security of IoT based Smart Cities, RERUM Deliverable D3.1, April 2014

[RE07] A. Rezgui, and M. Eltoweissy, "TARP: A Trust-Aware Routing Protocol for Sensor-Actuator Networks," Mobile Adhoc and Sensor Systems, 2007. MASS 2007. IEEE International Conference on , vol., no., pp.1,9, 8-11 Oct. 2007

[RERUM] EU FP7 ICT2013-SMARTCITIES-RERUM. Resilient, Robust and Secure IoT for smart city applications, http://www.ict-rerum.eu

[RKL++05] Roman, Dumitru, et al. "Web service 146odelling ontology." Applied ontology 1.1 (2005): 77-106.

[RP03] Raskin, Rob, and Michael Pan. "Semantic web for earth and environmental terminology (sweet)." Proc. Of the Workshop on Semantic Web Technologies for Searching and Retrieving Scientific Data. 2003.

[RP99] Royer, E. M., & Perkins, C. E. (1999). Multicast operation of the ad-hoc on-demand distance vector routing protocol. In Proceedings of the 5th annual ACM/IEEE international conference on Mobile computing and networking (MobiCom'99) (pp. 207–218).

[S06] T. Schmid, "GNU Radio 802.15.4 En- and Decoding," 2006.

[S09] G. Santucci, “From Internet of Data to Internet of Things,” in International Conference on Future Trends of the Internet, Jan. 2009.

[S76] Shafer G. A Mathematical Theory of Evidence. Princeton University Press: Princeton, New Jersey, USA, 1976.

[SA12] M. Z. Shakir and M.-S. Alouini “On the area spectral efficiency improvement of heterogeneous network by exploiting the integration of macrofemto cellular networks,” IEEE International Conference on Communications, pp. 5695-5700, Jun. 2012.

[SB98] R. Sutton and A. Barto, “Reinforcement learning: An introduction.” The MIT Press. [Online]

[SBL08] Byun, Sang-Seon, Ilangko Balasingham, and Xuedong Liang. “Dynamic spectrum allocation in wireless cognitive sensor networks: Improving fairness and energy efficiency.” Vehicular Technology Conference, 2008. VTC 2008-Fall. IEEE 68th. IEEE, 2008.

[SBPW04] J. Si, A. Barto, W. Powell, and D. Wunsch, Handbook of learning and approximate Dynamic Programming. Wiley-IEEE Press, 2004.

[SCK10] Song, S., Choi, B. Y., Kim, D. (2010). MR. BIN: Multicast routing with branch information nodes for wire- less sensor networks. In Proceedings of the 19th international conference on computer communications and networks (ICCCN 2010) (pp. 1–6).

[SCPR+08] Sá, Silva J., Camilo, T., Pinto, P., Ruivo, R., Rodrigues, A., Gaudêncio, F., et al. (2008). Multicast and ip multicast support in wireless sensor networks. Journal of Networks, 3(3), 19–26.

[SDR14] Alexander Shapiro, Darinka Dentcheva, Andrej Ruszczynski. Lectures on stochastic programming: Modeling and Theory. S.l. : MOS-SIAM, 2014. ISBN-13: 978-1611973426.

[SF02] Sentz, Kari, and Scott Ferson. Combination of evidence in Dempster-Shafer theory. Vol. 4015. Albuquerque, New Mexico: Sandia National Laboratories, 2002.

[SF96] Standard, Federal. "1037C." Department of Defence Dictionary of Military and Associated Terms in support of MIL-STD-188 (1996).

[SHB14] Shelby, Z., Hartke, K., Bormann, C., “The Constrained Application Protocol (CoAP)”, RFC 7252, June 2014

[SHC++04] Simulating the power consumption of large-scale sensor network applications. Victor Shnayder, Mark Hempstead, Bor-rong Chen, Geoff Werner Allen, Matt Welsh. Baltimore, MD, USA : Proceedings of the 2nd international conference on Embedded networked sensor systems ACM, 2004.

[SHL08] Y.L. Sun, Z. Han, and K. J. Ray Liu, “Defense of Trust Management Vulnerabilities in Distributed Networks”, IEEE Communications Magazine, Vol. 25, No.2, February 2008, pp. 112-119.

