Quality Of Service and MObility driven cognitive radio Systems

FP7-ICT-2009-4/248454

QoSMOS

D6.7

Integrated final functional specification of spectrum management framework and procedures

Contractual Date of Delivery to the CEC: 31/12/2012

Actual Date of Delivery to the CEC: 31/12/2012

Editor(s): János Bito, László Csurgai-Horváth, Péter Bakki (BME)

Internal Reviewers: Atilio Gameiro, Maziar Nekovee

External Reviewers: Uwe Beutnagel-Buchner, Alexandre Kholod, 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), Keith Briggs, Santosh Kawade (BT), Ulrico Celentano (UOULU)

Workpackage:	WP6
Est. person months:	10PM
Security:	PU
Nature:	P
Version:	1.0
Total number of pages:	109

Abstract:

This deliverable specifies the complete spectrum management framework including cognitive functions, security functions and test functions and elaborates on the capacity and limits of QoSMOS cognitive spectrum management.

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
AP	Access Point
BER	Bit Error Rate
BLER	BLock Error Rate
CCDF	Complementary Cumulative Distribution Function
CDF	Cumulative Distribution Function
CIR	Carrier-to-Interference Ratio
CM	Cognitive Manager
CM-RM	Cognitive Manager – Resource Manager
CM-SM	Cognitive Manager – Spectrum Manager
CN	Core Network
CNR	Carrier-to-Noise Ratio
CPFR	Common Portfolio Repository
CPOR	Common Policy Repository
CQI	Channel Quality Indicator
CSPC	Common Spectrum Control
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
IU	Incumbent User
LPFR	Local Portfolio Repository
LSPC	Local Spectrum Control
LTE	Long Term Evolution
MAC	Medium Access Control
MSC	Message Sequence Chart
MT	Mobile Terminal
NC	Network Cognition
OU	Opportunistic User

PDF	Probability Density Function
PHY	Physical Layer
PMF	Probability Mass Function
PMSE	Program Making and Special Events
QAM	Quadrature-Amplitude Modulation
QoS	Quality of Service
RAN	Radio Access Network
RAT	Radio Access Technology
RB	Resource Block
RM	Resource Management
RM	Resource Manager
RRM	Radio Resource Management
SAN	Spectrum Analyser
SINR	Signal-to-Interference plus Noise Ratio
SLA	Service Level Agreement
SM	Spectrum Management
SM	Spectrum Manager
SNR	Signal-to-Noise Ratio
SPI	Spectrum Efficiency Index
SPRR	Spectrum Provider Repository
SPRR	Spectrum Provider Repository
SS	Spectrum Sensing
SSE	Spectrum Selector
SSE	Spectrum Selector
TCP	Transmission Control Protocol
TDMA	Time Division Multiple Access
TTI	Transmission Time Interval
TVWS	TV White Space
UE	User Element
VoIP	Voice over IP
WAP	Wireless Application Protocol
WLAN	Wireless Local Area Network

Table of content

List of figures

List of tables

Executive Summary

This deliverable specifies the spectrum management framework including cognitive and opportunistic functions, and elaborates on the capacity and limits of QoSMOS cognitive spectrum management. It builds on deliverable [D6.5] that provides description of cognitive and opportunistic functions and self-learning capabilities of the spectrum management framework and on [D6.4] elaborating on trust, security, privacy, reliability and robustness of cognitive capacity. This document gives complete picture of the framework together with [D6.6] where implementation specific topics and evaluation of cognitive functions are discussed.

The Cognitive Manager – Spectrum Manager (CM – SM) reference model is derived and further elaborated from the specification given in deliverables [D2.1], [D2.2] and [D2.3], and it describes in detail the architecture, the procedures, the functional entities, their interactions and the interfaces connecting them inside CM-SM and to other QoSMOS entities outside. The document focuses on the QoSMOS scenarios as specified in [D1.2] and discuss scenario specific aspects of the framework.

The collaboration of the cognitive engines CM-SM and CM-RM (Cognitive Manager – Resource Manager) is one of the key aspects of spectrum management; therefore it is presented in detail here; however CM-RM specification is given in [D5.1], [D5.2] and [D5.3], where the description of the final structure of the cognitive manager for resource management is published.

1 Introduction

1.1 Scope and Objectives

This deliverable specifies the complete spectrum management framework including cognitive functions, security functions and test functions and elaborates on the capacity and limits of QoSMOS cognitive spectrum management.

QoSMOS WP6 gives here the final specification of the spectrum management framework and provides solid foundation for implementation and proof of concept demonstrations.

1.2 Organization of the document

The document is organized into chapters according to the functional decomposition of the framework that is presented in chapter 2 in detail. CM-SM reference model describes the functions, interfaces and the basics of their communication. A CM-SM entity is built of several concurrent threads implementing state machines. The internal states and procedures are discussed in chapter 2.1. Communication inside a CM-SM entity is specified by chapter 2.2 in a layered structure,

Chapter 3 specifies the collaboration between the two QoSMOS cognitive managers CM-SM and CM-RM that is one of the key aspects of spectrum management operation. The procedures and interfaces are specified in detail and mapped to architecture defined in [D2.2]. Groups of procedures are formed for different areas of spectrum management tasks related to CM-RM. The two cognitive managers are connected via direct communication channel (accessed over CM1 interface), that enables effective communication and clear separation of the functions of cognitive managers using the possibility of modular architecture.

Chapter 4 presents specification of repositories where spectrum management related information is collected, stored, and the aggregated context is accessed by CM-SM functions. CM-SM reasoning and decisions are based on the information retrieved from these repositories, that are designed to form a hierarchical and function oriented structure. The coexistence domain repositories are built on different databases and their services could be accessed via adaptation layer (AL) interfaces from coordination domain. Databases are organized according regulatory, geographical and functional areas; they store and give access to information needed for proper operation of coordination domain entities.

The following chapters (5 and 6) give the specification of the coordination and networking domain spectrum control operation and detailed analysis of various methods implemented in common (CSPC) and local (LSPC) spectrum control entities. The CSPC is implemented for wider geographical area served by an operator and it supports communication between coexistence and networking domain entities and provides centralized spectrum control functions; while LSPC is responsible for the spectrum control operation of a localized part of the operator’s RAT network. Ad-hoc and femtocell scenarios are discussed separately because they might be independent of operators, therefore different problems may arise, so different solutions are favourable.

Chapter 7 introduce optional elements of the spectrum management framework that could be realized in networking domain as functions of CM-SM END entity. Spectrum analyser (SAN) and spectrum selector are communicating directly with a local spectrum control (LSPC) or a local portfolio repository (LPFR) entity and they are responsible for quick deployment of spectrum portfolio to CM-RM and spectrum measurement data composition to spectrum portfolio data structure. Performance evaluation of spectrum selector (SSE) functions is performed for an environment where wireless microphones (WM) are present.

2 CM-SM Reference Model

In a nutshell, the role of the cognitive manager for spectrum management (CM-SM) [D6.2] is to manage (generate, deploy, maintain and revoke) spectrum portfolios and to provide them to the cognitive manager for resource management (CM-RM) [CelEtal2011]. In order to acquire the radio context, the CM-SM also interacts with the spectrum sensing (SS), as shown in Figure 2-1. Communication among remote functional blocks (except for the CM-RM/CM-SM interworking) happens through the adaptation layer (AL).

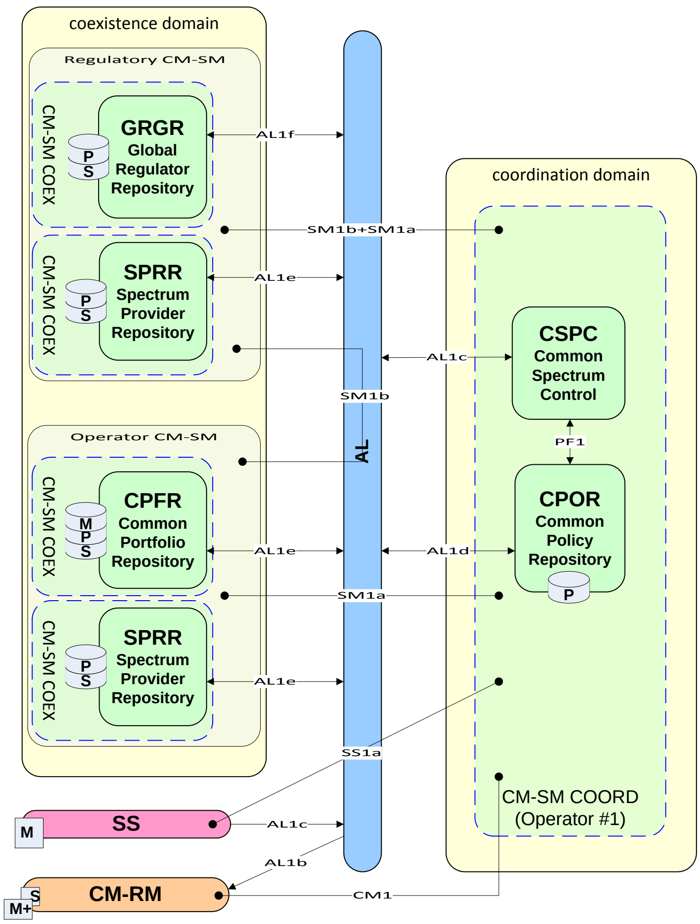

Figure 2-1: The CM-SM interacting with the other functional blocks and external repositories.

The global regulatory repository (GRGR) provides all the constraints and supporting information as required by the regulations. These, opportunely filtered and processed, are integrated by the CM-SM with radio context and the performance metrics received from the CM-RM to contribute to the common portfolio repository (CPFR). The information exchanged through the spectrum provider repository (SPRR) is exploited for spectrum trading and is also used for spectrum portfolio generation.

The QoSMOS architecture is designed to be flexibly adapted to diverse scenarios [D2.4]. In order to ease the design reuse, the internal functional blocks are assigned to topological domains [CelEtal2011] instead of rigidly associate them with network nodes. The split according to topological domains is shown in Figure 2-2.

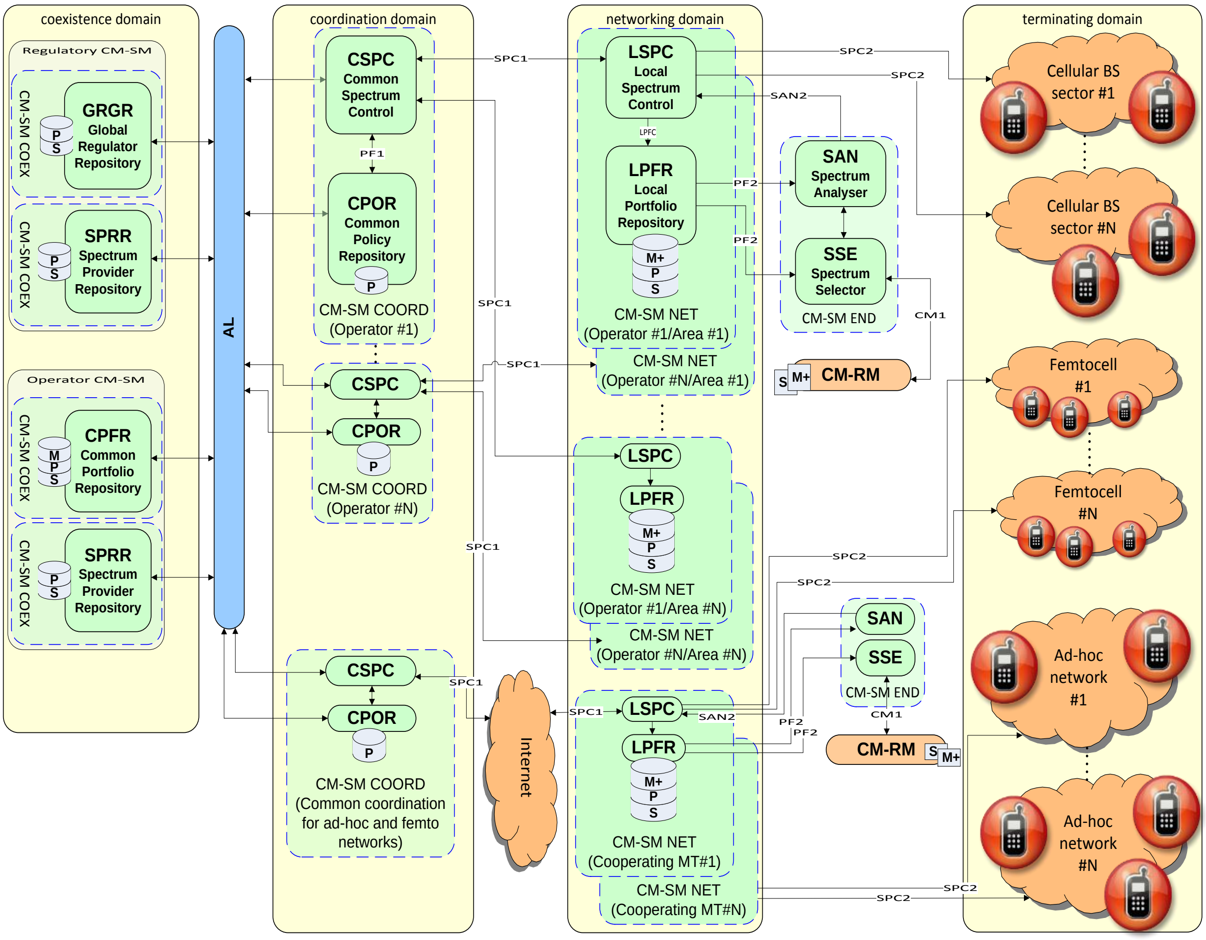

Figure 2-2: QoSMOS reference model of the CM-SM

The functions, which will be described shortly below, are divided into three groups. The repository access (CM-SM COORD) resides at the coordination domain (e.g., the core network in the cellular case) and accesses the repositories at the coexistence domain (e.g., the Internet). The local spectrum management (CM-SM NET) and the spectrum selection (CM-SM END) are placed at the network domain (e.g., at the base station in the above cellular example). The CM-SM NET is in charge of eliminating information not relevant locally and adding instead what peculiar at the networking domain, whereas the CM-SM END works in close interaction with the CM-RM.

The common spectrum control (CSPC) is in charge of accessing the repositories it manages and therefore of merging/splitting the relevant information. The common policy repository (CPOR) controls all spectrum usage constrains not covered at regulatory or spectrum provider side and that are needed at the radio access network in question. The CSPC also interfaces, through the AL, possible remote CM-SM's (or equivalent and compatible entities) in neighbouring networks.

The local spectrum control (LSPC) manages the local spectrum and is therefore accessing the local portfolio repository (LPFR), which, in short, caches the spectrum portfolio information and may trigger actions at the spectrum selector (SSE). Measurements from the CM-RM are gathered by the spectrum analyser CM-SM END and used at LPFR for local decision and sent further up to the CM-SM COORD through the LSPC for exploitation at a broader level (coordination and coexistence domain). The LSPC has the duty of all (networking-coordination domain) interactions towards the CM-SM COORD.

2.1 CM-SM state machines

A CM-SM entity implementation consists of multiple concurrent threads each implementing state machines responsible for handling a number of clients in parallel. The discussion of internal states of a CM-SM in the following refers to the totality of these. A particular thread may not implement all of the details as presented here but only those of relevance for a particular role and its current task in scope of that role. The notion of a role is introduced temporarily in the scope of this section to emphasize that a certain implementation of a CM-SM may focus on a certain set of capabilities. For example, a CM-SM may provide spectrum towards a CM-RM but may in parallel depend on another CM-SM instance for obtaining spectrum it needs to serve as a spectrum provider. A CM-SM instance that does not request spectrum from other CM-SM entities but only provides spectrum is said here to take the role of a spectrum provider (`SP`) while a CM-SM both requesting and providing spectrum is said to act as a spectrum manager (`SM`). A main difference between these is that a spectrum manager needs to enable both the portfolio release and portfolio revoke states, while a spectrum provider needs to enable only the portfolio release state since it never requests spectrum portfolios from another CM-SM instance.

Throughout the discussion of CM-SM states the following roles are assumed (see [D6.5] for naming details):

- `SM` – the spectrum manager. This role implements the full functionality of the CM-SM by implementing all states discussed below in a single entity. It can deploy and revoke spectrum portfolios, and it can request a spectrum portfolio from a spectrum provider such as the `SP` or `GDB`. A CM-SM of the networking or coordination domains is acting as an `SM`, for example.

- `SP` – a spectrum provider. This role implements the capacity to deploy and revoke spectrum portfolios. In contrast to the `SM` it has no `RM` clients but only `SM` clients. That is, it cannot request spectrum portfolios from another CM-SM and it cannot receive spectrum portfolios from a CM-RM potentially created through a spectrum sensing process, such as in ad-hoc scenarios. A CM-SM of the coexistence domain is acting as an `SP`, for example. In consequence, the communication between `SM` and `SP` is realized by the SM1a interface

- `GDB` – a Geolocation database realized by a CM-SM instance and at least one instance of an authority's Geolocation database. This role implements the capacity to respond to a portfolio request by an `SM`. It provides spectrum portfolios representing underutilized spectrum in a certain geographical area. The `GDB` role is realized for example in the Regulatory CM-SM of the coexistence domain that encapsulates an authority's database. In general a `GDB` is realized by a CM-SM instance acting as an `SP` and relying in this on a proprietary interface to an authority's Geolocation database. It translates between the Geolocation database's description of a spectrum opportunity and the QoSMOS spectrum portfolio based description. In that, the `GDB` may also extend the functionality of the authority's Geolocation database or may coordinate between distinct instances of an authority's Geolocation database. The communication between `GDB` and `SM` is realized by the SM1a/b interfaces. Note that a Geolocation database client may obtain a time-limited lease for a portfolio, which implies a portfolio revocation when the lease expires.

- `RM` – a resource manager. This role implements the capacity to request and release portfolios. It may provide spectrum portfolios generated from local sensing. Note that this is not an `SP` role since the `SM` involved here is not considering that spectrum as a resource provided but as context information. Although the spectrum portfolio interface data structure generated by an `RM` from sensing information may provide sufficient information to utilize the spectrum opportunity described, a CM-SM entity in its `SM` role must be involved in the process to provide, refine or certify spectrum utilization policies prior to deploying that spectrum portfolio. A realization of an RM is in the CM-RM of the networking domain. The communication between RM and SM is realized by the CM1, SPC1 and SPC2 interfaces.

- SS – a spectrum sensing sub-system. This role implements the capacity to provide spectrum portfolios generated from local sensing. A realization of an SS is in the SS entities of the networking domain reporting to a CM-SM or to a CM-RM. Since the SS role must include functions to create spectrum portfolios, this functionality might require involvement of a CM-SM or CM-RM providing that functionality. The communication between SS and SM is realized by the CM1 and SS1b interfaces and may involve the RM as a proxy.

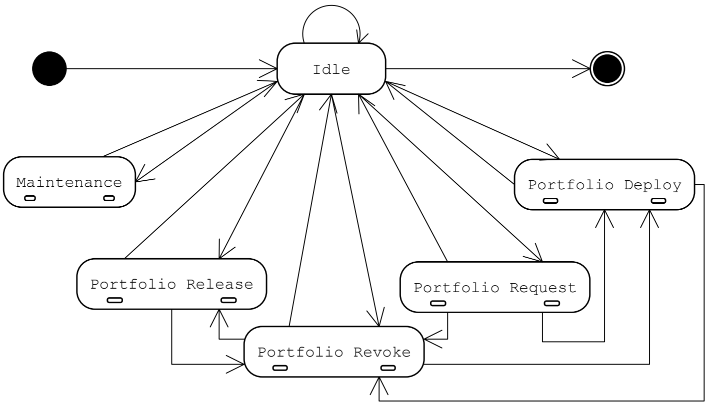

State	Valid for Role
Idle	SM,SP,GDB,RM,SS
Maintenance	SM,SP,GDB,RM,SS
Portfolio Release	SM,SP,GDB
Portfolio Revoke	SM,RM
Portfolio Request	SM,SP,GDB,SS
Portfolio Deploy	SM,RM

Figure 2-3: CM-SM top-level state diagram

Figure 2-3 depicts the top-level state diagram of the CM-SM. All states or sub-states shown can be realized as distinct concurrent processes or threads in an implementation. State transitions thus coincide with threat creation or termination. In the following CM-SM states are detailed further.

Idle state – Upon start-up a CM-SM is entering the Idle state. In this state it responds to requests from clients or it reacts on timed events. In consequence it enters one of the Maintenance, Portfolio Release, Portfolio Revoke, Portfolio Request or Portfolio Deploy states upon reception of a corresponding request from a client. It may remain in the Idle state in case no pending actions from earlier requests remain or in case of certain error conditions. Upon shut-down of the CM-SM, the Maintenance state is first reached for issuing requests to revoke or release portfolios earlier deployed or obtained, respectively, prior to the finally terminating the CM-SM.

Maintenance state – The Maintenance state is entered upon request of an internal timer event, or during start-up and shutdown procedures. In the Maintenance state the CM-SM is performing sanity checks on its internal threads and data structures, is checking current status and authorization of its clients and is executing internal performance monitoring tasks. A major task credited to the Maintenance state is in detecting terminated or disconnected clients and to initiate the release of spectrum portfolios deployed earlier to such client (potentially revoking portfolios from clients if possible) utilizing the functions of CM-SMs internal management interface. The Maintenance state is left before the next internal timer event. Tasks initiated while in Maintenance state are processed in parallel (by a dedicated thread) or by re-entering the Idle state and then another suitable state depending on the task scheduled.

Portfolio Release state – The Portfolio Release state is entered when a CM-SM (in its SM role) receives a request to release a spectrum portfolio from local use or from client use (i.e. RM or GDB) thus freeing up resources associated with this portfolio. This may involve communication with another CM-SM (in its SP role) when the SM is requesting to release portfolios in turn. The

Portfolio Release state is left after freeing up the portfolios referenced and confirming to the RM or GDB if required. Then Idle state is reached. The Portfolio Release state is entered also in consequence of a task initiated in Maintenance state, potentially subsequent to leaving the Portfolio Revoke state, which in general happens during regular CM-SM shutdown.

Portfolio Revoke state – The Portfolio Revoke state is entered by an SM in consequence of the request of a CM-SM in its SP role, or by an RM in consequence of the request of a CM-SM in its SM role. The Portfolio Revoke state may be included as a sub-state in several procedures such as programmed portfolio updates (initiated by the SM or by the RM). The Portfolio Revoke state is left after freeing up the portfolios referenced and confirming to the SM or to the SP. The Portfolio Revoke state also may be entered as a sub-state of the Portfolio Release state in case a programmed decision of the CM-SM demands to free up certain spectral resources during regular operations.

Portfolio Request state – The Portfolio Request state is entered by an SP upon request of an SM, by an SM upon request of an RM, or by a GDB upon request of an SM. The Portfolio Request state is left after deploying the requested portfolio to the requesting client, which includes issuing a request to the client (that in turn reaches the Portfolio Deploy state) before re-entering the Idle state. The Portfolio Request state in general is not a sub-state to any other state since it is initiated only by a protocol-driven request. The Idle state is reached from the Portfolio Request state immediately only if an error occurs and no portfolio can be deployed.

Portfolio Deploy state – The Portfolio Deploy state is entered by an SM or by an RM in consequence of receiving a portfolio from an SP, SM, SS, or GDB respectively in consequence of an earlier request or revoke procedure. The Portfolio Deploy state may be reached also upon a local programmed decision when populating a local spectrum portfolio repository. The Portfolio Deploy state thus mostly appears as a sub-state of other states. The Portfolio Deploy state is left after issuing a confirmation to the portfolio provider. The Idle state is reached from a Portfolio Deploy state only in the most basic procedures, usually used during initialization of a client when the client is populated with a default set of portfolios.

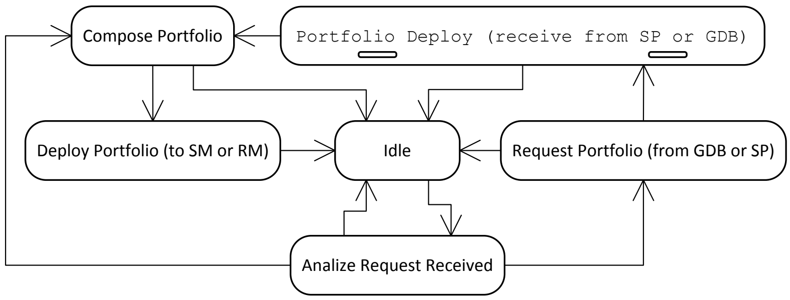

Figure 2-4: Portfolio Request state diagram

The Portfolio Request state is reached from the Idle state upon reception of a client requesting a spectrum portfolio. A first check is done on the request, to identify if this request can be satisfied from the local spectrum portfolio repository, or if additional spectrum must be acquired to compose a portfolio for deployment towards the requesting client. This is outlined by Figure 2-4. The Portfolio Deploy state is entered in order to wait for a portfolio request response and for integrating a portfolio received into the local portfolio repository.

The Idle state is reached upon all error conditions (i.e. request failures, portfolio compose failures or upon a portfolio request failure response).

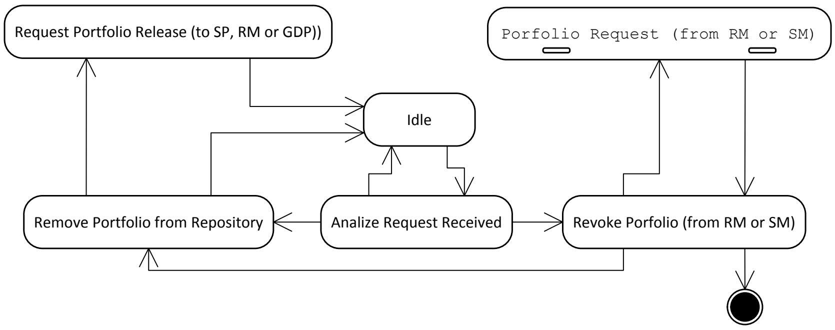

Figure 2-5: Portfolio Release state diagram

Figure 2-5 further details the `Portfolio Release` state. It is entered upon receiving a request from a client that wants to free up spectrum resources allocated earlier. If portfolios to be released are only of local relevance, the request can be satisfied by removing the portfolios from the local repository of the CM-SM instance. Local relevance is given if a portfolio has not been obtained earlier from an `SP`, `SM` or `GDB` (see role descriptions above) or has not been deployed towards an `SM` or `RM` (or is part of a deployed portfolio).
If a portfolio has been obtained earlier, a release request may be issued towards the portfolio provider, or the portfolio may be retained for further use until revoked by the provider.

If a portfolio has been deployed earlier (or the portfolio released is part of a portfolio deployed earlier), a revocation request needs to be issued towards the client. Such revocation request may involve a portfolio update procedure, replacing the portfolio revoked by a different portfolio prior to removing to ensure seamless spectrum handover, or it may cause a portfolio request by the client detecting that it has insufficient spectrum resources to continue operation.

Note that under certain conditions a revocation may fail leaving the CM-SM in a deadlock where it cannot release spectrum any more. This may be an unrecoverable exception similar to a 'panic' state of the CM-SM where automated recovery procedures likely will fail. In such cases the only option may be to terminate CM-SM operation. In practice, the CM-SM then first enters `Maintenance` state to determine potential next steps.

The `Idle` state is reached upon all error conditions except that it must not be reached through sub-states `Portfolio Deploy` and `Portfolio Request` since any error handling must be included with the previous sub-state attained before entering these states (see the discussion of the `Portfolio Request` state above).

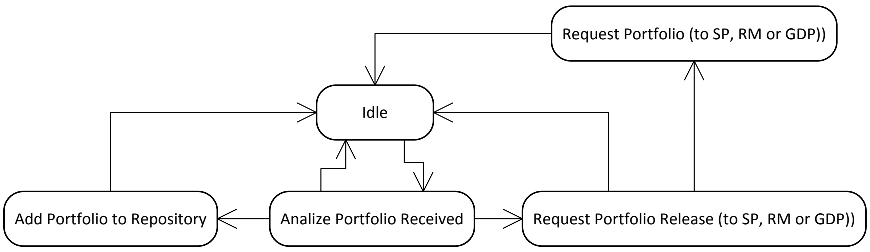

Figure 2-6: Portfolio Deploy state diagram

The `Portfolio Deploy` state is outlined by Figure 2-6. It is reached from the `Idle` state upon reception of a deployable portfolio from an `SP`, `SM`, or `GDB` possibly in response to an earlier portfolio request. In general, the portfolio received is added to a local repository and subsequently utilized as a spectrum resource.

