



Quality Of Service and MObility driven cognitive radio Systems

FP7-ICT-2009-4/248454

QoS MOS

D6.5

Specification of cognitive and opportunistic functions of the spectrum management framework

Contractual Date of Delivery to the CEC: 30/04/2012	
Actual Date of Delivery to the CEC: 08/06/2012	
Editor(s): Bernd Bochow, Marc Emmelmann (Fraunhofer)	
Internal Reviewers:	Santosh Kawade (BT), Klaus Moessner (UNIS)
External Reviewers:	Uwe Beutnagel-Buchner, Shyamalie Thilakawardana
Author(s):	Bernd Bochow (Fraunhofer), Ingo Karla (ALD), Peter Bakki, László Csurgai-Horváth (BME), Ramiro Samano Robles (IT), Pål Gronsund (TEL), Youngwook Ko, Miguel López-Benítez (UNIS)
Workpackage:	WP6
Est. person months:	40
Security:	PU
Nature:	R
Version:	Rev. 1.0
Total number of pages:	74

Abstract:

This deliverable provides the description of cognitive and opportunistic functions of the spectrum management framework enabling QoS and mobility. It builds on deliverable D6.3 initially defining scope, goals and limits of cognitive functions within the framework, and is complemented by deliverable D6.4 which elaborates on trust, security, privacy and reliability and robustness. This deliverable concludes the description of the CM-SM and provides an informal specification of cognitive functions and self-learning capabilities of the framework.

Keyword list: Cognitive spectrum management, Specification of cognitive functions, Opportunistic functions, Self-learning, CM-SM functions, interfaces and protocols



Abbreviations

ACK	Acknowledge
AL	Adaptation Layer
BER	Bit Error Rate
BLER	BLock Error Rate
CCDF	Complementary Cumulative Distribution Function
CM-RM	Cognitive Manager – Resource Manager
CM-SM	Cognitive Manager – Spectrum Manager
CPFR	Common Portfolio Repository
CIR	Carrier-to-Interference Ratio
CNR	Carrier-to-Noise Ratio
CSPC	Common Spectrum Control
DLC	Data Link Control
DTV	Digital TeleVision
DUR	Desired-to-Undesired power Ratio
FTP	File Transfer Protocol
GRGR	Global Regulator Repository
HTML	Hypertext Markup Language
HTTP	Hypertext Transfer Protocol
IEEE	The Institute of Electrical and Electronics Engineers
IPTV	Internet Protocol Television
LSPC	Local Spectrum Control
LPFR	Local Portfolio Repository
LTE	Long Term Evolution
MAC	Medium Access Control
MSC	Message Sequence Chart
PMSE	Program Making and Special Events
QoS	Quality of Service



RAN	Radio Access Network
RAT	Radio Access Technology
SAN	Spectrum Analyser
SINR	Signal-to-Interference plus Noise Ratio
SPRR	Spectrum Provider Repository
SNR	Signal-to-Noise Ratio
SPI	Spectrum Efficiency Index
SSE	Spectrum Selector
TCP	Transmission Control Protocol
TDMA	Time Division Multiple Access
TVWS	TV White Space
VoIP	Voice over IP
WAP	Wireless Application Protocol
RM	Resource Manager
SM	Spectrum Manager
SSE	Spectrum Selector
WLAN	Wireless Local Area Network



Table of content

1	INTRODUCTION.....	9
1.1	SCOPE AND OBJECTIVES.....	9
1.2	ORGANIZATION OF THE DOCUMENT	9
2	FUNCTIONAL DECOMPOSITION OF THE CM-SM REFERENCE MODEL.....	10
3	OVERVIEW OF CM-SM AND CM-RM COLLABORATION.....	12
3.1	INTRODUCTION	12
3.2	CM-SM TO CM-RM INTERWORKING.....	12
3.3	INFORMATION EXCHANGE	13
4	REPOSITORIES.....	16
4.1	GLOBAL REGULATOR REPOSITORY (GRGR)	16
4.1.1	<i>Functions</i>	16
4.1.2	<i>Interfaces</i>	17
4.2	COMMON PORTFOLIO REPOSITORY (CPFR).....	17
4.2.1	<i>Functions</i>	17
4.2.2	<i>Interfaces</i>	18
4.3	SPECTRUM PROVIDER REPOSITORY (SPRR)	18
4.3.1	<i>Functions</i>	18
4.3.2	<i>Interfaces</i>	19
4.4	COMMON POLICY REPOSITORY (CPOR)	20
4.4.1	<i>Functions</i>	20
4.4.2	<i>Interfaces</i>	20
4.5	LOCAL PORTFOLIO REPOSITORY (LPFR).....	21
4.5.1	<i>Functions</i>	21
4.5.2	<i>Interfaces</i>	22
5	COMMON SPECTRUM CONTROL (CSPC).....	23
5.1	FUNCTIONS IN CELLULAR SCENARIOS	23
5.2	FUNCTIONS IN AD-HOC AND FEMTOCELL SCENARIOS	26
5.3	OPPORTUNITY DETECTION FUNCTIONS IN THE CSPC.....	28
5.4	MULTI-OBJECTIVE PORTFOLIO OPTIMIZATION IN CSPC	29
5.5	INTERFACES	29
6	LOCAL SPECTRUM CONTROL (LSPC)	32
6.1	LSPC FUNCTIONS (OPERATOR LSPC)	32
6.2	LSPC FUNCTIONS (USER-EQUIPMENT LSPC).....	36
6.3	OPPORTUNITY DETECTION AND SPECTRUM PORTFOLIO MANAGEMENT FUNCTIONS IN THE LSPC	40
6.4	DISTRIBUTED SELF-ORGANIZING COGNITIVE-RADIO SPECTRUM MANAGEMENT	42
6.4.1	<i>Challenges for self-organized cognitive spectrum management</i>	42
6.4.2	<i>Distributed SON for cognitive radio spectrum managers</i>	43
6.4.3	<i>Prediction model</i>	46
6.4.4	<i>Next Steps</i>	50
6.5	COGNITIVE SPECTRUM UTILIZATION FOR STABLE, DENSE INDOOR FEMTOCELLS	50
6.6	INTERFACES	50
7	SPECTRUM ANALYSER AND SELECTOR	53
7.1	SAN FUNCTIONS	53



7.2	SSE FUNCTIONS	55
7.3	INTERFACES	56
8	SUMMARY AND CONCLUSIONS.....	59
9	REFERENCES	61
A	OPPORTUNITY SELECTION	63
A.1	SELECTION OF DTV BANDS FOR LTE UPLINK EXTENSION.....	63
A.1.1	INTRODUCTION	63
A.1.2	CONSIDERED SCENARIO AND MOTIVATION	63
A.1.3	CONFORMANCE METRICS	64
B	USER ACTIVITY MODELS	66
B.1	INTRODUCTION	66
B.2	LONG-TERM ACTIVITY MODEL FOR INCUMBENTS AND OPPORTUNISTIC USERS	66
B.3	ON/OFF ACTIVITY BASED MODEL.....	67
B.3.1	ON/OFF MARKOV-CHAIN MODEL FOR INCUMBENTS ACTIVITY	67
B.3.2	ESTIMATING FROM SPECTRUM SENSING.....	67
B.3.3	ESTIMATING FROM PACKET TRAFFIC OBSERVATION	68
B.3.4	AGGREGATION OF THE INCUMBENTS ACTIVITY	69
B.3.5	ON/OFF MODEL FOR OPPORTUNISTIC USER ACTIVITY.....	69
B.3.6	APPLICATION OF THE MODEL IN THE CM-SM DECISION-MAKING SYSTEM	70
B.3.7	AN EXAMPLE OF USING THE ON/OFF MODEL IN THE IEEE 802.22 STANDARD	70
B.4	SHORT TERM USER ACTIVITY BASED MODEL.....	71



List of figures

Figure 2-1: QoS MOS reference model of the CM-SM	10
Figure 2-2: Interfaces between coexistence and coordination domain entities over adaptation layer ..	11
Figure 3-1: High level perspective of CM-SM and CM-RM tasks, functions and responsibilities	13
Figure 3-2: Signalling between CM-SM and CM-RM.....	14
Figure 3-3: Signalling between CM-SM and CM-RM across domains	15
Figure 4-1: Accessing the GRGR and sample MSC	16
Figure 4-2: Accessing the CPFR and sample MSC.....	18
Figure 4-3: Accessing the SPRR and sample MSC.....	19
Figure 4-4: Accessing the CPOR and sample MSC.....	20
Figure 4-5: Accessing the LPFR and sample MSC.....	21
Figure 5-1: Accessing the CSPC in cellular scenarios and sample MSC.....	25
Figure 5-2: Accessing the CSPC in femtocell scenarios and sample MSC (portfolio update for cellular, portfolio deployment for femtocell)	27
Figure 5-3: Accessing the CSPC in ad hoc scenarios and sample MSC (role switch from ad-hoc node A to ad-hoc node B)	28
Figure 6-1: Accessing an operator LSPC and sample MSC (responding to a portfolio change request by a coordination domain CM-SM)	36
Figure 6-2: Accessing a user-equipment LSPC and sample MSC (fast update of shared spectrum portfolio by requesting the SSE, involving CM-RM, SAN and SS for context acquisition)	39
Figure 6-3: Schematic example of interactions and couplings in modern cognitive radio systems.	42
Figure 6-4: Schematic illustration of the distributed CM-SM architecture.....	43
Figure 6-5: Schematic drawing of a cellular layout, with one example of the “local area”.	44
Figure 6-6: Illustration of the signalling around the CM-SM	44
Figure 6-7: Sketch of the kind of information which are exchanged via signalling messages between the different distributed CM-SMs	45
Figure 6-8: Illustration of virtual cell internal sub-areas.....	47
Figure 6-9: Simplified operation of the virtual scheduler to predict the result a possible candidate parameter set.....	48
Figure 7-1: Accessing a SAN and sample MSC (SAN configured to create a portfolio from raw sensor data for storing to LPFR).....	54
Figure 7-2: Accessing an SSE and sample MSC (initialization and update of an SSE-deployed spectrum portfolio involving SS and CM-RM context updates).....	58
Figure A-1: Considered scenario.....	64
Figure A-2: Relation between CNR distance and DUR distance: a) proper operation of all DTV receivers, b) improper operation of some DTV receivers.	65
Figure B-1: Mapping the user activity model to QoS MOS CM-SM reference model.....	66
Figure B-2: One-day spectrum measurement.....	67



Figure B-3: Signal strength and ON/OFF activity	68
Figure B-4: Network traffic measurement	68
Figure B-5: Network traffic and ON/OFF activity	68
Figure B-6: Aggregated OFF length duration distribution for different number of users.....	69
Figure B-7: The ON/OFF model for opportunistic users	69
Figure B-8: CCDF of the aggregated opportunistic user ON activity for different user numbers	70
Figure B-9: Opportunity detection in CSPC	70
Figure B-10: Media access timing assumed	71
Figure B-11: User model data traffic generator	72
Figure B-12: Terminology: link utilization and access opportunities	73
Figure B-13: Multiplexing shared media access	73
Figure B-14: Traffic analysis by observing spectrum utilization and identifying traffic patterns from observation.	74
Figure B-15: Simplified algorithm to estimate user activity from traffic patterns observed.....	74



Executive Summary

This deliverable provides a description of cognitive and opportunistic functions of the spectrum management framework. It builds on deliverable [D6.3] initially defining scope, goals and limits of cognitive functions within the framework, and is complemented by deliverable [D6.4] elaborating on trust, security, privacy and reliability and robustness of cognitive capacity. This deliverable concludes the description and specification of cognitive functions and self-learning capabilities of the framework.

Starting from the QoS MOS reference model as specified by deliverables [D2.1], [D2.2] and [D2.3] and elaborating further on the reference model of the Cognitive Manager – Spectrum Manager (CM-SM) an informal specification of internal functional entities of the CM-SM and their interaction in a distributed environment is given. This description respects the QoS MOS scenarios as specified in [D1.2] and thus also provides a discussion on realization options for some specific scenarios and on performance issues. This discussion will be concluded in the upcoming deliverables D6.6 (Spectrum management framework integration and implementation report) and D6.7 (Integrated final functional specification of spectrum management framework and procedures).

The informal specifications given include an architectural view of interacting CM-SM entities, their individual functional capacities foreseen, and their (most simplified) interaction on the interface protocol level. Background details on functions and algorithms are provided by a number of annexes for better understanding the approach taken and to realize the complexity of cognitive spectrum management in the context of QoS MOS.

The discussion of the CM-SM architectural model, its functional entities and their interaction in this deliverable first presents the roles and functions of the various repositories storing and acting upon policies and spectrum portfolios. Next the domain model comprising coexistence, coordination and networking scopes is introduced. Functional entities then are discussed within their specific allocation to a domain while elaboration on functions provided to other domains and functions required from these. Different configuration options are discussed, which allows targeting the specific requirements of each of the various QoS MOS scenarios.

Although formal specifications of interface primitives and message formats have been prepared up to a level required for a proof of concept with respect to the most important functions of the CM-SM, they have not yet been included here. Message sequence charts provided as an example in this deliverable are derived and simplified from these and are currently tested in a reference implementation.



1 Introduction

1.1 Scope and Objectives

This deliverable concludes the description of cognitive and opportunistic functions of the spectrum management framework. It grounds upon deliverable [D6.3] initially defining scope, goals and limits of cognitive functions within the framework, and is complemented by deliverable [D6.4] elaborating on trust, security, privacy and reliability and robustness of cognitive capacity. This deliverable provides a description and informal specification of cognitive functions and self-learning capabilities of the framework.

The description and specification is based on the functional decomposition of the CM-SM (Cognitive Manager – Spectrum Manager) reference model, which provides a co-location model for the cognitive functions studied. Hence cognitive functions are described in the context they are used within, and in their interaction in a distributed environment. The latter is detailed further in the scope of interfaces involved in the exchange of context and control information between distributed instances, which is complementing the interaction through a shared environment as addressed earlier by [D6.3] and [D2.3].

1.2 Organization of the document

This document is organized into a main part and a number of annexes discussing further the approaches considered for realising the cognitive and opportunistic functions of a distributed Cognitive Manager – Spectrum Manager (CM-SM):

First, an overview of the generic interaction between QoS MOS cognitive managers (CM-SM and CM-RM) and interfaces involved is given. The decomposition of the QoS MOS cognitive spectrum manager (CM-SM) into its internal functional modules and related interfaces is presented, giving a concise picture of the CM-SM reference model as discussed in the scope of [D6.2] and [D2.2] here focusing on the cognitive capacity of the CM-SM.

Next, an informal specification of the cognitive capacity of the CM-SM with respect to functional modules given by the reference model and the role of related interfaces, as well as the exchange of information across these interfaces is elaborated upon with more detail.

The specification first details the databases of spectrum portfolios and spectrum policies, their functional role in the context of the CM-SM architecture, their internal functionality and the content they manage. In particular, functionality that goes beyond mere database functionality is elaborated in more detail.

The cognitive spectrum management functionality co-located with coordination and networking domain is presented next. This specification considers interaction between entities of the two domains within and across domains. In that it considers the main QoS MOS scenarios regarding cellular, femtocell and ad-hoc configurations with respect to their impact on the cognitive decision-making functions and strategies, context considered and output produced.

Further detail on the concepts, approaches and solutions is provided in the Annex, which forms the grounds for specifications given by this document and helps to picture the intricacies only briefly addressed in the scope of specifying functions and interfaces.

The document concludes by providing a brief summary and outlook towards a proof-of-concept realization touching the issue of performance metrics and testing and assessment of the cognitive functions of the QoS MOS CM-SM.



2 Functional Decomposition of the CM-SM Reference Model

The QoS MOS CM-SM reference model describes the topology and functionality of the QoS MOS cognitive spectrum management system. The architecture is kept modular to suit the scenarios defined, allowing for future expansions to support upcoming licensed and license-exempt radio and network technologies.

A functional decomposition of the reference model is shown by Figure 2-1, providing an overview of the relation between QoS MOS functional entities distributed to coexisting networks for the scenarios specified. Coexistence here refers to coexistence between cellular (wide-area and femtocell-based) and ad-hoc network applications.

The reference model defines different domains dedicated to providing functions to support coexistence in shared spectrum (coexistence domain), coordinating between shared spectrum users (and between networks of those, considering also coordination with networks following a different architectural approach), and managing infrastructures of wireless communication systems (networking domain) as well as wireless access networks and end-systems (terminating domain). Thus, domains have co-location, functional, topology and stakeholder aspects.

Extensibility of the system is maintained through distributing few functional entities across those domains, which determines the functions that must be provided by the specific functional entity regarding their role and capacity. Spectrum portfolio repositories, for example, may serve a dedicated stakeholder (e.g. a regulator, certification authority, operator) or may serve as a dedicated function (e.g. as a local spectrum pool or as a spectrum trader's database).




Figure 2-1: QoS MOS reference model of the CM-SM



Extensibility of the system is also maintained through well-defined logical interfaces between entities and domains. Two distinct types of interfaces must be considered here: Interfaces between functional entities and the Adaptation Layer (AL) and Interfaces between functional entities. This document is focusing on the functional entities of the CM-SM reference model and the interfaces between those functional entities. The AL is described in more detail in Deliverables [D2.1] and [D2.2] – Figure 2-2 is emphasizing on this structure.




Figure 2-2: Interfaces between coexistence and coordination domain entities over adaptation layer



3 Overview of CM-SM and CM-RM collaboration

3.1 Introduction

The QoS MOS system architecture (as documented and specified in [D1.2], [D2.1], [D2.2], [D2.3]) defines two main cognitive entities: the CM-RM mainly acting as a radio resource manager, and the CM-SM focusing on dynamic spectrum management. The CM-RM is operating with knowledge obtained from spectrum sensing and from the state of the wireless access system (i.e. the entities of the terminating domain and the networking domain). It is able to respond to state changes and resource requests within a short timeframe. The CM-SM operates on knowledge about spectrum utilization, spectrum efficiency and spectrum availability. As a dynamic spectrum management system it is responding within a much larger timeframe than the CM-RM.

Although no experimental results are available yet it is reasonable to assume that the CM-RM will operate in the sub-100ms range, while the CM-SM will show characteristic response times of 100ms...5s (networking domain), 1s...60s (coordination domain) and above 1min. up to days or weeks (coexistence domain). These figures currently are guestimates derived from simulations and are up to be confirmed by proof-of-concept experiments. For example, in a hypothetical scenario which comprises of fixed users (incumbents) and mobile users (opportunistic users) in the same UHF frequency band, the CM-SM is responsible for control of the spectrum allocation and the CM-RM is guaranteeing operation of the incumbents by collecting, storing and processing information and performing decision processes on the spectrum usage of the opportunistic users. Given that an opportunistic user moves with the speed of 60 km/h and the simulation area has the size of 35*35 km with 4 transmitters for incumbents, then a 1 minute joint CM-RM and CM-SM response time is sufficient for mal-usage detection and reaction (e.g. by revoking a spectrum portfolio). Further results will be addresses in upcoming deliverables D6.6 and D6.7.

The CM-SM thus relies upon context information provided by the CM-RM and responds to requests of the CM-RM to provide an amount of spectrum for consideration in radio resource management. In the general case, a CM-SM responds to spectrum requests by multiple CM-RMs and is responsible to optimize the deployment of spectrum to multiple radio access systems.

Hence, the CM-RM is a resource management entity focusing on the immediate demand of wireless access systems for spectrum resources, while the CM-SM is planning spectrum utilization across wireless access systems immediately regarding the policies given by stakeholders such as operators, regulators and spectrum traders.

In the following, a short overview of the interaction between CM-SM and CM-RM is given.