[SKC09] Song, S., Kim, D., & Choi, B. Y. (2009). AGSMR: Adaptive geo-source multicast routing for wireless sensor networks. In Proceedings of the wireless algorithms, systems, and applications, Lecture notes in computer science (Vol. 5682, pp. 200–209). Springer:Berlin/Heidelberg.

[SKS01] C. Schurgers, G. Kulkarni, M.B. Srivastava, Distributed assignment of encoded mac addresses in sensor networks, in: ACM MobiHoc’01,2001, pp. 295-298.

[SL++10] Q. Shen, Y. Liu, Z. Zhao, S. Ci, H. Tang, Distributed hash table based id management optimization for internet of things, in: IEEE IWCMC’10, 2010, pp. 686–690

[SMP15] M. Surligas, A. Makrogiannakis, and S. Papadakis, Empowering the IoT Heterogeneous Wireless Networking with Software Defined Radio, Workshop on Heterogeneous Networking for the Internet of Things, VTC Spring 2015, Glasgow, May 11-14 2015.

[SMPR+12] Sanchez, J. A., Marin-Perez, R., & Ruiz, P. M. (2012). Beacon-less geographic multicast routing in a real-world wireless sensor network testbed. Wireless Networks, 18(5), 565–578.

[SMY12] H. A. Suraweera, D.S. Michalopoulos and C. Yuen "Performance analysis of fixed gain relay systems with a single interferer in Nakagami-m fading channels," IEEE Trans. Veh. Technol., vol. 61, no. 3, pp. 1457-1463, Mar. 2012.

[SPI] Serial Peripheral Interface: http://en.wikipedia.org/wiki/Serial_Peripheral_Interface_Bus

[SPS99] R. Sutton, D. Precup, and S. Singh, "Between mdps and semi-mdps: A framework for temporal abstraction in reinforcement learning," 1999.

[SR14] Primary radio user activity models for cognitive radio networks: A survey. Yasir Saleem, Mubashir Husain Rehmani. 2014, Journal of Network and Computer Applications, pp. 1-16.

[SRLS07] Sanchez, J. A., Ruiz, P. M., Liu, J., & Stojmenovic, I. (2007). Bandwidth-efficient geographic multicast routing protocol for wireless sensor networks. IEEE Sensors Journal, 7(5), 627–636.

[SS05] Sabater, J. and Sierra, C.: Review on Computational Trust and Reputation Models. Artif. Intell. Rev. 24, 1, pp. 33-60 (Sep. 2005)

[SS09] L. Sørensen and K. E. Skouby (eds.), "User scenarios 2020 – a worldwide wireless future," OUTLOOK – Visions and research directions for the Wireless World, Wireless World Research Forum, no.4, July 2009.

[SS10] Sabater, J. and Sierra, C.: REGRET: A reputation model for gregarious societies. In Proceedings of the 4th Int. Workshop on Deception, Fraud and Trust in Agent Societies, in the 5th Int. Conference on Autonomous Agents (AGENTS'01), pp. 61-69, Montreal (2001)

[SS90] Skiena, S. "Dijkstra's Algorithm." *Implementing Discrete Mathematics: Combinatorics and Graph Theory with Mathematica, Reading, MA: Addison-Wesley* (1990): 225-227.

[SSS07] T. Schmid, O. Sekkat, and M. B. Srivastava, "An experimental study of network performance impact of increased latency in software defined radios," in Proceedings of the second ACM international workshop on Wireless network testbeds, experimental evaluation and characterization. ACM, 2007, pp. 59–66.

[STT14] Stamatakis, George, Elias Z. Tragos, and Apostolos Traganitis. "Energy efficient collection of spectrum occupancy data in wireless cognitive sensor networks." Wireless Communications, Vehicular Technology, Information Theory and Aerospace & Electronic Systems (VITAE), 2014 4th International Conference on. IEEE, 2014.