It may happen that the portfolio received does not match the requested portfolio characteristics (e.g. if the portfolio provider was short on spectrum resources) or does not satisfy the current demand (e.g. because the demand has changed between issuing a request and receiving the portfolio). In that case the portfolio may be released immediately and the request is re-issued (potentially with parameters having changed), or an additional complementing portfolio may be requested. The `Idle` state is reached next including upon all error conditions.

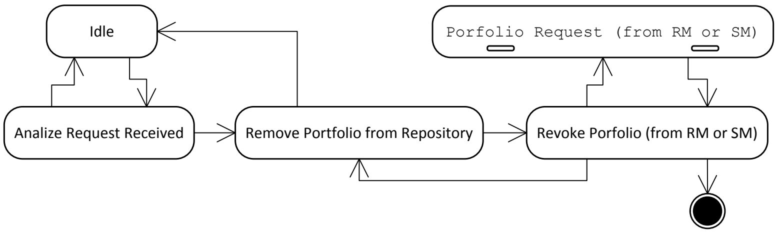

Figure 2-7: Portfolio Revoke state diagram

The `Portfolio Revoke` state as detailed by Figure 2-7 is reached upon reception of a portfolio revocation request from an `SM` or `SP`. Since portfolios requested to be removed may have been deployed earlier to clients of this CM-SM instance, a subsequent portfolio revocation request may be needed in turn, potentially causing subsequent portfolio requests from those clients (see the discussion of the `Portfolio Release` state above) before reaching the `Idle` state again. The `Idle` state is reached upon all error conditions except that it must not be reached through sub-states `Portfolio Deploy` and `Portfolio Request` since any error handling must be included with the previous sub-state attained before entering these states (see the discussion of the `Portfolio Request` state above).

Note that under certain conditions a revocation may fail leaving the CM-SM in a deadlock where it cannot release spectrum any more. In such cases the only option may be to terminate CM-SM operation. In practice, the CM-SM then enters `Maintenance` state (see discussion of the `Portfolio Release` state above).

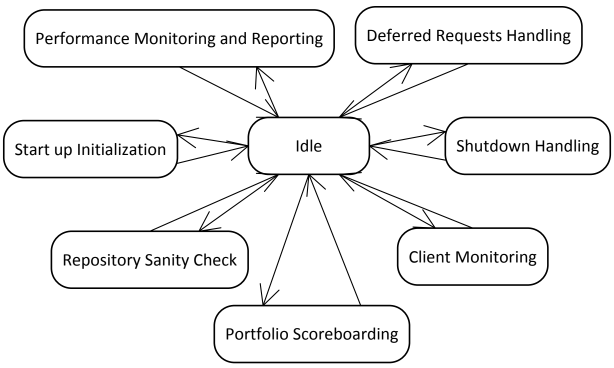

Figure 2-8: Maintenance state diagram

Details of the `Maintenance` state as given by Figure 2-8 have been already discussed in the scope of other CM-SM states. Most of the procedures executed in that state are considered management and administration processes that increase robustness and ensure resilience of the CM-SM. These are of no great relevance for the discussions made throughout this document and may change in future realizations.

In addition, the `Maintenance` state includes some performance optimization procedures that may be invoked depending on the state of distinct internal performance counters. As an example Figure 2-8 references a state denoted as `Portfolio Scoreboarding`. Under this term a number of methods are summarized that evaluate the utilization of spectrum portfolios in the local repository according to certain performance metrics such as “use count”, “area coverage”, “utilization history”, “interference vulnerability” and similar. Those metrics are continuously updated and are utilized to decrease the access time to the portfolio repository when selecting one out of many suitable portfolios for deployment in response to a client request.

2.2 Overview of CM-SM Service Primitives

Subsequent sections describe the service primitives based on the model shown in Figure 2-9. This layering allows separating adaptation to a communication subsystem (denoted as ‘interface primitives’, enabling access to the subsystem assumed to implement service primitives defined by the QoSMOS project) and adaptation to an application (denoted as ‘application layer primitives, i.e. enabling access of a CM-SM or CM-RM instance to a communication sub-system). The main benefit of such layering is in gaining flexibility regarding different (e.g. standardized or proprietary) communication sub-system implementations, and in increasing reusability of code across different application entities. In the terminology used reference is made to [D6.5].

The increase in flexibility is created from using a generic ‘setter/getter’ paradigm that allows choosing different implementations for the communication sub-system. CORBA, REST, XML-RPC, ETSI M2M are some of those that comply with the paradigm and have been evaluated for specific configurations and specific interfaces of the QoSMOS framework.

The example given in Figure 2-9 assumes a communication sub-system as defined through QoSMOS message sequences, protocol primitives and service access points (SAPs) implementing the protocol primitives as defined in [D2.3]. For example, a CM-SM instance calling the *revoke_object_request(PortfolioId)* application layer primitive is mapped to a call to *set_request(PortfolioId, Null)* protocol layer primitive, consequently generating an *SM1_portfolio_revoke.req(PortfolioId)* protocol message as defined for the QoSMOS communication sub-system. In contrast, a similar call to *set_request(PortfolioId, PortfolioObject)* would create an *SM1_ portfolio_deploy.req(PortfolioId, PortfolioObject)* message. The SAP instance used determines which QoSMOS interface is used. Hence, calling the identical primitive for a SAP towards a CM-RM would generate a *CM1_portfolio_deploy.req(PortfolioId, PortfolioObject)* message instead.

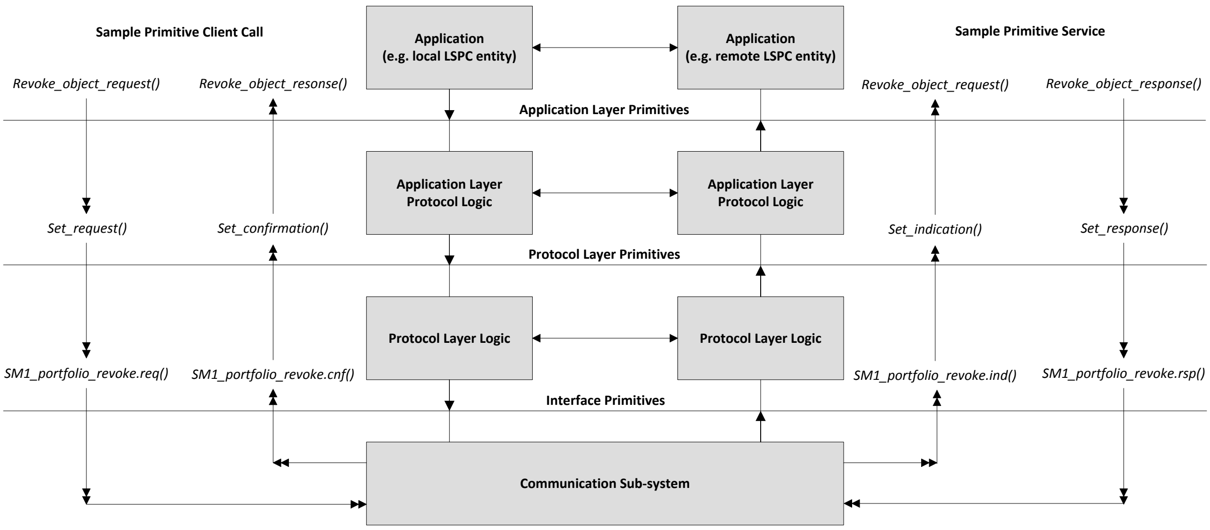

Figure 2-9: Layering of service primitives

The gain in reusability results from the option to streamline the application layer protocol logic to the needs of a specific QoSMOS entity. For example, CSPC and LSPC demand for the most complex implementation since these entities both have to implement most of the primitives and both client and server side functionality. All other entity-specific implementations can be created by omitting functionality of a CSPC and LSPC application layer protocol logic implementation and adding minor changes based on these for functions unique for a certain QoSMOS CM-SM functional entity in particular.

2.2.1 Application Layer Primitives

On the application layer a number of primitives are provided to interact with remote entities. For the current implementation the application layer SAP is common for all entities and roles while the implementation of the primitives is specific for a distinct role. The primitives given by the following table reflect the current status of the implementation. It does not reflect all suitable options as described by deliverables [D6.1] through [D6.5] but only those required to realize a fully functional framework. Subsequent implementations will likely not introduce new primitives. Functional enhancements in future implementations are expected to require only changes to the parameter lists of existing primitives.

Application layer primitives implement access to the end-to-end application layer protocol. In general, an application layer primitive utilizes one or more protocol layer primitives to realize certain functionality. Application layer primitives are realized at the Service Access Point (SAP) denoting the interface between application logic and underlying protocol support.

The term ‘object’ is used throughout subsequent tables and paragraphs as a generic term and can be read as either ‘deployable spectrum portfolio’, ‘empty (null) portfolio’, ‘spectrum portfolio model’, ‘context information data set’, or ‘control or configuration data set’ depending on the context of the call to the primitive and the capacities and role of entities exchanging information.

Table 2-1: Summary of Application Layer Primitives

Application Layer Primitives		
Primitive Name	**Primitive Description**	**Implemented by Entity**
Get_object_request	*Called by the local application to request a spectrum portfolio or context information from a service provider. The object requested might be referenced by a globally unique id if available, or by a description or model if referencing one of a kind.* *In general, a spectrum user requests a spectrum portfolio from a spectrum provider entity such as a spectrum manager or a Geolocation Database by calling this primitive.*	*CM-RM* *CM-SM (LSPC, CSPC, SAN, SSE)*
Get_object_response	*Called by a remote application in response to a get_object_request. Indicates success and returns the object requested if any. Also returning the globally unique identifier of the object returned.*	
Get_object_failure	*Called by a remote application in response to a get_object_request. Indicates failure of the request and returns the reason code and the reference to the object, which either is the globally unique identifier of the object requested or the description received with the get_object_request.*	
Set_object_request	*Called by the local application to request setting the value or one or more attributes of an object according to the parameters of the call. The object is referenced by a description when referencing all of a kind or by a globally unique identifier if addressing a distinct object.* *In general, a spectrum provider is calling this primitive to provide a spectrum portfolio and to request a spectrum user to consider the spectrum portfolio provided for subsequent use.*	*CM-SM (SPRR, CPOR, LSPC, CSPC, SAN)*
Set_object_response	*Called by a remote application in response to a set_object_request. Indicates success of the request and returns a globally unique identifier of the object referenced.*	

<table>
<tr><th colspan="3">Application Layer Primitives</th></tr>
<tr><th>Primitive Name</th><th>Primitive Description</th><th>Implemented by Entity</th></tr>
<tr><td>Set_object_failure</td><td>Called by a remote application in response to a set_object_request. Indicates failure of the request and returns the reason code and the reference to the object, which either is the globally unique identifier or the description received with the set_object_request.</td><td></td></tr>
<tr><td>Revoke_object_request</td><td>Called by a local application to revoke an object from use. The object must be referenced by its globally unique identifier.
This primitive in general is issued by an entity owning a spectrum portfolio towards an entity using the spectrum and unambiguously identifies the object to revoke from use.
Note – a pair of set_object_request and revoke_object_request is needed to replace a spectrum portfolio at a spectrum user.</td><td rowspan="2">CM-SM (GRGR, CPFR, SPRR, CPOR, LSPC, CSPC, SAN)</td></tr>
<tr><td>Revoke_object_response</td><td>Called by a remote application in response to a revoke_object_request. Indicates success of the request and returns the globally unique identifier for the object referenced. May indicate failure if the object referenced does not exist in the scope of the remote application.
In general, this primitive applies to spectrum portfolios.
Note – the main effect of revoke_object_response primitive is that the object is removed from the scope of the remote application and is returned to the owner. Hence, there is no further failure response since this could not be resolved by protocol means.</td></tr>
<tr><td>Release_object_request</td><td>This primitive is called by a local application to indicate that it is no longer using the object referenced by the globally unique identifier provided with the primitive.
Note – the main effect of this primitive is that the object is removed from the scope of the local application and is returned to the owner.</td><td rowspan="2">CM-RM
CM-SM (CFPR, SPRR, LSPC, CSPC, SSE)</td></tr>
<tr><td>Release_object_response</td><td>Called by a remote application in response to a release_object_request. Acknowledges the request, indicates success of the primitive or may indicate failure if the object referenced does not exist in the scope of the remote application.
Note – the main effect of release_object_request primitive is that the object is removed from the scope of the local application and is returned to the owner. Hence, there is no further failure response.</td></tr>
<tr><td>Register_notify_request</td><td>Called by a local application to request receiving indication from a remote application whenever an object referenced by the globally unique id provided along with the primitive is changing in the scope of the remote application in whole or regarding the attributes referenced along with the primitive.
In general, reasons for object changes might be internal to the remote application (e.g. as a result of an optimization process) or might be external to the remote application (e.g. in consequence of a context information change).</td><td rowspan="3">CM-RM
CM-SM (CPFR, CPOR, LSPC, CSPC, SAN, SSE)</td></tr>
<tr><td>Register_notify_response</td><td>Called by a remote application in response to a register_notify_request. Acknowledges the request, indicates success of the primitive or may indicate failure if the object referenced does not exist in the scope of the remote application. In addition, the register_notify_response primitive returns the globally unique identifier of the object referenced and the lease time.
Note – renewing the request through subsequent calls to register_notify_request might extend the lease time.</td></tr>
<tr><td>Notify_object_request</td><td>The notify_object_request primitive is called by a remote application whenever a change of an object or some of its attributes is detected by the remote application and an earlier register_notify_request issued by the local application indicates the willingness of the local application to receive notifications about changes of that object. The primitive provides the globally unique identifier of the object and the portions of the object data that require update in the scope of the local application.</td></tr>
</table>

2.2.2 Protocol Layer Primitives

On the protocol layer a number of primitives are provided to implement the application layer (end-to-end) protocol and to realize the application layer primitives as detailed in section 2.2.1. These primitives are implemented through message exchange mechanisms as documented in deliverables [D2.2], [D2.3] and [D6.5].

Protocol layer primitives implement access to the point-to-point communication protocol. In general, a protocol layer primitive utilizes one or more communication layer primitives to realize certain functionality. Protocol layer primitives are realized at the Service Access Point (SAP) denoting the interface between an end-to-end protocol logic instance and an underlying communication sub-system.

Note that some primitives are directed. That is, in a peer-to-peer communication only one peer needs to implement the *get_request* and *get_confirmation* primitives, while the other has to implement the *get_indication* and *get_response* primitives, for example.

Table 2-2: Summary of Protocol Layer Primitives

<table>
<tr><th colspan="3">Protocol Layer Service Primitives</th></tr>
<tr><th>Primitive Name</th><th>Primitive Description</th><th>Applies to Interface</th></tr>
<tr><td>Get_request</td><td>Requests to get a specified object from a remote entity when called. The remote entity is implicitly addressed by the SAP instance used in this call.The primitive must be called with a reference to the object requested. If successful, the primitive returns a transaction identifier to the caller which is unique in the scope of the SAP used.</td><td rowspan="4">Bi-directional:
CM1, SM1, SPC1, PF2

Unidirectional:
PF1, LPFC, SS1, SAN2, QS1, SPC2</td></tr>
<tr><td>Get_indication</td><td>If received by the remote entity, indicates that a request to get a specified object has been issued. The transaction identifier accompanying this primitive must be used in subsequent replies to this indication. The parameters of the indication are identical with that used by the corresponding get_request.</td></tr>
<tr><td>Get_response</td><td>If called by the remote entity, generates a response message for a corresponding get_indication. The primitive must be called with the transaction identifier received along with the corresponding get_indication and a return result code, and may be called with an object or object reference if applicable.</td></tr>
<tr><td>Get_confirmation</td><td>If received by the local entity, indicates that a request to get a specified object has been acknowledged. The transaction identifier accompanying this primitive is the same as returned by the call to the corresponding get_request. The parameters of the confirmation are identical with that used by the corresponding get_response.</td></tr>
<tr><td>Set_request</td><td>Requests to set a specified object at a remote entity when called. The remote entity is implicitly addressed by the SAP instance used in this call. The primitive must be called with parameters providing the object and a globally unique object reference. If successful, the primitive returns a transaction identifier to the caller which is unique in the scope of the SAP used.</td><td rowspan="4">Bi-directional
SAN2

Unidirectional
CM1, SM1, PF1, PF2, SPC1, LPFC, SS1, QS1, SPC2</td></tr>
<tr><td>Set_indication</td><td>If received by the remote entity, indicates that a request to set a specified object has been issued. The transaction identifier accompanying this primitive must be used in subsequent replies to this indication. The parameters of the indication are identical with that used by the corresponding set_request.</td></tr>
<tr><td>Set_response</td><td>If called by the remote entity, generates a response message for a corresponding set_indication. The primitive must be called with the transaction identifier received along with the corresponding set_indication and a return result code, and may be called with an object reference that in general matches the one used in the corresponding set_request.</td></tr>
<tr><td>Set_confirmation</td><td>If received by the local entity, indicates that a request to set a specified object has been acknowledged. The transaction identifier accompanying this primitive is the same as returned by the call to the corresponding set_request. The parameters of the confirmation are identical with that used by the corresponding set_response.</td></tr>
<tr><td>Register_request</td><td>Requests to register a local entity as a listener for remote changes of a specified object. The remote entity is identified implicitly through the SAP instance. The local entity must provide a globally unique reference to the remote object it intends to observe. It may provide data that narrows the set of events causing a change indication later on. If successful, the primitive</td><td>Bi-directional
CM1, SM1, SPC1, PF2</td></tr>
</table>

Protocol Layer Service Primitives		
Primitive Name	**Primitive Description**	**Applies to Interface**
	returns a transaction identifier to the caller which is unique in the scope of the SAP used.	*Unidirectional PF1, LPFC, SS1, SAN2, QS1, SPC2*
Register_indication	*If received by the remote entity, indicates that a request to register as an observer for an object in the scope of the remote entity has been issued. The transaction identifier accompanying this primitive must be used in subsequent replies to this indication. The parameters of the indication are identical with that used by the corresponding register_request.*	
Register_response	*If called by the remote entity, generates a response message for a corresponding register_indication. The primitive must be called with the transaction identifier received along with the corresponding register_indication and a return result code, and may be called with an object reference that in general matches the one used in the corresponding register_request. In addition, the remote entity may respond by narrowing or relaxing the conditions applicable for specifying the set of events causing a change indication later on. Furthermore the remote entity may specify a lease time that limits the duration of the request to a suitable value.*	
Register_confirmation	*If received by the local entity, indicates that a request to set a specified object has been acknowledged. The transaction identifier accompanying this primitive is the same as returned by the call to the corresponding set_request. The parameters of the confirmation are identical with that used by the corresponding set_response.*	
Notification_request	*Requests to notify a remote entity registered as a listener for changes of a specified object. The remote entity is identified implicitly through the SAP instance. The local entity provides a globally unique reference to the local object. It may provide an updated object or portions thereof. The primitive is called in the scope of the parameters set by the corresponding register_request. It does not generate any acknowledgment or confirmation from the remote entity and has no transaction identifier. Hence, the notification_request primitive always succeeds.*	
Notification_indication	*Indicates a change of an object when received by the remote entity. The parameters of the indication are identical with that used by the corresponding notification_request.*	
Error_indication	*Asynchronous error notification. Received by the local protocol implementation if an error condition was encountered while processing a primitive. This primitive is issued for all error conditions not related to a particular transaction.*	*All*

2.2.3 Primitive Parameters

The following table summarizes the primitive parameters used both in a call to an application layer primitive and to a protocol layer primitive. The following tables refer to the TLV encoded information elements. In-core representations may differ from that and depend on the programming language used.

Table 2-3: Summary of Parameter Types

Symbolic Type	Description
T_OCTET	*A primitive data type having 8 bits of information that can define up to 256 distinct states.*
T_OCTETSTRING	*A simple data type storing a sequence of octets. A string is a special use of a one-dimensional array or vector. An octet string may store information of type character but, in contrast to a character string, is not affected by character conversions.*
T_CHAR8	*A primitive data type suitable to represent a grapheme (i.e. the smallest semantically distinguishing unit in a written language). Note that common binary representations impose length restrictions. In this specification, the length of a character parameter is defined as 8 bits (e.g. [ISO8859-1]).*
T_STRING8	*A simple data type storing a sequence of characters. A string is a special use of a one-dimensional variable length vector.*
T_XMLSTRING	*A simple data type storing a sequence of octets representing an XML formatted document.* *The symbolic type of a XML string parameter is T_XMLSTRING. This specification requires implementing T_XMLSTRING as an octet string parameter T_OCTETSTRING. All implementations shall reference xml formatted strings only via T_XMLSTRING.*

Symbolic Type	Description
T_INT16, T_INT32	*A primitive integral data type representing natural numbers and their negatives. Note that common binary representations limit the number range due to machine word length restrictions. In this specification, the integer length is defined as 16 or 32 bits.*
T_FLOAT32	*A primitive data type storing real numbers, usually as floating-point numbers. Floating-point number representations as defined in IEEE Std 754™-2008 [IEEE754] can be taken as an example. In this specification, the length of a float parameter is defined as 32 bits.*
T_DSM_RESULT_CODE	*Enumeration of return result codes. An enumeration is a listing of elements of a set in a way that maps to an index set consisting of natural numbers. That is, each element of the set is unambiguously represented by an ordinal. The base type of T_DSM_RESULT_CODE is T_INT32.*
T_GUID	*The GUID information element is a structure consisting of type (DICTIONARY), scope (SCOPE_BASELINE) and dictionary (T_GUID) identifiers followed by an T_OCTETSTRING of length 16. It encodes a globally unique identifier (UUID) according to RFC4122.*
T_WGS84	*The WGS84 information element is a structure consisting of type (DICTIONARY), scope (SCOPE_BASELINE) and dictionary (T_WGS84) identifiers followed by three values of the geographical location (i.e. longitude, latitude and elevation) encoded as 32 bits signed integer values.*
T_PF_DEPLOY_IE	*The PF_DEPLOY information element is a complex structure consisting of type (DICTIONARY), scope (SCOPE_DSM) and dictionary (T_PF_DEPLOY) identifiers followed by an ordered set of information elements (SP_BLOCK_IE, AREA_IE, SP_OWNER_IE, PF_CERT_IE, SP_USAGE_RIGHTS_IE, SP_USER_IE, SP_QUALITY_IE, SP_POLICY_IE).* *Note that NULL portfolios do not contain any of the above and can only utilized to clear a portfolio data structure remotely..*
T_PF_MODEL_IE	*The PF_MODEL information element is a complex structure consisting of type (DICTIONARY), scope (SCOPE_DSM) and dictionary (T_PF_MODEL) identifiers followed by an unordered set of one or more information elements from SP_BLOCK_IE, AREA_IE, SP_OWNER_IE, PF_CERT_IE, SP_USAGE_RIGHTS_IE, SP_USER_IE, SP_QUALITY_IE, SP_POLICY_IE.*
T_PF_SENSED_IE	*The PF_SENSED information element is a complex structure consisting of type (DICTIONARY), scope (SCOPE_DSM) and dictionary (T_PF_SENSED) identifiers followed by an ordered set of information elements. Mandatory information elements are SP_BLOCK_IE, AREA_IE. Optional elements are SP_OWNER_IE, PF_CERT_IE, SP_USAGE_RIGHTS_IE, SP_USER_IE, SP_QUALITY_IE, SP_POLICY_IE.*
T_PF_CONTEXT_IE	*The PF_CONTEXT information element is a complex structure consisting of type (DICTIONARY), scope (SCOPE_DSM) and dictionary (T_PF_CONTEXT) identifiers followed by an ordered set of information elements from (ffs).*
T_PF_SENSING_IE	*The PF_ SENSING information element is a complex structure consisting of type (DICTIONARY), scope (SCOPE_DSM) and dictionary (T_PF_ SENSING) identifiers followed by an ordered set of information elements from (ffs).*
T_PF_CONTROL_IE	*The PF_ CONTROL information element is a complex structure consisting of type (DICTIONARY), scope (SCOPE_DSM) and dictionary (T_PF_ CONTROL) identifiers followed by an ordered set of information elements from (ffs).*

Table 2-4: Summary of Primitive Parameters

Primitive Parameters		
Parameter	**Parameter Description**	**Parameter Type**
Protocol Message Information Element	*A generic parameter containing a sequence of T_OCTET for further parsing into one or more information elements.*	*T_OCTET_STRING*
Transaction Identifier	*Unique identifier for a transaction. Unique in the scope of a given instance of a SAP.*	*T_INT16*
Result Code	*Return parameter of a primitive indicating success or failure of a certain transaction or of a call to a primitive in general.*	*T_DSM_RESULT_CODE*
Globally Unique Object reference	*A globally unique identifier (UUID) according to RFC4122*	*T_GUID*
Position	*A geographical position according to the geodetic reference system specified by NIMA TR8350.2.*	*T_WGS84*
Model Reference	*Character string representation to an object of a kind matching a 'Model' that has no distinct globally unique identifier.*	*T_STRING8*
Portfolio	*A spectrum portfolio data structure that can be deployed to a spectrum user and can be applied without further modification.*	*T_PF_DEPLOY_IE*

Primitive Parameters		
Parameter	**Parameter Description**	**Parameter Type**
Portfolio Model	*An incompletely specified spectrum portfolio that can be used as a model for completion by, e.g., context information, policy information, or by another spectrum portfolio.*	*T_PF_MODEL_IE*
Sensed Portfolio	*A portfolio data structure carrying a deployable spectrum portfolio that was generated from spectrum sensing. Can be used both as a model and as a deployable spectrum portfolio depending on the context.*	*T_PF_SENSED_IE*
Context	*A generic information element capable to describe any kind of context information in a given scope.*	*T_PF_CONTEXT_IE*
Sensing Information	*A specific context information element describing information obtained from spectrum sensing. Forms a superset of a T_PF_SENSED_IE since it may include context information not representable as a spectrum portfolio (e.g., spectrum occupation maps .*	*T_PF_SENSING_IE*
Control information	*A set of control and status information, or a set of configuration parameters. In general used for configuration and management of QoSMOS entities if suitable.*	*T_PF_CONTROL_IE*

3 CM-SM and CM-RM collaboration

3.1 Introduction

Collaboration between the CM-SM and CM-RM takes place through the CM1 interface, which provides a direct communication link between the two QoSMOS cognitive managers, enabling a simple and efficient communication without the need of an Adaptation Layer (AL). The split of responsibilities for spectrum and radio resource management between two cognitive managers operating at different time-scales and frequency granularities adds flexibility and efficiency to the management, selection and allocation of spectral and radio resources to the QoSMOS entities. However, this requires a simple communication mean for an efficient operation, which is provided by the CM1 interface. The CM1 interface allows a modular architectural design with a clear split of responsibilities between the two cognitive managers while at the same time enabling a close interaction between them for a more efficient operation and improved system performance. The existence of a direct communication interface between the cognitive managers removes the need for an intermediate communication entity (i.e., the AL), which results in a simple and fast communication and therefore in a more efficient interaction and operation.

The specification of this interface includes procedures to support the reporting of spectrum usage and spectrum performance, the management of spectrum portfolios, the management of spectrum policies, and the operation of the cognitive managers, which are designed to support the required interworking and collaboration between CM-RM and CM-SM entities. Such procedures are described in Section 3.3 of [D6.6] and are also included here in the next section (Section 3.2) in order to make this document self-contained. A formal description of the CM1 interface between the CMs, including the Message Sequence Charts (MSCs) and involved primitives for these procedures is presented in Section 3.3.