3.2 CM-SM to CM-RM interworking

Figure 3-1 depicts the interworking between CM-SM and CM-RM: The CM-RM manages spectrum and radio resources close to the physical layer on a comparatively short time scale. In a cellular network the CM-RM could be close to the cell resource scheduler, for example, having access to detailed short term information about the current situation in the cell. In that the CM-RM manages and operates within the spectrum resources and associated constraints given by the spectrum portfolio allotted by a collaborating CM-SM (cf. e.g. sect. 4).

The CM-SM is composing a spectrum portfolio based on context information obtained (among other sources) from the CM-RM in response to a spectrum resource request issued by the CM-RM. The spectrum portfolio deployed in response to such a request provides information about frequency bands for disposal to the CM-RM along with usage constraints (e.g. in form of policies) regarding, for example, acceptable transmission power and adjacent band emission limits.

The CM-SM obtains averaged, filtered context information from the CM-RM and takes into account external constraints such as information from spectrum sensing or Geolocation databases when



composing a spectrum portfolio. Decision-making in the scope of the CM-SM in consequence operates on a much longer term than the CM-RM. For a cellular network CM-SM could be related to a Self-Organizing Network (SON) entity, or to an Operation and Maintenance (O&M) centre, for example.

The CM-RM operates close to the actual resource assignment of the radio channel, such as the cell resource scheduler, on a time scale of milliseconds typically in the range of a few to 100ms. The CM-SM operates on longer time scales that depend on the particular use case, on the network operator's strategy preferences and on the domain the CM-SM is situated. Typical time scales could be in the range of seconds to several hours.

In addition to periodic operation procedures, the CM-SM can also be triggered by certain events. A change of traffic load may trigger the CM-RM to urgently request additional spectrum, or a change in spectrum availability indicated by spectrum databases may trigger the CM-SM to revoke and re-organize spectrum allocation. The CM-SM may need to react quickly on those triggers and may need to provide immediate resolution to an upcoming congestion situation. Acceptable response time upper limits and suitable resolution strategies strongly depend on the specific event since maintaining QoS for mobile users may demand for a (nearly) seamless handover between spectrum portfolios.

3.3 Information exchange

Basic signalling between a CM-SM and a CM-RM is outlined in Figure 3-2. Signalling across domains is more detailed by Figure 3-3 emphasizing the propagation of context from coordinating to terminating domains through multiple instances of CM-SM and CM-RM.




Figure 3-1: High level perspective of CM-SM and CM-RM tasks, functions and responsibilities




Figure 3-2: Signalling between CM-SM and CM-RM

Collaboration between CM-RM and CM-SM relies on the exchange of both status information and configuration information. Status information mainly flows from CM-RM to CM-SM: The CM-RM provides status information that enables the CM-SM to reason and decide on a suitable spectrum portfolio for this CM-RM. In a cellular system, for example, the CM-RM could provide information about cell load and how well a certain part of the spectrum was utilized.

Configuration information may originate from different sources such as network management or local control and management applications, but is conveyed mainly from a CM-SM to the CM-RM. Since there is no direct configuration command involved in this communication, the CM-RM derives its configuration from the information and constraints included in a portfolio (e.g. transmission power and adjacent band emission allowed). If conveyed towards a reasoning engine, a portfolio constitutes a set of facts provided by the cognitive engine of the CM-SM to collaborate with the cognitive engine of a CM-RM.

The exchange of facts between CM-SM and CM-RM further enables collaborative decision-making. A CM-RM may suggest strategies or may provide hints to the cognitive engine of the CM-SM. In case of increasing spectrum utilization, for example, the CM-RM may consider to request more spectrum resources to satisfy its resource demands. This could be done by explicitly (i.e. actively) requesting to enlarge its spectrum portfolio, or by continuously providing information about the level of spectrum utilization (i.e. its current spectral load) for having the CM-SM to choose a different spectrum portfolio composition strategy for the requesting CM-RM that leverages higher safety margins in spectrum allocation and eliminates the need for rapid requests of additional spectrum. Vice-versa, a CM-SM may request that behaviour to enable learning and planning capacities in its own cognitive processing.

Interworking between CM-SM and CM-RM mainly takes place in the networking domain (see Figure 2-1). Hence, Figure 3-2 focusses on the interaction of CM-SM and CM-RM in the networking domain. Signalling across domains (see Figure 3-3) takes place in case a local spectrum request cannot be satisfied (e.g. cannot be provided by local repositories LPFR). Resource requests originating from a CM-RM then need to be redirected towards the coordination domain or up to the coexistence domain. While requesting spectrum portfolios from a CM-SM instance in the coordination domain may be satisfied by a spectrum management procedure involving only the operator, or potentially also involves a spectrum trader, a request to the coexistence domain may not even result in deploying a spectrum portfolio but rather may result in a request to the regulator that there is a need to reconsider policies that limit spectrum utilization. Such a request may be forwarded to the management of a global repository (e.g. that of a coexistence domain CM-SM or of a GRGR, cf. sect. 4.1) and may



cause an automated or manual regulatory response. These are considered external interfaces out of scope for this deliverable and may be further discussed in the scope of QoS MOS WP1.




Figure 3-3: Signalling between CM-SM and CM-RM across domains



4 Repositories

4.1 Global Regulator Repository (GRGR)

4.1.1 Functions

The Global Regulator Repository (GRGR) is associated with a single regulatory domain (e.g. continent, country, city and city district) and provides information about spectrum availability and spectrum usage constraints regarding geographical areas within this domain. It is usually restricted in its extent regarding frequency bands addressed. A GRGR may instantiate as a table providing current spectrum regulations in a machine-readable format, as well as a database that can be queried actively (e.g. a TV white space Geolocation database). The GRGR may also instantiate as a single entity or in a distributed way where a CM-SM may access the GRGR via one out of many dedicated service access points. Some of these architectural and topological options may be subject to local regulations, for example, demanding a specific hierarchical organisation having a regulator's database controlling third party databases. The GRGR then may be instantiated in form of one of these databases or as a proxy or gateway to a distributed database infrastructure.

A CM-SM may access more than one GRGR entity simultaneously when operating across regulatory domains to support mobility between areas associated with different regulatory domains or different regulatory authorities.

When querying a GRGR a CM-SM must provide the geographical area and frequency band its request applies to. When responding a GRGR is expected to provide information about at least one contiguous frequency band within the frequency band queried along with following information:

- Responsible authority and applicable geographical area;
- Current spectrum licensee (incumbent technology, standard or stakeholder, channelization);
- Usage constraints (power constraints, spectrum mask, duty cycle, technology, policies);

Since realizations (i.e. vendor specific implementations) of a GRGR may behave different or may implement proprietary interfaces depending on local decisions of the operator or provider of a GRGR, a CM-SM of the coexistence domain is required to access the GRGR and to convert the response of a query to the GRGR into a spectrum portfolio representation. A gateway function is required to abstract the access to the specific GRGR implementation. It is up to the specific implementation if this gateway is considered a function of the GRGR or of the 'enclosing' CM-SM (Figure 4-1).

The abstraction of the GRGR implementation allows certain scenarios where requesting multiple GRGR instances, requesting GRGR and SPRR in parallel, or having the GRGR querying the SPRR on its own in advance of deploying a valid (i.e. qualified) spectrum portfolio is feasible to simplify procedures to merge the information retrieved from the GRGR and from the SPRR (Spectrum Provider Repository, cf. sect. 4.3). Potential communication between GRGR and SPRR is considered private and will not be addressed by this document.




Figure 4-1: Accessing the GRGR and sample MSC



4.1.2 Interfaces

The **SM1 interface** is used to exchange spectrum portfolios between a CM-SM entity and its associated spectrum portfolio repositories. The SM1 interface is provided by a lightweight implementation of a CM-SM wrapping one or more potentially proprietary implementations of a repository and performing necessary adaptations (e.g. by utilizing a gateway function). This interface splits between SM1a and SM1b. While SM1a is realized between CM-SM and its portfolio repositories for all instances of the CM-SM, SM1b is available in addition to SM1a only for instances of the CM-SM that are realized for the coexistence domain and for those CM-SM instances, that are in trusted collaboration with regulatory CM-SM instances. SM1a is dedicated to the exchange of credentials between CM-SM instances while SM1b is dedicated to the exchange of portfolios optionally containing credentials linked with that portfolio.

The **AL1 interface** is used to exchange information between distributed CM-SM entities and the QoSMOS adaptation layer. The AL1 interface splits between AL1a through AL1f. It is used as a management and control interface in this context to support the exchange of spectrum portfolios across domain boundaries. Its main purpose here is to identify, associate and exchange information with entities (e.g. suitable policy repositories, [1900.5],[1900.5.1]) based on some selection criteria such as applicable regulatory domain. It is an interface of the QoSMOS reference model.

The **AL1f** (GRGR-AL) control interface supports read-only access to GRGR contents.

4.2 Common Portfolio Repository (CPFR)

4.2.1 Functions

The Common Portfolio Repository (CPFR) is a dynamic (potentially distributed) database providing spectrum portfolios in the process of deploying spectrum to spectrum users, or consuming spectrum portfolios after revoking spectrum from spectrum users. Its main function is to provide spectrum portfolios to a CM-SM instance for further processing such as deploying spectrum to one or more spectrum users or performing split and merge operations prior to deployment. It keeps track of portfolios already deployed to spectrum users, which enables a CM-SM to revoke spectrum from spectrum users and to make it available to other spectrum users.

In certain ad hoc scenarios, a CPFR may serve as a temporary storage for exchanging spectrum portfolios between spectrum users sharing spectrum or for spectrum trading. In addition, it may aggregate information obtained from spectrum measurements to support a CM-SM in creating spectrum portfolios from spectrum sensing information.

The CPFR is the main repository that provides spectrum portfolios upon request of operators CM-SM entities from the coordination or networking domain. The CPFR receives spectrum portfolios from a co-located SPRR entity or from other CM-SM entities (e.g. a regulatory CM-SM or spectrum trader's CM-SM). In contrast to the SPRR, the CPFR is dynamic in nature since it reflects the current spectrum utilization context in form of spectrum portfolios deployed, spectrum portfolios currently not in use and spectrum utilization context derived from spectrum sensing regarding spectrum portfolios in use (e.g. interference situation).

A CPFR may store portions of spectrum portfolios across different databases to support efficient database implementations (e.g. using dedicated databases for frequency band descriptions, usage constraints, policies, licensee information, financial information, and geographical areas applicable). The CPFR may need to store portfolios already deployed in complete (potentially in a dedicated physical database) for various reasons:

- A portfolio may contain credentials tightly linked with the other information contained in a portfolio when composed and deployed to a spectrum user (e.g. certificates validating authority, serial number, lease time, amount of spectrum and spectrum mask), which is a coordination domain CM-SM from the perspective of the CPFR.



- A portfolio may be linked with a specific spectrum user potentially becoming a protected user by obtaining spectrum usage rights in form of a portfolio (e.g. PMSE devices utilizing TV white space and, depending on local spectrum regulations, attaining incumbent status through their operator's incumbent status), or it may be linked with specific technologies potentially including a relaxation of usage constraints for a specific technology.
- When revoking a portfolio, the portfolio under consideration must be referenced by some unique identifier used in communication with a spectrum user for technical reasons (e.g. reducing communication overhead) or for legal reasons (e.g. to implement non-repudiation).

It should be noted that spectrum portfolio revocation bears some timing considerations. Revoking a portfolio usually is a response either to an administrative (or regulatory) action or to an exceptional situation such as upturning malicious users or defective devices. While the former usually is a planned action that can be aligned with timing constraints, the latter requires applying de-escalating strategies. One option is to deploy a (sub-optimal) spectrum portfolio having a strictly limited lease time before revoking the existing portfolio, and before deploying a new spectrum portfolio. This allows mitigating the impact of a portfolio revocation and potentially avoids idling or shutting down infrastructure nodes (e.g. switching down base stations or putting them into maintenance mode).




Figure 4-2: Accessing the CPFR and sample MSC

4.2.2 Interfaces

The **SM1a interface** is used to exchange spectrum portfolios between a CM-SM entity and its associated spectrum portfolio repositories. Since the CPFR is associated with an operator's CM-SM the SM1b interface is not provided. That is, an operator's CM-SM does not provide credentials but can obtain credentials from a regulatory CM-SM by using the SM1b interface provided by a regulatory CM-SM.

The **AL1 interface** is used to exchange information between distributed CM-SM entities and the QoSMOS adaptation layer. The AL1 interface splits between AL1a through AL1f. It is used as a management and control interface in this context to support the exchange of spectrum portfolios across domain boundaries. Its main purpose here is to identify, associate and exchange information with entities (e.g. suitable policy repositories, [1900.5], [1900.5.1]) based on some selection criteria such as applicable regulatory domain. It is an interface of the QoSMOS reference model.

The **AL1e (CPFR-AL)** control interface provides read-only access to CPFR contents.

4.3 Spectrum Provider Repository (SPRR)

4.3.1 Functions

The Spectrum Provider Repository (SPRR) is a trusted entity either situated in the scope of a regulator, operator or spectrum trader. It is a database providing spectrum portfolios to a CM-SM instance for further processing.



Regulators will want to co-locate SPRR entities with GRGR entities to integrate policies or other usage constraints information with a portfolio prior to deploying such that spectrum users can be obliged to respect regulatory constraints when utilizing a spectrum portfolio obtained from a regulatory CM-SM instance. Spectrum portfolios coordinated through a regulatory SPRR are considered to represent temporary spectrum usage rights. A regulator in consequence may limit spectrum portfolios to specific technologies, licensees or further usage constraints.

Operators and spectrum traders will want to co-locate SPRR entities with CPFR entities for enabling fine-grained spectrum management. Spectrum portfolios coordinated through an operator's or spectrum trader's SPRR follow requirements set by network management and (dynamic) spectrum management systems of an operator in that they allocate and distribute spectrum portfolios as requested by entities of the networking and terminating domains.

An SPRR is a supporting entity utilized by a CM-SM. It usually does not realize the SM1 interface on its own but through its associated CM-SM. It may be implemented as a proprietary database storing complete spectrum portfolios, or may be implemented in form of a distributed database storing parts of spectrum portfolios. An SPRR, for example, may store frequency band descriptions, spectrum mask descriptions and policies across dedicated databases. A CM-SM then may follow a certain strategy to compose a spectrum portfolio from related portions according to operator's rules in that respecting regulator's constraints.

In spectrum trading scenarios the SPRR also stores and provides financial information about spectrum usage as well as spectrum usage rights constraints such as geographical area applicable, lease times, spectrum owners and subscribers or licensees.




Figure 4-3: Accessing the SPRR and sample MSC

4.3.2 Interfaces

The **AL1 interface** is used to exchange information between distributed CM-SM entities and the QoSMOS adaptation layer. The AL1 interface splits between AL1a through AL1f. It is used as a management and control interface in this context to support the exchange of spectrum portfolios across domain boundaries. Its main purpose here is to identify, associate and exchange information with entities (e.g. suitable policy repositories [1900.5], [1900.5.1]) based on some selection criteria such as applicable regulatory domain. It is an interface of the QoS MOS reference model.

The SPRR is not accessible via a public interface. Communication between GRGR and SPRR as well as between CPFR and SPRR is implemented through proprietary interfaces and through adaptation layer communication via the AL1e and AL1f interfaces.

The **AL1e** (SPRR-AL) control interface provides read-only access to the contents of the SPRR.



4.4 Common Policy Repository (CPOR)

4.4.1 Functions

The Common Policy Repository (CPOR) is associated with an operator's CM-SM of the coordination domain. It is used to store spectrum usage constraints complementing those included with spectrum portfolios obtained from coexistence domain entities such as regulators, spectrum traders or operators. In addition it keeps track on those spectrum portfolios deployed that have been amended by CPOR functions.

The CPOR enables spectrum sharing scenarios by amending spectrum portfolios. Its main function is to further constraint policies included with spectrum portfolios to enable sharing in the spatial, temporal or spectrum domains. When receiving a spectrum portfolio from an associated CM-SM it applies one or more policies stored to this portfolio. In that it adds further usage constraints to the portfolio. Policies to apply are selected by the CM-SM along with its request to modify a portfolio. In addition, the CPOR may implement reasoning capacities to ensure non-conflicting modifications to policies that may cause policy enforcement to intervene when utilizing a spectrum portfolio later on (see also [1900.5]).

Policies stored in the scope of a CPOR relate to entities of the networking and terminating domains. They will be implemented by those entities as a means to enable dynamic spectrum management across heterogeneous access networks and technologies. An operator may want to implement policies through a CPOR that increase spectrum efficiency (e.g. through spatio-temporal spectrum reuse), service-specific spectrum utilization (e.g. through scheduling mobile users to dedicated spectrum), or balance co-existence (e.g. through spatial interference mitigation).

Usage constraints introduced with actions of the CPOR may include but are not limited to restricting lease times, limiting frequency bands, valid geographical areas, technologies or spectrum masks, and requirements for spectrum sensing and incumbent protection (e.g. eviction delay when an incumbent is detected, or an obligation to query a Geolocation database prior to spectrum access).




Figure 4-4: Accessing the CPOR and sample MSC

4.4.2 Interfaces

The **PF1 interface** is used to exchange policies between the Common Spectrum Control (CSCP) function and the Common Policy Repository (CPOR) [1900.5]. It is an CM-SM internal interface of coordination domain entities. The data structures exchanged over the PF1 interface are spectrum portfolios consisting only of policies and related information required to determine the scope in that those policies apply (e.g. area, time or frequency band). Operator-managed spectrum usage constraints can be retrieved from the CPOR or stored to the CPOR via this interface.

The **AL1 interface** is used to exchange information between distributed CM-SM entities and the QoSMOS adaptation layer. The AL1 interface splits between AL1a through AL1f. It is used as a management and control interface in this context to support the exchange of spectrum portfolios across domain boundaries. Its main purpose here is to identify, associate and exchange information with entities (e.g. suitable policy repositories [1900.5], [1900.5.1]) based on some selection criteria such as applicable regulatory domain. It is an interface of the QoSMOS reference model.



The **AL1d** (CPOR-AL) control interface provides read-only access to CPOR contents.

4.5 Local Portfolio Repository (LPFR)

4.5.1 Functions

The Local Portfolio Repository (LPFR) is associated with a CM-SM of the networking domain. Networking domain CM-SM entities are co-located with dedicated management nodes or with network controllers such as a cellular base station controller or a WLAN access point. The LPFR hence is considered a local storage keeping track of spectrum portfolios obtained from CM-SM entities of the coordination domain. Spectrum portfolios kept by the LPFR are upon request deployed to CM-RM entities in the networking domain that in turn implement portfolios through their associated entities of the terminating domain.

In Addition, the LPFR stores information obtained from spectrum sensing and from associated CM-RM entities in form of spectrum portfolios for the purpose of keeping track of context (i.e. the radio scene) of the environment spectrum portfolios have been deployed to. This context information supports cognitive functions of local spectrum management (i.e. reasoning and decision-making as well as learning) and eventually is forwarded to the coordination domain enabling to evaluate and potentially revise earlier decisions of the cognitive spectrum management of each domain.




Figure 4-5: Accessing the LPFR and sample MSC



The LPFR is distributed by nature since portfolios (including policies and context related to portfolios) are stored in a scope of local relevance. That is, an LPFR co-located to a certain network controller may have access to topologically neighbouring entities (e.g. to base stations of geographically neighbouring cells) but not to the full infrastructure associated with a coordination or coexistence domain entity. When asked to deploy a spectrum portfolio to an associated CM-RM a CM-SM may utilize topological information about terminating domain entities controlled by this CM-RM as well as spectrum utilization information of portfolios deployed to neighbouring CM-RMs to optimize spectrum efficiency and interference metrics. This context is maintained by the LPFR through tagging portfolios deployed.