[STT15] Stamatakis, George, Elias Z. Tragos, and Apostolos Traganitis, Periodic collection of spectrum occupancy data by energy constrained cognitive IoT devices, IWCMC 2015 Coop-Cognitive Workshop (submitted, under review)

[STT15b] Stamatakis, George, Elias Z. Tragos, and Apostolos Traganitis, A two-stage Power and QoS aware Dynamic Spectrum Assignment scheme for Cognitive Wireless Sensor Networks, QoMex 2015 (submitted, under review)

[SWEET] Sweet ontology: Semantic Web for Earth and Environmental Terminology: http://sweet.jpl.nasa.gov/

[SZ10] C. Song and Q. Zhang, "Intelligent dynamic spectrum access assisted by channel usage prediction," in INFOCOM IEEE Conference on Computer Communications Workshops , 2010, March 2010, pp. 1–6.

[SZ10] C. Song and Q. Zhang, "Intelligent dynamic spectrum access assisted by channel usage prediction," in INFOCOM IEEE Conference on Computer Communications Workshops , 2010, March 2010, pp. 1–6.

[TA++11] K. Tan et al., “Sora: High-Performance Software Radio using General-Purpose multi-core Processors,” Communications of the ACM, vol. 54, no. 1, 2011.

[T++11] E. Z. Tragos, et al., “The Impact of Interference on the Performance of a Multi-path Metropolitan Wireless Mesh Network,” in Prof. of IEEE ISCC, pp.199-204, 2011

[T++14] E. Tragos et al., “Enabling reliable and secure iot-based smart city applications,” in IEEE PerCity 2014, 2014.

[TA14] F. P. Talebi. An information centric networking approach to context-aware dissemination of services and information, 2014, Ph. D. Thesis, UBC, Vancouver.

[T14] Thubert, P., The IPv6 Flow Label within a LLN domain, draft-thubert-6man-flow-label-for-rpl-05, August, 2014

[TA13] E. Tragos and V. Angelakis, “Cognitive Radio Inspired M2M Communications (Invited Paper),” in IEEE Global Wireless Summit 2013.

[TFAS11] E. Z. Tragos, A. Fragkiadakis, I. Askoxylakis, V. A. Siris, "The impact of interference on the performance of a multi-path metropolitan wireless mesh network", ISCC, 2011, 2013 IEEE Symposium on Computers and Communications (ISCC), 2013 IEEE Symposium on Computers and Communications (ISCC) 2011, pp. 199-204, doi:10.1109/ISCC.2011.5983840

[TZFS13] E. Tragos, S. Zeadally, A. Fragkiadakis, and V. Siris, “Spectrum assignment in cognitive radio networks: A comprehensive survey,” IEEE Communications Surveys and Tutorials, vol. 15, no. 3, 2013.

[UCO12] G. S. Uyanik, B. Canberk, S. Oktug, “Predictive spectrum decision mechanisms in Cognitive Radio Networks,” in proc. IEEE Globecom, pp.943-947, Dec. 2012.

[V++11] O. Vermesan, etal., “Internet of Things Strategic Research Roadmap”, IERC Whitepaper, 2011.

[VC04] Vida, R., Costa, L., Multicast Listener Discovery Version 2 (MLDv2) for IPv6, RFC 3810, June 2004

[VF14] O. Vermesan, P. Friess, Internet of Things: From Research and Innovation to Market Deployment. River Publishers, 2014.

[VMB++04] S. Victor, H. Mark, C. Bor-Rong, A. G. Werner, and W. Matt, “Simulating the power consumption of large-scale sensor network applications,” in Proceedings of the 2Nd International Conference on Embedded Networked Sensor Systems, ser. SenSys ’04. New York, NY, USA: ACM, 2004, pp. 188–200. [Online

[VIT67] Viterbi, Andrew J. "Error bounds for convolutional codes and an asymptotically optimum decoding algorithm." Information Theory, IEEE Transactions on 13.2 (1967): 260-269

[W++11] D1.2 – Initial Architectural Reference Model for IoT. IoT-A Deliverable.

[W++12] J. Wan, et al., "M2M Communications for Smart City: An Event-Based Architecture," in proc.IEEE 1 Computer and Information Technology (CIT), pp.895,900, 27-29 Oct. 2012.