3.2 Procedures requiring interactions

Several elementary procedures showing the interactions among the functional blocks of the QoSMOS reference model were identified in Section 6.1 of [D2.2]. The corresponding MSCs were provided in [D2.2] as well. Some of these procedures require certain interaction between the CM-RM and the CM-SM functional blocks in the form of basic communication operations, and can be summarised as follows:

- *Spectrum usage and performance reporting*. This procedure is related to the provisioning of reports on spectrum usage and performance from the CM-RM to the CM-SM. These reports are used by the CM-SM as an input to elaborate and keep up-to-date with the spectrum portfolios.
- *Spectrum portfolio management*. This set of procedures is related to the management of the spectrum portfolio by the CM-SM, especially the gathering of the information required for building-up the spectrum portfolio and the provision of the portfolio to the CM-RM. The interaction between the CM-SM and the CM-RM for spectrum portfolio management is described by a number of procedures:
    - *Portfolio deployment*. The purpose of this procedure is to deploy a spectrum portfolio either upon request of an associated CM-RM or upon request of a next-layer CM-SM entity. Deployment communication flow is always downwards within the architectural hierarchy, i.e. from next-layer CM-SM to CM-SM to CM-RM (e.g. from the coordination domain to the networking domain).
    - *Portfolio revocation*. The purpose of this procedure is to revoke a spectrum portfolio either upon request of an associated CM-RM or upon request of a CM-SM or next-

layer CM-SM entity. A spectrum portfolio may be revoked because it is not used any longer or it is for some reason reallocated (deployed) to another entity.

- *Portfolio update or modification*. The purpose of this procedure is the modification of a spectrum portfolio either upon request of an associated CM-RM or upon request of a CM-SM or the next-layer CM-SM entity. The modification of portfolios is an alternative to replacing portfolios. It is applicable to CM-RMs when extending or reducing their portfolios in use (e.g., when requesting backup channels). It is also applicable to CM-SMs managing portfolio repositories more efficiently. Basic modification operations on portfolios are split and merge.

- *Spectrum policy management procedures.* This set of procedures is related to the management of the spectrum usage policies and constraints associated to the deployed spectrum portfolios.
  - *Policy set.* The purpose of this procedure is to set up a policy for a corresponding spectrum portfolio. The policy determines the set of rules that must be followed and the conditions under which a deployed spectrum portfolio can be exploited by a QoSMOS entity. The final set of rules and conditions set for a particular spectrum portfolio is determined by the regulatory policies of a frequency regulation agency and the operator's own policies. The setting of a new policy is initiated by the entity in charge of deploying the spectrum portfolios (i.e., the CM-SM).
  - *Policy revocation.* The purpose of this procedure is to revoke a previously set spectrum policy for an already deployed spectrum portfolio. A policy revocation request is initiated by a CM-SM and results in the revocation of the corresponding policy in the requested CM-RM. After a policy revocation, the associated set of rules and conditions do not apply any longer.
  - *Policy update or modification*. The purpose of this procedure is the modification of a previously set spectrum policy for an already deployed spectrum portfolio. A policy update or modification request is initiated by a CM-SM and results in the modification of the corresponding policy in the requested CM-RM. After a policy modification, the new set of rules and conditions for the exploitation of the spectrum portfolio applies.
- *Cognitive managers operation.* These procedures are related to the management of CMs that start/cease their operation or reconfigure their operation parameters.
  - *Entity registration*. The purpose of this procedure is to register a CM-RM with the closest CM-SM and requesting an initial spectrum portfolio. If the closest CM-SM cannot resolve the portfolio request from its local portfolio repository, then the "next layer" CM-SM is contacted in order to resolve the request, though CM-RM interacts directly only with its closest CM-SM.
  - *Entity reconfiguration*. The purpose of this procedure is to trigger the reconfiguration of a CM or inform another CM that a reconfiguration has been performed. The reconfiguration of a CM-SM can be triggered internally or by a next-layer CM-SM, but not by a CM-RM. When this occurs, the CM-SM informs the associated CM-RMs that a reconfiguration has been performed. On the other hand, the reconfiguration of a CM-RM can be decided and initiated either internally by the CM-RM or by the CM-SM. In both cases the reconfiguration is informed/requested and confirmed. Reconfiguration procedures may occur for example in terms of operation distance from the primary system, transmission powers and load levels/number of users operating over a certain spectrum band.
  - *Entity deregistration*. The purpose of this procedure is to deregister a CM-RM or CM-SM entity before a controlled shutdown. When an entity needs to shut down (i.e., an instance of a CM-RM or CM-SM), the associated CM-SM(s) or CM-RM(s) need to be

informed about its intention. A CM-SM then must revoke or invalidate deployed spectrum portfolios before the requesting CM-RM or CM-SM is allowed to shut down. Usually, some safety measures have to be taken in order to avoid inconsistencies in the CM-SM repository in case a CM-RM or CM-SM disconnects unexpectedly. This case and its associated restart procedure applied after an unsolicited disconnect is considered out of scope for this procedure description.

It is worth noting that there are other more specific procedures not included in the categorisation above that may require interactions between the CMs, as it the case, for example, of the detection of the Radio Access Technologies (RATs) available in the vicinity of a QoSMOS system or the update of the transmission parameters. Such specific procedures are out of the scope of the level of detail of this chapter, and due to the clarifying nature of this chapter, they are therefore not treated here. A detailed description of such procedures along with the corresponding message sequence charts can be found in [D2.2] and [D2.3].

3.3 Interface specification

3.3.1 Interface overview

The CMs need to communicate in order to coordinate their decision processes. A bidirectional interface, designated as *interface CM1*, is defined between the CM-RM and CM-SM. In the CM-RM side, the CM1 interface is connected to the Networking Cognition (NC) sub-entity of the CM-RM, which is in charge of gathering information on the environment in the networking domain for its use and exploitation by other sub-entities of the CM-RM. In the CM-SM side, the CM1 interface is connected to the Spectrum ANalyser (SAN) and Spectrum SElector (SSE) sub-entities of the CM-SM. The SAN accumulates all spectrum information from the CM-RM, the Local PortFolio Repository (LPFR) and additional measurements from sensing devices. It is then able to combine all input data and generate new spectrum policies that will improve the operation of the network entity. The SSE is charged to prepare a pool of portfolios based on the portfolio and policy context stored at the LPFR, makes a pre-calculation on suitable spectrum for opportunistic usage and serves CM-RM spectrum portfolio and policy requests. A more detailed description of the SAN and the SSE modules can be found in Section 5.3 of [D6.3] and chapter 7 in this document.

The interface CM1 is used for exchanging the following, from the CM-RM to the CM-SM [D2.2]:

- Spectrum portfolio request.
- Spectrum portfolio update.
- Performance and usage reports of the active channels, including spectrum sensing measurement results. The CM-RM may send this information periodically, upon a request from the CM-SM or upon the occurrence of a particular event.
- Registration request/response.

The interface CM1 is used for exchanging the following, from the CM-SM to the CM-RM [D2.2]:

- Spectrum portfolio, including spectrum usage information and spectrum usage policies. The CM-SM may send this information periodically, upon a request from the CM-RM or upon the occurrence of a particular event.
- Specification of requested spectrum usage and performance reports.
- Request of a spectrum usage and performance report. The request may be sent for a punctual report or to initiate a periodic reporting.
- Registration request/response.

Depending on the particular procedure and configuration, the exchange of information through the CM1 interface is triggered by certain events, sent as a response to a request or sent on a periodic basis. Event-triggered messages are sent whenever a change is done in any relevant parameter. For example, as the CM-SM is subscribed to the different portfolios and data bases, the CM-RM is subscribed to changes in the CM-SM it is registered with. There is no polling process. Instead, when some change occurs, an event is triggered and all the entities subscribed to it are immediately updated. The CM-SM is subscribed to the different repositories and when they are updated they forward those changes to the subscribed or registered CM-RMs. Also, it is possible to carry out these actions through a simple query process initiated by the CM-RM, which is followed by a response from the corresponding CM-SM informing of performed changes, if any. In addition, it is worth noting that the possibility exists that these exchanges between CMs are made on a regular basis, following a prescribed period, so that at any moment they both would be aware of the current activity in the other CM. A more detailed discussion of the messages and information exchanged between the CM-RM and the CM-SM through the CM1 interface is provided in the following sections.

3.3.2 Naming convention for primitives and their parameters

A common naming convention for the messages exchanged between the CM-RM and the CM-SM (primitives) has been adopted as follows:

InterfaceName_PrimitiveName.MessageClassification(Parameters)

where

InterfaceName: Name of the interface over which the primitive is sent (CM1).

PrimitiveName: A unique descriptive name identifying the primitive.

MessageClassification: A suffix to be added to the primitive name in order to precise the primitive type. The primitives can be classified into four categories:

- *Request* (REQ). Used by one entity to request another entity to perform a given process.
- *Response* (RSP). Used by an entity to respond to a request received from another entity.
- *Confirmation* (CNF). Used by an entity to confirm it has received the request and started processing it. The actual response will be sent later.
- *Indication* (IND). Used by an entity to provide some information to another entity without having received a request from this entity.

Parameters: A list of parameter values sent along with the primitive. Some common parameters are listed below:

- *rmId*: Identification of the CM-RM sending the message or the CM-RM the message is addressed to, depending on whether the message is sent by or addressed to a CM-RM.
- *smId*: Identification of the CM-SM sending the message or the CM-SM the message is addressed to, depending on whether the message is sent by or addressed to a CM-SM.
- *rmCert*: Certificate of a CM-RM provided for authentication purposes.
- *smCert*: Certificate of a CM-SM provided for authentication purposes.
- *pf*: Identification of a spectrum portfolio.
- *pfNew*: Identification of the most recent spectrum portfolio in procedures where two or more spectrum portfolios are active simultaneously for some period of time.
- *pfParams*: Parameters or attributes of a spectrum portfolio.

- *pfNewParams*: Parameters or attributes of the most recent spectrum portfolio in procedures where two or more spectrum portfolios are active simultaneously for some period of time.
- *pl*: Identification of a spectrum policy.
- *plNew*: Identification of the most recent spectrum policy in procedures where two or more spectrum policies are active simultaneously for some period of time.
- *rpParams*: List of parameters to be provided in a spectrum usage and performance report such as the type of measurements to be provided (fraction of busy/idle channels, experienced channel quality, throughput, error rates and so on), the reporting periodicity, etc.

As an example, a CM-RM registers with a serving CM-SM by making use of the message or primitive CM1_Registration.REQ(smId, rmId, rmCert), indicating the CM-SM to which the request is addressed (smId), identifying the CM-RM which is sending the request (rmId) and providing a certificate (rmCert) to establish a trusted peer-to-peer communication channel.

3.3.3 Message sequence charts and primitives for basic procedures

The following tables provide the set of primitives associated to each of the basic procedures defined in Section 3.2, indicating their chronological order of occurrence, their direction of transmission (from to CM-RM to CM-SM, or from CM-SM to CM-RM) along with a brief description.

Table 3-1 provides the Message Sequence Chart (MSC) for spectrum usage and performance reporting.

Table 3-1: MSC for spectrum usage and performance reporting.

Spectrum usage and performance reporting			
No.	**Primitive**	**Direction**	**Description**
1	CM1_ReportSpec.IND(rmId, smId, rpParams)	RM←SM	Fixes conditions in which reports will be transmitted (type of measurements, reporting periods, etc.)
2	CM1_ReportSpec.CNF(smId, rmId)	RM→SM	Confirms the specification of reporting conditions
3	CM1_MeasurementReport.IND(smId, rmId)	RM→SM	Provides report with specified information (periodically)

Table 3-2 provides the MSC for spectrum portfolio management (portfolio deployment). When initiated by the CM-RM, the procedure starts at message #1. When initiated by the CM-SM or a next-layer CM-SM, the procedure starts at message #2.

Table 3-2: MSC for spectrum portfolio management (portfolio deployment).

Spectrum portfolio management – Portfolio deployment			
No.	**Primitive**	**Direction**	**Description**
1	CM1_PortfolioDeploy.IND(smId, rmId, pfParams)	RM→SM	Requests a spectrum portfolio placing its requirements in a describing parameter set
2	CM1_PortfolioDeploy.REQ(rmId, smId, pf)	RM←SM	Provides a portfolio to be deployed by the CM-RM
3	CM1_PortfolioDeploy.RSP(smId, rmId, pf)	RM→SM	Acknowledges deployment of the portfolio

Table 3-3 provides the MSC for spectrum portfolio management (portfolio revocation). When initiated by the CM-RM, the procedure starts at message #1. When initiated by the CM-SM or a next-layer CM-SM, the procedure starts at message #2.

Table 3-3: MSC for spectrum portfolio management (portfolio revocation).

Spectrum portfolio management – Portfolio revocation			
No.	**Primitive**	**Direction**	**Description**
1	CM1_PortfolioRevoke.IND(smId, rmId, pf)	RM→SM	Informs CM-SM that a portfolio is no longer in use
2	CM1_PortfolioRevoke.REQ(rmId, smId, pf)	RM←SM	Requests CM-RM to delete previously provided portfolio
3	CM1_PortfolioRevoke.RSP(smId, rmId, pf)	RM→SM	Acknowledges portfolio revoking request

Table 3-4 provides the MSC for spectrum portfolio management (portfolio update/modification). When initiated by the CM-RM, the procedure starts at message #1. When initiated by the CM-SM or a next-layer CM-SM, the procedure starts at message #2.

Table 3-4: MSC for spectrum portfolio management (portfolio update/modification).

Spectrum portfolio management – Portfolio update/modification			
No.	**Primitive**	**Direction**	**Description**
1	CM1_PortfolioUpdate.IND(smId, rmId, pf, pfNewParams)	RM→SM	Informs CM-SM about the operating channel(s) in order to request a modification of its current portfolio
2	CM1_PortfolioDeploy.REQ(rmId, smId, pfNew)	RM←SM	Provides a portfolio to be deployed by the CM-RM
3	CM1_PortfolioDeploy.RSP(smId, rmId, pfNew)	RM→SM	Acknowledges deployment of the portfolio
4	CM1_PortfolioRevoke.REQ(rmId, smId, pf)	RM←SM	Requests CM-RM to delete previously provided portfolio
5	CM1_PortfolioRevoke.RSP(smId, rmId, pf)	RM→SM	Acknowledges portfolio revoking request

Table 3-5 provides the MSC for spectrum policy management (policy set).

Table 3-5: MSC for spectrum policy management (policy set).

Spectrum policy management – Policy set			
No.	**Primitive**	**Direction**	**Description**
1	CM1_PolicySet.REQ(rmId, smId, pl)	RM←SM	Provides spectrum usage policies relative to a deployed spectrum portfolio, to be followed by the CM-RM
2	CM1_PolicySet.RSP(smId, rmId, pl)	RM→SM	Acknowledges policy set request

Table 3-6 provides the MSC for spectrum policy management (policy revocation).

Table 3-6: MSC for spectrum policy management (policy revocation).

Spectrum policy management – Policy revocation			
No.	**Primitive**	**Direction**	**Description**
1	CM1_PolicyRevoke.REQ(rmId, smId, pl)	RM←SM	Requests CM-RM to delete a previously provided policy
2	CM1_PolicyRevoke.RSP(smId, rmId, pl)	RM→SM	Acknowledges policy revoking request

Table 3-7 provides the MSC for spectrum policy management (policy update/modification).

Table 3-7: MSC for spectrum policy management (policy update/modification).

Spectrum policy management – Policy revocation			
No.	**Primitive**	**Direction**	**Description**
1	CM1_PolicySet.REQ(rmId, smId, plNew)	RM←SM	Provides a new policy to be followed by the CM-RM
2	CM1_PolicySet.RSP(smId, rmId, plNew)	RM→SM	Acknowledges deployment of the policy
3	CM1_PolicyRevoke.REQ(rmId, smId, pl)	RM←SM	Requests CM-RM to delete previously provided policy
4	CM1_PolicyRevoke.RSP(smId, rmId, pl)	RM→SM	Acknowledges policy revoking request

Table 3-8 provides the MSC for CMs operation (entity registration). The entity registration procedure is always initiated by the CM-RM. Therefore, the procedure starts always at message no. 1.

Table 3-8: MSC for CMs operation (entity registration).

CMs operation – Entity registration			
No.	**Primitive**	**Direction**	**Description**
1	CM1_Registration.REQ(smId, rmId, rmCert)	RM→SM	Requests registration to serving CM-SM and establishes a trusted peer-to-peer communication channel
2	CM1_Registration.RSP(rmId, smId, smCert, regInfo)	RM←SM	Returns result of registration request and provides registration information
3	CM1_Registration.IND(smId, rmId)	RM→SM	Acknowledges registration result and information sent by the CM-SM
4	CM1_PortfolioDeploy.IND(smId, rmId, pfParams)	RM→SM	Requests a spectrum portfolio placing its requirements in a describing parameter set
5	CM1_PortfolioDeploy.REQ(rmId, smId, pf)	RM←SM	Provides a portfolio to be deployed by the CM-RM
6	CM1_PortfolioDeploy.RSP(smId, rmId, pf)	RM→SM	Acknowledges deployment of the portfolio

As mentioned in Section 3.2, both the CM-RM and CM-SM may undergo a reconfiguration process for various reasons. The reconfiguration of a CM-SM can be triggered internally or by a next-layer CM-SM (but not by a CM-RM). When this occurs, the CM-SM informs the associated CM-RMs that a reconfiguration has been performed. Table 3-9 provides the MSC for CMs operation (entity reconfiguration) when the reconfiguration takes place in the CM-SM.

Table 3-9: MSC for CMs operation (CM-SM reconfiguration).

CMs operation – CM-SM reconfiguration			
No.	**Primitive**	**Direction**	**Description**
1	CM1_Reconfiguration.IND(rmId, smId, pfParams)	RM←SM	Informs that a reconfiguration has taken place

On the other hand, the reconfiguration of a CM-RM can be decided and initiated either internally by the CM-RM or by the CM-SM. In both cases the reconfiguration is informed/requested and confirmed. Table 3-10 provides the MSC for CMs operation (entity reconfiguration) when the reconfiguration takes place in the CM-RM. When the reconfiguration is triggered by the CM-RM, the procedure starts at message #1. Otherwise, the procedure starts at message #2.

Table 3-10: MSC for CMs operation (CM-RM reconfiguration).

CMs operation – CM-RM reconfiguration			
No.	**Primitive**	**Direction**	**Description**
1	CM1_Reconfiguration.IND(smId, rmId, pfParams)	RM→SM	Used by a CM-RM to indicate an internal reconfiguration
2	CM1_Reconfiguration.REQ(rmId, smId, pf)	RM←SM	The CM-SM requests/accepts a CM-RM reconfiguration
3	CM1_Reconfiguration.RSP(smId, rmId, pf)	RM→SM	Acknowledges internal reconfiguration

Table 3-11 provides the MSC for CMs operation (entity de-registration). When initiated by the CM-SM, the procedure starts at message #1. When initiated by the CM-RM, the procedure starts at message #2.

Table 3-11: MSC for CMs operation (entity deregistration).

CMs operation – Entity deregistration			
No.	**Primitive**	**Direction**	**Description**
1	CM1_DeRegistration.IND(rmId, smId)	RM←SM	Notifies CM-RM to initiate deregistration procedure
2	CM1_DeRegistration.REQ(smId, rmId)	RM→SM	Requests deregistration from the CM-SM
3	CM1_PortfolioRevoke.REQ(rmId, smId, pf)	RM←SM	Requests CM-RM to delete previously provided portfolio
4	CM1_PortfolioRevoke.RSP(smId, rmId, pf)	RM→SM	Acknowledges portfolio revoking request
5	CM1_DeRegistration.RSP(rmId, smId)	RM←SM	Acknowledges the deregistration of the CM-RM

As mentioned in Section 3.2, a detailed description of some more complex and specifically targeted procedures along with the corresponding MSCs and primitives can be found in [D2.2] and [D2.3].

4 Repositories

4.1 Global Regulator Repository (GRGR)

4.1.1 Functions in GRGR

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 geo-location database). The GRGR may also instantiate as a single entity or in a distributed way. In the latter case it serves CM-SM 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 one or more contiguous channels 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 from a specific use case implementation allows certain scenarios that merge information retrieved from the GRGR and from the SPRR (Section4.3) prior to deploying a valid (i.e. qualified) spectrum portfolio. For example, requesting GRGR and SPRR (Spectrum Provider Repository) in parallel, or having the GRGR to query the SPRR on its own in a local process can simplify procedures to merge the information retrieved from the GRGR and from the SPRR. Potential communication between GRGR and SPRR is considered private and will not be addressed by this document.

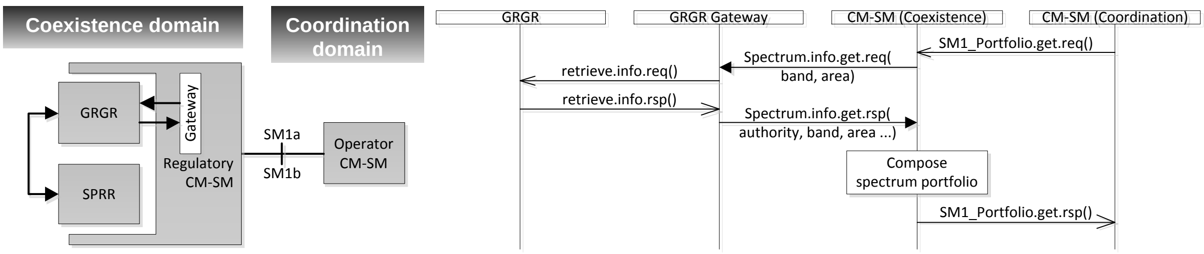

Figure 4-1: Accessing the GRGR and sample MSC

4.1.2 Interfaces in accessing GRGR

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 as explained in references [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.

4.1.3 TV whitespace geo-location database functional description and evaluation

One of the key requirements to enable opportunistic spectrum access is efficient automated spectrum management. One basic spectrum management method is a geo-location database approach, which has been adopted in the recent rules of Federal Communications Commission in the US to allow opportunistic spectrum access to use TV whitespace. Geo-location database is the preferred approach by the UK regulator Ofcom in accessing TV white space. In the FCC scenarios, the devices of the opportunistic systems identify their locations by using a geo-location capability and then query the database to determine which TV channels they can use at their locations. To verify that the reuse of the TV channel does not cause harmful interference to the incumbent receiver, the database checks whether the signal-to-interference ratio of the incumbent receiver can be kept at a required level. A transmit power which satisfies the signal-to-interference ratio requirement corresponds to an allowable transmit power of the opportunistic system.

QoSMOS has implemented a prototype TV whitespace database for demonstration purposes in the UK, and this database is also used by the QoSMOS demonstration spectrum manager. A demonstration web-browser interface is available at http://www.ict-qosmos.eu/project/demos.html. For full details of how this database is constructed and of the interfaces to it see [D6.6 and D6.7 section 4.1]

4.2 Common Portfolio Repository (CPFR)

4.2.1 Functions in CPFR

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).

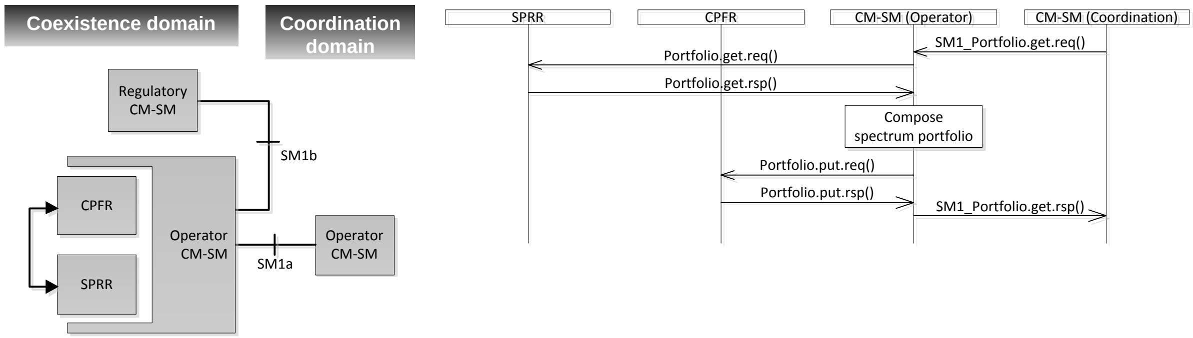

Figure 4-2: Accessing the CPFR and sample MSC

4.2.2 Interfaces in accessing CPFR

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 deployment. In this way spectrum users can be obliged (if the regulator so decides) 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.

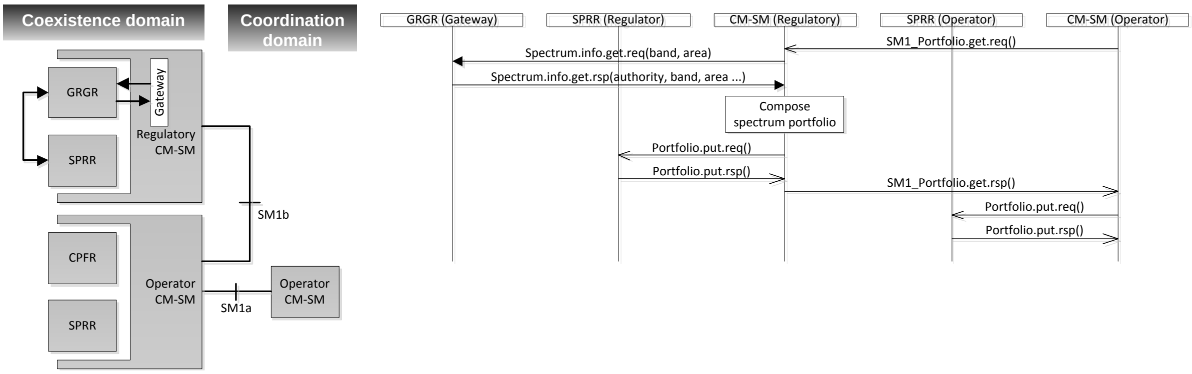

Figure 4-3: Accessing the SPRR and sample MSC

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 QoSMOS 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).

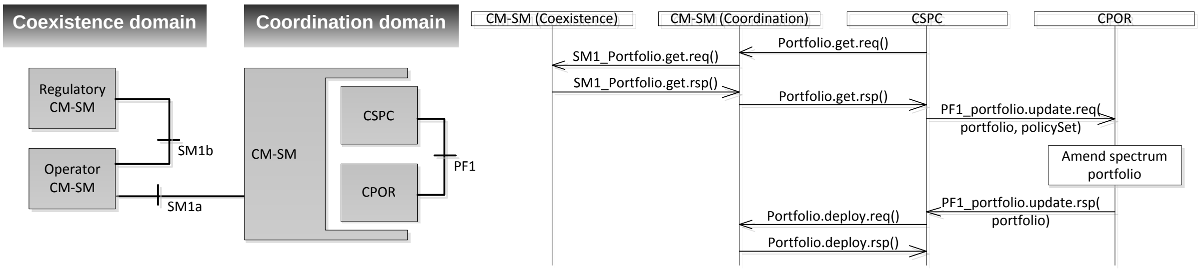

Figure 4-4: Accessing the CPOR and sample MSC

4.4.2 Interfaces

The **PF1 interface** is used to exchange policies between the Common Spectrum Control (CSPC) function and the Common Policy Repository (CPOR) [1900.5]. It is a 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.

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), 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.