Communication with the LPFR takes place through an LSPC entity of the associated CM-SM (see section 6.1), except when co-located with a network controller. For this special flavour of a CM-SM (denoted as CM-SM END) the LPFR provides portfolios to an SSE entity (see section 7) and obtains context information from a SAN entity (see section 7.1). SSE and SAN are detached functions that can greatly enhance the performance of the LPFR by local caching consequently lowering significantly the response time to a CM-RM portfolio request.

4.5.2 Interfaces

The **LPFC interface** is used to exchange spectrum portfolios between the Local Portfolio Repository (LPFR) and the Local Spectrum Control (LSPC). It is an CM-SM internal interface of networking domain entities. Except for CM-SM END entities the LPFC interface is the only way to access the LPFR for storing and retrieving deployable spectrum portfolios.

The **PF2 interface** is used to exchange spectrum portfolios between the portfolio processors Spectrum Analiser (SAN) and Spectrum Selector (SSE), and the Local Portfolio Repository (LPFR). It is an CM-SM internal interface of networking domain entities and applies to CM-SM END entities only.

The **PF2 (LPFR-SSE)** interface is used to retrieve spectrum portfolios (i.e. the LPFR deploys spectrum portfolios to an SSE entity).

The **PF2 (LPFR-SAN)** interface is used to store spectrum portfolios (i.e. the LPFR obtains spectrum portfolios from a SAN entity).



5 Common spectrum control (CSPC)

5.1 Functions in cellular scenarios

The Common Spectrum Control (CSPC) is associated with an operator's CM-SM in the coordination domain which, in the cellular case, is situated in an operator's core network. A single CSPC instance is responsible for a network or part of the network sharing the same context. It acts as a centralized spectrum management entity supported by one or more localized LSPC entities in the networking domain. A minimum of one CSPC instance per operator is assumed. In consequence of a network design and management decision, there may exist additional CSPC instances dedicated, for example, specifically to one operator's RAT or RAN. In that case, the interaction between those instances regarding spectrum management towards the networking domain should be kept on a minimum level. An operator may benefit from multiple CSPC instances if spectrum allotment or spectrum usage rights are valid for a wide area or have been made technology independent (e.g. in re-farming or pooling scenarios).

The CSPC implements a number of functions for manipulating spectrum portfolios including at least:

- Interfacing with coexistence domain entities via the SM1a/b interface.
 - Request spectrum portfolios, policies and spectrum information from coexistence domain entities via the SM1a interface
 - Request or provide credentials via the SM1b interface (mandatory if mutual authentication is required, otherwise optional).
- Interfacing with networking domain entities via the SPC1 interface.
 - Deploy spectrum portfolios to networking domain entities upon request of those networking domain entities or upon request of coexistence domain entities.
 - Revoke spectrum portfolios from networking domain entities in consequence of earlier deploying updated spectrum portfolios or upon request of coexistence domain entities.
- Interfacing with an instance of the CPOR via the PF1 interface.
 - Request a CPOR to apply operator's policies to a spectrum portfolio prior to deploying this portfolio to networking domain entities.
 - Add or remove operator's policies to/from a CPOR.
- Cognitive functions to compose spectrum portfolios according to requests of networking domain entities and to the constraints set by coexistence domain entities prior to request a CPOR to apply operator policies.
- Collaboration and cooperation functions with other instances of coordination domain CM-SM instances of the same or of other operator's for the purpose of sharing or trading spectrum through coexistence domain entities (e.g. operator, for intra-operator coordination, or spectrum trader, for inter-operator coordination).

Optionally, the CSPC may choose to forward spectrum information requests received from networking domain entities towards coexistence domain entities if the information requested is not available at the CSPC, or it may decide to forward (based upon operator's policies) measurement information obtained from networking domain entities and indirectly also from terminating domain entities to the coexistence domain. This mediator function is required since networking domain entities cannot directly communicate with coexistence domain entities, and since coordination domain entities cannot communicate directly among each other (except using proprietary interfaces). This is due to the



requirements for trusted association (i.e. authentication) and communication of entities in the coexistence domain to accept input from other domains entities.

In the case of cooperation between CM-SM instances of the coordination domain, a CSPC is also involved when conveying information between networking domain entities of different operators, such as for exchanging spectrum measurements. An exchange of policies may take place between CSPC entities of the same operator in case a new CM-SM entity is introduced or if a CM-SM was temporarily disabled (e.g. for maintenance reasons). For example, if a CM-SM in the coordination domain becomes active initially or after some downtime, it requires an update of operator's policies. Keeping in mind that the policy management and utilization mainly is a reasoning process, it might be more convenient to synchronise policies between distributed CSPC entities rather than managing policies in a central location – even if policy rules are static on their own, their salience depends on utilization history and other cognitive processes that would require continuous synchronisation. In addition, only CSPC entities may know exactly which spectrum portfolio is utilized under which policy by which entity of the networking domain. In particular this applies to policies for shared spectrum (e.g. for back-off channels shared across access network cells).

The cognitive capacity, potentially including robustness enhancing measures as outlined by [D6.4] (cf. D6.4 sect. 5 on robust decision-making in spectrum management), of the CSPC includes

- Reasoning on context in the process of context filtering, and decision-making when selecting suitable context parameters to consider as context for the general reasoning process. This process is considered to utilize low complexity pre-determined rule sets and deterministic algorithms operating on context parameters selected to create facts to consider further. Available context parameters are described in more detail in [D6.2] and [D6.3].
- Reasoning on facts obtained to further infer facts suitable as an input to decision-making. This process is considered to utilize an expert system realizing a suitable reasoning engine (e.g. based on logical reasoning, case-based reasoning, instance-based reasoning, or similar). Its purpose is to obtain facts that enable a decision engine to select a suitable course of action which usually is not possible considering context parameters or derived facts directly.

Context parameters and derived facts basically describe a region in the state space. That is, they describe what can be observed. For decision-making, facts need to describe a target that must be achieved. That is, they describe a desire. For example, it can be observed how many users are sharing a certain frequency band, but a-prior knowledge about interference characteristics is needed to conclude that additional users may be assigned to that frequency band.

- Decision-making derives a reasonable set of actions (e.g. on the composition rules for spectrum portfolios) from facts. Assuming that facts generated by a reasoning engine either may describe desires or knowledge, a decision engine may have available a set of pre-defined rules that result in a certain configuration of a spectrum portfolio when triggered. A desire then may trigger an action (e.g. request spectrum), and knowledge selects the way how to implement that action for a given context (e.g. the amount of bandwidth to request from a coordination domain entity).

For example, the DARPA XG [XGL2004] as an early approach described policy rules through triplets of 'selector', 'opportunity' and 'usage constraints'. It has been shown that this approach suits the basic requirements of policy radios. For clarification, XGL here is assumed as a special application of the QoSMOS cognitive spectrum management approach and, if suitable to achieve equivalent functionality, also can be understood as functional validation the QoSMOS approach.

The XGL 'selector' describes the characteristics of a frequency band (e.g. issuing authority, bandwidth, time limits, applicable technology, and similar). The 'opportunity' can be seen as a context that can be observed for the spectrum described by the 'selector' and is characterizing the conditions



that must be met to consider that spectrum as a potential opportunity. The ‘usage constraints’ describe the limits (or policies) that apply to spectrum usage (e.g. to the device configuration) if the spectrum described by the ‘selector’ would be utilized if an opportunity is observed. QoS MOS spectrum portfolios form a superset of DARPA XGL policy rules.

All information in an XGL ‘selector’ also is present in a QoS MOS spectrum portfolio. In practice, it contains accumulated information obtained from coexistence and coordination domain entities as a static description of an amount of frequency spectrum. If a CSCP needs to find suitable spectrum to satisfy the request of a CM-RM, it will search available spectrum portfolios for exactly those parameters until a best-match is achieved, or it will request a spectrum portfolio from coordination domain entities using those parameters as a requirements description. Since in spectrum management a ‘best match’ is depending on context (e.g. spectrum may be used or denied depending on accumulated interference) the search for a best match already is a cognitive process including a planning for future use of spectrum requested and obtained. That is, when querying spectrum the CSCP may not request spectrum exactly according to a CM-RM’s requirements, but may alter parameters to increase re-usability of spectrum requested upon knowledge from earlier requests of the same kind.

The information contained in the XGL ‘opportunity’ is available in a QoS MOS spectrum portfolio through the policies and usage constraints set by the issuing CM-SM entity of the coexistence domain. A description of an opportunity consequently is derived from reasoning on those policies and usage constraints resulting in a set of facts that can be compared with observations (or vice-versa converting observations to parameters of a policy). In the QoS MOS CM-SM architecture these observations may be obtained from querying a CM-RM or by querying spectrum sensors. Thus, a spectrum portfolio determines which parameters and parameter values describe an opportunity, a CM-RM or spectrum sensor provides the observation, and the CSCP performs the reasoning required to infer comparable facts from both. In consequence, the CSCP can decide if a certain spectrum portfolio satisfies the request of a CM-RM by selecting a portfolio based on its static description and by comparing if it matches the current context.




Figure 5-1: Accessing the CSCP in cellular scenarios and sample MSC

The information given by XGL ‘usage constraints’ is directly included in a QoS MOS spectrum portfolio in form of its policies and usage constraints parameters embedded. In addition the CSCP



includes operator policies from the CPOR as applicable and deploys the resulting spectrum portfolio to the requesting CM-RM.

In consequence, a spectrum portfolio composed as described above (i.e. selected based on its static description, matching the current context as far as considered, and embedding regulatory and operator's usage constraints) and deployed to a networking domain entity contains all information required to realize a policy based system to the extent of requirements as given earlier by the DARPA XG.

5.2 Functions in ad-hoc and femtocell scenarios

In a **femtocell scenario** co-location and functionality of the CSPC is equivalent to the cellular case. For associating a CSPC with a femtocell infrastructure two options seem feasible:

- A local femtocell infrastructure is maintained and managed by the same operator as that of a surrounding wide area cellular infrastructure.
- A local femtocell infrastructure is coexisting with a surrounding wide area cellular infrastructure and with neighbouring femtocells but is either unmanaged or is managed by different operators.

Thus a decision is needed if a single CSPC instance shall manage both networking domain entities associated with cellular and femtocell entities in the terminating domain, or if multiple CSPC instances will collaborate with dedicated cellular and femtocell control points. A single CSPC per operator is close to the centralized spectrum management approach, multiple CSPC instances that collaborate in the scope of one operator are close to a distributed spectrum management scenario.

Multiple CSPC instances per operator in the cellular case may be feasible in case a CSPC associates with different networking domain entities controlling the infrastructure of the same RAN or RAT, and operator manages different RANs or RATs concurrently. That is, it should be considered to achieve a balance between coordination effort between CSPC instances and scalability and performance issues for a single CSPC instance. If different RANs don't share spectrum or different RATs operate in allotted spectrum it might be feasible to foresee dedicated CSPC instances.

In case of femtocells associated with the same operator, a single CSPC may control both wide-area cellular and femtocells infrastructures if they share the same geographical area and the same frequency bands. Alternatively a femtocell infrastructure may be considered a dedicated RAT sharing spectrum with a surrounding wide area cellular infrastructure. The latter enables a functional splitting of the CSPC: one CSPC instance coordinates among femtocells while the other coordinates femtocells with wide-area cellular control points, which seems a reasonable trade-off between complexity of cognitive functions and collaboration overhead.

A topological decision thus affects scalability, communication overhead, spectrum efficiency and complexity of reasoning and decision-making. The main benefit of an approach involving multiple CSPC instances is in the lower complexity of cognitive functions (e.g. in terms of rules to consider) while a single CSPC instance enables more balanced spectrum utilization and offloading gain potentially increasing spectrum efficiency when sharing spectrum between wide-area cellular and local femtocells (due to less interaction between distributed cognitive engines through the controlled environment as outlined in [D2.3] and [D6.3]).

In an **ad hoc scenario** the cognitive functionality of the CSPC is equivalent to the cellular case except that connectivity of the CSPC in case an ad-hoc network is not connected with an infrastructure may aggravate collaborative functions. Although different architectural options exist, cognitive functions mandate a CSPC situated at the coordination domain to communicate with coexistence domain entities and to maintain its trust relationship in this communication. In consequence, a CSPC cannot be collocated with a mobile ad-hoc node as long as this node does not provide sustained (potentially also reliable) connectivity with a network infrastructure.




Figure 5-2: Accessing the CSpC in femtocell scenarios and sample MSC (portfolio update for cellular, portfolio deployment for femtocell)

In case of occasionally connected ad hoc networks CSpC and LSPC may collaborate more closely to overcome connectivity problems. In consequence, CSpC functions may be allocated temporarily to LSPC entities situated in the networking domain and being collocated with mobile ad-hoc nodes. In particular, an LSPC may take responsibility for implementing a reliable SPC1 interface and for managing spectrum portfolios autonomously within the limits set.

The LSPC in conjunction with the LPFR then may take responsibility for part of the functionality implemented through the collaboration of CSpC and CPOR. In case of being disconnected from the network infrastructure operator's policies may be applied to spectrum portfolios deployed earlier by the CSpC by the LSPC when there has been a connection available that implemented the SPC1 interface. During disconnected periods, the CSpC will not obtain information on context changes from the networking domain and will not be able to deploy or revoke portfolios. Networking domain entities will have to take responsibility for proper reactions to context changes thus.

In consequence the CSpC must allow deploying operator's policies to networking domain entities in addition to spectrum portfolios. Since there is no trust relation between networking domain entities and coordinating domain entities the CSpC has to ensure by proper pre-processing of spectrum portfolios that networking domain entities when taking decisions in response to a local context change do not violate regulatory or operator's policies and spectrum usage constraints. Moving cognitive functionality temporarily over from a coordination domain entity to a networking domain entity thus



in general demands for policy enforcement functions in ad hoc configurations in the networking and terminating domain.

From the discussion above it follows that in ad hoc scenarios the CSPC has to provide additional cognitive capacities:

- Earlier decisions taken by networking domain entities while disconnected from the infrastructure need to be considered prior to deploying new portfolios to enable prediction or planning portfolio modification to expect from networking domain entities. Hence, case based reasoning may play a stronger role in ad hoc scenarios than for cellular ones.
- Spectrum portfolios deployed are not considered final but will be modified in advance of utilization by network domain entities depending on the specific context encountered. As such the CSPC may generate and deploy a set of alternative spectrum portfolios along with policies that determine under which conditions to utilize them. Hence, predicting user behaviour may play a stronger role in ad hoc scenarios than for cellular ones.

In consequence, the CSPC in ad hoc scenarios must be enhanced for implementing decision-making under uncertainty




Figure 5-3: Accessing the CSPC in ad hoc scenarios and sample MSC (role switch from ad-hoc node A to ad-hoc node B)

5.3 Opportunity detection functions in the CSPC

In composing a suitable spectrum portfolio the CSPC applies a spectrum user model to estimate and potentially predict spectrum utilization in shared spectrum. The CSPC receives measurements from associated networking domain entities and combines information from several locations of the geographical area covered by the terminating domain entities associated with networking domain entities. From the spectrum user activity observed, the CSPC can estimate the utilization of spectrum for the area covered by the networking domain entity (usually a network controller such as an access point or base station) requesting a suitable amount of spectrum. The estimate obtained can be used to determine the amount of spectrum required for achieve a certain (i.e. predicted) interference level likely to be experienced by additional spectrum users in a shared band. According to the usage constraints in its spectrum portfolios available for deployment the CSPC may also decide upon the expected eviction rate of spectrum users if incumbent protection is required. The modelling schemes as detailed in Annex B thus provide the context for a pre-selection of suitable frequency bands in composing a spectrum portfolio for the CSPC and the policies to apply by networking domain entities utilizing the spectrum portfolios.

When composing a spectrum portfolio, the CSPC cannot depend on statistical properties only if shared spectrum usage constraints demand for incumbent protection. The most prominent use case here is TV white space spectrum. In detecting opportunities the CSPC then requires a-prior knowledge about



incumbent's position and transmit characteristics as well as radio propagation conditions between incumbents, victim devices and secondary spectrum users. Information on incumbents and signal estimations based on propagation models applicable to incumbents are foreseen to be provided by TV white space Geolocation databases upon discretion of local regulations.

A drawback of the Geolocation database approach yet is in its inefficiency if multiple spectrum users (e.g. a RAN cell or an ad hoc network) and, in particular, mobile users need to be considered. In such cases the CSPC relies upon its own interference models to determine the geographical area for that a database needs to be queried. In collaboration with networking domain entities (e.g. by deploying suitable operator's policies along with spectrum portfolios) the CSPC provides spectrum portfolios with location-dependent spectrum usage constraints to ensure that regulatory interference thresholds can be met for all areas the spectrum portfolio deployed applies to. The methods considered for the QoS MOS CM-SM are further detailed in Annex A.

5.4 Multi-objective portfolio optimization in CSPC

In this section the topic of multi-objective portfolio optimization is initially covered only. It is seen as a functionality of the CSPC which allows creating portfolios to satisfy initial requests of a spectrum user. Further studies and results on multi-objective spectrum portfolio optimization will be detailed in the scope of the upcoming deliverable D6.7.

When composing spectrum portfolios at the CSPC optimization fusing both pricing (economical) and technical (radio and load) parameters needs to be addressed with multi-objective optimization techniques. Since this spectrum management approach is crucial in multi-cell or multi-access point scenarios, multi-objective optimization will be carried out mainly in the CSPC using input from all relevant repositories with information for the spectrum portfolio composition. Depending on the pricing scheme and the network architecture, the multi-objective optimization of economical and radio parameters can be done either in a joint manner or in an independent fashion. If the pricing of the spectrum is fixed, then the optimization of the economic terms can be done even offline, and then be incorporated into the optimization of the radio parameters in other entities. This means that depending on the network architecture and the pricing scheme, we can have different optimization schemes. In addition, in multi-objective optimization problems, there is no unique solution to a given optimization problem, but instead a group of optimal solutions can be derived. Operators will have to decide a given trade-off between the objective functions they may want to evaluate. For example, in the simplest scenario, a trade-off between revenue given a particular spectrum allocation and the risk generated by the use of different radio interface in an opportunistic manner should be agreed between the operator and the user, which can be used to calculate the optimum solution that complies with the given trade-off. Another consequence of this issue is that different solutions to the optimization problem can be dynamically selected according to the scenario and pricing scheme. Therefore, both the type of solution and the trade-off requirements of the different objective functions can also be potentially included as part of the spectrum portfolio information.

At the local side, multi-objective portfolio optimization can also be implemented for spectrum allocation and radio resource management, depending on the scenario addressed and pricing scheme to be used. For example, if the pricing scheme depends heavily on load and radio conditions (i.e. the pricing becomes increasingly dynamic), then some parts of the multi-objective optimization can be carried out in the local controller. The signalling bandwidth required to exchange the parameters of the pricing information from the spectrum portfolio repositories over the interfaces must be also estimated to achieve a good trade-off performance.

5.5 Interfaces

The Common Spectrum Control (CSPC) accesses the **SM1a interface** to request spectrum portfolios from coexistence domain entities such as regulatory or operator CM-SMs. A CSPC may also request operator's policies through the SM1a interface. If a trusted relationship with coexistence domain



entities is required, a CSpC may access the **SM1b** interface to exchange credentials needed given that the coexistence domain entity is implementing the SM1b interface and is making it accessible to coordination domain entities.

A CSpC entity may also provide spectrum portfolios to coexistence domain entities through the SM1a interface to communicate, for example, spectrum measurements or policy data (i.e. piggy-backing measurements in spectrum portfolio interface data structures).