[W3C11] W3C. (2011, W3C SSN Incubator Group Report. Available: http://www.w3.org/2005/Incubator/ssn/wiki/Incubator_Report

[WM07] Lessons learned from an extensive spectrum occupancy measurement campaign and a stochastic duty cycle model. Matthias Wellens, Petri Mähönen. Orlando : International conference on Cognitive Radio oriented Wireless Networks and Communications, Crowncom, 2007.

[WRM09] Empirical Time and Frequency Domain Models of Spectrum Use. Matthias Wellens, Janne Riihijarvi, Petri Mahonen. s.l. : Elsevier Physical Communications, 2009, Vol. 2.

[WTBH+12] Winter (editor), T., Thubert (editor), P., Brandt, A., Hui, J., Kelsey, R., Levis, P., et al. (2012). RPL: IPv6 Routing Protocol for Low power and Lossy Networks. RFC 6550.

[WW11] J. Wamicha and S. Winberg, “IEEE 802.11 OFDM Software Defined Radio beacon frame transmission,” in AFRICON, 2011. IEEE, 2011.

[X++11] G. Xilouris, et al., FP7 ICT project, Deliverable D2.1: "ALICANTE Overall System and Components Definition and Specification", September 2011, http://www.ict-ALICANTE.eu/

[XJL08] Optimal Bandwidth Selection in Multi-Channel Cognitiver Radio Networks: How much is too much? Dan Xu, Eric Jung, and Xin Liu. s.l. : Dynamic Spectrum Access Networks (DYSPAN'08), 2008.

[XJL12] Efficient and Fair Bandwidth Allocation in Multichannel Cognitive Radio Networks. Dan Xu, Eric Jung, and Xin Liu. s.l. : IEEE Transactions on mobilie computing , 2012.

[YB++07] Allocating dynamic time-spectrum blocks for cognitive radion networks. Yuan Yuan, Paramvir Bahl, Ranveer Chandra, Thomas Moscibroda, Yunnan Wu. s.l. : Mobihoc'07, 2007.

[YCZ08] Proactive channel access in dynamic spectrum networks. Yang L., Cao L., Zheng H. s.l. : Physical Communication, 2008.

[YK10] M. Yuriyama, T. Kushida, Sensor-cloud infrastructure—physical sensor management with virtualised sensors on cloud computing, in 13th International Conference on Network-Based Information Systems (NBiS) (,Takayama (Japan), 2010), pp. 1–8. 2010

[YKD06] Z. Yao, D. Kim, and Y. Doh, “PLUS: parameterized and localized trust management scheme for sensor networks security,” inProc. 2006 IEEE Int. Conf. Mobile Adhoc Sensor Syst., pp. 437–446.

[YLH10] Spectrum allocation algorithm in cognitive ad hoc networks with high energy-efficiency. L. Yu, C. Liu, and W. Hu. s.l. : Green Circuits and Systems (ICGCS), International conference, 2010.

[Z++12] Y. Zhang, et al., “Cognitive machine-to-machine communications: visions and potentials for the smart grid,” IEEE Network, vol.26, no.3, pp.6-13, May 2012.

[Z1] Zolertia Z1 Platform. [Online. Available: http://zolertia.com]

[Z99] Zacharia, G.: Collaborative Reputation Mechanisms for Online Communities. Master’s thesis, Massachusetts Institute of Technology (1999)

[ZSD12] G. Zhan, W. Shi, and J. Deng, "Design and Implementation of TARF: A Trust-Aware Routing Framework for WSNs," *Dependable and Secure Computing, IEEE Transactions on* , vol.9, no.2, pp.184,197, March-April 2012

[ZYI04] Q. Zhang, T. Yu, K. Irwin: A classification scheme for trust functions in reputationbased trust management, in: Proceedings of ISWC Workshop on Trust, Security, and Reputation on the Semantic Web (2004)

[ZZ11] D. Zhao and X. Zhou, “Spectrum sensing using prior probability prediction,” in Wireless Communications, Networking and Mobile Computing (WiCOM), 2011 7th International Conference on, Sept 2011, pp. 1–4.