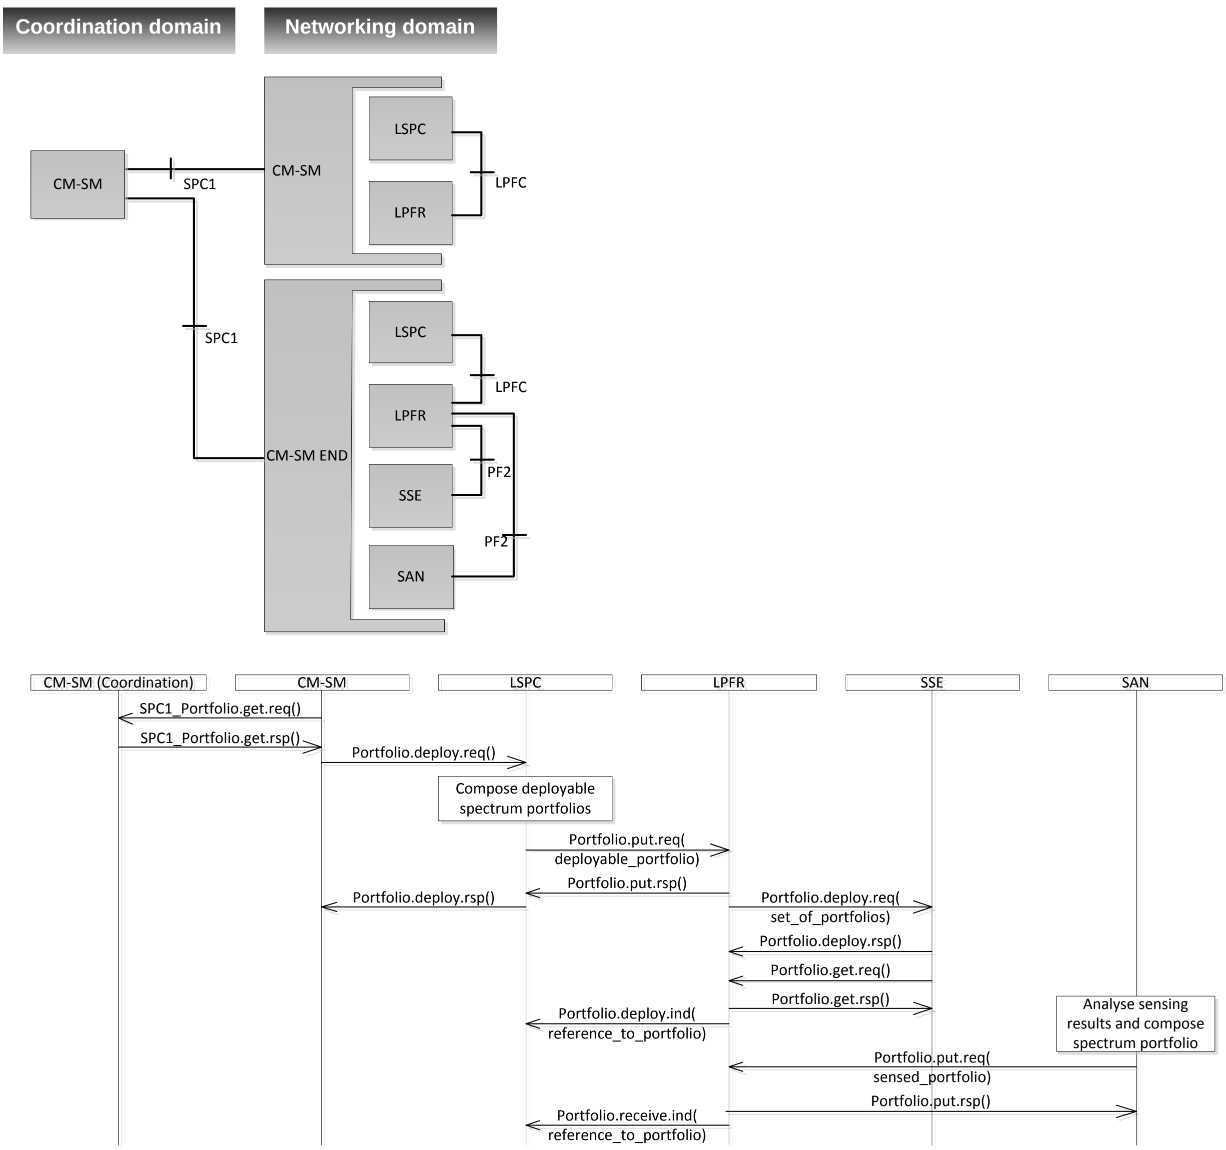

Figure 4-5: Accessing the LPFR and sample MSC

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 Analyser (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. QoSMOS spectrum portfolios form a superset of DARPA XGL policy rules.

All information in an XGL ‘selector’ also is present in a QoSMOS 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 CSPC 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 CSPC 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 QoSMOS 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 QoSMOS 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 CSPC performs the reasoning required to infer comparable facts from both. In consequence, the CSPC 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.

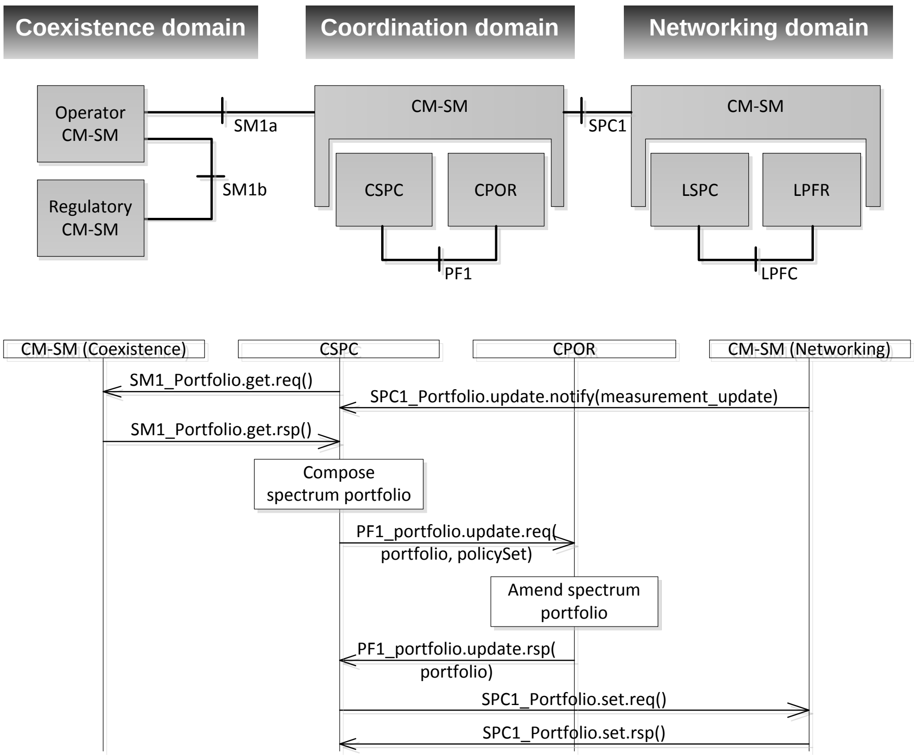

Figure 5-1: Accessing the CSPC in cellular scenarios and sample MSC

The information given by XGL ‘usage constraints’ is directly included in a QoSMOS spectrum portfolio in form of its policies and usage constraints parameters embedded. In addition the CSPC includes operator policies from the CPOR as applicable und 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 as 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.

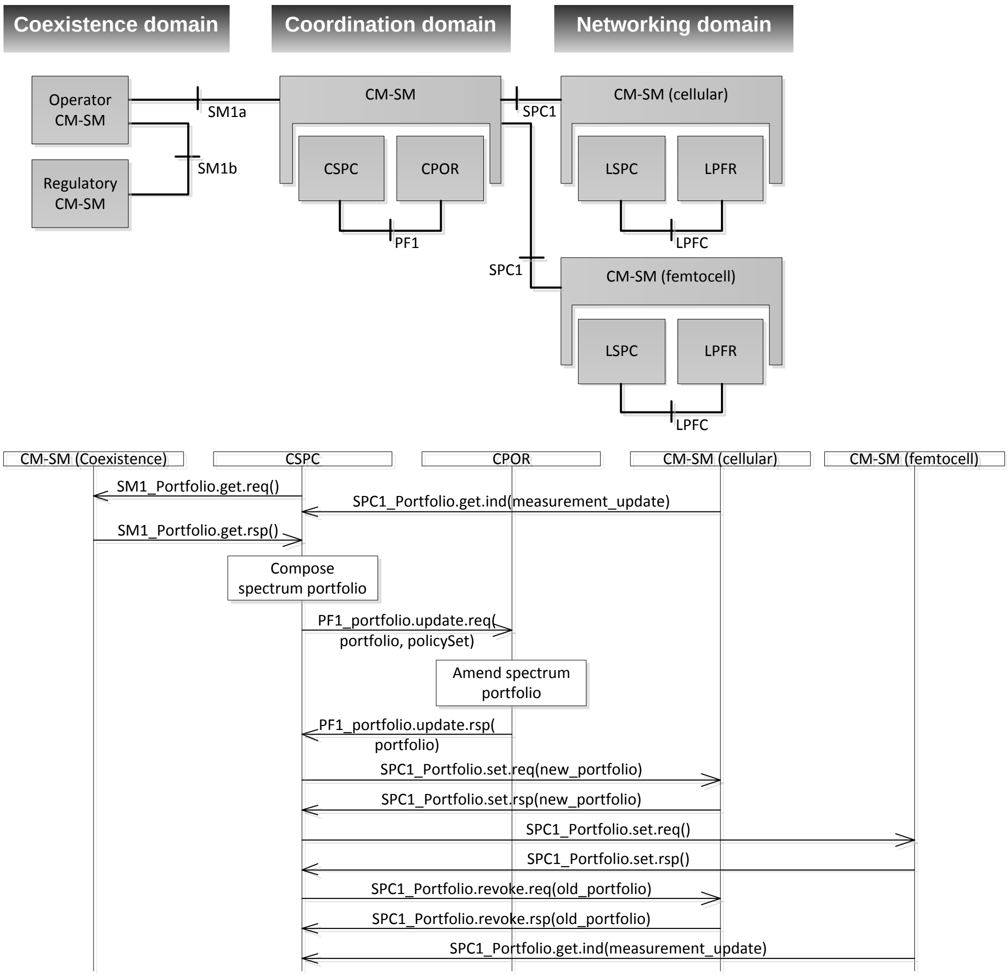

Figure 5-2: Accessing the CSPC in femtocell scenarios and sample MSC (portfolio update for cellular, portfolio deployment for femtocell)

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.

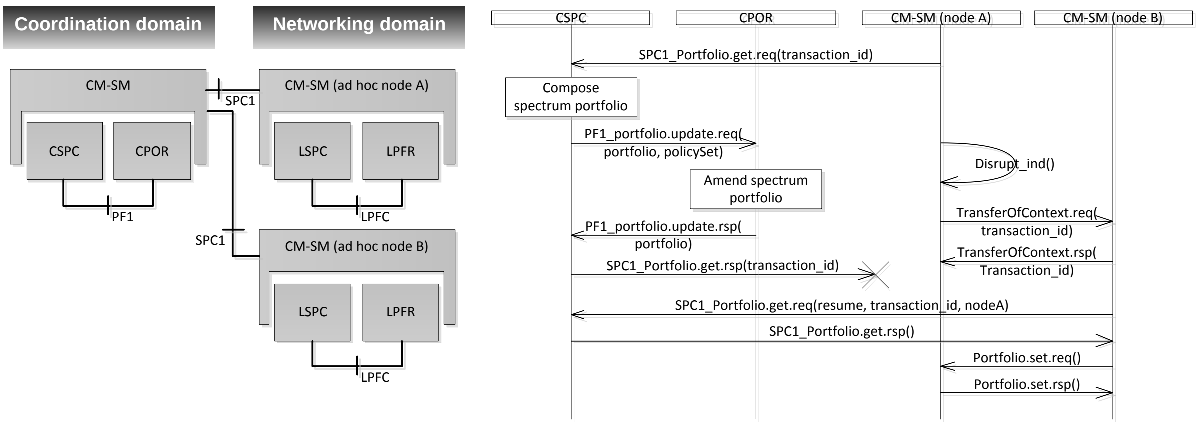

Figure 5-3: Accessing the CSPC in ad hoc scenarios and sample MSC (role switch from ad-hoc node A to ad-hoc node B)

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.

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 where the spectrum portfolio is deployed.

5.4 Advanced multi-objective portfolio optimization for spectrum selection/aggregation in cognitive radio

Multi-objective portfolio optimization has been used in the problem of frequency selection/aggregation of cognitive radio systems in deliverable [D6.6]. The main conclusion was that technical and economic target functions can be conveniently reached in a cognitive radio network by means of a multi-objective port-folio optimization problem given an appropriate balance between return and risk components and certain values of economic parameters of the available frequency bands. However, the approach considered includes low level of interaction between CM-RM and CM-SM. In this section we provide an extension of these previous developments that will allow a tighter interaction between both entities and that will allow for the deployment of more dynamic pricing schemes which depend also on network conditions and that manage a finer granularity at the radio resource unit level.

5.4.1 Joint frequency selection, resource allocation, and user scheduling based on multi-objective portfolio optimization

As observed in the previous deliverables, the theory of portfolio optimization can be applied in the problem of frequency selection. It was also observed that some modifications to the objective functions are required in order to consider some crucial network parameters such as interference and network load. This convergence of network and economic parameters has also a consequence on the level of cooperation between network entities dedicated to these two functionalities, i.e. the CM-SM and the CM-RM. In this section we take this level of cooperation to the limit, which actually means that both entities should be hosted in the same location. In this section we propose the study of a complete fusion between the CM-SM and the CM-RM functionalities.

Consider the deployment scenario depicted in Figure 5-4 with two WiMAX-based networks, one indoor and one outdoor, operating in two different frequency bands denoted by f_{in} and f_{out}, respectively. The outdoor network consists of a base station (BS) located at the centre of the deployment. The indoor network operates inside a rectangular building with a horizontal footprint of dimensions denoted by width *a* and depth *b*, and its BS is located at the centre of the building which is located at the coordinates denoted by $x_{\text{in}}, y_{\text{in}}$. The analysis is focused on the downlink considering a bandwidth of 10MHz using the transmission parameters described in [WIXSLS] for TDD WiMAX systems. The resource management entity (CM-RM) of the two networks is centralized at the BS of each network. A cognitive radio is assumed to be operating at the BS of each network. Each BS is able to allocate all the time and frequency resources of its operating frequency band and, in addition, allocate some users in an opportunistic manner in the neighbour frequency band of the adjacent network. The propagation model for outdoor is defined by the WINNER B1 model [WIND112].

Indoor propagation model is a B5 WINNER model [WIND112]. A building penetration loss of 8 dB is used in all simulations. Figure 5-5 displays the architecture of system-level simulator used for the evaluation of the proposed multi-objective optimization framework for cognitive radio. Note that the entities for CR-RM and CM-SM have a close interaction in the simulation architecture. Figure 5-6 depicts the OFDMA super-frame for resource allocation in a multi-cell scenario with different frequency bands for primary and secondary transmissions. In the general case it is assumed that each frequency band is licensed to a different operator. However, each operator can access other frequency bands licensed to other operators in an opportunistic manner with different economic parameters.

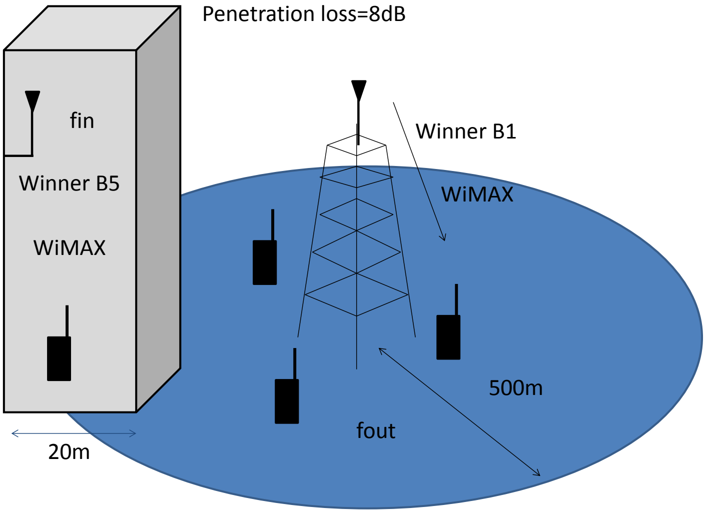

Figure 5-4: Testing scenario

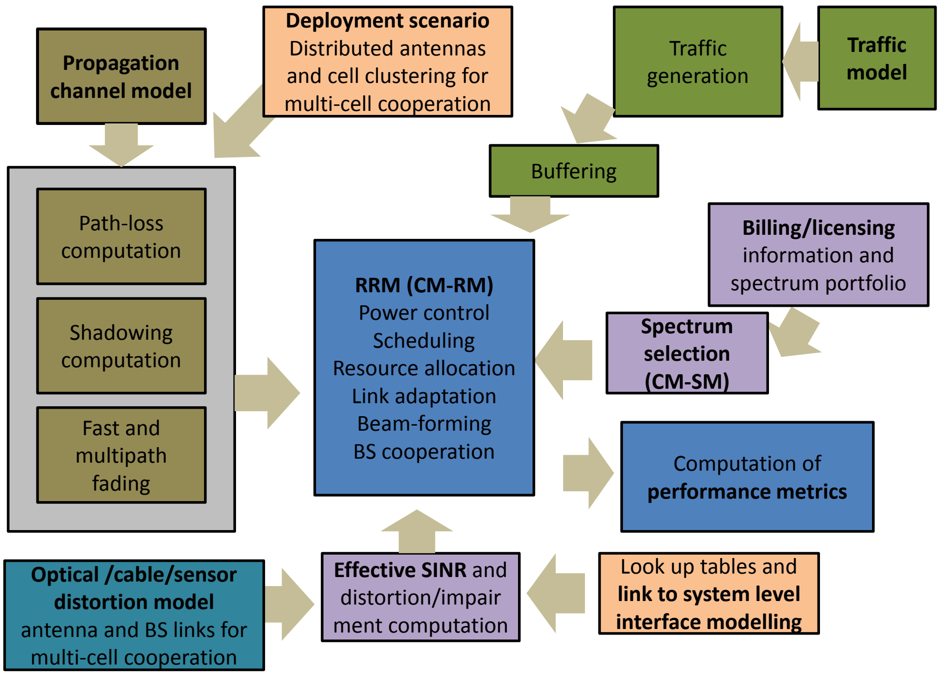

Figure 5-5: Block diagram of the system level simulator for testing multi-objective optimization for frequency selection of cognitive radio networks.

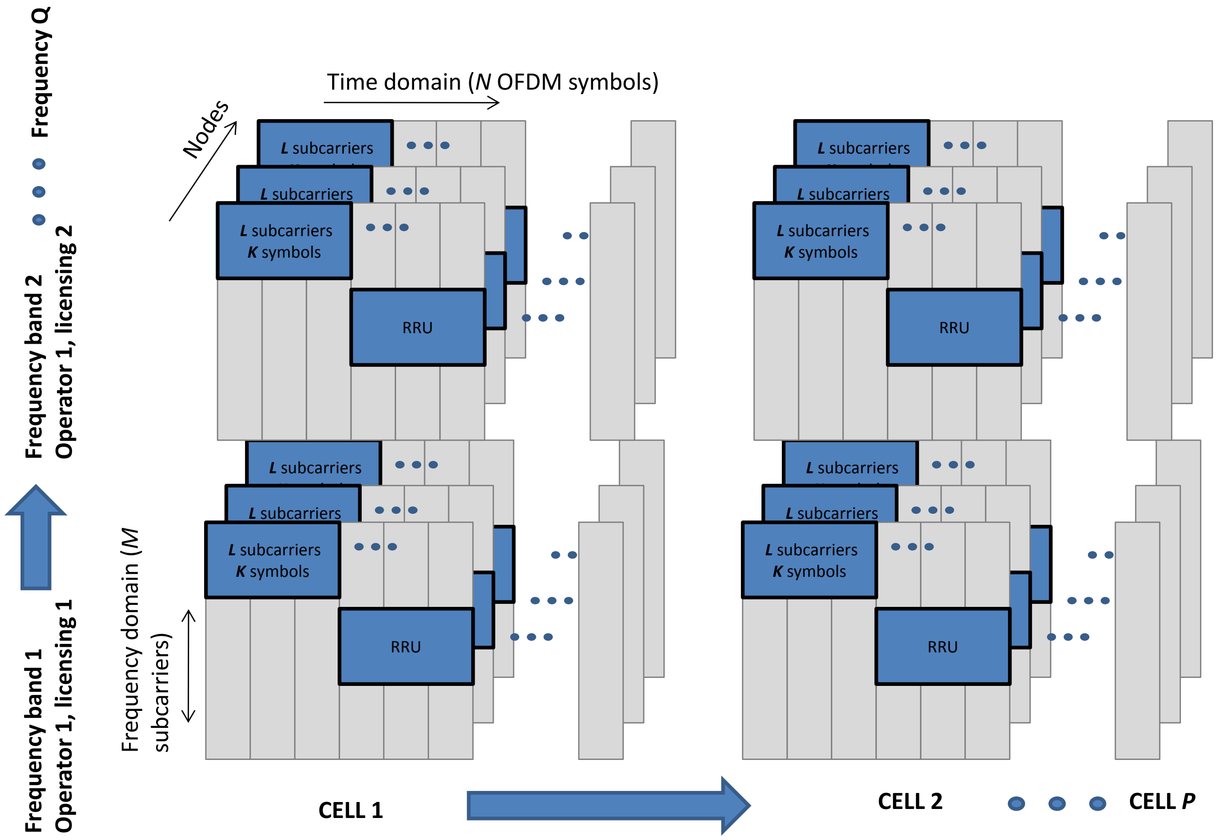

Figure 5-6: Frame definition for OFDMA resource allocation in cellular wireless networks

Table 5-1: System level simulation parameters

Parameter	Explanation
Layout	Manhattan grid
Building width	20 m
Channel model	WINNER B1
Antenna radiation pattern	Omnidirectional
BS antenna gain	15 dB
Mobile antenna gain	0 dB
Wall penetration loss	8 dB
Simulator mode	Combined snapshot
Traffic model	Full queue
N. freq. bands	2
Bandwidth	10 MHz
Subcarriers per symbol	1024
Data subcarriers per symbol	720
Frame length	20 symbols
Radio resource unit	10x720
Frame duration	5 ms
CQI feedback delay	10 ms
LSLI compression	EESM

For each user in the network, we calculate an expected throughput performance for each one of the radio resources in the network. The expected performance has to consider the measured interference and load from the adjacent network. The next step is to calculate a multi-objective function for each user *j* allocated to a different radio resource *i* of the frequency band *k*:

$$f_{i,j,k} = p_{i,j,k}\overline{T}_{i,j,k} - \mu C_{i,j,k}\overline{T}^{2}_{i,j,k}$$

where $p_{i,j,k}$ is the return when user *j* transmits on radio resource *i* of the frequency band *k*, $\overline{T}_{i,j,k} = E\left[T_{i,j,k}\right]$ is the expected throughput when user *j* transmits on radio resource *i* of the frequency band *k*, $C_{i,j,k}$ is the risk when user *j* transmits on radio resource *i* of the frequency band *k*, and $\overline{T}^{2}_{i,j,k}$ is the variance of the throughput when user *j* transmits on radio resource *i* of the frequency band *k*

The problem can be stated then as follows:

$$\{U,R,F\}_{\mathrm{opt}} = \arg\max_{\{U,R,F\}} \sum f_{i,j,k}$$

where U, R, F are the sets of users, resources and frequency bands, respectively, to be allocated. The optimization can also include a proportional fairness indicator in the risk part of the objective function. The joint optimization can be carried out in different ways. The most obvious one is simply to sort all user-frequencies-resource combinations and select the set that maximizes the objective function. Figure 5-7 shows the advantages of the joint optimization approach for frequency selection and radio resource management over the conventional management approach used in deliverable [D6.6], where frequency selection (CM-SM) and radio resource management (CM-RM) functionalities have a small degree of interaction. The figure shows the total throughput considering primary and secondary user transmissions versus the total transmit power normalized to the noise variance of the receiver (in dBs). Gains are obtained for all values of transmit power used.

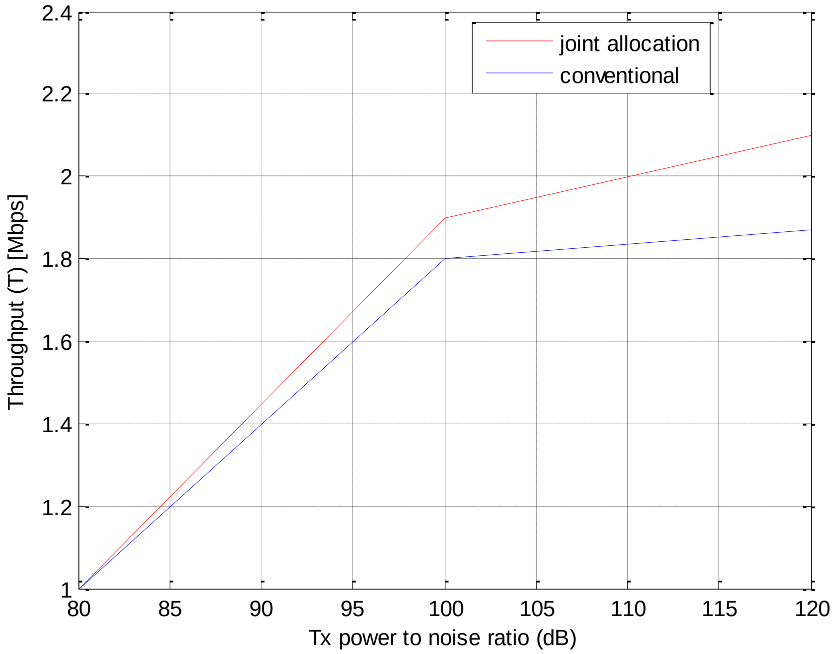

Figure 5-7: Throughput of primary+secondary users indoor network using joint resource allocation, user scheduling, and frequency selection based on multi-objective portfolio optimization.

5.4.2 Joint frequency selection, resource allocation, user scheduling and beam-forming for distributed antenna systems based on multi-objective portfolio optimization

In this section we extend the proposed algorithm of the previous subsection to a distributed antenna system with beam-forming. The scenario is depicted in Figure 5-8, where the distributed antenna system is positioned outdoor. The system is allowed to work in one frequency and opportunistically in a secondary band that is being used inside the buildings of the scenario. The objective is to select a different user attached to a different antenna in the system using the criteria explained before:

$$f_{i,j,k}(W) = p_{i,j,k} T_{i,j,k}(W) - \mu C_{i,j,k} T_{i,j,k}^{2}(W)$$

where W is the set of beam-forming vectors to be optimized. The problem can now be stated then as follows:

$$\{U, R, F, W\}_{\text{opt}} = \arg \max_{\{U,R,F,W\}} \sum f_{i,j,k}$$

This optimization problem can be split into different stages. Each node is assigned with the users that maximize the objective function. Then, for each resource, a weighted least squares optimization is initialized in order to control the interference created between users assigned to the same resource but on different antenna. The weighted least squares optimization targets both power levels and beam-forming vectors with the aim to satisfy a SINR level of the chosen modulation and coding scheme for each scheduled user. The weights of each residual error term are given by the portfolio multi-objective optimization. In this way we include the economic information in the joint power and beam-forming optimization. If the solution does not reach an error function equal to zero, then the modulation and coding schemes or the set of scheduled users are modified. Another least squares optimization stage that reuses the result of the previous stage is initialized with the updated set of modulation and coding schemes and scheduled users. The process is repeated until the error function is equal to zero, which means that the users are allowed to transmit with the given modulation and coding scheme and with the transmit power levels and the beam-forming vectors calculated in the latest least squares optimization stage. Figure 5-9 shows the results of the improved approach using beam-forming and the proposed resource allocation scheme with frequency selection in cognitive radio systems. Figure 5-10 displays the multicell Manhattan deployment scenario for system level evaluation, and Figure 5-11 shows the flowchart diagram of the proposed algorithm with joint frequency selection, user scheduling, beam forming and resource allocation.

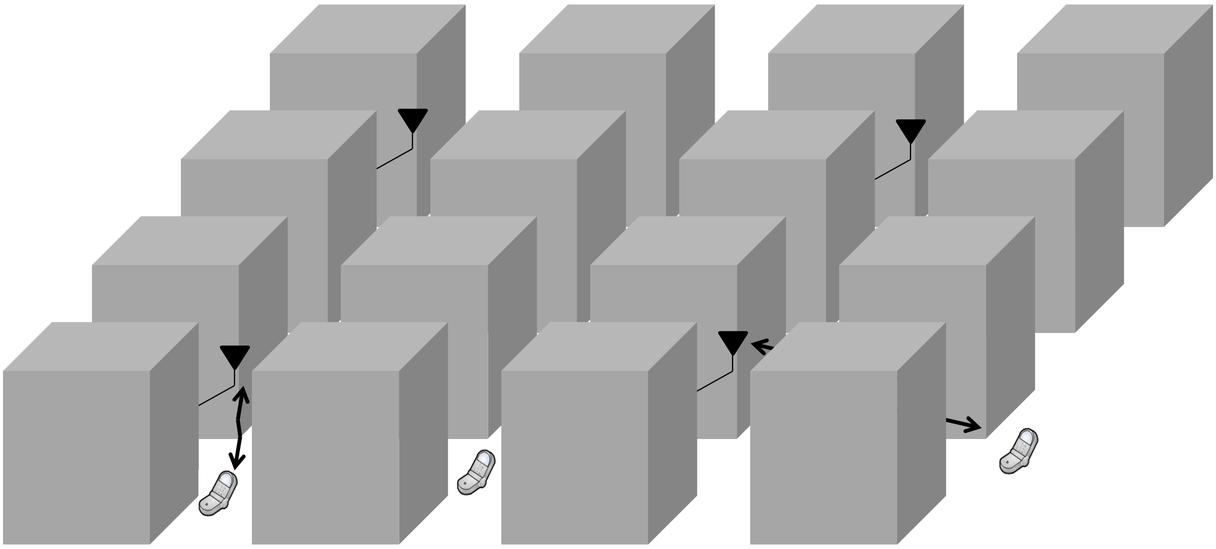

Figure 5-8: Manhattan scenario for testing of cognitive radio algorithms.