In addition a CSpC instance may request a coexistence domain entity to convey a spectrum portfolio to other coordination domain entities. In case both source and destination entities in such exchange rely on a trusted relationship with the conveying coexistence domain entity, this may take place by simply forwarding a spectrum portfolio signed by the source entity. In case at least one of the source or destination entities is not in a trusted relationship with the conveying coexistence domain entity, a transfer of trust (e.g. by having the conveying coexistence domain entity to sign the spectrum portfolio conveyed prior to forwarding it to the destination) is needed. Since both the coexistence domain and coordination domain entities in untrusted relation may be owned by the same operator, there may exist other methods to verify the trustworthiness of the source entity that allow the coexistence domain entity to take responsibility for the trustworthiness of the information conveyed by the spectrum portfolio under consideration.

The CSpC accesses the **SPC1 interface** for deploying spectrum portfolios to networking domain entities. In case of an ad hoc scenario a sub-set of operator's policies matching the utilization methods foreseen for spectrum portfolios provided is deployed in addition through this interface. In practice this interface is utilized mainly for deploying spectrum portfolios from a spectrum management entity (e.g. an operator's central CM-SM) to spectrum users (e.g. network control points of the same operator's RANs).

Networking domain entities may also utilize communication through the SPC1 interface for due coordination among each other and for communicating measurements and policies from the networking domain to the coexistence domain. In consequence, the primitives provided for the SPC1 interface must provide the same functionality as those for the SM1a interface. Specifications for the SPC1 interface thus are a sub-set of SM1a interface specifications.

In the case of ad hoc scenarios the SPC1 provides the same functionality as for the cellular case but the management of this interface and its realizations is different since ad hoc networks may or may not have connectivity with a network infrastructure. Physical connections realizing that interface may be disruptive, connection endpoints may be chosen dynamically and opportunistically, and connection up-time may be random. Thus, a networking entity situated in the ad hoc network needs to establish and control the physical connection to an associated CSpC, while in the cellular case the CSpC realization is controlling the connection. Special consideration may be needed to ensure that transactions disrupted (e.g. the request of a spectrum portfolio and its response carrying the portfolio deployed) can be recovered across some connection tear-down event.

The **QS1 interface** supports some scenarios where interworking of core network management entities and spectrum management entities will be required. This interface splits between QS1a and QS1b. While QS1a is realized between core network management entity and CM-SM, QS1b is realized between core network management entity and CM-RM.

The procedures associated to this interface are similar to those described in 3GPP TS 36.413 [TS36.413].

The QS1 interface has been introduced to realize, for example, a centralized management of distributed CSpC instances of a single operator. It allows to exchange management and control information between CSpC instances and an operator's core network management system. It is not supporting the exchange of spectrum portfolios or policies for the purpose of cognitive spectrum management but can be used to initialize repositories, for establishing associations between coordination and coexistence



domain entities, and for maintaining connections between those entities that realize the interfaces discussed so far.

The primitives provided by this interface allow to control the operation of CM-SM entities and in particular to set strategies how the CSPC shall split and merge spectrum portfolios for subsequent deployment towards networking domain entities. Strategies deployed to the CSPC herein reflect the operator's network management paradigms and expectations with respect to shared spectrum utilization, efficient use of frequency resources across managed RANs and RATs, interference situation handling, handling handover and offloading situations and the grade of QoS to provide in these situations. Additionally, strategies set determine conditions (including policies) when to deploy, revoke or modify spectrum portfolios, and how to react upon regulatory changes or, in general, on context changes that can not be observed by the cognitive functions of the CSPC. The latter includes test and training situations that are especially set-up for optimizing performance of the CSPC and its collaborating entities.

This interface actually does not fit into the QoS MOS domain model. It has been allocated to the coordination domain because it provides a means to coordinate between entities of the QoS MOS reference model and those aside this model.

The **AL1 interface** is used to exchange information between distributed CM-SM entities and the QoS MOS adaptation layer. The AL1 interface splits between AL1a through AL1f. It is used as a management and control interface in this context to support the exchange of spectrum portfolios across domain boundaries. Its main purpose here is to identify, associate and exchange control information with entities (e.g. suitable policy repositories [1900.5], [1900.5.1]) based on some selection criteria such as applicable regulatory domain. It is an interface of the QoS MOS reference model.

The **AL1a** (CSPC-AL) control interface provides communication with other coordination domain and coexistence domain entities.



6 Local spectrum control (LSPC)

The local spectrum control (LSPC) entity complements the functions of the common spectrum control (CSPC). It is situated in the networking domain and has two distinct flavours depending on its co-location with an infrastructure-based control point (denoted ‘operator LSPC’) or with local area network control point or a cooperating mobile device (denoted ‘user-equipment LSPC’). This distinction corresponds to the CM-SM NET and CM-SM END instances foreseen for the networking domain. Regarding QoS MOS scenarios, the operator LSPC is mostly dedicated to infrastructure-based configurations while the user-equipment LSPC may be utilized preferably for ad hoc and unmanaged operation, including co-location with mobile nodes in disruptive networks.

An LSPC instance communicates with CM-SM entities in the coordination domain and communicates with a CM-RM through a CM-SM END entity. An LSPC instance may communicate directly with a CM-RM entity for certain scenarios that demand for tight coupling of end system and infrastructure spectrum management such as for those TV white spaces applications demanding that end systems register directly with a Geolocation database.

6.1 LSPC functions (operator LSPC)

The LSPC when co-located with an infrastructure-based network control point such as a cellular base station or an access point (including a managed femtocell) implements spectrum management for associated networking domain entities such as CM-RM entities. Since a single LSPC instance in general manages multiple spectrum users (e.g. a cellular base station serving a certain geographical area and a number of mobile terminals within an operator’s RAN of a certain RAT), it has two main tasks:

1. Collecting spectrum portfolio requests from its associated networking domain entities, computing the accumulated spectrum demand and requesting a spectrum portfolio from its associated coordination domain CM-SM that can satisfy the accumulated spectrum demands.
2. Receiving spectrum portfolios from an associated coordination domain CM-SM instance, composing individual spectrum portfolios and responding to networking domain entities’ requests for spectrum by deploying individual spectrum portfolios along with suitable policies to the networking domain (i.e. its associated CM-RMs).

Cognitive decision-making of the LSPC is characterized by highly dynamic context such that an LSPC always operates on uncertain knowledge (assuming that context changes are random or correlate in a chaotic way). In consequence, an LSPC has to find a balance between overprovisioning spectrum and risking interference among spectrum users.

On the other hand an LSPC instance is conveying measurements between networking domain entities and coordination domain entities and thus has more accurate (e.g. timelier and more detailed) context information in a local scope than a CSPC, which has less accurate context but a more global scope (see section 5.2).

For its main tasks as stated above, the LSPC can query network domain entities by providing a spectrum portfolio that defines the area of interest of the requesting LSPC in terms of frequency bands, location or technologies by utilizing the information elements of a spectrum portfolio data structure as a descriptor. In the optimal case – depending on the sensing capacities of devices deployed – it may obtain in response to querying networking domain entities the following context information (with increasing complexity):

- Presence detection results for spectrum users in the frequency bands observed (including spectrum users of a previously specified, a-prior known or of an unknown technology);
- Temporal and spectral statistics on spectrum utilization for frequency band observed (including in-band and out-of-band detections) potentially in form of averaged duty cycle



(activity vs. silence periods) and variance of duration of active periods (i.e. the distribution of spectrum access periods and quiet periods observed);

- Aggregated, filtered and pre-processed information potentially omitting irrelevant measurements such as detections found below the interference thresholds set for the frequency bands of interest;
- The position of spectrum users associated directly or indirectly with the requesting LSPC and their local radio scene measurement, which comes closest to a static RF environment map (see [1900.1a], [1900.6a]). Indirect association here may refer to topological neighbourhood (e.g. through collaboration with network control points serving geographically neighbouring areas, different RANs or RATs or non-overlapping frequency bands.
- Temporal changes of the information above potentially parameterized in a suitable mobility model characterized by speed, direction, sojourn times or similar parameters.

Since only user equipment LSPC instances may obtain context from spectrum sensors directly (see section 6.2) this information is mainly obtained from CM-RM entities or CM-SM END entities in collaboration with CM-SM or CM-RM entities associated with the requesting LSPC. The LSPC in turn can make this information available to other CM-SM entities (in the coordination domain as well as in the networking domain) to support cognitive processes implemented by these entities. The information can be provided in form of context information or in form of policies (e.g. coexistence policies) generated by the LSPC from this context information in a separate decision-making process. Since this information is encoded into one or more spectrum portfolios, it is closely related to a radio environment map (see [1900.1a]).

The LSPC implements a number of functions for manipulating spectrum portfolios including at least:

- Interfacing with coordination domain entities via the SPC1 interface.
 - Request spectrum portfolios, policies and spectrum information from coordination domain entities via the SPC1 interface.
 - Provide measurement information obtained from associated networking domain entities to coordination domain entities via the SPC1 interface upon request of a coordination domain entity.
- Interfacing with networking domain entities via the PF2 or CM1 interface.
 - Deploy spectrum portfolios to networking domain entities upon request of networking domain entities or upon request of coordination domain entities via the PF2 or CM1 interface.
 - Revoke spectrum portfolios from networking domain entities in consequence of earlier deploying spectrum portfolios updates or upon request of coordination domain entities via the PF2 or CM1 interface.
 - Receive context information (e.g. measurements) from other networking domain entities via the PF2 interface (if the source is a CM-SM instance) or via the CM1 interface (if the source is a CM-RM instance).
- Interfacing with an instance of the LPFR via the LPFC interface.
 - Store and retrieve spectrum portfolios along with related status, utilization and history information (i.e. if unused, deployed or revoked, to which networking domain entity it has been deployed, which spectrum portfolios obtained from other networking domain entities for which purpose or objective, and a reference to its parent if it has been derived from another portfolio, and similar).



- Retrieve, modify and store modified portfolios in the course of composing spectrum portfolios upon request of other associated networking domain entities or upon request of coordination domain entities.
- Cognitive functions to compose spectrum portfolios according to requests of other networking domain entities and to the constraints set by coordination domain entities considering current context as provided by the requesting networking domain entity or from the LPFR.
- Collaboration and cooperation functions with other instances of networking domain CM-SM instances for the purpose of collaborative decision-making and context exchange.

At any point in time an instance of the LSPC can decide to forward context information towards associated coordination domain entities or to request context information from coordination domain entities if its decision-making processes encounters situations where additional context may reduce uncertainty or risk (e.g. by requesting to add some redundancy, see [D6.4]).

For its cognitive decision-making process the LSPC strongly relies on the LPFR (see section 4.5 Local Portfolio Repository (LPFR)). Since this repository records spectrum portfolios available as well as portfolios deployed along with spectrum utilisation experienced earlier for deployed portfolios it is storage for a-prior knowledge, ontology for a case-based reasoning process, as well as a training data repository for self-learning capacities. That is, all context information obtained from other networking domain entities must be seen in relation to the information kept through the LPFR since these resulted from decisions that have been made earlier and have been recorded through the LPFR.

To ensure a short response time to spectrum requests an LSPC has to apply more sophisticated predictive methods. In consequence the reasoning engine of LSPC instance may need to evaluate alternative courses of actions concurrently and mitigate decision upon availability of context at a given deadline. In particular an LSPC may need to decide in a first step based on different objectives and strategies (e.g. on interference minimization vs. optimization of spectrum utilization) selected from current risk factors (e.g. risk of creating interference) while in a second step a “quick decision” based on most recent context has to be taken on the preference on several similar courses immediately in advance of deploying a spectrum portfolio. The LSPC thus requires an optimization regarding the timeliness of decisions made much more than for the CSPC.

The cognitive capacity, potentially including robustness enhancing measures as outlined by [D6.4], of the LSPC includes

- Reasoning on context in the process of context filtering, and decision-making when selecting suitable context parameters to consider as context for the general reasoning process (through low complexity pre-determined rule sets and deterministic algorithms comparable to the functionality of a CSPC, see section 5).
- In addition the LSPC context filtering must identify context suitable to be forwarded to other networking domain entities or to coordination domain entities (e.g. by selecting parameters with reasonable change rates). In the course of communicating context, the LSPC may decide on further fusion of context parameters. The process may involve both pre-determined rule-sets and reasoning of higher complexity on the communication of context when determining parameters to forward and their respective update frequency. Cognition here may support estimating the relevance of context to associated entities.
- Reasoning on facts obtained from context evaluation to further infer facts suitable as an input to decision-making similar to the corresponding CSPC functions. In contrast to the CSPC, LSPC decision-making is more dynamic regarding timeliness and concurrency of requests.
- Preparation of alternatives (e.g. potential decisions to choose from) following more than one objective at a time, enabling simplified and rapid decision-making in a final conclusive step.



This corresponds to an emphasis on the planning phase in an OODPA loop (see Mitola in [Fette06]).

In general an LSPC has to respond to a very limited set of possible requests originating from a coordination domain CM-SM, from a networking domain CM-SM, or from a networking domain CM-RM including:

- A request to deploy an initial spectrum portfolio.
This request is satisfied by reasoning upon the context provided (e.g. amount of frequency spectrum requested and desired spectrum attributes) and context a-prior known (e.g. amount of spectrum available). Potential decisions of the LSPC would be
 - to provide a spectrum portfolio that satisfies the request as given from its local LPFR,
 - to provide a spectrum portfolio allotting more spectrum than requested but not satisfying requested attributes,
 - to provide less spectrum than requested but providing proper attributes.

The latter two options can be seen as temporary decisions and may occur in conjunction with requesting additional spectrum portfolios from an associated coordination domain CM-SM which may take some time in order of seconds to weeks depending on the measures that need to be taken to obtain new spectrum (which may involve spectrum auctioning or reorganization of already allotted spectrum).

Spectrum portfolio optimization criteria may be the price of spectrum, lease times, load factors (e.g. number of request or amount of spectrum already deployed) or number and kind of spectrum users for the frequency bands considered. A-prior knowledge such as request success rate, response time of coordination domain CM-SM entities or attributes of the requesting entity (e.g. serving highly relevant users, areas, events, or services) also influences LSPC decisions.

- A request to change or to extend a spectrum portfolio.
A change may be required in consequence of a coordination action (e.g. resulting from a network management request), from coexistence issues arising, or from increasing or decreasing spectrum demands of spectrum users (e.g. due to traffic load changes during daytime and overnight). A change request is satisfied by first deploying a new spectrum portfolio and then revoking the spectrum portfolio deployed previously, or by deploying a spectrum portfolio complementing the existing one.
 - The first option is very similar to deploying an initial spectrum portfolio except that revoking a spectrum portfolio later on may compensate some of the optimization criteria when seen as a single transaction. In consequence this is a trading situation and could be handled by the LSPC as such. Since context may have changed since the spectrum was deployed originally, a spectrum portfolio may become more or less valuable at the time it is revoked.
 - The second option may result in a quicker response time and higher spectrum availability but may lead to higher fragmentation of spectrum depending on the availability of contiguous frequency bands. In addition, a spectrum user (i.e. a CM-RM in this case) has to be prepared to operate on multiple spectrum portfolios. If this is experienced as a drawback depends on the specific situation. A CM-RM may request extension of its spectrum portfolio, for the purpose of offloading mobile terminals from its main spectrum portfolio or needs to handle specific handover or connectivity situations, where a complementing spectrum portfolio would be considered as beneficial.



- A request to revoke a spectrum portfolio.
Spectrum revocation may be required in consequence of a coordination action (e.g. resulting from a network management request), or upon request of a spectrum user in response to diminishing its operation (e.g. prior to a power-down or switching into a maintenance mode). The latter may happen in scenarios where wide-area cells are switched off temporarily in favour of a more power-efficient femtocell service. An LSPC may decide to reserve the spectrum portfolio revoked for later use by the same spectrum user for some time and upon request, or may decide to handover this spectrum portfolio to a different spectrum user. Revoking a spectrum portfolio for the purpose of deploying it to a different user may become a common use case for power efficient wireless access assuming that frequent system reconfigurations due to a change of spectrum used may unnecessarily increase power consumption of infrastructure as well as mobile terminals.




Figure 6-1: Accessing an operator LSPC and sample MSC (responding to a portfolio change request by a coordination domain CM-SM)

6.2 LSPC functions (user-equipment LSPC)

The LSPC when co-located with a local area network control point (e.g. an access point or an unmanaged femtocell) implements spectrum management for associated networking domain entities. In contrast to an operator LSPC, a user-equipment LSPC may serve only few CM-RM instances. In particular an instance of the user equipment LSPC may be co-located with an SSE, a SAN and a CM-



RM in a single mobile terminal in an ad-hoc network. A user equipment LSPC is deploying spectrum portfolios to a spectrum selector entity (SSE) and is receiving spectrum portfolios from a spectrum analyser (SAN) entity.

As outlined above (see section 6.1) an LSPC may implement a decision-making strategy where a reasoning engine develops alternative courses of action and a concluding less complex decision-making engine picks the most suitable from those alternatives based on most recent context. In a user-equipment LSPC the concluding decision-making is located at the SSE which then acts as a rapid decision-engine and spectrum portfolio cache. In conjunction with a SAN and SSE entity, an LSPC may quickly respond to context changes triggered by spectrum sensors without involving potentially time-consuming reasoning processes.

For ad-hoc scenarios a user-equipment LSPC has to implement a role-handover strategy since connectivity with an infrastructure may be disrupted frequently and, in consequence, communication with an associated coordination domain CM-SM may fail. Hence, the LSPC could be co-located with multiple mobile ad-hoc terminals in a network, while only one of these instances associates with a coordination domain CM-SM (see also section 5.2) at a given time. Such strategy has both a protocol and cognitive aspect:

- A protocol between LSPC instances must exist that allows exchanging the context of an ongoing transaction between LSPC and CSPC. In case of disrupted communication a different LSPC entity should be able to conclude a transaction without loss of information on both ends. This protocol may be proprietary and thus is not addressed further in this deliverable. This may include synchronizing between instances of the LPFR if there is a one-to-one association of LSPC and LPFR has been selected as a design choice.
- The cognitive engine of an LSPC instance may be utilized also to optimize role handover in an ad-hoc scenario. Context information about spectrum utilization in an ad-hoc scenario is in any way available at all instances of an LSPC and adding context about connectivity of nodes within the ad-hoc network and towards a fixed infrastructure is likely possible. Hence the LSPC may plan communication with a coordination domain CM-SM both on the availability of relevant context updates for the coordination domain and upon availability of a communication link, which may include multi-hop and store-and-forward strategies that involve potential role-handover candidates to reduce protocol overhead as a side-effect of multi-hop communications.

A co-location of LSPC, SAN and LPFR allows creating portfolios from spectrum observation. It enables data fusion of spectrum observations obtained directly from spectrum sensors and from CM-RM entities providing additional context information obtained from terminating domain entities (e.g. spectrum sensors co-located with access points, base stations or mobile terminals [1900.4], [1900.4a]). Actually, a co-location is not mandatory but rather preferable to realize short response times in communication between the three entities. Low delay communication increases correlation between raw sensor data and fused data provided by CM-RM entities and enhances the timeliness of decisions based on this information exchange.

The SAN entity is creating a spectrum portfolio data structure from spectrum observations and forwards this to the LSPC which may utilize the spectrum portfolio obtained in several ways:

- The LSPC may decide to utilize the spectrum portfolio data structure obtained from a SAN entity as pure context information. The spectrum portfolio data structure is processed by the LSPC as any other context information. It may be forwarded to other networking domain or coordination domain entities as such and it may be stored by the local LPFR instance. If stored locally, it may be referenced subsequently by the SAN for updating partially or in whole. Usually, a stored spectrum portfolio data structure will be removed at a certain time after its last update.