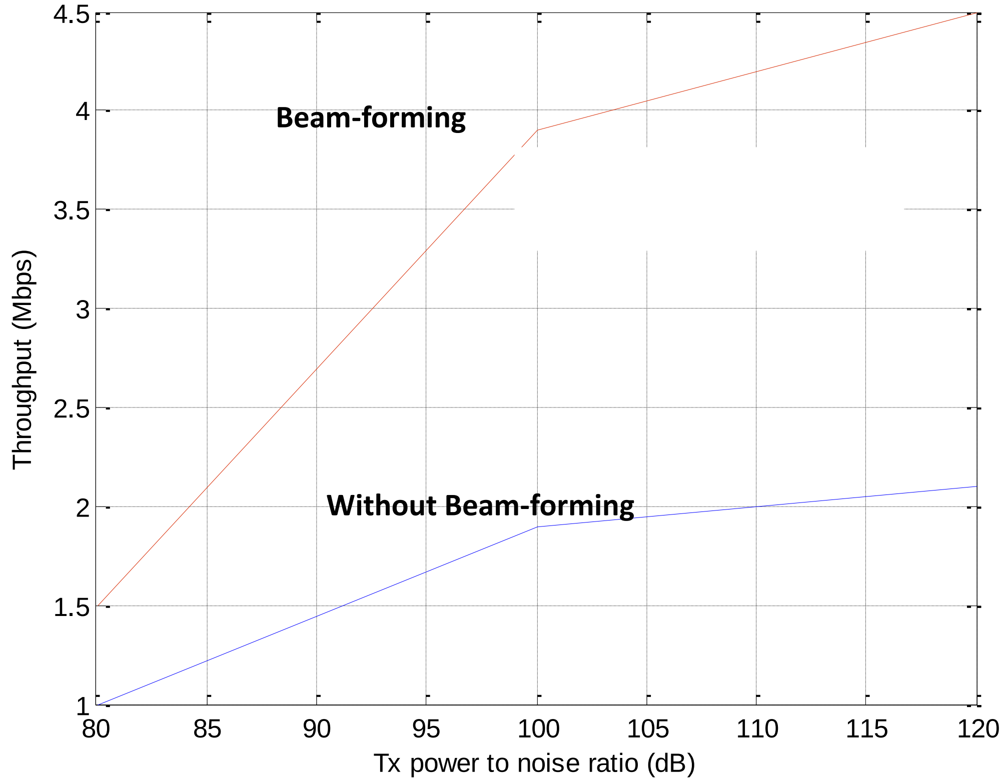

Figure 5-9: Throughput of primary+secondary users indoor network using joint resource allocation, user scheduling, frequency selection and beam-forming based on multi-objective portfolio optimization in the test scenario of Figure 5-8.

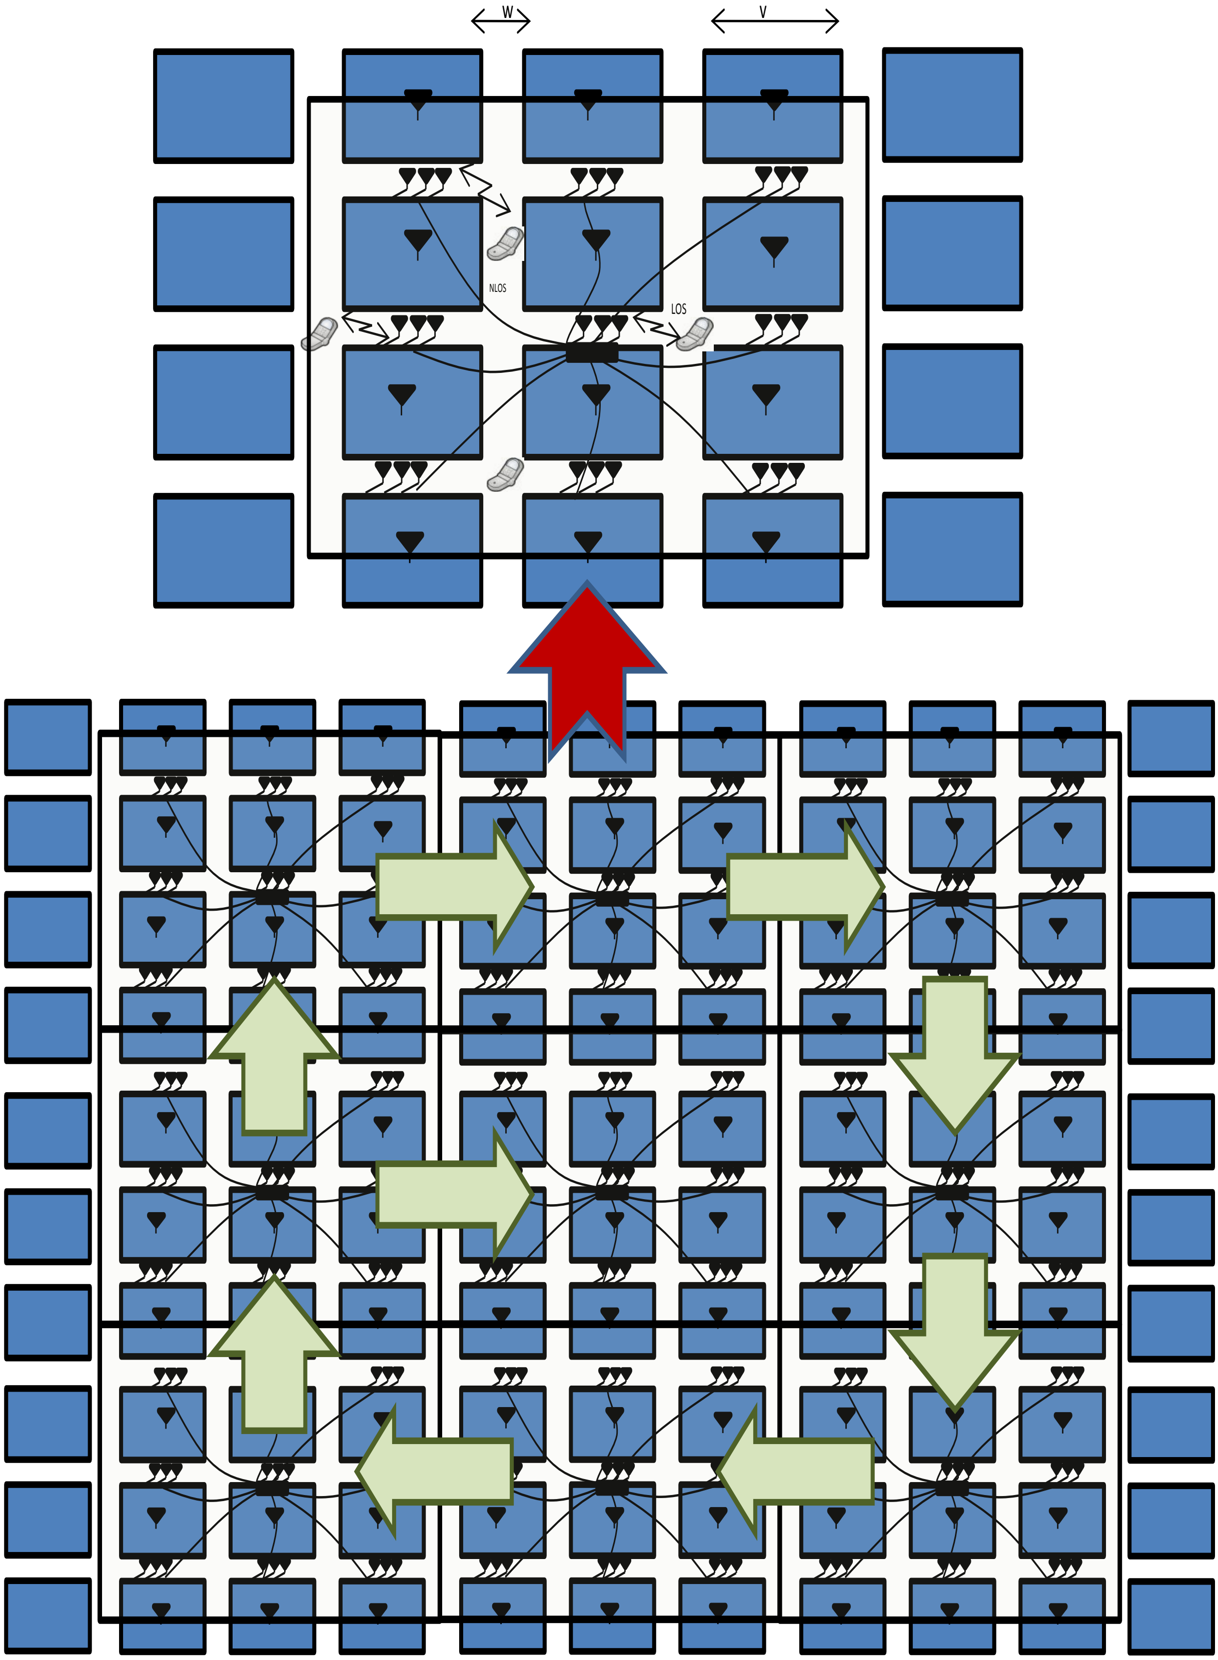

Figure 5-10: Manhattan multi-cell scenario for system-level evaluation of distributed antenna systems with beam-forming and cognitive radio.

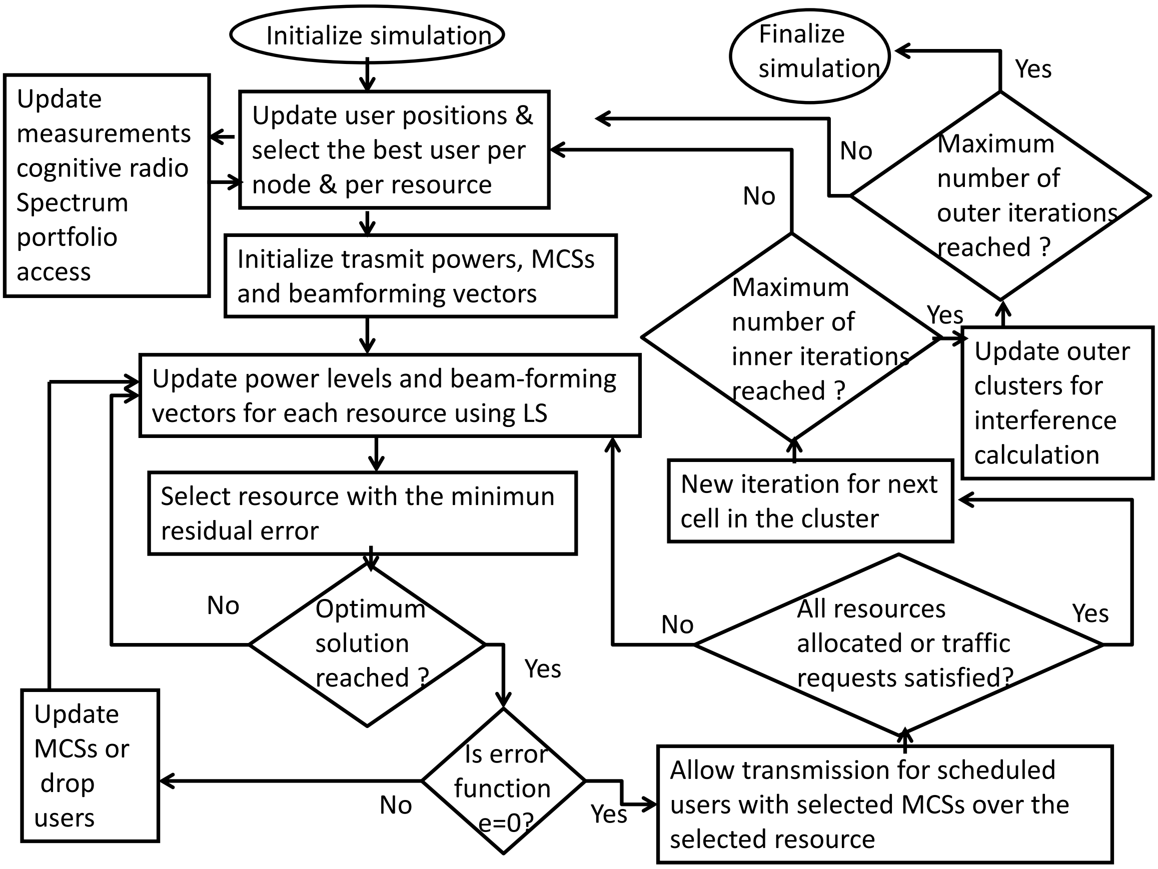

Figure 5-11: Flowchart describing the algorithm for radio resource management and frequency spectrum selection for DAS with beam-forming and cognitive radio using multi-objective portfolio optimization.

5.4.3 Throughput and Return-Risk Trade-off Regions of Cognitive Radio Systems with Random Transmission Control

The allocation of radio resources in a modern cognitive radio network with heterogeneous licensing/billing agreements, as well as different radio conditions must reflect the complex trade-offs between the radio and economic worlds. This problem can be conveniently expressed as a multi-objective portfolio optimization, which has been investigated in detail in the fields of economics and finance theory.

This section attempts to partially address this issue by optimizing the usage of spectrum bands for primary and secondary transmissions in terms of economic indicators such as return and risk, in addition to the conventional network performance metrics such as throughput and load. The section proposes a multi-objective function that attempts to optimize the random transmission probability on each available resource of each frequency band such that the total return is maximized and the risk (or variance of the return) is minimized or controlled. This optimum transmission probability is compared to the transmission policy that maximizes the throughput region of the system. In this work it is assumed that both primary and secondary transmissions are uncoordinated, and thus resemble a random access network. To facilitate analysis, a packet reception model for Rayleigh channels is also proposed that allows for calculation of correct packet reception statistics in closed-form expressions. The boundaries of both the throughput and risk-return regions are derived in a non-parametric closed-form expression, which allows for a geometrical interpretation of the resource allocation problem. By inspection and by analytical work, conditions are derived for simultaneously achieving maximum throughput and maximum economic performance. This means that operators can simultaneously maximize revenue and network resource efficiency. A third transmission policy is further proposed that maximizes return while controlling risk and ensuring a level of quality of service for primary and secondary users. This means that operators can maximize revenue and network efficiency while minimizing risk and ensuring some fairness between primary and secondary users. The trade-off performance regions of this transmission policy are characterized and discussed. The results in this section also give some indication of the economic parameters in terms of return and risk that will provide operators with optimum network and revenue performances, which can be also useful in the design of billing schemes.

5.4.3.1 System model

Consider the deployment scenario in Figure 5-12 with $L=2$ networks, operating in $K=2$ different frequency bands f_1 and f_2, respectively. The two networks are assumed to have a total of M orthogonal radio resources available for allocation on each of the frequency bands, one for primary (licensed) user transmissions and the other one for secondary (opportunistic) transmissions. For simplicity we consider all the resources to be statistically identical. Let us focus on a single radio resource: a primary user will experience a channel with its serving base station denoted by h_p and a channel with the potential source of interference form a secondary transmission denoted by h_{sp}. Similarly, a secondary user will experience a channel with its serving base station denoted by h_s and a channel with the potential source of interference from a primary transmission denoted by h_{ps}. All channels will be modelled as circular complex Gaussian random variables with zero mean and variances denoted, respectively, by: σ_p^2, σ_s^2, σ_{sp}^2, and σ_{ps}^2. Primary transmissions are regulated by a randomized Bernoulli process with parameter r_p. Secondary transmissions are regulated by another independent Bernoulli random process with parameter r_s. It is assumed that BSs always have information ready to be transmitted to the users (full queue). Now, consider that primary and secondary users receive packets with N symbols, denoted here, respectively, by $\mathbf{x}_p=\left[x_p(0),\ldots,x_p(N-1)\right]^T$, and $\mathbf{x}_s=\left[x_s(0),\ldots,x_s(N-1)\right]^T$, where $\mathbf{x}_s=(\cdot)^T$ is the vector transpose operator. Let us assume that packets comply with the unitary power constraint $E\left[\mathbf{x}_p^H\mathbf{x}_p\right]=1$ and $E\left[\mathbf{x}_s^H\mathbf{x}_s\right]=1$, where $E[\cdot]$ is the statistical average operator. The signal received by primary and secondary users in the absence of interference can be written, respectively, as follows:

$$\mathbf{y}_p=h_p\mathbf{x}_p+\mathbf{v}_p \qquad \text{and} \qquad \mathbf{y}_s=h_s\mathbf{x}_s+\mathbf{v}_s$$

where $\mathbf{v}_p=\left[v_p(0),\ldots,v_p(N-1)\right]^T$ and $\mathbf{v}_s=\left[v_s(0),\ldots,v_s(N-1)\right]^T$ are the Gaussian noise vectors modelled as a complex circular Gaussian variables with zero mean and variance σ_v^2. The signals of primary and secondary transmissions in the presence of interference from another user are thus given, respectively, by:

$$\mathbf{y}_{p|s}=h_p\mathbf{x}_p+h_{sp}\mathbf{x}_s+\mathbf{v}_p \qquad \text{and} \qquad \mathbf{y}_{s|p}=h_s\mathbf{x}_s+h_{ps}\mathbf{x}_p+\mathbf{v}_s$$

The instantaneous signal-to-noise ratio (SNR) for primary and secondary transmissions in the absence of interference can be written, respectively, as follows:

$$\gamma_p=\frac{\left|h_p\right|^2}{\sigma_v^2} \text{ and } \gamma_s=\frac{\left|h_s\right|^2}{\sigma_v^2}$$

Finally, the instantaneous signal-to-interference-plus-noise ratio (SINR) for primary and secondary transmissions in the presence of interference can be written, respectively, as:

$$\gamma_{p|s} = \frac{|h_p|^2}{\sigma_v^2 + |h_{sp}|^2} \quad \text{and} \quad \gamma_{s|p} = \frac{|h_s|^2}{\sigma_v^2 + |h_{ps}|^2}$$

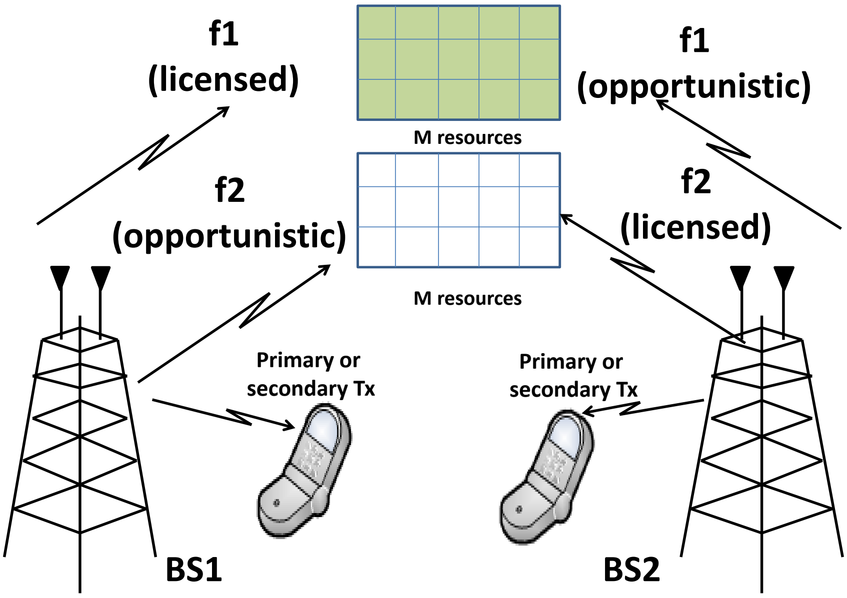

Figure 5-12: System deployment scenario.

5.4.3.2 Packet reception model

This section presents the packet reception model for primary and secondary transmissions in the absence or presence of interference. The correct reception probability of packet transmissions towards a primary and secondary user in the absence of interference, denoted here, respectively, by q_p and q_s, are defined as the probability that the instantaneous SNRs surpass a reception threshold β:

$$q_p = \Pr\{\gamma_p \geq \beta\} \quad \text{and} \quad q_s = \Pr\{\gamma_s \geq \beta\}$$

Since all channels are Rayleigh-distributed, it can be proved that the SNRs are exponentially distributed with parameters given, respectively, by $E[\gamma_p] = \hat{\gamma}_p = \frac{\sigma_p^2}{\sigma_v^2}$ and $E[\gamma_s] = \hat{\gamma}_s = \frac{\sigma_s^2}{\sigma_v^2}$. Therefore, the reception probabilities can be written as the complementary cumulative distribution function (CCDF) of the exponential distribution valued at β, which can be written as:

$$q_p = e^{-\frac{\beta}{\hat{\gamma}_p}} \quad \text{and} \quad q_s = e^{-\frac{\beta}{\hat{\gamma}_s}}$$

The correct reception probabilities of packet transmissions towards primary and secondary users in the presence of interference, denoted here, respectively, by $q_{p|s}$ and $q_{s|p}$, are defined as the probability that the instantaneous SINRs surpass a reception threshold β:

$$q_{p|s} = \Pr\{\gamma_{p|s} \geq \beta\} \quad \text{and} \quad q_{s|p} = \Pr\{\gamma_{s|p} \geq \beta\}$$

By substituting the SINR in the previous expression for the correct reception probability $q_{p|s}$, we obtain:

$$q_{p|s} = \Pr\left\{ \frac{|h_p|^2}{\sigma_v^2 + |h_{sp}|^2} \geq \beta \right\}$$

which can be rewritten as follows

$$q_{p|s} = \Pr\left\{ |h_p|^2 \geq \beta\left(\sigma_v^2 + |h_{sp}|^2\right) \right\} = \Pr\left\{ |h_p|^2 - \beta|h_{sp}|^2 \geq \beta\sigma_v^2 \right\}$$

By using the change of variables $z_p = |h_p|^2$, $u_p = -\beta|h_{sp}|^2$, and $w_p = z_p + u_p$, then the previous expression the previous equation becomes:

$$q_{p|s} = \Pr\left\{ z_p + u_p > \beta\sigma_v^2 \right\} = \Pr\left\{ w_p > \beta\sigma_v^2 \right\}$$

Therefore, it is now possible to calculate the statistics of w_p to derive an analytical expression for $q_{p|s}$. Let us first consider that the probability density and the characteristic functions of z_p and u_p are given, respectively, by

$$f_{z_p}(z_p) = \frac{1}{\hat{z}_p} e^{-\frac{z_p}{\hat{z}_p}} \quad \text{and} \quad \Psi_{z_p}(i\omega) = \frac{1}{1 - i\omega\hat{z}_p}, \quad z_p > 0$$

where $\hat{z}_p = E[z_p] = \sigma_p^2$ and

$$f_{u_p}(u_p) = \frac{1}{\hat{u}_p} e^{-\frac{u_p}{\hat{u}_p}} \quad \text{and} \quad \Psi_{u_p}(i\omega) = \frac{1}{1 - i\omega\hat{u}_p}, \quad u_p < 0$$

where $\hat{u}_p = E[u_p] = \beta\sigma_p^2$. Since z_p and u_p are statistically independent, the characteristic function of their sum $w_p = z_p + u_p$ is given by the product of their individual characteristic functions:

$$\Psi_{w_p}(i\omega) = \Psi_{z_p}(i\omega) \qquad \Psi_{u_p}(i\omega) = \frac{1}{1 - i\omega\hat{z}_p} \frac{1}{1 + i\omega\hat{u}_p}$$

which can be rewritten by partial fraction expansion (PFE) as:

$$\Psi_{w_p}(i\omega) = \Psi_{z_p}(i\omega) \qquad \Psi_{u_p}(i\omega) = \frac{A_p}{1 - i\omega\hat{z}_p} + \frac{B_p}{1 + i\omega\hat{u}_p}$$

where $A_p = \left(1 + \frac{\hat{u}_p}{\hat{z}_p}\right)^{-1}$ and $B_p = \left(1 + \frac{\hat{z}_p}{\hat{u}_p}\right)^{-1}$.

The back-transform of the previous expression provides a CCDF of w_p given by:

$$1 - F_{w_p}(w_p) = A_p e^{-\frac{w_p}{\hat{z}_p}}$$

Finally, the correct reception probability of a primary transmission in the presence of interference from a secondary user can be calculated as the CCDF in the previous expression valued at $w_p = \beta\sigma_v^2$:

$$q_{p|s} = 1 - F_{w_p}\left(\beta\sigma_v^2\right) = A_p e^{-\frac{\beta\sigma_v^2}{\hat{z}_p}}$$

Following the lines of the derivation of $q_{p|s}$, the expression for correct reception probability of a secondary transmission in the presence of interference from a primary user is given by:

$$q_{s|p} = A_s e^{-\frac{\beta\sigma_v^2}{\hat{z}_s}}$$

where $A_s = \left(1 + \frac{\hat{u}_s}{\hat{z}_s}\right)^{-1}$, $\hat{z}_s = E[z_s] = \sigma_s^2$, $\hat{u}_s = E[u_s] = \beta\sigma_s^2$

5.4.3.3 Throughput and Return-risk regions

5.4.3.3.1 Throughput region

The main network performance metric to be used in this section is throughput, which can be defined as the long-term ratio of total number of correctly transmitted packets to the total number of time-slots used in the measurement. In our context, the total throughput for primary user transmissions can be calculated as the summation of all possible cases of correct packet reception either in the absence of interference from secondary transmissions with probability $r_p\bar{r}_s$, where $\overline{(\cdot)} = 1 - (\cdot)$, or in the presence of such interference with probability $r_p r_s$. This can be mathematically written as:

$$T_p = r_p \bar{r}_s q_p + r_p r_s q_{p|s}$$

which, for convenience, can be rewritten as:

$$T_p = r_p q_p - r_p r_s Q_p$$

where $Q_p = q_p - q_{p|s}$. Similarly, for secondary user transmissions we obtain:

$$T_s = r_s q_s - r_p r_s Q_s$$

where $Q_s = q_s - q_{s|p}$. Let us now define the concept of throughput region. For this purpose, let $\mathbf{T} = [T_p \quad T_s]^T$ be the vector of stacked throughput values of primary and secondary users, and $\mathbf{r} = [r_p \quad r_s]^T$ be the vector of stacked transmission probabilities. The throughput region C_T is the union over all possible realizations of transmission probabilities:

$$C_T = \left\{\tilde{\mathbf{T}} \mid \tilde{T}_p = T_p(\mathbf{r}), \tilde{T}_s = T_s(\mathbf{r}), 0 \le r_p \le 1, 0 \le r_s \le 1\right\}$$

which can be simply considered as the region where all the possible values of user throughput can be found.