- The LSPC may utilize the spectrum portfolio data structure obtained as a self-learned spectrum portfolio.
In a local context (e.g. in ad-hoc scenarios) an LSPC may learn about spectrum availability by sensing spectrum for incumbent or other spectrum user activities. In order to utilize this spectrum opportunistically it must have obtained a spectrum portfolio from a coordinating domain CM-SM at an earlier point in time (which by intention includes “obtained at manufacturing and certification time”). The self-learned portfolio must respect the policies set by the certified portfolio and must operate within its authoritative limits if utilized as a spectrum portfolio.
- The LSPC may forward the spectrum portfolio data structure to another LSPC (e.g. from a user-equipment LSPC to an operator LSPC).
For local spectrum coordination and coexistence reasons an LSPC may forward the spectrum portfolio data structure to another (e.g. geographically or topological neighbouring CM-SM or CM-SM END) instance. The originating LSPC may modify (e.g. fuse, filter or average) its contents as needed when forwarded as context information. If utilized locally as a spectrum portfolio and subsequently forwarded as context information, the LSPC must restrict the information contained in the spectrum portfolio data structure to the authoritative limits set by the enclosing spectrum portfolio under that it operates. In addition, it must sign the forwarded spectrum portfolio data structure as the originator and user of this spectrum portfolio. A receiving LSPC then may utilize the context obtained as “the quiet situation” since it does not reflect the spectrum utilization caused by the originating LSPC (i.e. by the networking domain entities associated with the originating LSPC) utilizing this as a spectrum portfolio.
- The LSPC may forward the spectrum portfolio data structure to a coordination domain CM-SM.
For information and coordination purposes the LSPC may decide to forward a spectrum portfolio data structure to a coordinating domain CM-SM as context information. In that it may flag the spectrum portfolio as in-use under the authoritative spectrum portfolio obtained earlier. There is no need to modify the contents of the portfolio data structure here since coordinating domain CM-SM and networking domain CM-SM END are in an implicit trust relationship through association and the CM-SM END may even have obtained the authoritative spectrum portfolio from exactly that CM-SM which may use the context to validate the spectrum portfolio.
- The LSPC may forward the spectrum portfolio data structure to an SSE entity.
If the LSPC decides to utilize the spectrum portfolio data structure obtained from its associated SAN entity as a self-learned spectrum portfolio it may immediately forward this to an SSE entity for utilization by associated CM-RM entities. This may allow reacting rapidly to changes in the observed environment (e.g. if multiple LSPC entities apply similar strategies for the same geographical area). The LSPC can quickly decide upon forwarding but is ‘out of the loop’ afterwards unless it revokes the spectrum portfolio. Hence, this strategy is of high relevance for local use but is much too restrictive for collaboration in a larger (managed) environment.

Cognitive methods of the user-equipment LSPC will decide upon a strategy how to cooperate with SSE and SAN entities as the main exchange between terminating and coordinating domains. Applicable methods jointly forming suitable cognitive methods have been described earlier in the scope of [D6.1], namely genetic algorithms (focus on optimization), neural networks (instance-based reasoning) and game theory (focus on performance assessment and validation). In that it may select dynamically one or more of the strategies described above. For this the LSPC will need a number of operator policies to guide such decision, which can be realized through a straight-forward rule-set with few fact evaluations necessary to conclude.




Figure 6-2: Accessing a user-equipment LSPC and sample MSC (fast update of shared spectrum portfolio by requesting the SSE, involving CM-RM, SAN and SS for context acquisition)



If no operator policy exists (which usually already considers current context or guides through providing alternatives for certain context situations) such that most suitable desires can be inferred from, higher complexity decision-making is required and the LSPC must be ‘kept in the loop’ for optimizing both the strategy and the spectrum portfolio in parallel.

In consequence, the LSPC needs to implement an iterative ‘trial-and-error’ process first deploying a self-learned spectrum portfolio obtained from the SAN to the SSE, then observing its impact on the environment and modifying the spectrum portfolio according to the feedback of the SAN. If progress in terms of predefined metrics occurs, forwarding to cooperating LSPC entities may stabilize this ‘evolution’ in case of competitive opportunistic spectrum users. When reaching a stable state, a spectrum portfolio data structure may be forwarded to a coordination domain CM-SM, which in turn may formulate a ‘self-learned operator policy’ from this context information.

Although this approach may be realized as an algorithm (e.g. as a genetic algorithm) it may lead to the formulation of a suitable case-based reasoner or, in particular, to a set of standardized case descriptions based on numerical values that enable instance-based reasoning. An application of instance-based reasoning for the LSPC here is much less complex than general solutions and even may be downloaded to the SSE which then may implement concluding decision-making as discussed above based upon instance based reasoning closest to the spectrum user.

A suitable case-based reasoner (which is here assumed as a function of the LSPC) will monitor the decision loop created by SAN, SSE and LSPC and will derive suitable case descriptions from this observation (i.e. references to spectrum portfolios and related context that led to the deployment of this portfolio). Further looking at the SAN, monitoring the use of selected spectrum portfolios (i.e. appropriateness), and on the reports of the CM-RM, monitoring utilization of spectrum by spectrum users (i.e. efficiency), will allow to tag portfolios created and used earlier by a salience or precedence parameter for later selection (potentially including moderate modifications) as a deployable spectrum portfolio. It is expected that this will speed-up significantly the response time to CM-RM spectrum portfolio requests.

6.3 Opportunity detection and spectrum portfolio management functions in the LSPC

In composing a suitable spectrum portfolio the LSPC utilizes similar models as the CSPC (see sections 5.3 and 5.4). In contrast to the CSPC the scope for spectrum utilization optimizations by the LSPC is rather limited to the scope set through the spectrum portfolios obtained from a CSPC (i.e. the coordination domain CM-SM it is associated with). The LSPC can assume that the CSPC already performed a global optimization across RATs, RANs and associated operator’s infrastructures. In most scenarios the LSPC thus can focus its operation on a single technology, a limited geographical extend or a narrow set of frequency bands and spectrum access and sharing strategies.

While the CSPC is optimizing spectrum portfolios based on complex and rather long-term user models or spectrum pricing models, the LSPC performs rather quick scheduling tasks that even may have to interoperate with a certain technology’s inherent spectrum utilization optimization such as LTE subcarrier multiplexing – potentially not on a time-scale that an CM-RM must be aware of but with distinct knowledge about the impact of its optimization process on such technology specific optimization strategies.

The LSPC thus focuses on the construction of spectrum portfolios from spectrum opportunities it is aware of for the spectrum portfolios including usage constraints, regulatory constraints and operator’s policies obtained from a coordination domain entity by performing a number of tasks (potentially concurrently) that include:



- Maintenance of the LPFR to ensure its consistency with corresponding repositories at the coordination and coexistence domain.

When receiving an update of a spectrum portfolio received earlier from its associated coordination domain CM-SM, it has to evaluate the impact of this change on its spectrum portfolios deployed earlier to other networking domain entities. To ensure consistency the LSPC has to take decisions which spectrum portfolio is affected and which networking domain entity must be addressed for updating or revoking spectrum portfolios obtained earlier. This process involves cognitive functions that have to enable incremental decisions, which is actually a matter of context filtering and managing a priori knowledge. In particular, decision-making follows different inference rules for evolving within a locally changing state space.

- Compose spectrum portfolios according to the requests of other networking domain entities (i.e. CM-RM entities).

Starting from a spectrum portfolio obtained, the LSPC applies the very same strategies and algorithms as the CSPC when composing a spectrum portfolio. Since the amount of resources available (i.e. the input set of frequency bands a spectrum portfolio can be constructed from) is more limited and the policies and usage constraints are more restrictive than for CSPC decisions the LSPC will likely have fewer alternatives available to select from when composing spectrum portfolios. Although this will speed up decision-making in one way, it also may increase the risk for decisions or the potential for not being able to come to a decision at all. The LSPC thus may need to consider the robustness issues discussed in [D6.4] more closely than the CSPC.

When composing spectrum portfolios the LSPC needs a certain degree of awareness about the technology of terminating domain entities associated with CM-RMs it is deploying spectrum portfolios to. For example, their reconfiguration capacity, RF bandwidth and granularity of bandwidth, transmission power limits, and similar may assist the LSPC in optimising its selection of context parameters to consider in decision-making. In addition, knowledge regarding the characteristics of the incumbent, if any, may be needed (e.g. channelization) as well as about spectrum sensors. This kind of awareness helps to categorize context parameters available according to their relevance and accuracy in robust decision-making.

- Preparing spectrum portfolios for later use by the SSE (CM-SM END only).
- Obtaining spectrum portfolios from the SAN (CM-SM END only) for updating the LPFR with context information from spectrum sensing, for adding portfolios, or for merging with existing spectrum portfolios.

Besides cognitive decision-making on which way to consider a spectrum portfolio received from a SAN the LSPC here may need to realize self-learning capacities.

- A cognitive process may be needed to categorize the spectrum portfolio received if it must be considered as a set of context parameters (i.e. a set of spectrum measurements) or if it could be recognized as a spectrum opportunity (i.e. that it does not conflict with other spectrum portfolios or their policies and usage constraints). This decision cannot be taken by the SAN since operator's policies are available to the LSPC only due to its management role towards the LPFR.
- Self-learning may be required to decide if a spectrum portfolio received from a SAN describes a spectrum opportunity and if it is beneficial to select that opportunity. The goal of self-learning here is in optimizing the decision parameters and rules according to the benefit of earlier decisions in this scope (which may be seen as a more sophisticated trial and error strategy).

LSPC decisions based on spectrum user observations directly impact spectrum utilization and may produce harm to incumbents or other spectrum users in the presence of malicious users in a way tampering spectrum observations. This includes the option to force a rejection of



spectrum portfolios obtained from coordination domain CM-SM entities due to contradicting observations and bears the risk of conveying attacks to the coexistence domain.

Continuous observation of the LSPC cognitive decision-making and self-learning for outlier detection will increase robustness of managing spectrum portfolios as shown in [D6.4].

6.4 Distributed self-organizing cognitive-radio spectrum management

6.4.1 Challenges for self-organized cognitive spectrum management

Modern Cognitive Radio Systems become more and more diverse, in terms of heterogeneity, cell layouts with a multitude of different cell sizes, quickly varying and inhomogeneous traffic, as well as various spectrum possibilities and certain interferences on certain parts of the spectrum.

Figure 6-3 does schematically illustrate such a cognitive radio scenario, where for each cell or base station it has to be decided which part of the frequency spectrum (illustrated by the rainbow-coloured-row) it shall best use and which transmission power it shall use; the resulting coverage range is schematically illustrated in the figure by the circles and the arrows indicate its modification with transition power changes. The mobile phones represent different traffic densities in certain areas which need to be considered when choosing the best suited cognitive radio settings.




Figure 6-3: Schematic example of interactions and couplings in modern cognitive radio systems.

One major challenge in Cognitive Radio and in particular for the QoS MOS Cognitive-Manager Spectrum-Manager (CM-SM) is how to organize and to decide which radio access entity (e.g. base station, cell) is using which part of the spectrum and with which power. This spectrum plus power organisation, configuration and optimisation challenge is subject to strong interactions between the nodes, such as interferences and interactions in their “coverage areas”. Furthermore, the cognitive radio system does constantly need to be adapted and re-optimized when the situation is changing, such as e.g. different user-traffic load, altering interference situations and updated external constraints such as modified spectrum database entries.

As this issue is far too complex for manual handling, powerful self-organizing networks (SON) techniques are required to solve this configuration, adaptation and optimization challenge for Cognitive Radio and for the QoS MOS spectrum management (CM-SM) in particular. These Self-X techniques do especially need to be able to handle and resolve the complex interactions of the highly coupled parameters, also among different nodes.



For various reasons, centralized CM-SM approaches are not suitable for this optimization challenge which thus requires distributed solutions for this complex optimization challenge. These reasons include that centralized solutions cannot anymore handle (well) a large area, that there are limitations in the (or non-existent) cooperation between different kind of vendors or systems on the same spectrum. Furthermore, distributed spectrum management entities are much better suited to ensure robustness and stability, as outlined in detail in [D6.4].

6.4.2 Distributed SON for cognitive radio spectrum managers

6.4.2.1 Fully distributed CM-SM architecture

There is a distributed CM-SM architecture, where each cognitive node has its own CM-SM which in a simple form is illustrated by Figure 6-4. The CM-SM functionality can be realized as a single individual CM-SM attached to a particular cognitive node, or as a CM-SM entity controlling a multitude of CM-RMs. In the latter case, the CM-SM instance creates an individual virtual instance for each of its controlled nodes then running an individual instance of this CM-SM's evaluation and decision engine in the scope of that particular node. The CM-SM architecture is designed to be generic. It thus can handle any kind of cells, including macro, metro and femto cells. In particular, this concept is also sufficiently powerful to manage and optimize a heterogeneous network with a large diversity of cell types on a cell individual basis, including a large amount of small and femto cells.

This CM-SM decides on a “longer” time scale, e.g. semi-static time scale, which part of the spectrum portfolio (which part(s) of the bandwidth part(s), which part(s) of the frequencies, which part(s) of the spectrum(s)) and which other relevant configuration parameters (i.e. transmission power) the resource manager CM-RM is allowed to use. The CM-RM then operates on a shorter time scale (e.g. dynamic) within the parts of the spectrum portfolio and within the configuration constraints set by the CM-SM.

The distributed individual spectrum managers can communicate and do arrive to find together optimized configuration parameter settings for the whole system as is specified in the following chapters.




Figure:

Figure 6-4: Schematic illustration of the distributed CM-SM architecture.

Within the QoS MOS framework, these techniques developed here for spectrum and parameter configuration and optimization are located at the spectrum selector functionality in the LSPC functions.

6.4.2.2 Distributed SON operation on a “local area”

This SON approach uses the concept of “local areas”, a cluster of some cells within the neighborhood around the particular CM-SM in which this distributed algorithm is running. Each CM-SM is optimizing a “local area”, this means it is optimizing the spectrum portfolio and the relevant



parameters (such as i.e. transition power) for itself, and for other “neighboring” CM-SMs within a “local area”.

Figure 6-5 schematically draws a cellular layout with one example for a local area. Each cell has its own, distributed, CM-SM. The CM-SM in the dark red “centre” cell is capable to influence also parameter settings of the SM-SMs in the first (light-red) tier of neighbouring cells, while having a knowledge of the situation even in a larger area, e.g. also of ones of the CM-SMs two (or more) tiers away.

Due to the interactions and due to the interferences the spectrum- and power settings of neighbouring CM-SM entities are highly coupled, they cannot be individually optimized, and during the parameter finding process, the situation, setting, interactions with and from neighbouring entities have to be considered. The local area contains that group of CM-SMs which need (or should) be considered as there are directly interactions with the “centre” CM-SM.




Figure 6-5: Schematic drawing of a cellular layout, with one example of the “local area”.

6.4.2.3 Distributed CM-SM SON entity optimization procedure

Each spectrum manager has attached or included one SON entity which runs independently its distributed SON algorithm on its particular local area. The flow chart in Figure 6-6 illustrates its main SON operation steps.




Figure 6-6: Illustration of the signalling around the CM-SM

The high level SON operation is as follows: In the first step, the CM-SM is creating knowledge about the current situation within its local area, i.e. the situation also of other CM-SM. The CM-SM can either use stored information from previously exchanged messages, and/or the CM-SM can sent out a signalling message requesting another node for information at that moment when the CM-SM needs to make a decision. A futher mechanism is that each nodes informs its neighbouring ones, whenever its own situaion changes by a relevant amount so that each node can assume that its stored information are always reasonaly recent and accurate. This information includes for example the currently used



configurations, e.g. the theoretical available and the currently used spectrum portfolio, its power settings, indicators about the quality of the spectrum, such as the interference situation on the particular parts of the spectrum, as well as information about the traffic load. This information may either be obtained when needed, and/or previously stored information may be used.

The SON entity/functionality of the “centre” CM-SM” is then evaluating possible candidate sets in order to find the best suited parameter combinations for the spectrum managers in the local area. These candidate parameter sets can be an intelligently chosen subset of parameter combinations out of the complete parameter space of all possible parameter combinations within the CM-SMs in the local area.

The simplest search algorithm would be to assess all options via brute force, but there are more intelligent and more runtime efficient search algorithms. Thereby the expected system performance and the expected energy consumption of each accessed particular parameter set is predicted via a “sufficiently well suited” prediction model, which calculates virtually the future system behaviour in the case that this particular parameter set would be installed. This prediction of the future network performance is very tricky, requires innovative novel approaches, and this solution will be described in more detail in [D6.4].

It shall be noted again, that this is an offline assessment of possible candidate parameter sets, without actually installing (testing, trying) these in the field. After having virtually evaluated a/the large set of candidate combinations, the SON entity then selects the best suited one and these found best suited parameter settings are then installed in the CM-RMs (/cells) within the local area.

In this way, the optimal -predicted- parameter set is found and installed for each CM-SM within the local area. As this local area optimization process did also consider the situation in the surrounding CM-SMs, it is unlikely, that the neighbouring CM-SM are not happy with the new settings which were calculated by a neighbouring CM-SM.

6.4.2.4 CM-SM signalling message exchange

This approach –as well as any other SON technique– does require some kind of information exchange and/or messages sent between nodes. There are some different variations of how the information is concretely embedded into existing or new signalling messages, but the following, kind of information exchange is related to this SON concept as illustrated in Figure 6-7.




Figure 6-7: Sketch of the kind of information which are exchanged via signalling messages between the different distributed CM-SMs



The concept here is flexible and is not restrict to a certain way how this information is exchanged, and via which architectural interface the signalling messages are exchanged. For example within the cellular use case and for the LTE-technology, the cognitive information messages may use the inter-eNBs X2 interface or could communicate via its S1 link via the core network.

The CM-SMs exchange the following kind of information:

- Information about their current configurations and settings, e.g. which part of the spectrum portfolio is assigned to use and parameter configurations such as e.g. transmission power.
- Information about –e.g. averaged values- about currently experienced (average) “radio and load conditions”, such as e.g. about their traffic load and about how much interference is observed on a particular part of the spectrum portfolio.
- Commands (suggestions) from one CM-SM to another CM-SM to use a certain part of the spectrum portfolio, and to use a certain configuration parameters, such as e.g. a certain transmission power.
- Optionally, direct trigger messages to initiate an action such as to start the local area evaluation + optimization procedure.

6.4.3 Prediction model

In order to evaluate the quality of considered new parameter sets, several other SON approaches install these candidate parameter settings in the field, let the system run for some time and then to observe the system feedback. However, this in the field testing is no longer suitable for complex and highly interacting parameter optimization challenges, there are too many parameter options, it takes too much time to assess these, the system performance would decrease while testing a not-good parameter set, and this single-node trying is not suitable for coordinated distributed SON operation of the system.

Therefore it is required to be able to carry out an offline calculation to evaluating the quality of potential new parameter sets in the local area. This offline calculation requires an internal system understanding, including all the interactions and couplings, in order to be able predict the quality of a potential candidate parameter set. In the following this prediction model and its internal tools are specified.

6.4.3.1 Generic classification of parameters according to their effects

As a well suited level of abstraction, the different cell-parameter variation techniques are described and modelled in a generically according to their main effect on the system and on the inter-node interactions.

- a) There is one group of parameters of which their variation affects the area, within which the mobiles are (resultingly) assigned to a particular cell. In this concrete case of cognitive radio spectrum managers, this is here the basestation transition power on a particular frequency or frequency band.
The size of this area affects the amount of offered input traffic which shall be served by that particular cell. For example via variation of the transition powers, traffic can be moved between different systems or access possibilities.
- b) The other group of parameters influences the amount or efficiency of the available resources for one cell. Here, these parameters include the amount of spectrum, which parts of the spectrum is used, and the interference situation on this spectrum. This involves all different kind of inter-cell and inter-system interference coordination and management. As a result it is required handling of inter-cell interference issues, how much one cell affects or is affected by the interference on a particular resource by the use (or non-usage) of that resource in a neighbouring cell.