5.4.3.3.2 Return-risk region

The main objective of this section is to optimize the transmission parameters of primary and secondary users according to network and economic metrics. To achieve this goal, we will borrow concepts from the theory of multi-objective portfolio optimization, which is commonly used in the fields of Economics and Finance theory. In a financial portfolio optimization problem, a portion of each financial asset from a given set of assets has to be optimized so as to maximize revenue or return and simultaneously control or minimize the risk or variance of the return. In this section we will consider each correct packet transmission as a financial asset. Therefore we will attempt to optimize the transmission probabilities of primary and secondary transmissions in the same was as optimum allocation weights for the different assets are calculated in a financial portfolio optimization problem. For this purpose, let us define some economic metrics such as the instantaneous return per correctly transmitted packet in primary mode as p_p, and in secondary mode as p_s. Since packet reception is a random process, let us define the instantaneous return of a primary user in the absence of interference as follows:

$$\text{Instantaneous return primary (no interference)} = R_p = p_p t_p$$

where t_p is a binary random variable that takes the value of $t_p = 1$ if the packet was correctly transmitted and the value of $t_p = 0$ if the packet was incorrectly transmitted. By averaging the previous expression we obtain:

$$\text{Average return primary (no interference)} = E[R_p] = \hat{R}_p = E[p_p]\Pr\{t_p = 1\} = \hat{p}_p \Pr\{\gamma_p \geq \beta\}$$
$$= \hat{p}_p q_p$$

Following the lines of the previous derivation, the remaining expressions for average return of primary and secondary packet transmissions in the absence and presence of interference can be written, respectively, as follows:

$$\hat{R}_s = \hat{p}_s q_s, \quad \hat{R}_{p|s} = \hat{p}_p q_{p|s}, \text{ and } \hat{R}_{s|p} = \hat{p}_s q_{s|p}$$

Let us now calculate the risk of a primary packet transmission in the absence of interference as the variance of the return:

$$\text{Risk primary (no interference)} = S_p = E\left[\left(R_p - \hat{R}_p\right)^2\right] = E\left[R_p^2\right] - \hat{R}_p^2$$
$$E\left[p_p^2\right]\Pr\{t_p = 1\} - \left(\hat{p}_p q_p\right)^2 = E\left[p_p^2\right]\Pr\{\gamma_p \geq \beta\} - \left(\hat{p}_p q_p\right)^2 E\left[p_p^2\right]q_p - \left(\hat{p}_p q_p\right)^2$$

Note that the terms $E\left[p_p^2\right]$ and $E\left[p_p\right] = \hat{p}_p$ are related to the financial fluctuations of the transmissions in the primary band, whereas the terms q_p and q_p^2 are related to the fluctuations of the channel conditions. Now, both types of fluctuations usually occur in different time-scales. Since we are interested in short-term fluctuations, particularly in the timescale of radio resource allocation schemes, it would be possible to ignore the terms related to financial fluctuations. However, in this section these terms will not be dropped as they will be used for assigning different priority levels to primary and/or secondary transmissions, thus allowing us to control in a better manner the proposed short-term radio resource allocation schemes. Following the lines of the derivation of S_p, the remaining expressions for the risk of primary and secondary packet transmissions in the absence and presence of interference can be written, respectively, as follows:

$$\hat{S}_s = E\left[p_s^2\right]q_s - \left(\hat{p}_s q_s\right)^2, \quad \hat{S}_{p|s} = E\left[p_p^2\right]q_{p|s} - \left(\hat{p}_p q_p\right)^2, \quad \text{and } \hat{S}_{s|p} = E\left[p_s^2\right]q_{s|p} - \left(\hat{p}_s q_{s|p}\right)^2$$

Finally, the total average return combining primary and secondary transmissions, denoted here by $\hat{R}$, can be calculated as the summation of all possible cases either in the absence of interference (with probability $r_p\bar{r}_s$ for primary transmissions and with probability $r_s\bar{r}_p$ for secondary transmissions) or in the presence of such interference with probability $r_s r_p$. This can be mathematically written as follows:

$$\text{Total average return} = \hat{R} = r_p\bar{r}_s\hat{R}_p + r_s\bar{r}_p\hat{R}_s + r_p r_s\left(\hat{R}_{p|s} + \hat{R}_{s|p}\right)$$

which for convenience can be rewritten as:

$$\hat{R} = r_p\hat{R}_p + r_s\hat{R}_s - r_p r_s\hat{U}_R$$

where $\hat{U}_R = \hat{R}_p + \hat{R}_s - \hat{R}_{p|s} - \hat{R}_{s|p}$. Using a similar approach, the total risk, denoted here by $\hat{S}$, can be calculated as follows:

$$\text{Total risk} = \hat{S} = r_p\bar{r}_s\hat{S}_p + r_s\bar{r}_p\hat{S}_s + r_p r_s\left(\hat{S}_{p|s} + \hat{S}_{s|p}\right)$$

which for convenience can be rewritten as

$$\hat{S} = r_p\hat{S}_p + r_s\hat{S}_s - r_p r_s\hat{U}_S$$

where $\hat{U}_S = \hat{S}_p + \hat{S}_s - \hat{S}_{p|s} - \hat{S}_{s|p}$. Having defined the total average return and the total risk, let us now define the concept of return-risk region. For this purpose, let $\mathbf{M} = \left[\hat{R} \quad \hat{S}\right]^T$ be the vector of stacked return and risk values. The return-risk region C_M is the union over all possible realizations of transmission probabilities:

$$C_M = \left\{\tilde{\mathbf{M}} \mid \tilde{R} = \hat{R}(\mathbf{r}), \tilde{S} = \hat{S}(\mathbf{r}), 0 \le r_p \le 1, 0 \le r_s \le 1\right\}$$

which can be simply considered as the region where all the possible values of return and risk can be found.

5.4.3.4 Optimum transmission policies

5.4.3.4.1 Throughput-region optimum

To derive the boundaries of the throughput region, a multi-objective optimization is here proposed, where all T's can be simultaneously optimized. For simplicity, let us address this optimization by maximizing an individual throughput function to a fixed constraint on the remaining throughput function:

$$\mathbf{r}_{\text{opt}} = \arg\max_{\mathbf{r}} T_p \quad \text{subject to } T_s = G_s$$

where G_s is the limit or bound on the constraint of the throughput of the secondary user. By changing the value of the constraint G_s it is possible to explore the boundaries of the throughput region, which can also be regarded as the Pareto optimal trade-off curve of the multi-objective optimization problem. The solution to this optimization problem can be proved, by using the method of Lagrange multipliers, to be equivalent to setting the following Jacobian determinant to zero:

$$\left|\mathbf{J}_t\right| = 0$$

where $|\cdot|$ denotes the determinant operator and $\mathbf{J}_t$ is the Jacobian matrix with elements given by:

$$\mathbf{J}_t = \begin{pmatrix} \frac{\partial T_p}{\partial r_p} = -r_s Q_p + q_p & \frac{\partial T_s}{\partial r_p} = -r_s Q_s \\ \frac{\partial T_p}{\partial r_s} = -r_p Q_p & \frac{\partial T_s}{\partial r_s} = -r_p Q_s + q_s \end{pmatrix}$$

By substituting these expressions back in the Jacobian determinant we obtain the following expression for the optimum transmission probabilities of primary and secondary packet transmissions

$$r_p \frac{Q_s}{q_s} + r_s \frac{Q_p}{qp} = 1$$

This expression and the throughput expressions in provide a parametric form (in terms of the transmissions probabilities) of the boundary of the throughput region. To obtain a non-parametric form, it is possible to obtain r_p and r_s from the previous expression and substitute them back in the expressions for throughput to obtain: $r_p = \sqrt{\frac{T_p q_s}{Q_s q_p}}$ and $r_s = \sqrt{\frac{T_s q_p}{Q_p q_s}}$. If these expressions are substituted back in the previous expression we obtain:

$$\sqrt{\frac{T_p Q_s}{q_s q_p}} + \sqrt{\frac{T_s Q_p}{q_s q_p}} = 1$$

5.4.3.4.2 Return-risk optimum: Financial Portfolio optimization

To derive the boundaries of the return-risk region, a multi-objective portfolio optimization is here proposed, where both return and risk functions can be simultaneously optimized. For simplicity, let us address this optimization by maximizing the return subject to a fixed constraint on the risk value:

$$\mathbf{r}_{\text{opt}} = \arg\max \hat{R} \quad \text{subject to } \hat{S} = \hat{S}_0$$

where $\hat{S}_0$ is the limit or bound on the risk function. By changing the value of the constraint $\hat{S}_0$, it is possible to explore the boundaries of the return-risk region, which can also be regarded as the Pareto optimal trade-off curve of the multi-objective optimization problem. The solution to this optimization problem can be proved to be equivalent to setting the following Jacobian determinant to zero:

$$|\mathbf{J}_e| = 0$$

where $|\mathbf{J}_e| = 0$ is the Jacobian matrix with elements given by:

$$\mathbf{J}_e = \begin{pmatrix} \frac{\partial \hat{R}}{\partial r_p} = -r_s \hat{U}_R + \hat{R}_p & \frac{\partial \hat{R}}{\partial r_s} = -r_p \hat{U}_R + \hat{R}_s \\ \frac{\partial \hat{S}}{\partial r_p} = -r_s \hat{U}_s + \hat{S}_p & \frac{\partial \hat{S}}{\partial r_p} = -r_p \hat{U}_s + \hat{S}_s \end{pmatrix}$$

By substituting these expressions back in the Jacobian determinant, we obtain the following expression for the optimum transmission probabilities of primary and secondary packet transmissions:

$$r_p \left(\hat{R}_p \hat{U}_s - \hat{U}_R \hat{S}_p \right) + r_s \left(\hat{U}_R \hat{S}_s - \hat{U}_s \hat{R}_s \right) = \hat{S}_p \hat{R}_s - \hat{R}_p \hat{S}_s$$

This expression together with the expressions for return and risk provide a parametric form (in terms of the transmissions probabilities) of the boundary of the return-risk region.

5.4.3.4.3 Optimum transmission policy with fairness

Having studied the boundaries of the two trade-off regions under investigation, it is clear that maximum throughput or economic performance comes at the expense of performance degradation of either primary or secondary users, particularly in scenarios with high interference. To address this issue, let us now derive a transmission policy that maximizes return and controls the risk, but which is also able to ensure a given throughput performance (quality of service) to both primary and secondary users. This can be expressed as follows:

$$\mathbf{r}_{\text{opt}} = \arg\max_{\mathbf{r}} \hat{R} \text{ subject to } \hat{S} \le \hat{S}_0,\ T_p \ge \alpha_p q_p,\ T_s \ge \alpha_s q_s$$

where α_p and α_s are the ratio constraints on primary and secondary throughput performance, respectively. In the following derivation we assume that the inequality constraints $T_p \ge \alpha_p q_p$ and $T_s \ge \alpha_s q_s$ can be simultaneously achieved. Let us now attempt to optimize the return in by considering the equality constraints $T_p = \alpha_p q_p$ and $T_s = \alpha_s q_s$, and the inequality constraint $\hat{S} \le \hat{S}_0$. This optimization can be proved to be identical to the optimization problem for the boundary of the throughput region previously derived. If the throughput region solution does not comply with this inequality constraint, the solution boils down to finding the intersection of one of the equality constraints $T_p = \alpha_p q_p$ and $T_s = \alpha_s q_s$, with the equality constraint $\hat{S} = \hat{S}_0$. The solutions for these two intersection points are given by the solutions for the following equations:

$$r_p^2\left(Q_p\hat{S}_p - \hat{U}_s q_p\right) + r_p\left(\hat{S}_s\hat{U}_s\alpha_p q_p - Q_p\hat{S}_0\right) - \alpha_p q_p \hat{S}_s = 0$$

$$r_s^2\left(Q_p\hat{S}_s\right) + r_s\left(\hat{S}_s\hat{U}_s\alpha_p q_p - Q_p\hat{S}_0\right) + q_p\hat{S}_s - \alpha_p q_p \hat{S}_p = 0$$

for the intersection of $\hat{S} = \hat{S}_0$ and $T_p = \alpha_p q_p$, and

$$r_s^2\left(Q_s\hat{S}_s - \hat{U}_s q_s\right) + r_s\left(q_s\hat{S}_p + \hat{U}_s\alpha_s q_s - Q_s\hat{S}_0\right) - \alpha_s q_s \hat{S}_p = 0$$

$$r_p^2\left(Q_s\hat{S}_p\right) + r_p\left(\hat{U}_s\alpha_s q_s - \hat{S}_p q_s - Q_p\hat{S}_0\right) + q_s\hat{S}_0 - \alpha_s q_s \hat{S}_s = 0$$

for the intersection of $\hat{S} = \hat{S}_0$ and $T_s = \alpha_s q_s$. The maximum of these two solutions is finally selected as the optimum transmission policy.

5.4.3.4.4 Results

This subsection presents the sketches of the two types of trade-off performance regions under the two transmission policies previously described. The results will be divided in two cases: one where the interference between secondary and primary users is high and the second one when this interference is relatively low. Primary users will be modelled with channel parameter of $\sigma_p^2 = 4$, while secondary users will use a parameter of $\sigma_s^2 = 4$. Interference parameters in the high interference scenario will be given by $\sigma_{sp}^2 = 10$ and $\sigma_{ps}^2 = 12$. Interference parameters in the low interference scenario will be given by $\sigma_{sp}^2 = 3$ and $\sigma_{ps}^2 = 2$. Reception threshold is set to a value of $\beta = 1$. Regarding the economic parameters, the secondary transmission will always be considered with a higher risk value with $E\left[p_s^2\right] - \hat{p}_s^2 = 10$, while the risk in primary transmission will be fixed too to a value of

$E\left[p_p^2\right]-\hat{p}_p^2=0.1$. In the case of the average return we will consider two cases: one where the return of the primary transmission is higher than the average return of the secondary transmission ($\hat{p}_p=5$, $\hat{p}_s=1$), and the second case where secondary transmissions experience higher average returns than primary transmissions ($\hat{p}_p=1$, $\hat{p}_s=5$). Figure 5-13 shows the throughput (left) and return-risk regions (right) of the throughput-optimum and the economic-optimum transmission policies in a system with high interference and where return in the primary is higher than the return in the secondary. Note that the throughput region is non-convex and that the projection of the economic-optimum transmission policy does not coincide with the throughout curve. The return-risk region shows a peculiar shape bounded by the points of maximum return with $r_p=1$ and $r_s=0$ and the point with maximum risk given by $r_p=0$ $r_s=1$. Figure 5-14 shows the results of a system with high interference and where return in the primary is lower than the return in the secondary. In comparison with the previous case we can observe that the return-risk region has a more complex shape. This is because now the economic parameters have a conflictive relation: the band with the higher risk has also the higher return. This causes the shape of the region to have complex crossings between the boundary conditions. The point of maximum return is now given by $r_p=0$ and $r_s=1$ and not in $r_p=1$ and $r_s=0$ as in the previous case. This point of maximum return at high values of risk has a conflict with the point of maximum return at low values of risk which is given by $r_p=0$ $r_s=1$. This means that the boundary conditions given by the lines characterized by $r_p=1$ and $r_p=0$ intersect or cross paths at some point the figure, whereas in the previous case of Figure 5-13 this intersection never happens inside the return-risk region. The result of this issue is a complex shape that is actually described as a hyperboloid coming out of the plane of the reader. We can observe too the in general the case in Figure 5-13 achieves higher levels of return and lower values of risk than the case in Figure 5-14. The area of the region is also larger in the first case than in the later one. This fact suggests that in cognitive radio is convenient to have a primary frequency band with higher return than a secondary frequency band. We remind the reader than in all cases we assume risk of the secondary transmission is higher than the risk of primary transmission. Figure 5-13 and Figure 5-14 also show the optimality region for the optimum transmission scheme that ensures a given level of quality of service for both primary and secondary users and thus a given level of fairness between them. The region is displayed in shaded colours both in the throughput and in the return-risk regions. We can observe that in Figure 5-13 the maximum return solution is given by the intersection of the risk equality constraint and the secondary throughput constraint. By contrast, in Figure 5-14 the optimum point is given by the intersection of the risk equality constraint and the primary throughput performance constraint. This follows from the complex shape of the return-risk region as previously discussed. The optimum policy with fairness shows that it is possible to maximize return while controlling risk and also ensuring a given level of quality of service for both types of users even in a high interference scenario. Figure 5-15 shows the results of a system with low interference and where return in the primary is higher than the return in the secondary. Figure 5-16 shows the results of a system with low interference and where return in the primary is lower than the return in the secondary. In the low interference regime we can observe that the shape of the throughput region has become convex. Therefore, any increase of primary or secondary user performance is not as performance expensive as it was in the high interference scenario. Another interesting feature is that both optimum transmission policies to achieve throughput and return-risk regions describe exactly the same region. This means that the boundaries of the throughput and return risk region coincide with the projection of one another in both domains. This result simply suggests that in a low interference scenario operators will be able to maximize revenue, reduce risk and at the same time achieve optimum network performance levels. Now, the optimum transmission policy with fairness shows an interesting solution in Figure 5-16, where the maximum return point is given actually by the intersection of the risk equality constraint with the point of the throughput region, as compared to the case with high interference where it was not possible to achieve maximum throughput region performance. The optimum random transmission policies derived in this section for purposes of spectrum management can be implemented in the QoSMOS architecture in the CM-SM entity. It is assumed that the CM-SM entity will be in charge of the calculation of the optimum transmission probabilities for each available frequency band (licensed and opportunistic) with the help of channel or interference levels information collected by the CM-RM entity. All the results calculate in this section assume knowledge of long-term channel and interference statistics, which means that the interaction between CM-SM and CM-RM can be kept at low levels.

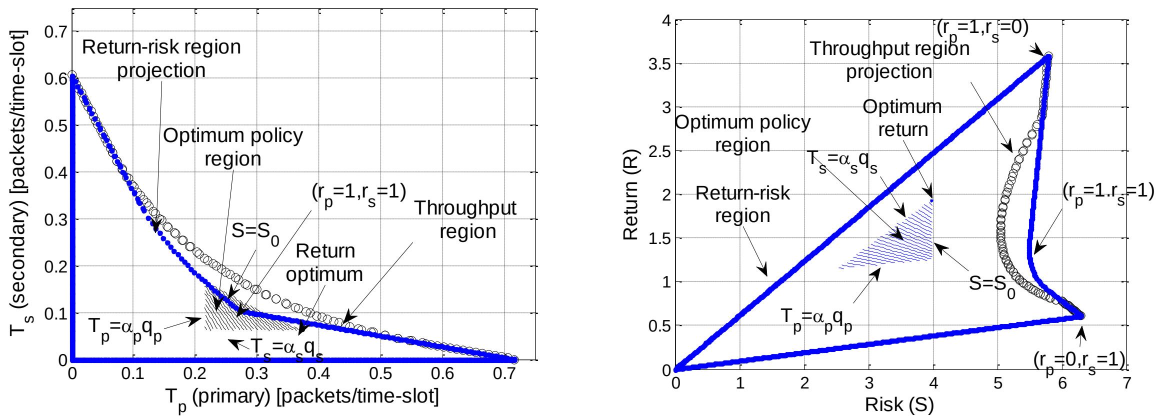

Figure 5-13: Throughput region (left) and return-risk trade-off region (right) of a cognitive radio system with high interference and with a primary frequency band transmission with higher return than secondary frequency band transmission.

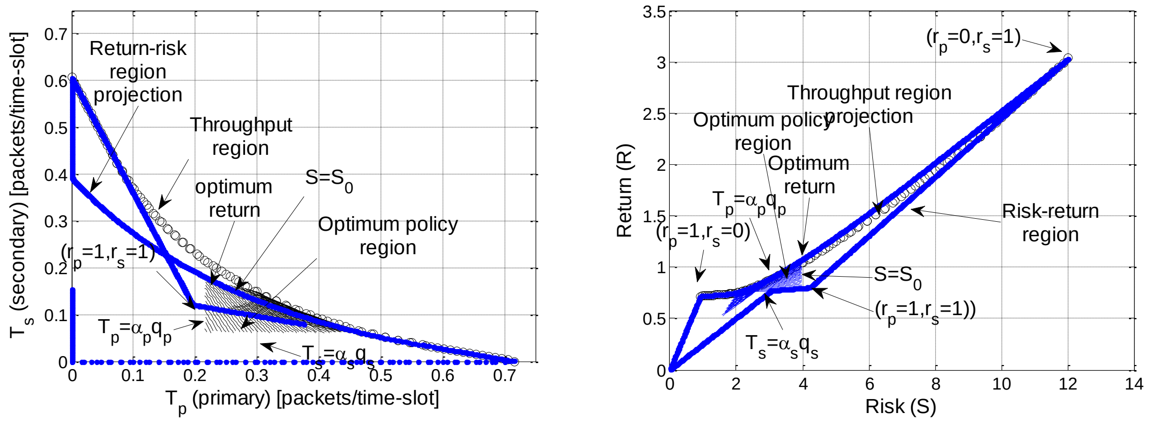

Figure 5-14: Throughput region (left) and return-risk trade-off region (right) of a cognitive radio system with high interference and with a primary frequency band transmission with lower return than secondary frequency band transmission.

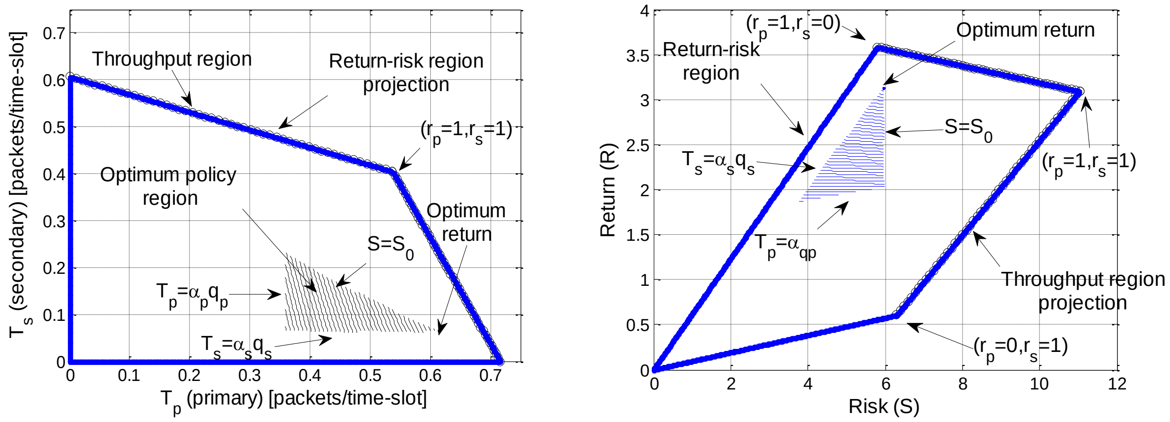

Figure 5-15: Throughput region (left) and return-risk trade-off region (right) of a cognitive radio system with low interference and with a primary frequency band transmission with higher return than secondary frequency band transmission.

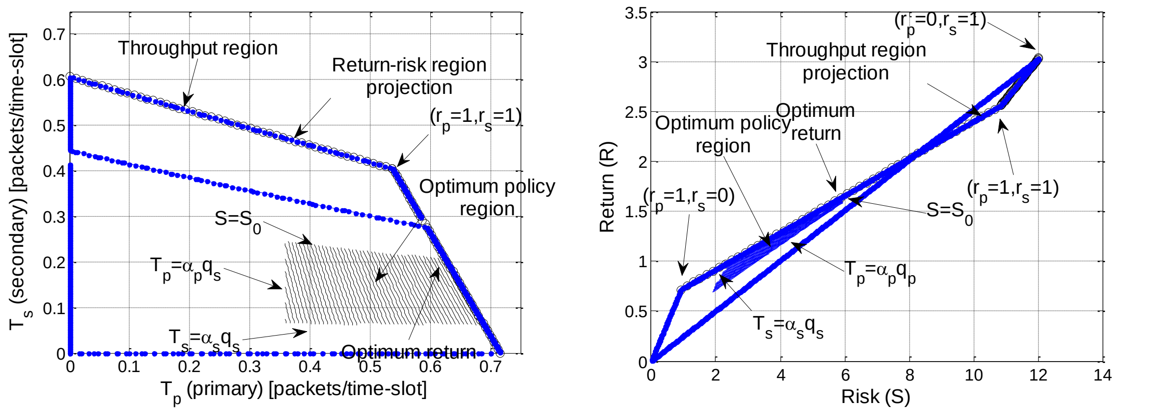

Figure 5-16: Throughput region (left) and return-risk trade-off region (right) of a cognitive radio system with low interference and with a primary frequency band transmission with lower return than secondary frequency band transmission.

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 QoSMOS 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).

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) 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 4.1).
- 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 rules-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.

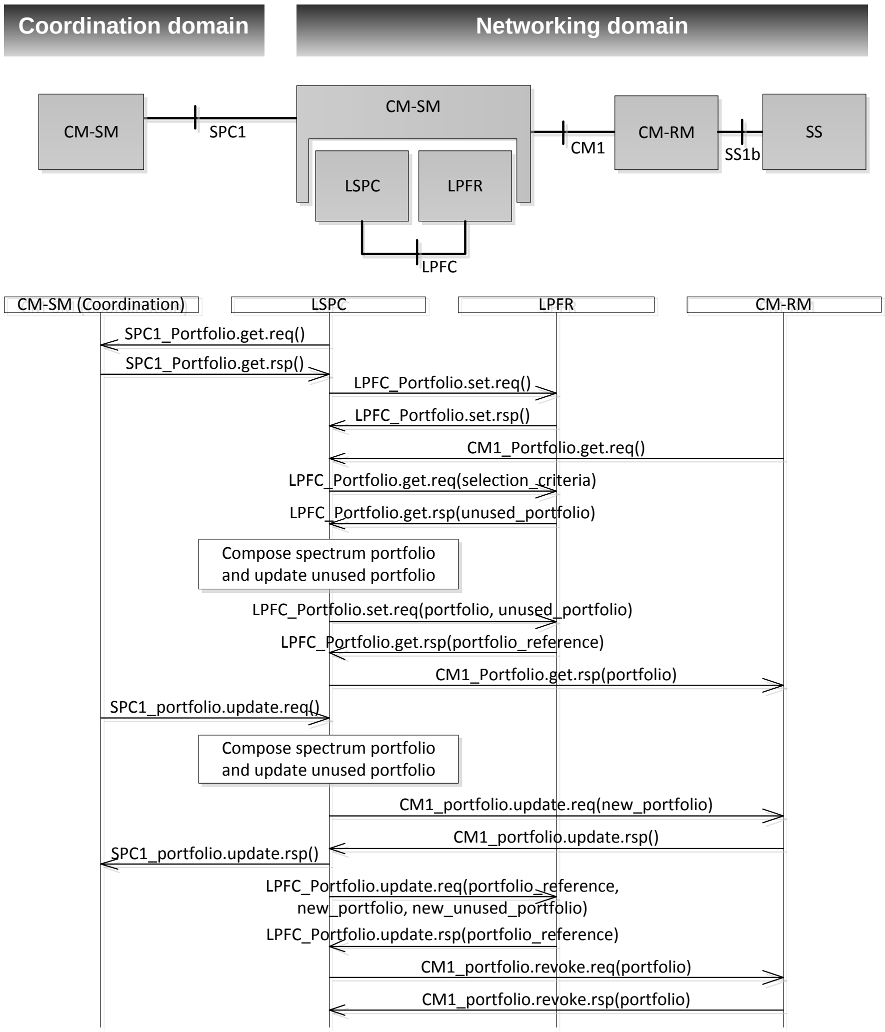

Figure 6-1: Accessing an operator LSPC and sample MSC (responding to a portfolio change request by a coordination domain CM-SM)

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 on-going 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.