6.4.3.2 Internal usage of ‘virtual sub-areas’

Cell internal, the cell area is virtually sub-divided into several smaller “Virtual Sub-Areas” as schematically illustrated in Figure 6-8. Each of these virtual sub-areas consists of a part of the complete cell area. There is one virtual centre area and one separate virtual sub-area towards each neighbouring cell.




Figure 6-8: Illustration of virtual cell internal sub-areas.

As one realisation example, a border sub-area of Cell A towards Cell B is that area within which the mobile terminals are served by cell A and within which the (e.g. pilot) radio channel from cell B is the strongest neighbour within a certain dB-signal strength window (e.g. the channel from cell B is by e.g. 2 dB weaker than the channel from the serving cell A). The cell scheduler knows the radio channel properties of its own mobiles and can thus calculate the (average) situation within these virtual sub-areas.

Within the internal calculations, the user traffic and the resources are treated separately for each virtual sub-areas (and are thereafter convoluted to obtain the full cell behaviour). Where possible, precise sub-area data can optionally be exchanged between cells, but it is also possible to use only the standard inter-cell exchanged information and to estimate the properties of the virtual sub-areas, e.g. by assuming average values when no precise information can be obtained.

When calculating the effects of a certain parameter variation, then the effect of this parameter variation is calculated (modelled) with respect to its impact to particular sub-areas. The sub-area model has large the advantage, that the effects of parameter variations are limited to few selected sub-areas only, while the rest of the cells remain (basically) non-affected.

A variation of the “cell area” affecting parameters, i.e. here the base station transmission power, shifts input traffic between the border sub-areas from one cell to the attaching sub-area of the neighbouring cell. This traffic shift is illustrated in the right part of the Figure 6-8 above, where the red marked area is shifted between these two cells. Thereby the amount of sifted traffic can be derived based on the amount of cell border shift, and based on the traffic density in the shrinking sub-area.

A variation of the “resource affecting parameters” alters the amount of resources and additionally –via inter-cell interferences– the “resource efficiency” as outlined in the next paragraph. The left part of Figure 6-8 illustrates the example that the red –centre– cell does not use a certain part of the frequency spectrum; as a result, the mobile users in the –green marked– border sub areas of the neighbouring cells do experience less inter-cell interference on this particular part of the spectrum which the red centre cell is omitting.

6.4.3.3 Virtual resource efficiencies as a generic describing tool

As a tool for the internal calculations, an internal variable “Resource Efficiency” is introduced. This variable describes, how well radio data can be delivered on a certain resource (e.g. on a certain frequency-part) and in a certain area (e.g. within a cell or within a certain sub-area of a cell). It shall be noted, that this variable does not (necessarily) need to have correct absolute values; it is sufficient if relative values are used describing a relation between different areas and between different resources, so that based on these relative values decisions can be made how to e.g. shift resources between cells.



Within the cell internal calculation, the cell creates values for virtual resource efficiencies for the following separate areas and resources:

- 1) Distinction whether or not a particular resource is “interference coordinated” or not; there are two different resource efficiencies for this resource in that area:
 - a) Firstly a Virtual Resource Efficiency for the case that the closest neighbour is creating inter-cell interference on that resource (under the currently present traffic load situation in this neighbouring cell).
 - b) Secondly, a Virtual Resource Efficiency for the case that the closest neighbour is not creating any inter-cell interference on that resource, because that neighbouring cell does not use this resource, the neighbouring cell is restricting it.
- 2) Spatial distinction, separate values for the cell internal virtual sub-areas.
These two above described resource efficiencies are created for each of the virtual sub-areas of the cell.

These “Virtual Resource Efficiencies” are a generic description, which allow reflecting in a generic way several characteristic aspects of a radio network cell:

- 1) These various virtual resource efficiency values are created for each cell individually, considering and reflecting the actual situation in and around that cell.
- 2) These resource efficiencies are reflecting the current traffic load and the current inter-cell interference situation.
- 3) Via the separate handling of inner-cell and outer-cell areas, these resource efficiencies reflect also the resulting cell scheduler policy, i.e. the fairness characteristics how the cell scheduler prioritizes the cell centre and cell border users.

6.4.3.4 Virtual scheduler for assessing a potential new parameter set

With the help of the above employed tools and pre-calculations, the virtual scheduler is then predicting, how well the cell would be able to handle a new situation under the assumption that a new set of configuration parameters (see above: different resources, different cell areas) would be installed in the real system. The main steps of this virtual scheduler are given in the flow-chart Figure 6-9.




Figure 6-9: Simplified operation of the virtual scheduler to predict the result a possible candidate parameter set



The Virtual Scheduler performs the following calculations:

- 1) Calculating the new traffic amount in all the sub-areas of a cell according to traffic shift between the border subareas of two neighbouring cells according to the variation of the cell “cell-border-modifying SON parameter (see above)”,
- 2) Assuming the new candidate resource distribution in the cells in the local area, determine then how many resources are available, and how many of these (and where) benefit from no direct neighbour interference.
- 3) Assume to serve virtually all the offered input traffic and calculate how many resources the cell would need to be able to fully serve all requested input traffic. For this resource wish calculation, (in the most simple form) the following approximating basic equation could be used:

$$\text{CarriedTraffic} = \text{ResourceEfficiency} * \text{NumberOfResources}$$

Thereby this cell-resource-wish calculation is first done for each sub-area separately and in that sub-area the “good” resources (those who do not suffer inter-cell interference) are taken first, and the remaining traffic –if any– is thereafter served by the other resources, by those resources which suffer inter-cell interference from the nearest neighbouring cell.

Then the virtual wishes from all sub-areas are added to obtain the total number of resources (e.g. LTE-PRBs) which would be needed by this cell to fully serve the new traffic amount under the given resource-distribution and the given traffic load situation in the local area.

It shall be noted, that this simple and generic scheduling approximation does not describe the system as precise as the real scheduler who operates on a much shorter time scale and uses much more information + complexity. However, this scheduling describes and predicts well the characteristic behaviour of the cells and to allow very well to compare different parameter options and to decide on optimized parameter sets.

6.4.3.5 Prediction of the system performance and of the energy consumption

The above virtual scheduler does offline calculates (predicts) for an assumed candidate parameter, how many (i.e. fine-granular) resources (e.g. how many LTE Physical Resource Blocks (PRBs)) would be needed to be able to fully serve all the requested input traffic. This then allows predicting the quality of a particular set of candidate parameters by calculating:

- 1) The averaged system throughput and average the user experienced quality of service; this calculation is based on comparing the actually available resources with the virtually wished amount of resources. Assuming e.g.
 - a) that all users can fully be served if enough resources are available,
 - b) that the service quality of the users is degraded by a certain amount (e.g. linear percentage) if the cell does not have enough resources.
- 2) The energy consumption of the system:

It is calculated which amount of resources are available in a cell, the (maximum) power level and the actual “resource usage percentage” (= which amount of the time an available resource is actually being used or scheduled on average). Combined with an energy consumption model, this information is then the basis to estimate the relative energy consumption of the system and to compare the energy consumption of different possible candidate parameter sets.

These performance and energy contributions can then be convoluted to a metric, e.g. to create a single combined rating value. The algorithm which offline varies and assesses a multitude of possible candidate parameter sets is then selecting the best rated parameter combination and will then initiate that this best parameter set will be configured in the system.



6.4.4 Next Steps

The SON approach discussed in previous sections is enabling the spectrum manager to select and optimize the spectrum portfolio while handling and resolving strong interactions and couplings between different parameters and effects. As a distributed CM-SM individual technique it automatically configures and optimizes the situation around each individual CM-SM instance according to its individual particular situation, such as traffic load, interference or spectrum opportunities available. In contrast to other techniques based on measuring the system feedback, a generic prediction model for cellular networks based on fast offline computations is used to quickly determine or improve spectrum portfolio and parameter configuration. It relies on precise modelling of individual cells for predicting in a very diverse and possibly quickly changing situation. Hence the prediction model needs to adapt itself to the currently present cell individual situation. This self-adaptation of the SON model itself is realized by several self-learning techniques as outlined in detail in [D6.4]. Concept and related self-learning techniques discussed will be documented in more detail in deliverable D6.7, including the results of a first assessment regarding capabilities and limitations.

6.5 Cognitive spectrum utilization for stable, dense indoor femtocells

The LSPC from the femtocell perspective is taken into account in the case of using the spectrum-limited orthogonal frequency division multiplexing access (OFDMA) femtocells application. Specifically, multiple indoor femtocells, each of which serves multiple femtocell users, are randomly deployed in a small area (e.g., enterprise environment) within the coverage of the existing macrocell network, and the co-channel deployment of the femtocells causes interference towards neighbouring macrocell users. For this case, the CM-SM functionality proposed in [D6.4] is focused on addressing decision making with respect to the LSPC functionality (i.e., a cognitive and spectrum management functionality) that from the femto perspective maintains the consistence of the LPFR with the corresponding repositories by identifying and developing opportunities for handling the spectrum usage information: the active spectrum capacity (in number of active sub-carriers) per femtocell user and the power allocation per sub-carrier.

For femtocells the LSPC take into account local spectrum management functions with respect to the spectrum capacity increase (in number of available sub-carriers) per femtocell, the active spectrum capacity increase (in number of active sub-carriers) per femtocell user and the power allocation per sub-carrier.

Once the spectrum usage information is given to the femtocell access points, along with a properly chosen power allocation level per sub-carrier, the selections of the active spectrum capacity increase (in the number of active sub-carrier) influences the energy usage balance between the signalling and the data transmission inherent at each femtocell. Decision on such energy usage balance through a selection of the active spectrum capacity increase is taken into account alongside a selection criterion of the power allocation to each sub-carrier under the co-channel interference requirement.

For composing spectrum portfolios in a femtocell scenario the LSPC is made aware of the number of femtocell users that are waiting to access the co-channel spectrum, the range information of which is assumed to be given from the spectrum sensing devices. The LSPC is also aware of the tolerated threshold level that can be used to limit the co-tier interference from femtocells towards the neighbouring incumbent receiver.

LSPC decision-making further needs to consider the joint local management of active spectrum capacity and energy consumption to maintain the reliable performance of the femtocells while guaranteeing the co-channel interference requirements.

6.6 Interfaces

The **CM1 interface** is used by the CM-SM and CM-SM END entities to exchange spectrum portfolio data structures with a CM-RM entity in the networking domain.



The Local Spectrum Control (LSPC) receives spectrum portfolio requests from a CM-SM along with further descriptors detailing the request (e.g. by giving number and desired attributes of spectrum portfolios requested) and with most recent context information if needed. In addition it deploys spectrum portfolios through this interface.

In the presence of a Spectrum Analyser (SAN) and Spectrum Selector (SSE) entity (i.e. for CM-SM END realisations) the LSPC is the first point of contact for a CM-RM. The LSPC entity then utilises the CM1 interface to exchange control information and negotiate requirements and configurations regarding the utilization of SSE and SAN entities in the course of receiving context information from the CM-RM and related spectrum sensing entities (via the SS1 interface) as well as deploying spectrum portfolios in response to requests made by the CM-RM.

In contrast to the CM-SM realisation, where the LSPC responds immediately with a spectrum portfolio data structure, the LSPC of a CM-SM END realisation responds by providing information about the pool of spectrum portfolios a CM-RM may request. This is to assist the CM-RM in requesting the most suitable portfolios on demand in a fast communication with the SSE entity that deploys selected spectrum portfolios for use by the requesting CM-RM through this interface.

The SAN entity utilizes the CM1 interface to obtain pre-processed spectrum sensing information and other context information from CM-RM entities.

The **LPFC interface** is used to exchange spectrum portfolios between the Local Portfolio Repository (LPFR) and the Local Spectrum Control (LSPC). It is an CM-SM internal interface of networking domain entities. Except for CM-SM END entities that implement the PF2 interface, the LPFC interface is the only way to access the LPFR for storing and retrieving deployable spectrum portfolios. The LPFC mainly is a database interface which allows to search for entries, to restrict this search to certain attributes (e.g. spectrum portfolios consisting of specific frequency bands) or to search for spectrum portfolios that satisfy a certain requirement (e.g. that provide a minimum contiguous bandwidth).

For CM-SM END entities that do not implement the PF2 interface, SSE and SAN may directly access the LPFR through an ‘empty’ LSPC entity. This shortcut allows trivial realizations of generating spectrum portfolios through observation and utilizing them immediately. That is, a SAN stores a spectrum portfolio in a degenerated LPFR which in turn is retrieved by an SSE entity. The LSPC then acts as a proxy of the rudimentary LPFR implementing only store and retrieve primitives of the LPFC interface, responding with an error indication for all other primitives.

The Local Spectrum Control (LSPC) entity utilizes the **SPC1 interface** for exchanging spectrum portfolios with remote CSPC entities. It is an CM-SM internal interface between networking domain and coordination domain CM-SM entities. In addition an LSPC may convey spectrum portfolio data structures through a CSPC entity towards remote LSPC entities.

The **SPC2 interface** is used to exchange spectrum portfolios or parts thereof directly between LSPC entities and proprietary control and management function situated in the radio access part without involving a CM-RM entity. It is a CM-SM internal interface of networking domain entities and external (proprietary) entities. The SPC2 interface is implementation and technology-dependent and may be proprietary or standardized in a different scope.

In certain scenarios the LSPC directly communicates local spectrum management decisions for femtocell access points, femtocell controllers and wireless access points through the SPC2 interface. Implementing the SPC2 interface then may involve additional gateway functions that can be seen as minimalistic CM-RM realizations.

The **SAN2 interface** is used to control and configure the Spectrum Analyser (SAN) entity and to exchange spectrum portfolios data structures between SAN and Local Spectrum Control (LSPC) entities in the networking domain. It is an CM-SM internal interface of networking domain entities. A spectrum portfolio data structure when issued by a SAN entity may carry context information or a self-



learned spectrum portfolio depending on the interpretation made by the LSPC's strategies regarding the utilization of a SAN and the pre-processed spectrum portfolio data structures as its outcomes.

The **SS1 interface** is used for the exchange of context information from spectrum sensing entities. The SS1 interface splits between SS1a and SS1b. While the SS1a is used in communication between spectrum sensing and the CM-SM, the SS1b is used between spectrum sensing and CM-RM. It is an interface of the QoS MOS reference model.

An LSPC entity may utilize the SS1 interface to obtain spectrum sensing information without relying upon a SAN or CM-RM entity to obtain spectrum sensing information, information about spectrum sensor capabilities and incumbent detection indications (see also [1900.6], [1900.6a]).

Since an LSPC does not 'own' spectrum sensors, it in turn does not implement spectrum sensor control functions. For configuration and control of spectrum sensors utilized by an LSPC a CM-RM must be involved for selecting and programming a suitable set of sensors avoiding conflicts with its own needs. Sensors in turn then register with the CM-SM END and send their measurement updates via the SS1 interface to the LSPC. For instructing the CM-RM the LSPC conveys a spectrum portfolio to the CM-RM through the CM1 interface, or indicates its need for spectrum measurements along with deploying a spectrum portfolio to the CM-RM for utilization. This approach allows the CM-RM to coordinate and plan utilization of sensors situated in the terminating domain while reducing communication overhead by duplicated information if a sensor is reporting to a CM-RM as well as to a CM-SM.

The **PF2 interface** is used to exchange spectrum portfolios between the portfolio processors Spectrum Analyser (SAN) and Spectrum Selector (SSE), and the Local Portfolio Repository (LPFR). It is an CM-SM internal interface of networking domain entities.

The **AL1a** (LSPC-AL) control interface provides communication with other networking domain and coordination domain entities.



7 Spectrum analyser and selector

The two CM-SM architectural entities spectrum analyser (SAN) and spectrum selector (SSE) are optional elements and are realized in the networking domain as a function of a CM-SM END entity only. They are directly interfacing with an LSPC (section 6.2) and LPFR (section 4.5) entity to enable

- a. Quick deployment of a spectrum portfolio to a spectrum user through an instance of the CM-RM and
- b. Compose spectrum measurements into a spectrum portfolio data structure for later use as a context parameter or as a self-learned spectrum portfolio for later deployment via an SSE entity.

The main purpose of the SSE is to provide a caching function for spectrum portfolios, which significantly reduces response times for users requesting spectrum through avoiding most of the cognitive decision processes of a fully featured CM-SM. In collaboration with an LSPC, an SSE entity may implement part of the decision-making by selecting from a set of spectrum portfolios prepared and pushed to the LPFR by the LSPC in a more complex cognitive process. The decision process performed by the SSE then selects a suitable spectrum portfolio potentially based on the spectrum analysis performed by the SAN entity. To support basic spectrum sharing scenarios the SSE may implement in addition simple spectrum portfolio split and merges.

A SAN / SSE pair would be able to collect spectrum measurements, to create a spectrum portfolio out of these and to store the portfolio in the LPFR. An SSE entity may request the LPFR and retrieve this as a spectrum portfolio to be deployed to a spectrum user.

The LPFR is involved in such scenario as a spectrum portfolio store only, that can be accessed also by the LSPC for management purposes such as deploying an initial spectrum portfolio to the SSE or for retrieving sensed context or spectrum portfolios for further processing. Herein the LPSC is involved mainly as a management entity not demanding for any cognitive capacity, but cannot be omitted in whole even for the most basic configuration since it takes responsibility of the control functions that allocates initial spectrum portfolios upon request of the CM-RM, which cannot be directed to the SSE.

The configuration described above enables a realization of spectrum management for a single node (e.g. a mobile terminal or an ad-hoc terminal) that may operate based on spectrum sensing in a network that only occasionally has connectivity with an infrastructure or where initial spectrum portfolios, usage constraints or policies are deployed only once at manufacturing time, for example.

Since most of the SSE and SAN functions are algorithmic, subsequent descriptions will focus on the functionality directly related to or interacting with other CM-SM functions that employ cognitive or opportunistic capacities.

7.1 SAN functions

The main functions of a SAN entity are that of a context filter and analyser. It receives spectrum measurement information from spectrum sensors and related pre-processed information from associated CM-RM entities. When receiving context from multiple spectrum sensors or CM-RM instances the SAN also performs context fusion algorithms. The major outcome of this process are one or more spectrum portfolio data structures consisting of descriptors of the frequency bands for that measurements have been performed as well as measured parameters and parameter values related to these frequency bands. Complementing those spectrum measurements a CM-RM may also provide context information that associates other or derived information such as data stream measurements to a certain frequency band. In particular, data rates, bit error rates, SNR or SINR, or data rates experienced on higher protocol layers (i.e. data link layer) may contribute to an overall 'quality of spectrum' metric for the purpose of evaluating suitability of a portfolio for a certain usage scenario prior to deploying a spectrum portfolio to the spectrum user.



Cognitive functions of the SAN are limited to context filtering and processing. Reasoning and decision-making hereby enables the SAN to handle dynamic spectrum portfolios as determined by the SSE and LSPC and their deployment strategies. That is, all changes of the spectrum portfolio which is communicated between SSE and CM-RM and utilized by spectrum users in the terminating domain will result in a more or less different set of context parameters observed (e.g. different, more or less frequency bands to observe).

In particular the SAN depends on a number of primitive decision rules that control composition of elementary operations on parameters (e.g. routing through processing elements, selection of fusion schemes, and configuration of time-domain interpolation or decimation if needed). In addition, robustness issues such as assessing accuracy, relevance and trust of parameters prior to establishing the details of processing may be needed (see [D6.4]).

The SAN may also utilize algorithms that allow detecting and classifying spectrum user activity of both incumbents and other spectrum users. When detecting a certain kind of incumbent, the SAN then may reconfigure context processing to control accuracy and adjust relevance of context parameters. When detecting, for example, a PMSE device in a TV band, the SAN may need to switch to a narrow band analysis scheme to decide if there still is a TV white space opportunity for neighbouring bands.

In order to utilize the SAN as a ‘versatile context processor’ the SAN must obtain an overall analysis strategy either from the LSPC or at manufacturing time. This strategy first of all determines the goal of the context analysis, which is either to provide context data to the LPFR for later use by the SSE in selecting a suitable spectrum portfolio from those stored in the LPFR, or to decide if an opportunity exists that will extend the choices available to the SSE. Both is a valid strategy and may be used in conjunction. In consequence the SAN emits spectrum portfolio data structures to the LPFR that need to self-describe as a spectrum portfolio or as a context parameter set. Both query functions of the SSE as well as LPFR database smart search functions need to make this distinction too.




Figure 7-1: Accessing a SAN and sample MSC (SAN configured to create a portfolio from raw sensor data for storing to LPFR)



So far the SAN may be able to create spectrum portfolios from spectrum observation, but it cannot create policies from observations. Although thinkable to create a policy from observing the behaviour (i.e. etiquette) of other spectrum users, such an approach is very limited since it requires sophisticated spectrum sensing and spectrum user classification as well as a complex reasoning and decision-making capacity. The trustworthiness of potential results thus is questionable. Hence, regulatory policies and other spectrum usage constraints cannot be replaced by an autonomous process and must be made available to the SAN and to the SSE prior to initially accessing spectrum.

Nevertheless it is reasonable to allow a SAN to infer usage policies (potentially resulting in self-learned operator's policies) from spectrum observations by evaluating the gain or benefit of a decision through applying suitable performance metrics, for example. That is, if the context obtained indicates that a recent change of spectrum portfolio by the SSE has led in consequence to, for example, an increase of efficiency of spectrum use, it may infer that the salience of the new portfolio is higher than the old one and may recommend its preference in form of a policy to the SSE. Inferred policies then may be communicated to the SSE through assistance of the LSPC in order to allow the SSE to increase performance of its local spectrum decisions in the future. In fact this strategy describes a collaborative and distributed reinforcement self-learning process.

7.2 SSE functions

The main functions of an SSE entity are that of a spectrum portfolio cache and intelligent selection function applying decision-making to select and deploy spectrum portfolios to spectrum users. It is accessing the LPFR to retrieve a spectrum portfolio from a pool of portfolios made available by coordination domain CM-SM entities or by a collaborating SAN entity through the LSPC. For certain scenarios (e.g. TV white space utilization) the LSPC may also consult a Geolocation database either directly or through a coordination domain CM-SM.

Smart search functions of the LPFR need to support locating spectrum portfolios for retrieval based on descriptive attributes of resources or usage constraints such as, for example

- Searching for a best match of frequency bands, (i.e. centre frequency, bandwidth, RF emission or duty cycle constraints ...), contiguity of spectrum (i.e. amount of contiguous space vs. frequency gaps), price of spectrum (i.e. price vs. lease time), minimum quality of spectrum (i.e. average activity of other spectrum users), or geographical areas (e.g. disjunctive vs. overlapping).
- Defining precedence for attribute matches such as setting a preference for a match in contiguity of frequency bands vs. quality of spectrum.
- Searching for groups of spectrum portfolios such as those having disjunctive frequency bands (e.g. for normal and back-off operations) or complementing frequency bands (e.g. for normal operations and off-loading or handover purposes).

Hence, the LPFR must be able to provide upon request multiple spectrum portfolios as a result of a single search operation, which can be considered state-of-the-art for both relational and object oriented databases. The SSE then caches those spectrum portfolios and delivers on demand of its associated CM-RMs.

The SSE deploys a spectrum portfolio upon request of a CM-RM entity. The selection of a suitable spectrum portfolio relies on context information stored by the LPFR in form of one or more spectrum portfolio data structures, and of context conveyed by the CM-RM along with its request. All spectrum portfolios obtained from the LPFR consist of a description of the spectrum opportunity including usage constraints and policies applicable. In addition, spectrum portfolio data structures obtained from



the LPFR may also provide context information for decision-making.¹ In this case policies included with the spectrum portfolio data structure must prohibit its use as a spectrum portfolio.

Decision-making by the SSE is time-constrained, which forbids complex cognitive processes for the time being. In consequence SSE decision-making might be algorithmic or heuristic in form of a neural network, for example. In fact, the collaboration of SSE and LPFR in spectrum portfolio selection forms a case-based reasoning process in that the SSE realizes reasoning and decision-making and the LPFR provides the ontology.

The SSE operates on a pool of spectrum portfolios stored in the LPFR that must be constructed in a suitable way for being deployed without further considerations.

- Due to timing constraints for SSE requests, spectrum portfolios must be deployable without change, or must require only minimal modifications (i.e. simple split and merge operations) before deployment towards a spectrum user. That is, the SSE must not be obligated to compose spectrum portfolios.
- The number of distinct spectrum portfolios stored by the LPFR must be adequate for a given purpose or scenario. The number of spectrum portfolios justifying the implementation of an SSE entity in a certain configuration depends on the number of spectrum users, the geographical area covered, the number and dynamics of incumbents and their interference protection requirements, for example.
- Since an SSE may serve more than one CM-RM at a time, different spectrum portfolios may be deployed towards different CM-RMs. In consequence, spectrum portfolios should be composed and grouped for certain goals such as mitigating interference by spectrum reuse over distance. That is, similar to conventional spectrum planning, spectrum portfolios may be composed for complementing each other in terms of lease time, coverage area and frequency, for example.

The LSPC is responsible for ensuring such constraints since it can apply more complex cognitive processes compared to the SSE. That is, composing of new spectrum portfolios satisfying above demands takes place in parallel with SSE operations and results in an LPFR update ‘in the background’. In addition, the LSPC needs to control SAN operations in a way that ineffective, unfavourable or conflicting spectrum portfolios will not be persistently stored in the LPFR.

In consequence the flexibility of an SSE is strictly limited which makes it a QoSMOS CM-SM entity that is optimized for a single purpose and for very few scenarios only.

7.3 Interfaces

The **PF2 interface** is used to exchange spectrum portfolios between the portfolio processors Spectrum Analiser (SAN) and Spectrum Selector (SSE), and the Local Portfolio Repository (LPFR). It is an CM-SM internal interface of networking domain entities and applies to CM-SM END entities only.

The **PF2** (LPFR-SSE) interface is used by the SSE to retrieve spectrum portfolios from the LPFR and the **PF2** (LPFR-SAN) interface is used by the SAN to store spectrum portfolios to the LPFR.

The **SAN2 interface** is utilized by the Spectrum Analyser (SAN) entity to forward spectrum portfolio data structures to the Local Spectrum Control (LSPC) entity for evaluation and further processing. The LSPC is utilizing the SAN2 interface for configuration and control of the SAN. The SAN may be configured by the LSPC to directly forward spectrum portfolio data structures to the LPFR by means of PF2 interface primitives. A spectrum portfolio data structure when issued by a SAN entity may carry context information or a self-learned spectrum portfolio depending on the interpretation made by

¹ In fact, spectrum portfolio data structure always can provide additional context, regardless if they are utilized as a spectrum portfolio or not.



the LSPC's strategies regarding the utilization of a SAN and of the pre-processed spectrum portfolio data structures as its outcomes.

The **CM1 interface** is used by the CM-SM and CM-SM END entities to exchange spectrum portfolio data structures with a CM-RM entity in the networking domain.

The LSPC entity utilizes the CM1 interface to exchange control information and negotiate requirements and configurations regarding the utilization of SSE and SAN entities in the course of receiving context information from the CM-RM and related spectrum sensing entities (via the SS1 interface) as well as deploying spectrum portfolios in response to request made by the CM-RM.

The SSE entity utilizes the CM1 interface to receive spectrum portfolio requests from a CM-SM along with further descriptors detailing the request (e.g. by giving number and desired attributes of spectrum portfolios requested) and with most recent context information if needed. Furthermore, the SSE deploys selected spectrum portfolios for use by the requesting CM-RM through this interface.

The SAN entity utilizes the CM1 interface to obtain pre-processed spectrum sensing information and other context information from CM-RM entities.

The **SS1 interface** is used for the exchange of context information from spectrum sensing entities. The SS1 interface splits between SS1a and SS1b. While the SS1a is used in communication between spectrum sensing and the CM-SM, the SS1b is used between spectrum sensing and CM-RM. It is an interface of the QoS MOS reference model.

The SAN entity receives spectrum sensing information through this interface directly from spectrum sensing entities.




Figure 7-2: Accessing an SSE and sample MSC (initialization and update of an SSE-deployed spectrum portfolio involving SS and CM-RM context updates)



8 Summary and Conclusions

This deliverable concludes the refined specification of cognitive and opportunistic functions of the spectrum management framework. It grounds upon deliverable [D6.3] initially defining scope, goals and limits of cognitive functions within the framework, and is complemented by deliverable [D6.4] elaborating on trust, security, privacy and reliability and robustness of cognitive capacity. This deliverable provides a description and informal specification of cognitive functions and self-learning capabilities of the framework.

This deliverable provided a coarse overview of the interaction between QoS MOS cognitive managers (CM-SM and CM-RM) and elaborated further on interfaces involved in this interaction from a specific perspective of cognitive spectrum management. Then, based on a functional decomposition of the QoS MOS cognitive spectrum manager, CM-SM internal functional modules and related interfaces are described focussing on cognitive capacities. This description is complementing [D6.2] (context filtering, aggregation and communication) and [D6.4] (flexibility, robustness and cognition) and concludes [D6.3] (initial description of functions of the spectrum management framework). It will form the basis for upcoming deliverables that will provide a final integrated specification and will highlight implementation aspects.

The specification details on the databases of spectrum portfolios and spectrum policies, their functional role in the context of the CM-SM architecture, their internal functionality and the content they manage. In particular, functionality that goes beyond mere database functionality is elaborated in more detail. Cognitive spectrum management functionality co-located with coordination and networking domain is presented and their interaction across domains is discussed. In that it considers the main QoS MOS scenarios regarding cellular, femtocell and ad-hoc configurations with respect to their impact on the cognitive decision-making functions and strategies, context considered and output produced. A number of annexes complete the specification by further detailing on the concepts, approaches and solutions for spectrum user modelling and opportunity detection and selection that form the functional basis for cognitive spectrum management.

Upcoming deliverables D6.6 (Spectrum management framework integration and implementation report) and D6.7 (Integrated final functional specification of spectrum management framework and procedures) will conclude the specifications of the cognitive spectrum management framework by further elaborating on function details, assessment and testing of functions, on interface primitives and data structures, on protocols as well as on the assessment of the framework in whole.

The concepts and solutions discussed in this deliverable will be forwarded to the proof-of-concept evaluation (prospectively in form of executable computer programs) planned for QoS MOS WP7. From the perspective of a CM-SM development this integration will have to concentrate on three focus topics. The proof-of-concept evaluation then will complement the framework assessment by providing key performance indicators and measurement results:

- It should validate the interfaces between CM-SM entities, and between CM-SM and other entities provided by other work packages.
- It should provide a proof-of-concept for cognitive functions in spectrum management focussed on the functions that have to be present for demonstration of key scenarios.
- It should clearly identify key performance indicators and verify which parameters and metrics discussed so far prove suitable to evaluate the performance of a spectrum management system based on cognitive and self-learning capacity.

A **validation of interfaces** in general requires a reference implementation including implementations for major application scenarios. Clearly, the QoS MOS project does not have sufficient resources available to implement all or even the main scenarios in whole as set forth by the business evaluation of WP1. Fortunately, the design of interfaces as specified by this deliverable is based on the exchange



of spectrum portfolio data structures mainly. Interface primitives for the various interfaces of the CM-SM architecture are very similar to each other and thus may be seen as specific sub-sets of some (abstract) top-level interface that combines an interface-specific command set with a generic information element common to all interfaces. A validation of interfaces thus can be made cumulative by implementing and testing representative interfaces in the course of a proof-of-concept, which will at least connect CM-SM entities with other CM-SM entities (SPC1 interface) and with CM-RM entities (CM1 interface).

A **proof-of-concept of cognitive functions** is more difficult to achieve since it relies on the availability of theoretical solutions for a certain optimization problems. In addition, the result of a cognitive decision-making process very much depends on context and a-prior knowledge available to the cognitive engine. Furthermore, the QoS MOS CM-SM relies on algorithms, cognitive functions and collaborating cognitive engines at the same time. Their interaction thus must be considered when evaluating the outcome of a cognitive decision-making process, which may deviate from the optimum case due to a fuzzy behaviour of the collaboration of all these.

When operating within its training set, a cognitive function can be assumed as stable and optimal within its functional limits (which may be variable when considering self-learning capacity) if the results of decision-making are reproducible for a certain fixed set of static context parameters and if they are ideally closest to the optimum case regarding the specific optimization goal. In order to prove such behaviour, the environment in which a cognitive 'system-under-test' must operate is likely completely artificial. The uncertainty of real-world context will make it a very demanding task to evaluate correctness of cognitive decision-making since context is likely much less reproducible and may lead to a completely different internal behaviour, although its output (e.g. the optimum spectrum portfolio) will remain the same.

Beyond its training set a cognitive engine is expected to behave more robust and resilient than comparable algorithmic solutions. Suitable **performance parameters and metrics** are needed to evaluate this behaviour, preferably without requiring knowledge about an optimum solution to compare with. In general absolute metrics are required to compare performance with earlier test cases and relative metrics are required to evaluate the behaviour of a cognitive engine if no such test cases are available (which is often the case for operating outside training environment in real-world set-ups).

Given that the main output of cognitive spectrum management consists of one or more spectrum portfolios, absolute metrics are mostly static and can be evaluated by the CM-RM in terms of spectrum efficiency, interference level, number of eviction events, number of users supported in a certain portfolio, CM-SM response times and similar. Relative metrics are likely to be evaluated internally by the CM-SM and consist of, for example, the number of alternative actions for consideration to decision-making (which impacts the CM-SM response time to CM-RM requests), dynamics of spectrum portfolios (e.g. expressed in fragmentation, duplication, underflow rate of spectrum blocks available to satisfy a request), improvement rate over time or spatial distance (e.g. in terms of response time and spectrum portfolio dynamics as well as spatial reusability), and improvement over time of safety margins required (e.g. in terms of spectrum overprovisioning to satisfy CM-RM requests).

As can be seen from these examples most of the performance parameters and metrics are linked to each other and future efforts will have to address among other challenges the development of a set of (almost) orthogonal performance metrics. Such set of metrics maybe helpful for current optimization of the CM-SM but will be mandatory for comparing different CM-SM realisations and for setting up test specification for cognitive spectrum management in the future.



9 References

- [D1.2] “QoS MOS consolidated scenarios”, FP7-ICT-2009-4/248454 QoS MOS project deliverable D1.2, Dec. 2010.
- [D2.1] “Initial description of system architecture options for the QoS MOS system”, FP7-ICT-2009-4/248454 QoS MOS project deliverable D2.1, Apr. 2010.
- [D2.2] “System architecture options for the QoS MOS system”, FP7-ICT-2009-4/248454 QoS MOS project deliverable D2.2, Dec. 2010.
- [D2.3] “System specification and evaluation criteria”, FP7-ICT-2009-4/248454 QoS MOS project deliverable D2.3, Nov. 2011.
- [D6.1] “Initial description of spectrum management framework, requirements analysis and approach selected”, FP7-ICT-2009-4/248454 QoS MOS project deliverable D6.1, Jun. 2010.
- [D6.2] “Description of context filtering, aggregation and communication functions of the spectrum management framework”, FP7-ICT-2009-4/248454 QoS MOS project deliverable D6.2, Apr. 2011.
- [D6.3] “Initial description of cognitive and opportunistic functions of the spectrum management framework”, FP7-ICT-2009-4/248454 QoS MOS project deliverable D6.3, Nov. 2011.
- [D6.4] “Trust and security functions of the spectrum management framework”, FP7-ICT-2009-4/248454 QoS MOS project deliverable D6.4, Apr. 2012.
- [IEEE2003] “Traffic Models for IEEE 802.20 MBWA System Simulations”, IEEE P 802.20; Draft 802.20 Permanent Document, July 8, 2003
- [ORACLE2008] “Public Deliverable D1.3 – ORACLE System Model and Implementation Guide“, The ORACLE consortium, FP6-IST-2004 027965, www.ist-oracle.org, November 2008.
- [XGL2004] DARPA XG Working Group, “DARPA XG Policy Language Framework, Request for Comments”, v. 1.0, prepared by BBN Technologies, Cambridge MA, USA, 2004.
- [TS36.413] 3GPP TS36.413, “3rd Generation Partnership Project; Technical Specification Group Radio Access Network; Evolved Universal Terrestrial Radio Access Network (E-UTRAN); S1 Application Protocol (S1AP) (Release 10)”, Mar 2012
- [Aky08] I. F. Akyildiz, W-Y. Lee, M. C. Vuran and S. Mohanty, “A Survey on Spectrum Management in Cognitive Radio Networks”. IEEE Communications Magazine, vol. 46, no. 4, Apr. 2008, pp. 40-48.
- [1900.1a] IEEE Draft Std 1900.1a™ / D0.06; “Draft Standard for IEEE Standard Definitions and Concepts for Dynamic Spectrum Access: Terminology Relating to Emerging Wireless Networks, System Functionality, and Spectrum Management. Amendment: Addition of New Terms and Associated Definitions”; February 2012
- [1900.6] IEEE Std 1900.6™-2011; “IEEE Standard for Spectrum Sensing Interfaces and Data Structures for Dynamic Spectrum Access and other Advanced Radio”; April 2011
- [1900.6a] IEEE Draft Std 1900.6a™; “Draft IEEE Standard for Spectrum Sensing Interfaces and Data Structures for Dynamic Spectrum Access and other Advanced Radio – Amendment: Procedures, Protocols and Data Archive Enhanced Interface”; June 2012
- [Fette06] Fette, B.; “Cognitive Radio Technology”, Newnes, 2006.



-
- [1900.4] IEEE Std 1900.4™-2009; “IEEE Standard for Architectural Building Blocks Enabling Network-Device Distributed Decision Making for Optimized Radio Resource Usage in Heterogeneous Wireless Access Networks”; Feb 2009.
- [1900.4a] IEEE Std 1900.4a™-2011; “IEEE Standard for Architectural Building Blocks Enabling Network-Device Distributed Decision Making for Optimized Radio Resource Usage in Heterogeneous Wireless Access Networks – Amendment 1: Architecture and Interfaces for Dynamic Spectrum Access Networks in White Space Frequency Bands”; Sep 2011.
- [1900.5] IEEE Std 1900.5™-2011; “IEEE Standard for Policy Language Requirements and System Architectures for Dynamic Spectrum Access Systems”; Jan 2012
- [1900.5.1] IEEE Draft Std 1900.5.1™; “IEEE Standard Policy Language for Dynamic Spectrum Access Systems”; June 2012



A Opportunity selection

A.1 Selection of DTV bands for LTE uplink extension

A.1.1 Introduction

The CM-SM is responsible for building the spectrum portfolios based on a set of external constraints, such as regulatory and operator policies, and on spectrum sensing results. The CM-SM is responsible as well for the management of the spectrum portfolios, including cognitive spectrum management (decision making) methods to decide how to allocate portions of the available spectrum to the requesting entities or spectrum users (e.g., the base stations of a cellular system) [Aky08]. The SSE is the module responsible for this particular task. The CM-SM operates over relatively wide blocks of spectrum, at medium/long time scales and taking into account several licensed primary bands. When building up the spectrum portfolios and selecting candidate bands for secondary operation, the CM-SM needs to account for the potential consequences of selecting a certain band, not only for the primary system in terms of resulting interference levels, but also for the secondary system in terms of performance as well as the overall efficiency of spectrum utilisation. One of the candidate spectrum bands commonly considered for CR applications, and also within the framework of the QoS MOS project, is the Digital TeleVision (DTV) band. This constitutes indeed the “cellular extension in TeleVision White Spaces (TVWS)” scenario defined in [D1.2]. However, reusing a particular primary spectrum band by making use of a specific secondary technology has an impact on the operation and performance of both systems, thus requiring a careful and detailed study on the conditions under which the coexistence between primary and secondary systems in the considered scenario would be feasible along with the resulting technical implications. Both quantitative and qualitative reference results as well as some guidelines would be useful in order to help decision entities, i.e. the Spectrum Selector (SSE), to make decisions on the adequacy of selecting DTV bands for the extension of a cellular system and evaluate the expectable consequences in terms of protection of the primary DTV system, performance of the LTE cellular system and efficiency of spectrum utilisation. The following sections provide a more detailed discussion of this scenario and its motivation as well as several metrics to be considered by the SSE, regarding DTV bands, when preparing the pool of portfolios and making pre-calculations on suitable spectrum bands for opportunistic usage.