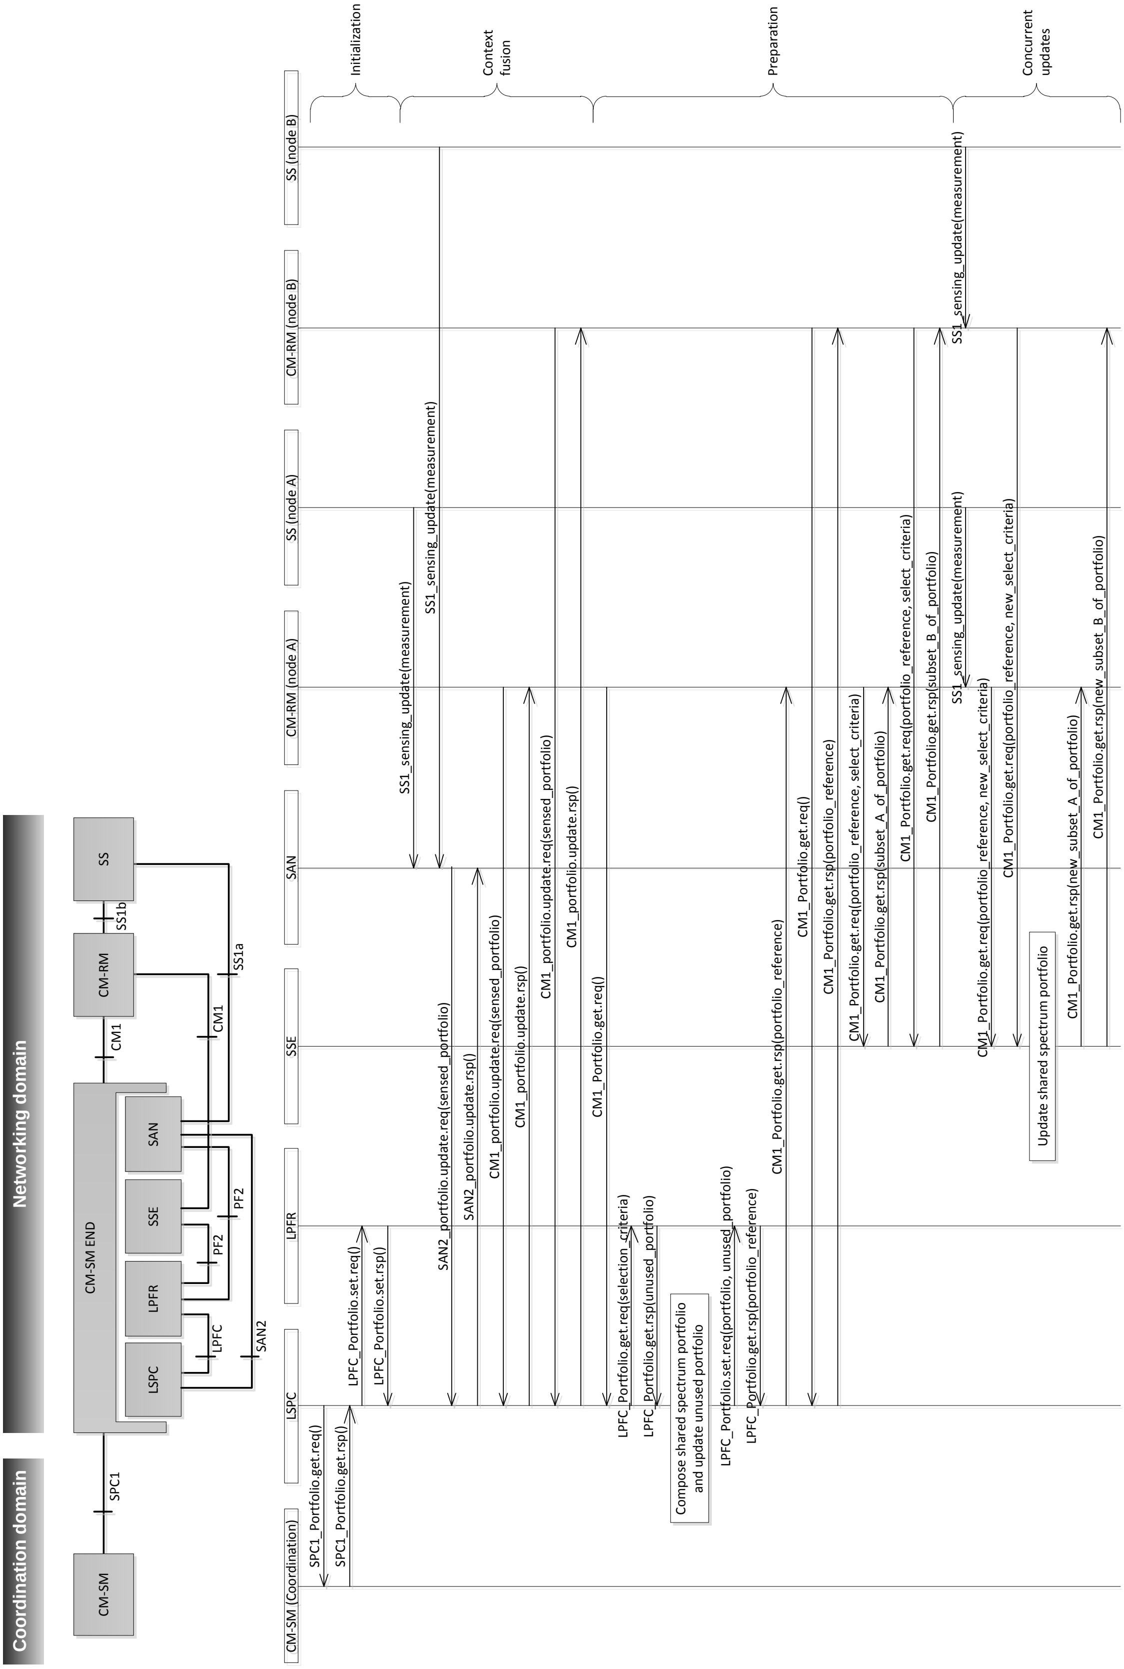

Figure 6-2: Accessing a user-equipment LSPC and sample MSC (fast update of shared spectrum portfolio by requesting the SSE, involving CM-RM, SAN and SS for context acquisition)

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.

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 than 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 User activity modelling for cognitive spectrum management

6.4.1 Introduction

In this section we will provide an overview of the role of long-term user activity in the cognitive spectrum management systems, and the algorithmic approach of the different calculations will be also outlined. The main elements of this method can be found in D6.4; in the following this will be briefly described and extended with algorithms.

The long-term observation the ON/OFF activity of incumbents and opportunistic users gives the large-scale overview of a cognitive system. The activity duration statistic of different users can be applied to build a model, express the distribution of the length of the activity and the activity-free periods. If the probability of incumbents channel utilization is less than the expected opportunistic user activity duration, an opportunity is detected for a cooperative operation. To calculate the required statistics a continuous updating of the activity parameters is necessary. This is the task of the CM/SM system; therefore computing and data storage capability is required in this entity.

An opportunity allows using the spectrum as a shared resource, allocating it both for incumbents and opportunistic users. Two scenarios were discussed, a general radio channel and a wireless IEEE 802.11 computer network, serving as a primary channel.

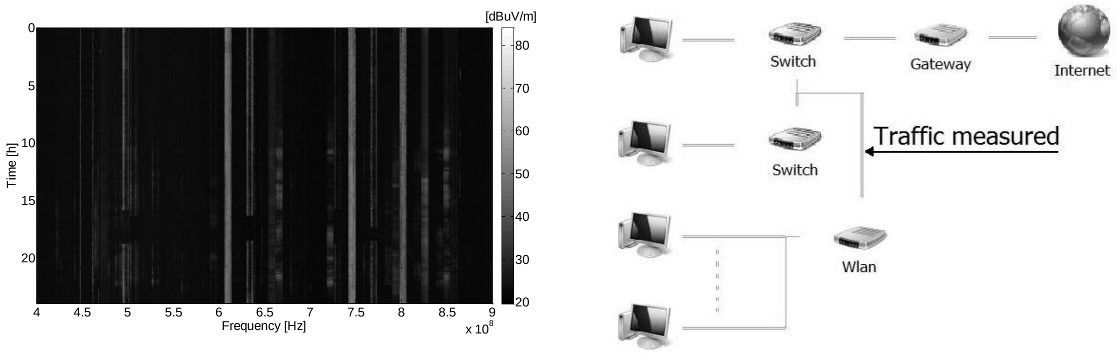

Figure 6-3: General radio channel and IEEE802.11 network measurement

The ON/OFF statistical properties for the channels are derived from real measurements. By using an aggregation method multiple incumbents' activity processes can be encapsulated, resulting a more sophisticated partitioned Markov model, fitted to the aggregated process.

The opportunistic users are considered as Internet users; this model is parameterized from latest internet usage statistics, extended with parameter variability functionality. It will be proved with simulations that the aggregated activity-free length distribution of the incumbents gives the possibility for opportunistic users to join the same network. This method is a general tool; it is intended to implement in the Cognitive Manager (CM) subsystem to support the decision mechanisms of the resource and spectrum management system. The results can be easily adapted for different operational environment, if training data is available for model parameterization.

6.4.2 Relevant scenarios

The main goal is to develop a statistical tool for cognitive systems to support the incumbents/primary and the opportunistic/secondary usage of the spectrum in a shared manner. The long-term observation of activities of incumbents exposes the statistical properties of the connections. Probability chains (Markov chains) will be parameterized with this information, where the different states of the chain represent the current status of the system. The state transition probabilities determine the most probable future status of the system. To model the channel availability, two state ON/OFF models provide a feasible solution.

The Markov modelling of incumbent’s activity presumes on the existence of a Cognitive Manager (CM) in the Cognitive Radio (CR) system that provides sufficient information on the user activity process. On the other hand, spectrum sensing technologies are also applicable in the parameterization process and the algorithms can be built into the CM as functionalities.

In order to model the incumbents ON/OFF activity, we apply two-state, discrete time Markov chain. Besides the simplicity of this model, it was selected as the bursty behaviour of the user activity that can be modelled easily with this probability chain.

The parameterization of the incumbents ON/OFF model is based on measured network traffic, which consists of the data-flow of a WLAN access point. The number of users served by this AP was varying between 1 and 15 during the measurements. Furthermore, with this simple model several users can be simulated by aggregation, and the distribution of the aggregated process provides appropriate data to parameterize a more complex partitioned Markov model. This results a general CR tool to simulate the availability periods of the investigated network.

In order to prove that opportunistic user activity can be fitted in the same channel when there is no activity of incumbents, a variable parameter-set ON/OFF model was developed. This model provides the statistics of the secondary activity duration and confirms that in the investigated environment a cooperative network could be a feasible solution.

6.4.3 Algorithms for long-term user activity modelling

This section describes a Markov model used to describe user activity. In addition the parameterization, algorithms, numerical and simulation results are presented here.

There are two main parts in the model; they area long-term incumbent’s model and an opportunistic user behaviour model. The description of the mathematical background for ON/OFF simulation is followed by the incumbent’s model parameterization methods and also introduces the aggregation scheme. Afterwards the ON/OFF model for opportunistic users is shown, extended with user behaviour dependency and the variable parameter set.

6.4.3.1 ON/OFF Markov-chain model for incumbents activity

In this section the incumbents ON/OFF activity model is detailed.

6.4.3.1.1 The base model

The ON/OFF activity of incumbents can be modelled with a 2-state discrete state, discrete time Markov chain with the transition matrix (6.4-1) where the transition interval was 1 sec:

$$\overline{\overline{P}} = \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix} \tag{6.4-1}$$

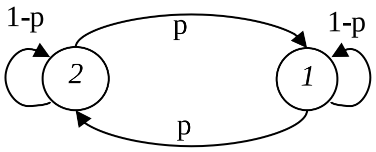

Figure 6-4: 2-state Markov model

State *1* stands for the OFF and state *2* for the ON activity, while p_{ij} are the probabilities to step from state *i* to *j*. The complementary cumulative distribution (CCDF) of the duration of the ON state can be expressed with 6.4-2),

$$F_{ON}^{C}(n) = \lim_{N\to\infty} \sum_{i=n}^{N} p_{21} p_{22}^{i} \quad , \quad n = 1,2,3,\ldots,N \tag{6.4-2}$$

for *n* discrete time slots up to $N\to\infty$. The function $F_{ON}^{C}(n)$ gives the probability that the user activity is ON for duration *n* or longer. Similarly, the CCDF for the OFF state duration is 6.4-3:

$$F_{OFF}^{C}(n) = \lim_{N\to\infty} \sum_{i=n}^{N} p_{12} p_{11}^{j} \quad , \quad n = 1,2,3\ldots N \tag{6.4-3}$$

denoting the probability $F_{OFF}^{C}(n)$ 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. At first let us perform the calculations and generate time series from a hypothetical transition matrix as it follows: $\overline{\overline{P}} = \begin{pmatrix} 0.99999 & 0.00001 \\ 0.01000 & 0.99000 \end{pmatrix}$.

Algorithm: generating ON/OFF time series

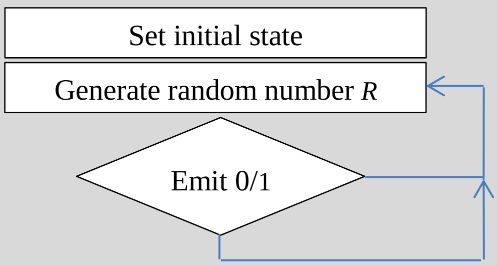

In Figure 6-5 a realization of the ON/OFF time series and the CCDF of the ON length are depicted, as calculated from the synthetic time series and with the use of equation 6.4-2.

.

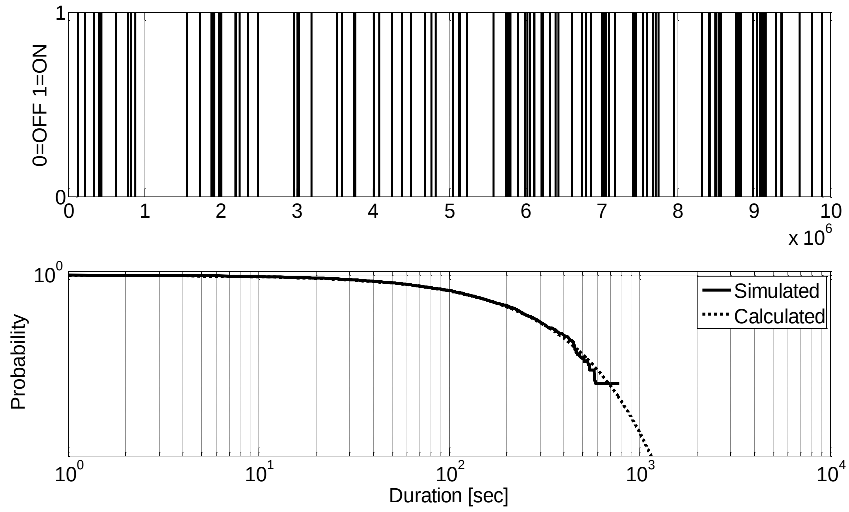

Figure 6-5: ON/OFF time series and the CCDF the ON length

The figure depicts the duration in [sec] resolution for the simulated time series and for the CCDF as well. The length of the simulated time series influences the precision of the simulation. As the number of the long ON events is relatively low, this explains the lower precision of the simulation at high durations.
For the further simulations, with the use of realistic results the Markov model can be parameterized from real measurements as shown in the next sections.

6.4.3.2 Parameterization of the incumbents ON/OFF model - based on spectrum measurement

To determine the transition probabilities for the ON/OFF model, another feasible solution is observing signal strength levels from a spectrum measurement. By scanning and recording the radio band where the incumbents are communicating, valuable information is provided about the user activity as the function of time and frequency.
In 2010 the Budapest University of Technology, in cooperation with the Hungarian National Media and Electronic Communications Authority, performed a 24-hour measurement in a medium size city in Hungary with the parameters listed in **Error! Reference source not found.**:

Table 6-1: Spectrum measurement parameters

Location	Frequency band [MHz]	Resolution [kHz]	Measured parameter
Pécel, Hungary	400-900	30	field strength [dBµV/m]

The measurement was performed with the Rohde-Schwarz ESMB monitoring and test receiver with HE-309 passive broadband receiving dipole antenna. The typical noise figure of the receiver is 9 dB, and the field strength has been detected in dBµV/m units.
This frequency range includes in Hungary several air navigation frequency bands, emergency frequencies for mobile devices and terrestrial television broadcast frequencies.
In order to determine the incumbent activity at any selected frequency, the signal strength can be examined at a specific threshold (a well-defined level above the noise). The duration of the ON/OFF sequence is applied to parameterize a two-state Markov model and to determine the elements of equation 6.4-1.

Algorithm: Modelling the usage of context information, incumbents ON/OFF model

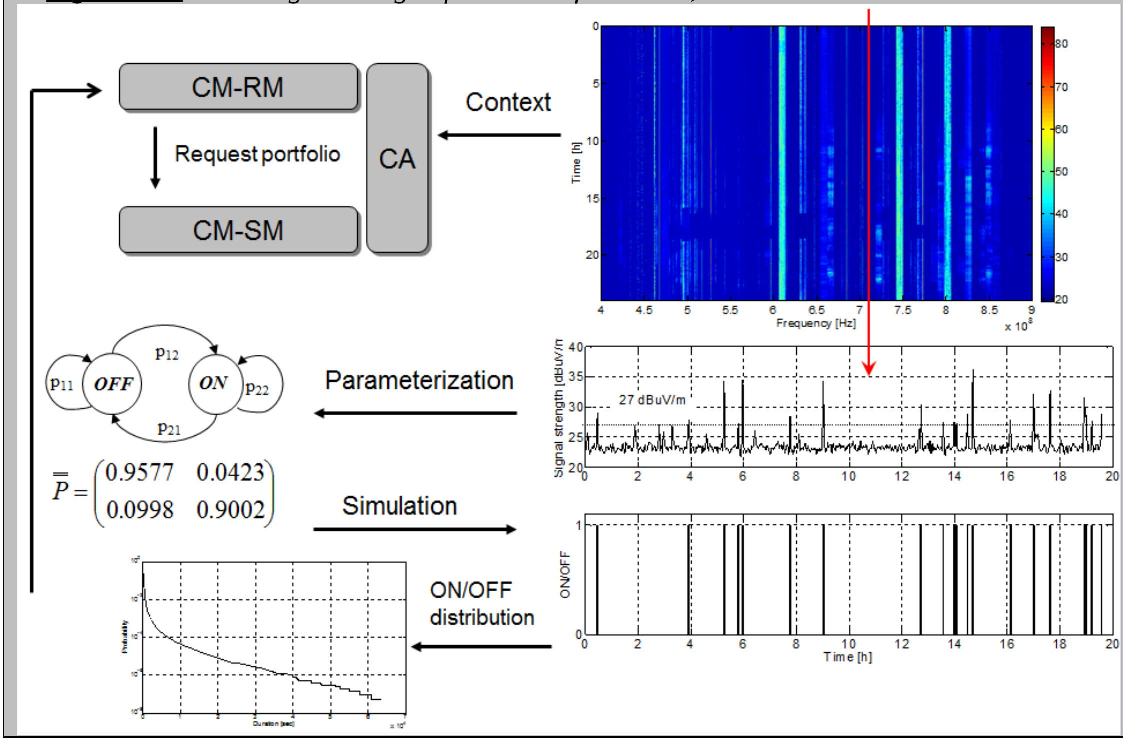

These calculations were created by selecting 433.98 MHz frequency and 27 dBμV/m level as the system noise floor. The ON→OFF transition probability $p_{21}=Pr(S_i+1=0|S_i=1)$ is a conditional probability where S_i is the i'th sample of the ON/OFF time series, while the OFF→ON transition probability is $p_{12}=Pr(S_i+1=1|S_i=0)$.

As the transition matrix in 6.4-1 is a stochastic matrix, p12 and p21 are determining the whole ON/OFF process and the corresponding Markov chain. The values of the transition matrix is $\overline{\overline{P}} = \begin{pmatrix} 0.9577 & 0.04323 \\ 0.0998 & 0.90020 \end{pmatrix}$, according to the measured time series, and it will be applied in further simulations.

Algorithm: determining the transition matrix parameters from measurements

The transition matrix is a stochastic matrix, therefore:

$$\overline{\overline{P}} = \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix} = \begin{pmatrix} 1-p_{12} & p_{12} \\ p_{21} & 1-p_{21} \end{pmatrix} \tag{6.4-4}$$

We can write the state vector z with the steady state probabilities at state n:

$$z(n) = [Z_{OFF}(n) \quad Z_{ON}(n)] = [1 - Z_{ON}(n) \quad Z_{ON}(n)] \tag{6.4-5}$$

The state vector z of the model can be calculated with considering a first order Markov process:

$$z(n+1) = z(n) \cdot \overline{\overline{P}} \tag{6.4-6}$$

By solving this equation (with Gauss-elimination), p_{12} can be calculated:

$$p_{12} = \frac{Z_F}{1 - Z_F} \cdot p_{21} \tag{6.4-7}$$

6.4.3.2.1 Parameterization of the incumbents ON/OFF model - based on Internet traffic measurement

In order to determine the transition probabilities for the ON/OFF model, we performed an analysis of network traffic over an IEEE 802.11 wireless access-point. By scanning and recording the number of data packets as the function of time, valuable information is provided about the incumbent's activity. The location of the measurements was the university campus and it was performed over a busy network node, using the Wireshark ver.1.4.3 network protocol analyser software.

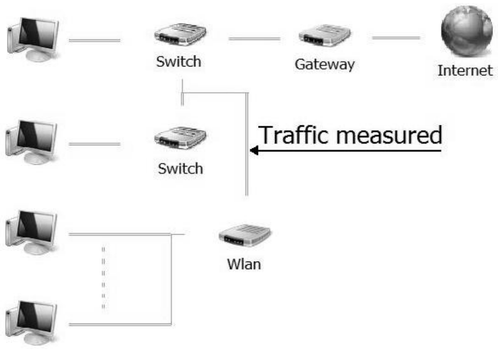

Figure 6-6: Network traffic measurement

In order to determine the incumbent activity, the number of packets was examined at a specific threshold as a two-state process. For the calculations the mean of the traffic speed was selected as a threshold (16 packet/sec). The time periods of the ON/OFF sequence (see Figure 6-7: One hour network traffic and ON/OFF activity) can be applied to parameterize a two-state Markov model and then to determine the elements of equation 6.4-1.

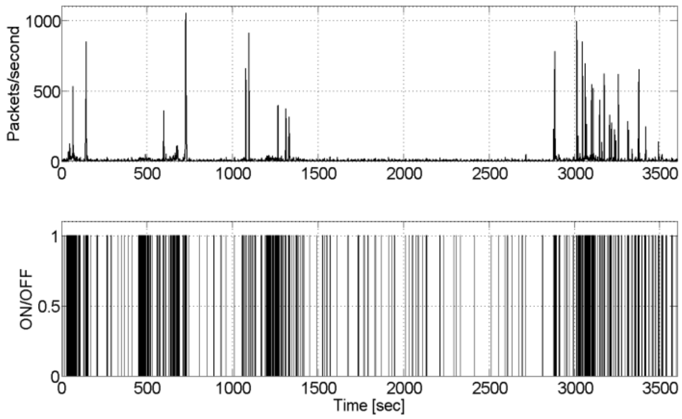

Figure 6-7: One hour network traffic and ON/OFF activity

As an example, the figure was created by selecting the average packet rate (16 packets/sec) as threshold level. Aggregating a number of such ON/OFF incumbent models are used to show the possible existence of opportunistic users in cognitive systems.

The ON→OFF transition probability $p_{21}=Pr(S_i+1=0|S_i=1)$ is a conditional probability where S_i is the i'th sample of the ON/OFF time series, while the OFF→ON transition probability is $p_{12}=Pr(S_i+1=1|S_i=0)$.

As the transition matrix is a stochastic matrix, p_{12} and p_{21} determine the whole ON/OFF process and the corresponding Markov-chain. According to the real measured time series the values of the transition matrix $\overline{\overline{P}} = \begin{pmatrix} 0.8866 & 0.1134 \\ 0.5309 & 0.4691 \end{pmatrix}$, and this is applied during simulations.

Algorithm: *The transition matrix can be determined similarly as it was described in the previous section. The equations 6.4-4 – 6.4-7 were used for the parameterization process.*

6.4.3.2.2 Aggregated incumbents activity

Considering more than one incumbents, an aggregated user activity model is needed.

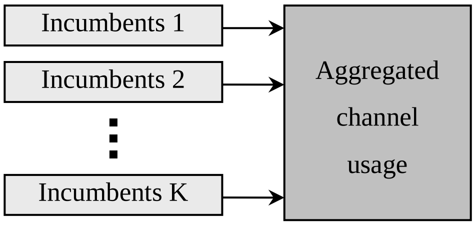

Figure 6-8: Aggregation of incumbents

Let us assume that the capacity of the central service provider is not limited by the number of incumbents. It means that the number of incumbents cannot exceed the capacity limit of the system.
Certainly, as the number of users is increasing, the percentage of the aggregated channel usage will be higher.
The *ON* activity at the aggregated channel is the logical *OR* conjunction between incumbents *1…K*, according to Figure 6-8 In a Cognitive Radio system, besides the statistics of the incumbents, the inter-user activity has also a great importance. The distribution of the OFF duration for aggregated users is a valuable tool for estimating the period, when the radio spectrum is free or unused and therefore available for opportunistic users.
In Figure 6-9 the result of a simulation can be seen, calculated from 10^3 realizations of time series, where each of them consists of 10^7 number of samples. The realization of a time series is an individual set of simulated data.

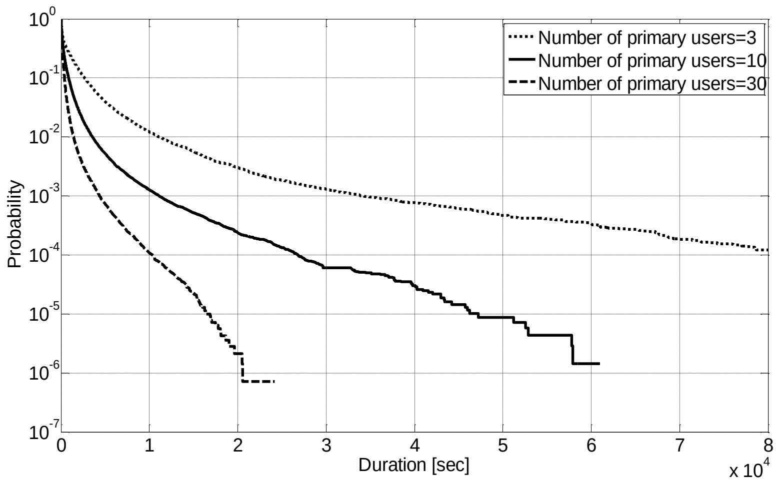

Figure 6-9: Aggregated OFF length CCDF for different number of users (probability of duration longer than the abscissa)

During the simulation the number of concurrent incumbent users was set to 3, 10 and 30, respectively. As the user number is increasing, the probability the OFF state is decreasing, as it was expected. It is worthy of note that according to the individual user Markov-models, the aggregated incumbents user traffic will have also a self-similar behaviour. The meaning of self-similarity in this context is that the process is statistically self-similar: the same statistical properties can be observed at many scales of the stochastic process.

A further result of the simulations was that we were able to estimate the percentage of channel usage. With the given simulation parameters the percentage of ON state for a single user was around 0.1% of total time, whilst the aggregated channel was occupied in 49% of total time, considering 10 concurrent incumbent users. This result shows that in a cognitive system a capacity can be assigned for secondary usage, in conjunction with appropriate predicting technology for spectrum occupancy.
The OFF sequences are appearing in the simulated time series as in a burst structure (see **Error! Reference source not found.**), that can be originated from the primary nature of the ON/OFF Markov model. Fritchman proposed a partitioned discrete state and discrete time Markov model to simulate burst-type errors on digital communication channels. The Fritchman model [FRI67] is an *N*-state Markov-chain with *k* errors in the first partition and *N*-*k* error-free states in the second partition. The distribution function of the state durations has been expressed by Fritchman and proved its applicability for bursty processes. As the bursty behaviour of the investigated channel can be well observed in Figure 6-7, the implementing of the Fritchman model seems to be reasonable.
As the aggregation process does not change the self-similarity behaviour of the single ON/OFF models, a Fritchman-model for the aggregated OFF length sequences is implemented within our work. The number of states in this model depends on the complexity of the modelled phenomena. In practical point of view the number of states is usually below or equal to 5; and can be proved by simulations.

6.4.3.2.3 Markov Model for the OFF process of the aggregated incumbents

The CCDF of simulated OFF sequences for the aggregated process can be applied to parameterize a Fritchman-model, using the gradient-method. This method is based on the linear regression of the log-survivor function (natural logarithm of the survivor/CDF function), intended to model with Markov-chain. The states in a simplified Fritchman-model partitioned in two parts, as it can be seen in **Error! Reference source not found.**, denoted as 4/1 model, a model with 5 states in all.

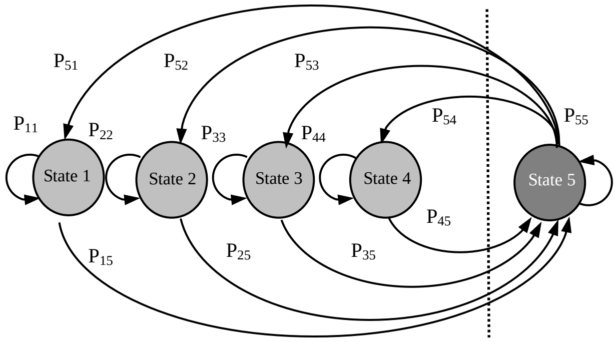

Figure 6-10: A 4/1 Fritchman-model

In our model-adaptation the states in the multi-state partition are represented with different OFF event durations, while the single state in the other partition stands for the ON duration. The CCDF of the 4-state partition can be expressed analytically with equation (6.4-8), where M denotes the number of states, where in this case M=5:

$$F_{OFF}^{C}(n)=\sum_{i=1}^{M-1}\frac{p_{Mi}}{p_{ii}}p_{ii}^{n} \qquad (6.4\text{-}8)$$

The linear regression is based on equation (6.4-9), which leads to the well-known equation of linear:

$$\ln\left(F_{OFF}^{C}(n)_i\right)=\ln\left(\frac{p_{Mi}}{p_{ii}}p_{ii}^{n}\right)=n\ln\left(p_{ii}\right)+\ln\left(\frac{p_{Mi}}{p_{ii}}\right) \qquad (6.4\text{-}9)$$

therefore the slopes and crossings with the vertical axis of the regression lines provide the transition matrix elements of the Fritchman-model.