A.1.2 Considered scenario and motivation

The considered scenario comprises a DTV broadcast link as the primary system and a LTE cellular network as the secondary system (see Figure A-1). As it can be appreciated in Figure A-1, the DTV station broadcasts a TV signal for the DTV receivers within a certain coverage area. However, this signal is also captured by the receivers of the LTE system, thus leading to some interference levels on the cellular network. Similarly, the signal of LTE transmitters leads to some undesired interference over the incumbent receivers.

The focus of the considered scenario is on the uplink of the LTE system. The benefits of exploiting primary bands for uplink transmissions are manifold. The path loss reduction due to a lower frequency of operation² results in increased battery life for the mobile terminals and coverage outage reduction since the uplink is more seriously power-limited. Moreover, the reuse of licensed bands for uplink transmissions enables the LTE system to place downlink transmissions, which are in general more bandwidth consuming, in part of the spectrum allocated to the uplink, thus leading to an increase in the overall system capacity. This scenario compliments previous studies of the same scenario performed in the context of QoS MOS, where the focus was on the downlink (Section 7.5 of [D6.1], Section 4.2 of [D6.2] and Section 4.3 of [D6.3]).

² The path loss reduction from the LTE band (2000 MHz) to the TV band (600 MHz) is around 18 dB according to the COST 231 Hata model and 10 dB according to the free space model (worst case).




Figure A-1: Considered scenario.

A.1.3 Conformance metrics

The feasibility of selecting a portion of a DTV band for opportunistic LTE uplink transmissions depends on the resulting interactions between the primary DTV and secondary LTE systems. Such interactions and the resulting performance of both systems can be analysed by means of three main groups of conformance metrics, aimed at analysing and quantifying the protection of the primary system, the performance of the secondary system and the efficiency of spectrum utilisation. Most of these metrics are outlined in or based on those described in [D2.1] and [D6.1].

The protection of the primary system can be analysed by means of the Carrier-to-Noise Ratio (CNR) and Desired-to-Undesired power Ratio (DUR), the latter following the same concept of the Carrier-to-Interference Ratio (CIR). While the CNR is independent of the secondary system and its interference, this parameter allows to determine the distance from the DTV transmitter at which the minimum CNR is satisfied and thus the intended coverage area of the primary transmitter. Within this coverage area, the aggregated interference generated by the secondary system (quantified by means of the DUR) must be lower than the maximum tolerable level. In other words, taking as a reference point the primary DTV transmitter, the distance at which the minimum required DUR is observed must be larger than the distance at which the minimum required CNR is experienced in order to guarantee an appropriate protection of the primary system. Notice that an appropriate operation of the DTV receivers requires not only a minimum CNR but also a minimum DUR to be met. This concept is illustrated in Figure A-2. As long as this protection criterion is met, the DTV band can be selected by the SSE for its secondary reutilisation by the LTE system.

The performance of the secondary system can be analysed mainly in terms of transmission rates such as the net data throughput. However, other metrics can be useful as well to quantify the performance of the secondary system in terms of error rates, such as the Block Error Rate (BLER) and Bit-Error Rate (BER), and the experienced channel quality in terms of common metrics such as the Signal-to-Interference plus Noise Ratio (SINR) experienced by the User Equipments (UEs) and the employed transmission powers. Not only the average values of these metrics but their distribution over the UEs of the system need to be considered for a more complete evaluation of the real performance.

The efficiency of spectrum utilisation can be quantified in terms of a Spectrum Efficiency Index (SPI) similar to that defined in [D2.1], which for the particular case of LTE can be expressed in terms of the quotient between the number of Resource Blocks (RBs) allocated in a sector/cell and the total number of RBs available in that sector/cell. For example, for a 5-MHz chunk of licensed spectrum, the number of RBs available in a LTE channel is equal to 25. If the average number of exploited RBs is 20, then $SPI = 20/25 = 0.8$. This definition of the SPI quantifies the efficiency of spectrum utilisation in terms



of the fraction of available licensed spectrum being exploited by the secondary system. However, it is worth noting that the SPI does not quantify how efficiently is being exploited the spectrum that is actually being used. For example, in some cases the scheduler may be configured to distribute all the available RBs among the requesting UEs. In such a case, the SPI would always be equal to one as long as there is at least one UE per sector/cell. However, the whole chunk of spectrum may be used by one or several UEs at various modulation and coding rates in order to achieve the desired data rate per user, thus leading to various efficiency levels for the same SPI. In such a case, the SPI would not be a representative metric of how efficiently the spectrum is actually being used. An alternative and more convenient metric to quantify the efficiency of spectrum utilisation is the *bandwidth utilisation*, defined as the quotient between the total data throughput in a sector/cell and the maximum achievable bit-rate at the highest modulation and coding rate. The main interest of this parameter lies in its ability to quantify the real efficiency of the spectrum utilisation in a single parameter by capturing the impact of many relevant aspects such as the overhead resulting from collisions, signalling messages, packet headers, back-off timer delays and any other network control data. Spectrum efficiency can also be evaluated in terms of the classical concept of data rate (bits per second) per bandwidth unit (Hz).




Figure A-2: Relation between CNR distance and DUR distance: a) proper operation of all DTV receivers, b) improper operation of some DTV receivers.

These conformance metrics will be used to analyse and evaluate, based on comprehensive system-level simulations, the adequacy of selecting DTV bands for the extension of the uplink component of an LTE cellular system as well as the expectable consequences in terms of protection of the primary DTV system, performance of the secondary LTE system and efficiency of spectrum utilisation. The final results along with some guidelines for the SSE will be provided in a subsequent deliverable.



B User activity models

B.1 Introduction

In CR systems the observation of activity for different user types plays important role. The observation is related both to the incumbents and to the opportunistic users. We do not deal here how these observations can be done physically, but we suppose that a cognitive system is prepared to acquire as many information about the system operation and about the environment, as it is possible. However, not all of this information is required for the different types of modelling activity. The context filtering mechanism supports the algorithms at the higher levels; it receives, sorts, ranks the available information and only the relevant data will be transferred to the decision algorithms.

The context information can be sensors data, spectrum sensing information, traffic and channel measurements on wired or wireless mediums. In this section we introduce different models that are candidates to be implemented in the spectrum management framework in order to support the cognitive decision making and enhance the efficient frequency allocation in the system.

One of the proposed cognitive decision making algorithms is related to the modelling of the long-term user activity.

B.2 Long-term activity model for incumbents and opportunistic users

The observation of the long-term incumbents and opportunistic user ON/OFF activity may lead to the large-scale overview of a cognitive system. The observed activity duration statistic can be used to build a model to express the distribution of the length of the active and activity-free periods. If the probability that the incumbents utilize the channel for a given period is less than the expected opportunistic user activity duration, an opportunity is detected for a cooperative operation. In order to calculate the required statistics, a continuous updating of the activity parameters is necessary. This could be the task of the CM/SM system; therefore a computing and data storage capability is required in these entities. The mapping of the task and functionalities to the QoSMOS CM-SM reference model is depicted in Figure B-1:




Figure B-1: Mapping the user activity model to QoSMOS CM-SM reference model



Two different scenarios have been investigated, a general radio channel and a wireless IEEE 802.11 computer network as a primary channel. The ON/OFF properties of the channels are derived from measurements. A model for aggregated incumbents' activity has been fitted to the ON/OFF process. The opportunistic users are taken in account as Internet users with behaviour specific parameter sets.

According to the simulations, the aggregated activity-free length distribution of the incumbents gives the possibility for opportunistic users to join to the same network. The results can be adapted as a general tool in the CM_SM system to support the decision mechanisms.

In section B.3 the algorithm will be detailed as a probability-chain based approach.

B.3 ON/OFF activity based model

This section details a probability-chain based approach for long-term user activity modelling. The two main sections are summarizing the long-term incumbents and the opportunistic user behaviour modelling.

B.3.1 ON/OFF Markov-chain model for incumbents activity

The base model for incumbents ON/OFF activity is a 2-state discrete time/state Markov chain.

State 1 represents the OFF (inactive) state, while state 2 the ON activity. The complementary cumulative distribution (CCDF) the ON state duration can be analytically expressed for n discrete time slots up to $N \rightarrow \infty$. This function gives the probability that the user activity is ON for duration n or longer. Similarly, the CCDF for the OFF state duration denotes the probability that the user activity is OFF for duration n or longer. This kind of user activity model is applicable as a generative model for a synthetic ON/OFF time series.

B.3.2 Estimating from spectrum sensing

To determine the transition probabilities for the ON/OFF model, a feasible solution is observing signal strength levels from a spectrum measurement. There are dedicated entities in the CM-SM model to perform this task. By scanning and recording the radio band where the incumbents are communicating, valuable information can be gathered about the user activity as the function of time and frequency, as Figure B-2 shows:




Figure B-2: One-day spectrum measurement

The incumbent ON/OFF activity at a selected frequency can be determined from the signal strength at a specific threshold. The timing of the ON/OFF sequence is applicable to parameterize the two-state Markov model and determine its transition matrix.




Figure B-3: Signal strength and ON/OFF activity

B.3.3 Estimating from packet traffic observation

In order to determine the ON/OFF model parameters an analysis of network traffic over an IEEE 802.11 wireless access-point can be also successful. By scanning and recording the number of data packets as the function of time, valuable information is provided about the incumbent's activity.




Figure B-4: Network traffic measurement

In order to determine the incumbent activity, the number of packets can be extracted and count for a specific duration. The timing of this ON/OFF sequence can be applied to parameterize a two-state Markov model and determine its transition matrix.




Figure B-5: Network traffic and ON/OFF activity



B.3.4 Aggregation of the incumbents activity

Considering more than one incumbents, the user activity can be aggregated. We calculated the OFF duration distribution for the aggregated ON/OFF process and a specific multiple-state, partitioned Markov chain (Fritchman model) has been fitted to this model. The bursty behaviour of the aggregated process proved that this well-known model will be applicable in the CM-SM architecture.




Figure B-6: Aggregated OFF length duration distribution for different number of users

B.3.5 ON/OFF model for opportunistic user activity

The opportunistic users would like to utilize the same radio spectrum as incumbents when no main (primary) activity can be foreseen. In our approach we specify the opportunistic users that are generating internet traffic (WEB browsing, using chat applications, email send/receive operations, transferring files, etc.). In order to calculate the ON/OFF model parameters, we selected a representative internet usage statistics for online computer use at different locations (office, home, other) extended with the knowledge of the average online durations. To refine this model, we assumed that the model parameters are varying and following a normal distribution around their mean values.




Figure B-7: The ON/OFF model for opportunistic users

After the simulations the number of synthesized opportunistic user ON/OFF time series can be calculated. Moreover, an aggregation system similar to that one that has been applied for incumbents' activity, the ON duration activity for opportunistic users can be determined:




Figure B-8: CCDF of the aggregated opportunistic user ON activity for different user numbers

The distribution function of the opportunistic user ON activity can be applied to estimate the probability of a specific duration, when the opportunistic user is generating network traffic.

B.3.6 Application of the model in the CM-SM decision-making system

The applicability of the long-term user activity model in the CM-SM system is quite clear. From the sensor data and other measurement results the context filtering entities are capable to transfer the most relevant information to the processing unit, where the model parameters for incumbents and opportunistic users can be determined. Afterwards the aggregated incumbent's OFF duration statistics and the aggregated ON duration statistics for the opportunistic users can be continuously calculated and adapted to the varying environment.

The decision making is simple the comparing of this distributions: if the probability that the incumbents are OFF for a specific duration is higher, than the probability that the opportunistic users are ON, there is an opportunity in the system to utilize the same channel/spectrum by the opportunistic users, that is originally dedicated to the incumbents.

The next figure depicts this opportunity detection algorithm. This functionality can be placed in the Common Spectrum Control system (CSPC) as part of the spectrum management framework.




Figure B-9: Opportunity detection in CSPC

B.3.7 An example of using the ON/OFF model in the IEEE 802.22 standard

An interesting usage of the ON/OFF model in state-of-the-art cognitive radio systems is to apply it for the established IEEE 802.22 standard. One concern for the IEEE 802.22 standard operating the TV bands is the appearance of wireless microphones as the primary users.



Wireless microphones will in many cases appear suddenly following unpredictable patterns. An example of such is the usage of wireless microphones to cover sudden media events, e.g. a car accident or a sudden crime. In this case, the ON/OFF model will be of limited value. However, there are other cases where wireless microphones are used following more predictable patterns, e.g. at specific times in churches every Sunday or at concerts every Wednesday, Friday and Saturday. These usage patterns might be monitored by sensing functionality in IEEE 802.22 terminals and stored in a historical database. By using these statistics, the IEEE 802.22 system can benefit from accessing channels with lower probability of being occupied both to increase system performance and to limit interference to the wireless microphones.

The IEEE 802.22 standard has already defined an internal spectrum management framework [D6.1, Section 9.1.4], but it could be useful to extend this by using the ON/OFF model for wireless microphones activity to prioritize the channels considered in the spectrum manager as described in [D6.3, Section 4.4.3].

B.4 Short term user activity based model

A short term user activity model is complementing the long term On-Off model for the purpose of estimating idle periods of multiple users accessing a shared medium in a TDMA strategy. In particular this model considers a more realistic behaviour of the protocol stack and a proper distribution of access delay and access duration controlled by MAC and DLC sub-layers. Additionally, it considers application and human user behaviour as well as service specific response times of remote application servers to model the link resource utilisation of a shared spectrum user. A probabilistic model assuming the findings of [IEEE2003] on top of a deterministic medium access and packet flow model as shown in Figure B-10 is used to model the user access and utilisation of a shared medium.




Figure B-10: Media access timing assumed

Protocol processing: this time interval summarizes any terminal-internal packet delay caused by processing and regular handling/buffering of data packets.



Access delay: this time interval summarizes the packet delay caused by any delay in accessing the shared media. This delay includes the time interval a packet must be buffered until a suitable idle time becomes available. If the link will not become available within the maximum delay allowed, the packet is recognized as not sent but may cause a delay for subsequent packets by delaying higher layer protocol processing.

Uplink/Downlink utilization: if a packet is placed successfully to the shared media, this parameter denotes the time interval the packet utilizes the physical link. This time interval depends on the current physical link characteristics. If the lower layer model simulates packet retransmits or packet fragmentation, the total time interval which is reported to the higher layer protocol includes any trial retransmit or generation of multiple packets due to fragmentation.

Server response time: this time interval summarizes the delay between a pair of request-response packets on the physical link caused by the remote peer. This includes all delays caused by border elements such as wireless access points, base-stations or gateways and all infrastructure routers between terminal and remote server.

The model actually was designed for modelling and observing the behaviour of a single wireless link [ORACLE2008]. In this context it is applied to model the utilization of the shared resource (i.e. a shared spectrum frequency band). This generalization is valid as long as all competing spectrum users access and utilize a certain frequency band in whole such that the frequency band under consideration can be seen as a single broadcast medium. In all other cases (e.g. narrow and wide-band users concurrently using the same shared frequency band) each fragment of a frequency band must be considered a dedicated link.




Figure B-11: User model data traffic generator

An individual traffic generator is realized for each type of application associated with a terminal and a user (Figure B-11). Multiple applications and multiple users on a single terminal share the higher layer protocol stack entities. Different terminals share only lower layer protocols in order to allow for modelling a distributed MAC or collaborating terminals. The application model includes server-side parameters (for example, specifies the probability density of server response times) and human user behavioural parameters, such as the distribution of HTML content read times. Currently models for



FTP, HTTP and VoIP are available – WAP is a special case of HTTP and IPTV is a generalization of the VoIP application model. A sequence of application layer protocol packets (packet stream) generated here includes response packets from the remote peer if needed. This allows setting individual server parameters for each application model instance.

The observed data traffic (see Figure B-12) for a single application, terminal, or user shows a typical interleave of times of utilized and idling link. Short idle periods usually show up between single packets (or between a data packet and its ACK packet for some wireless protocols) while longer idle periods occur between sections of application activity (denoted here as a “session”). The definition of “session” depends on the scope of the underlying application generating that type of traffic and could be, for example, the download of some data file or the page read of a Web page including multiple text, image and audio objects.




Figure B-12: Terminology: link utilization and access opportunities

The model can be configured for multiple users, terminals and applications by setting individual parameters. In general those parameters are distribution functions for protocols and applications utilized, number and sizes of objects and embedded objects of complex rich media context, response times of servers, applications and human users and similar, setting the relative weight in a mix of activities on the modelled link. The model both considers statistical and QoS-based multiplexing of streams to a shared link (see Figure B-13).

The model described so far can be applied to the problem of estimating shared spectrum user activity and link utilization if some basic parameters are known for competing spectrum users such as data rate achievable at the PHY for a terminal utilizing shared spectrum, number of terminals concurrently using the same frequency band, an estimate for the number and type of application processes hosted on each terminal, and multiplexing scheme applied to the packet streams generated by those applications. The model allows reflecting a dynamic change of allocated spectrum since data packets generated by an application receive a temporal dimension at the very moment when they are put onto the shared medium.




Figure B-13: Multiplexing shared media access

Currently, the complexity of the model does not allow to estimate shared spectrum user behaviour in real-time by direct application of the model. Hence, it is more feasible to generate training sequences from the model, to observe spectrum utilization in real-time and to identify traffic patterns that allow categorizing the data traffic observed and to estimate model parameters or specific utilization parameters such as link idle times from these. Depending on the overall optimization purpose, it might be reasonable to estimate accumulated spectrum utilization (e.g. for interference calculation), congestion probability (e.g. through media access collisions by concurrent spectrum users), efficient resource utilisation (e.g. through estimating the probability of link idle times), or QoS prediction (e.g. by estimating the probability density of link idle times).



For example, a sliding window traffic analysis (Figure B-14, Figure B-15) has been used to predict upcoming idle times of the wireless link. The method tries to recognize typical traffic patterns that usually precede an idling link, and from categorizing traffic patterns and ‘guessing’ traffic parameters according to the model, tries to predict the upcoming idle time. Subsequent processing then provides a probabilistic model of idle times (i.e. inter-arrival time and duration of idle times) of the shared medium, enabling to predict the data rate and access delay to expect from the current utilization of the shared frequency band.




Figure B-14: Traffic analysis by observing spectrum utilization and identifying traffic patterns from observation.




Figure B-15: Simplified algorithm to estimate user activity from traffic patterns observed