Algorithm: Transition matrix elements calculation from the regression line parameters

If the regression lines 1-4 are given with $f(n) = a \cdot n + b$, they parameters are applicable to calculate the transition matrix. According to the Fritchman-model properties, only the main diagonal and the last row and column consist of non-zero elements. According to the stochastic matrix features the last column is: $p_{i5}=1-p_{ii}$, similarly $p_{55}=1-p_{51}-p_{52}-p_{53}-p_{54}$.

An example, according to the following figures:

	a	B	$p_{ii}=\exp(a)$	$p_{i5}=1-p_{ii}$	$p_{5i}=p_{ii}\cdot\exp(b)$
Line 1.	-0.04123	-0.339	p_{44}=0.9596	p_{45}=0.0404	p_{54}=0.6837
Line 2.	-0.00129	-2.146	p_{33}=0.9987	p_{35}=0.0013	p_{53}=0.1168
Line 3.	-0.0001559	-4.724	p_{22}=0.9998	p_{25}=0.0002	p_{52}=0.0089
Line 4.	-0.0001114	-5.498	p_{11}=0.9999	p_{15}=0.0001	p_{51}=0.0041

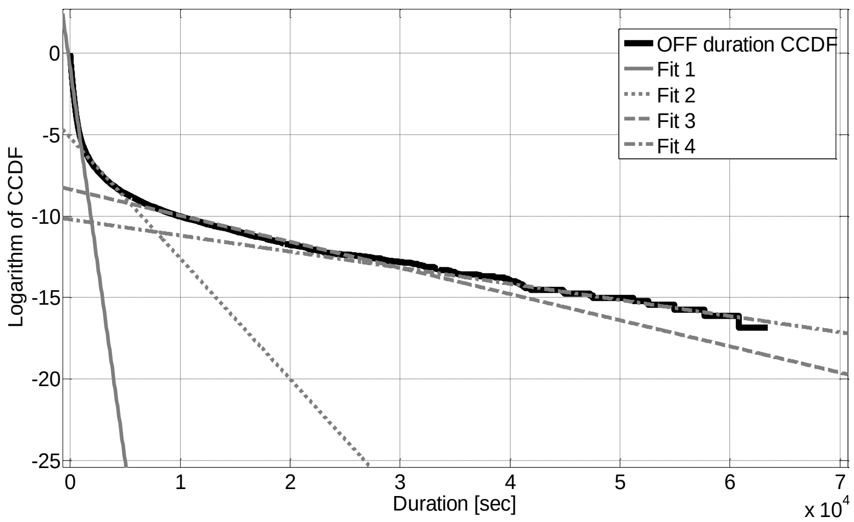

Figure 6-11: Linear regression for model parameterization

The number of states depends on the number of regression lines required to correctly approximate the original log-survivor function; four states were enough in the present case. In **Error! Reference source not found.** the linear regression of the aggregated OFF activity duration can be seen. The regression was performed for the simulation with 10 concurrent users, as it was depicted in Figure 6-9. By applying equation (6.4-9), the transition matrix of the Fritchman-model can be calculated.
With the transition matrix and by using equation (6.4-8) the Fritchman-model is a perfect tool to simulate the aggregated OFF duration distribution for CR systems, as it can be seen in Figure 6-12. This is a prediction tool to estimate the probability of the communication-free periods when multiple users are utilizing the same radio service.

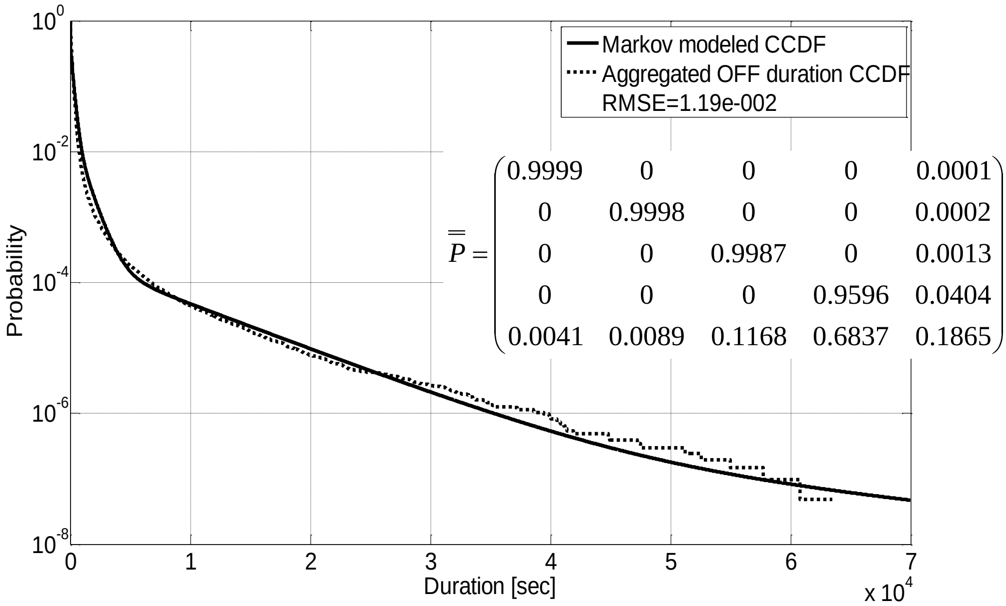

Figure 6-12: Markov modelled OFF duration CCDF

The transition matrix of the model is also given in Figure 6-12. The most important outcome of this work is that a simple partitioned Markov-model is applicable to determine the probability of the activity-free duration in a CR system, where several incumbents are communicating with different ON/OFF activity distributions.

6.4.3.2.4 ON/OFF model for opportunistic user activity

In the cognitive radio environment that we are modelling, the opportunistic users are intending to utilize the same radio spectrum as incumbents, when no main (primary) activity can be foreseen. In our study we specify this CR system for opportunistic users that generates internet traffic (WEB browsing, using chat applications, email send/receive operations, transferring files, etc.). In order to parameterize the ON/OFF model, we selected a different approach than as it was shown in section 6.4.3.2. A representative US internet usage statistics has been used to calculate the model parameters. This statistics was released by Harris Interactive Poll at the end of 2009 [HP09] and it contains the adult online computer use at different locations (office, home, other) and the average online hours/week is also given.
Applying this assessment, the parameter-set has been calculated for the ON/OFF model:

Table 6-2: ON/OFF model parameters for typical opportunistic user

p_{11}	p_{12}	p_{22}	p_{21}
0.999997	0.000003	0.99980	0.00020

The calculation based on the 13 hours/week average internet usage statistics. In order to refine this model, we assumed that the model parameters are varying and following a Normal distribution around their mean values (13 hours/week).
In this manner the variability of the user behaviour can be taken into account during the simulation. The mean of the distribution $\mu=p_{ij}$ is the transition probability, while the variance σ^2 can be calculated from the statistics, taking into account the standard deviation of the weekly internet usage statistics.

Algorithm: Modelling the usage of context information, adaptive opportunistic user activity model

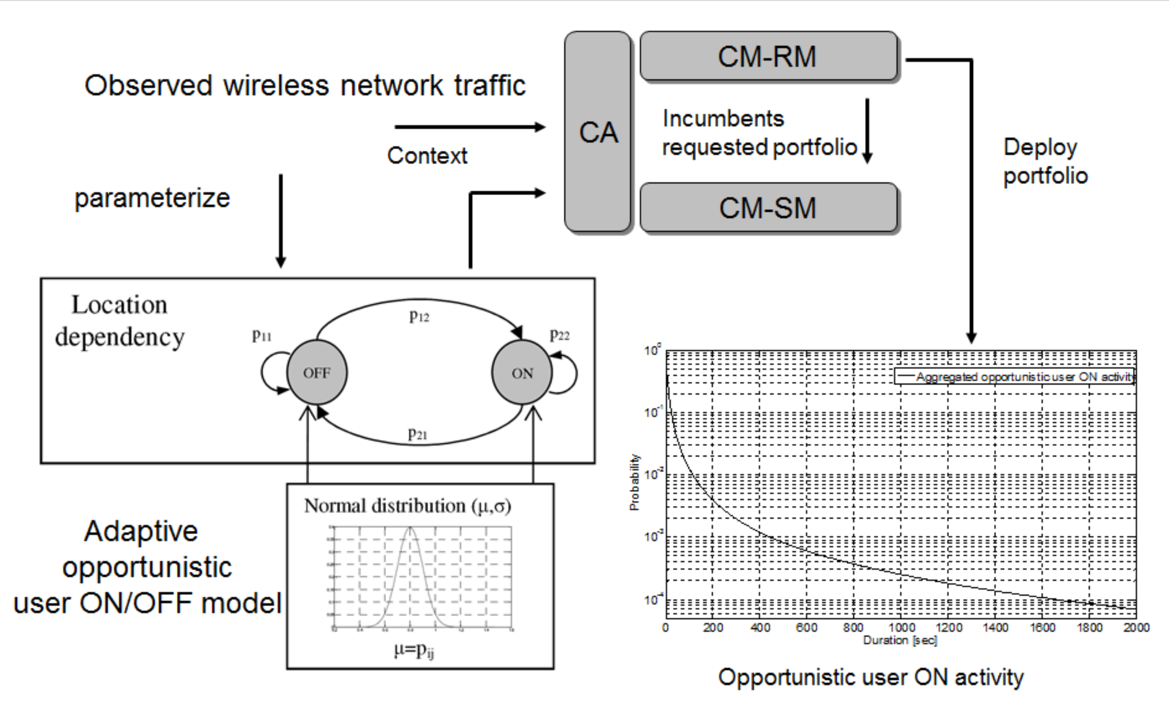

Simulation results can be seen in Figure 6-13. The number of simulated opportunistic user with ON/OFF time series of 10^3 is used, where each of them are consisting of 10^7 numbers of samples. The aggregation schema was performed according to Figure 6-9, assuming 3, 6 and 9 concurrent users with the above described Markov-model.

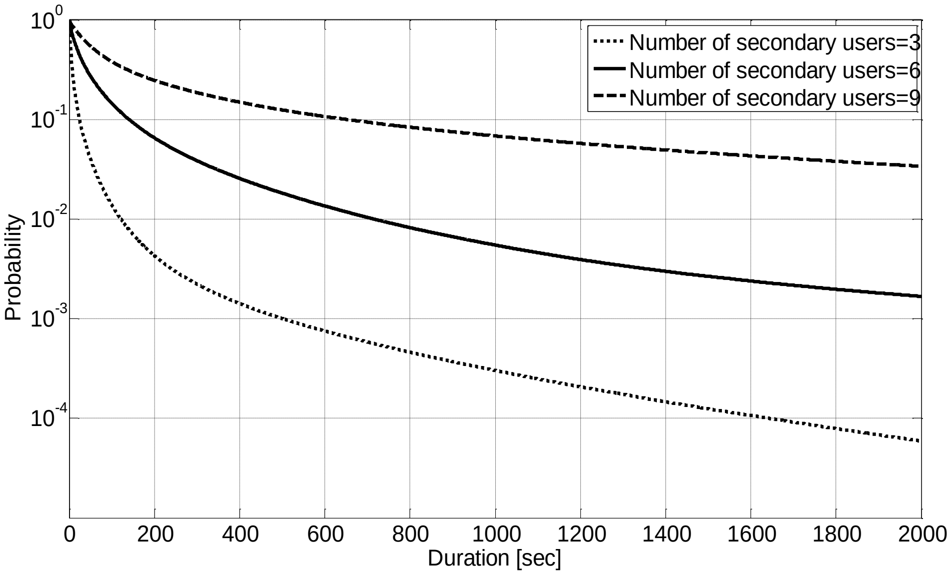

Figure 6-13: CCDF of the aggregated opportunistic user ON activity for different number of users

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. It can be observed, that in case of 3 concurrent opportunistic users the probability of a longer traffic period (more than 30 minutes) is lower than 10^{-4}. The simulated aggregated incumbents OFF duration probability for 10 concurrent incumbents at $2 \cdot 10^3$ seconds (duration in x axis in Figure 6-9) is $1.1 \cdot 10^{-3}$ (y axis in Figure 6-9). This value is one order of magnitude larger than the probability of the opportunistic user ON activity with similar duration. Therefore our simulations demonstrate that in the cognitive radio system considered here, the radio spectrum is not fully utilized by the incumbents; furthermore, the same spectrum can be assigned for secondary internet users.

6.5 Performance evaluation of distributed self-learning self-organized network (SON) spectrum management

6.5.1 Introduction and Solution Approach

One of the major challenges in cognitive radio spectrum management is to find the best suited spectrum portfolio and its power settings for each individual cognitive node. The current/future diverse and heterogeneous cellular telecommunication scenarios require being able to handle and resolve the large interactions and strong couplings between the different cognitive nodes, in addition these highly coupled parameters shall quickly be adapted and optimized to its present situation. In order to handle and resolve this spectrum management issue, an approach for distributed self-learning Self-Organizing-Network (SON) spectrum management has been developed as described in detail in [D6.4] [D6.5]. In the present study, this concept has been implemented in a radio system simulator for cellular wireless networks in order to validate this concept, and to evaluate its performance for cognitive radio as well as to characterize its capabilities and limitations.

6.5.2 Proof of Concept Simulation Scenario

This SON + self-learning concept (presented in [D6.4] and [D6.5]) has been implemented into a strongly simplified LTE-A system simulator which simulates many snapshots with random mobile user placements according to various traffic distributions and load levels. Thereby, the LTE cells could be considered as cognitive nodes for which the best suited parameters and spectrum parts needs to be found. The scheduler of the cells represent the cognitive radio resource manager and operate frequency (physical resource block (PRB)) selective with 10 groups or orthogonal frequency resources, which correspond to 10 different (parts of) frequency bands. In the current implementation, orthogonal physical resource blocks (PRBs) within the same frequency band are taken. The scheduler does also consider the current inter-cell interference on each scheduled resource (PRB) depending on the actual real time load and resource utilisation level in the other interfering cells.

In order to observe the effects as clearly as possible, for this proof of concept study, a small heterogeneous cellular playground was chosen, consisting of one large macro cell with 4 small metro cells inside. The SON algorithm uses the granularity of grouping the spectrum resources into 10 orthogonal resource groups (groups of 5 PRBs; 50 PRBs in total). The SON algorithm can either vary the cell coverage area by modifying the handover thresholds via the LTE-A Bias values in algorithms steps of +/- 1 dB or by modifying the cell transmission power via 9 discrete levels of 1.0, 1.4, 2.0, 2.8, 4.0, 5.6, 10.0, 14.0 and 20 watts. The graphs presented here show example results with modifying the Bias value. Simulation studies with power variations have shown that the general study results and the lessons learned about the SON operation are comparable and similar. To be able to observe and follow the development of the SON parameter adaptations, the presented simulation studies limit the parameter variation to steps of one single value within one simulation cycle. In other words, within each single SON algorithm step, a cell may alter its configuration parameters by +/– one resource group and by +/– 1 dB Bias value. Thereafter, the system simulation is running again to determine the new system performance with that improved parameter set which has been found by the SON algorithm. The x-axis in the Figure 6-14 and Figure 6-15 shows these algorithm iteration sequences of alternating system simulations and distributed SON parameter optimizations.

6.5.3 Simulation Study Results

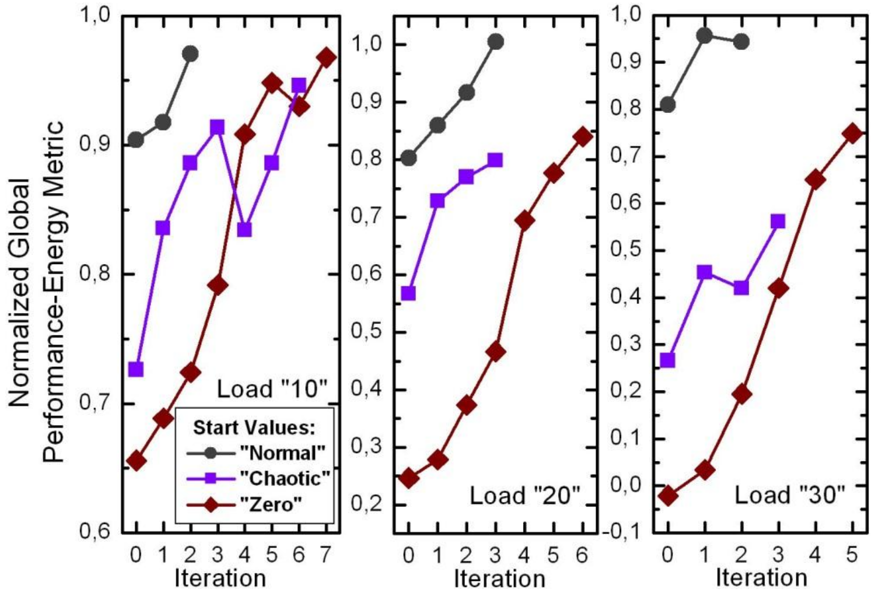

Figure 6-14: Performance increase of the SON-optimized system for several initial starting parameter configurations for different traffic loads

Figure 6-14 shows how the normalized global performance-energy metric is developing with an on-going SON algorithm operation, for different traffic load levels and for different starting configurations. Thereby, ‘Normal’ corresponds to cell configuration parameters as these would be expected by manual configuration as the typical cell parameter settings for these heterogeneous cells. ‘Chaotic’ corresponds to a little bit arbitrary starting values for the cell configuration parameter settings and ‘Zero’ corresponds to that all parameters are set to their starting values (i.e. all resources are allowed to be used and neutral cell borders).

It can be seen that the SON algorithms gradually improve in all situations, for different traffic load levels and for different starting values the performance-energy-metric, which is here taken as

“Metric = Throughput + UserQualityOfService – EnergyConsumption“, (6.5-1)

where all contributions are normalized to achieve a scenario and load independent metric. Starting from any scenario, and starting from any currently present initial parameter settings, the SON + self-learning technique does always gradually improve the cells’ parameter settings. In certain situations of other simulation scenarios, it may occur, that the observed metric-performance decreases a little bit within the first iteration step which is there the consequence of the random start without any starting knowledge of the system; then during the first iteration step, the SON algorithms learn themselves the properties and internal configuration values for their self-learning prediction model, and thanks to this initial self-learning, the model thereafter gradually improves its optimization metric. A more steep performance increase would be achieved when removing the restricting limit, that within one algorithm iteration step, any SON parameter can only be changed by one single parameter step.

Figure 6-15 presents in more detail the algorithm behaviour. The top graph displays, that, after the initial increase pretty good configuration parameter settings were found for the cells and the system performance-energy metric varies in a stable way within a certain range reasonably close to the theoretical maximum, but however does not actually reaches it.

The bottom graph illustrates how the individual contributions towards the total metric are composed ; the alternation of the values is caused by the change between competing alternative sets of parameter values. For understanding and analysing the pure behaviour and to identify occurring effects, this feature is intentionally shown in its pure form; future more complex algorithms could consider the parameter histories and resolve such alternating effects.

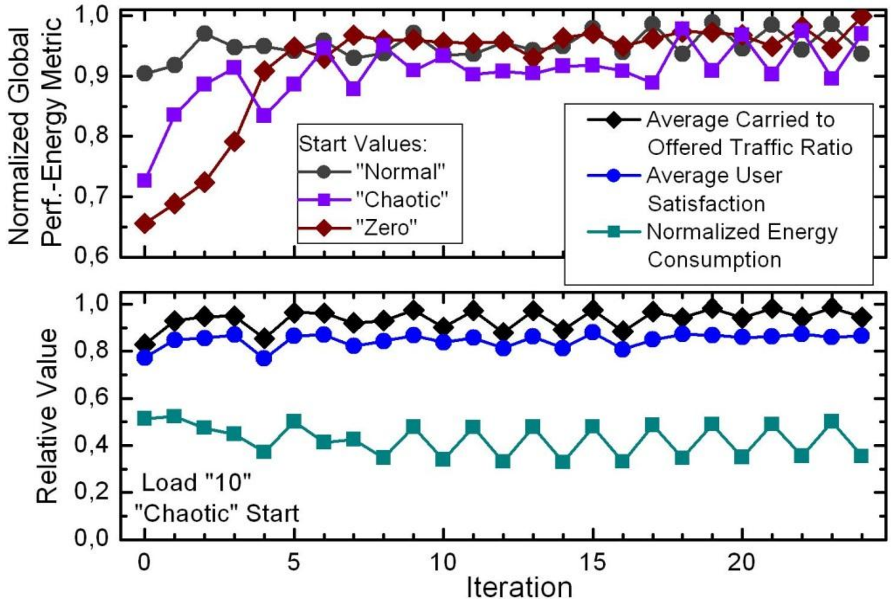

Figure 6-15: Stable operation and convergence of the SON algorithm with on-going simulation duration

6.5.4 Discussion about Capabilities and Limitations

These proof of concept simulation studies validate the SON + self-learning approach for cognitive radio spectrum management and the pure studies on an intentionally very simple scenario allow understanding its capabilities, its characteristic behaviour and effects which may occur and which may need to be considered and handled in more detail.

It can be seen that the self-learning adaptation of the SON system itself performs well; the distributed SON model is adapting itself to the particular individual situation of each single cognitive node, and the SON system then successfully finds suitable and optimized configuration parameters for each cell. In addition SON + self-learning operation can be rather fact, as it involves only offline computations without any feedback measurements for evaluating candidate parameter settings. The SON operation is capable to cope with different traffic load situations and it works for any initial starting parameter configuration values, although reasonable starting parameters improve the algorithm convergence speed.

But while the SON + self-learning system does always quickly improve any unsuitable parameter settings, it does only manage to reach “pretty good” performance results. As this simple and generic modelling cannot be as precise as the very diverse real system, this prediction model without any direct system feedback cannot be expected to reach the theoretical maximum. However it can be suggested to use first this fast offline SON model to find appropriate parameter settings for each individual cell. Thereafter another SON technique could follow, for example another SON technique for fine-tuning which may adapt only one single parameter, which may be slower and may use system feedback when testing some optional parameters on the real system.

It has been found that there is a kind of operation range within which this SON + self-learning technique for cognitive radio spectrum management is working well and for which it has been developed and adapted. Within this range of more or less reasonable parameter values and traffic loads, the implemented proof of concept model operates stably, and provides better results. But for a production system it is still required to handle certain special cases and to extend its parameter range and test it for some particularly extreme parameters settings. Basically, this approach is well suited for the operating range for which it has been developed, adapted and tested. If sometimes maybe unwanted behaviour should still occur, then it is easy and straightforward to detect and to block it, but actually resolving such aspects may then require suitable further SON model adaptations and maybe extensions to handle those particular cases.

An abbreviated comparative summary of the capabilities and strengths versus the limitations and risks of the considered SON model is illustrated in the following table (Table 6-3):

Table 6-3: Capabilities and strengths versus the limitations and risks of the considered SON model

Capabilities and Strengths	Limitations and Risks
Fully generic SON approach, working fine (in most situations) Generic building block for Cognitive Radio and other applications and basis for application specific fine tuning	Fine tuning does still require model extensions and special case handling which may result in that a concrete implementation becomes application specific (unless re-adaptation for another use case)
Achieves (pretty) good parameter settings for each individual node (cell)	Does (usually) not achieve the absolute theoretical performance maximum
Very fast (purely offline calculations) No parameter testing in the field Does not require system feedback Does not require any starting values; Simply activate the SON+Self-Learning entity and it works	There is a certain operation range, a range of more or less reasonable parameters for which the SON system has been designed and adapted for. Extending its range requires further model extensions and testing.

6.5.5 Discussion and further research

This novel solution approach for cognitive radio spectrum management has been shown to be working fine; the distributed SON entities in the spectrum managers configure and optimize many highly interacting configuration parameters via a self-learning prediction model which uses offline calculations without the need of direct system feedback. It is a major step towards the vision of simply put any cognitive nodes into a very diverse, heterogeneous environment, and then letting the SON entities simply adapt and optimize themselves automatically to the currently given situation and to quickly changing environments while thereby resolving the major inter-cell interactions and parameter couplings. The spectrum managers adapt to both, to changing needs from the resource managers as well as to modified external constraints such as spectrum databases. This distributed self-learning approach includes that the cognitive nodes are capable to cope and optimize themselves in the best possible way to externally given non-changeable situations, such as interfering nodes from other non-cooperating networks.

This fast SON concept achieves a pretty good performance and energy optimization, but this simple generic model is not perfect. It is proposed to first let this self-learning SON approach to find quickly reasonably good configuration parameters in a fast offline computation. Thereafter, another SON technique could be added for fine tuning, possibly a single-parameter only or single effect only SON technique which may also require more direct system feedback and which may be slower acting. For a production system, further adaptation work is still required, to extent the operation range and increase its accuracy via more detailed modelling and to add further special case handling as well as to cope with extreme out-of-range parameters and traffic load situations.

For the next steps, it is proposed to add this SON + self-learning entity to a larger more complex scenario and to add further details according to the concrete priorities of that particular SON use case and for that concrete investigated cognitive radio model.

This SON + self-learning entity can be considered as an independent, separate SON box or SON function, which can simply be attached to a spectrum manager or to another cognitive node. It provides a simple interface for integrating it into and for supporting the operation of existing and future cognitive manager and base station products.

6.6 Optimized Cognitive Spectrum Utilization for Stable, Dense Indoor Femtocells

This section considers the context of orthogonal frequency division multiplexing access (OFDMA) based indoor femtocells, overlaid to the existing macrocell by spectrum sharing. For given sub–channels available at the OFDMA femtocells. I It has been specified, in [D6.4] [D6.5], that the functionality of cognitive spectrum management can opportunistically enhance reliability of the spectrum management operation from the femtocell perspective by investigating the method for a joint energy and spectrum (i.e., sub–channels) control for dense femtocell channels. It was shown that a joint design of the energy and the spectrum to the finite, random sub-channels available at the dense femtocells can result in the increased outage capacity. Accordingly, overall energy and spectrum usage among femtocells was developed. In particular, we developed methods of (i) formulating the aggregate amount of the energy usage by taking into account the cost of both the sub–channels at the femtocells and the energy usage for channel feedbacks and data transmissions per femtocell, and (ii) restricting the maximum power allocation level per femtocell based on the aggregate interference rise (by the nearby femtocells) at the incumbent macrocell receiver. Based on these methods, the performance of the joint design has been analysed by deriving expressions for the aggregate energy usage, the interference rise at the nearest macrocell receiver and the outage capacity with limited energy usage at the femtocells.
In this chapter, we will complete this functionality to extend towards optimized cognitive spectrum management by investigating the optimization problem in which the achievable outage capacity per femtocell is maximized under realistic constraints. To that end, we find jointly the optimum energy usage as well as the optimum number of active sub–channels per femtocell. As per our optimum analysis and numerical results, it will be clearly found that in a dense femtocell downlink fading channels deployment, the CM-SM functionality, using the optimum energy–utilizing scheduling in the sense of the maximization of the outage capacity, should implement cognitive spectrum scheduling scheme that converges towards the round–robin scheduling at some extreme cases. This is the inverse of the traditional greedy approach that allocates the spectrum to only the best among the candidates and is known to be optimal without the cognitive spectrum and energy usage functionality.

6.6.1 Optimization problem formulation

Consider the case of using the outage–sensitive and the energy–limited femtocells application. In this case, we focus on maximizing the outage capacity $C_{out,l}$ of femtocell. The outage capacity used in this work is referred to as the maximum achievable rate *R* such that the probability that the sum rate C_l per femtocell is less than or equal to that rate (*R*) is less than or equal to a given threshold ε for $\varepsilon > 0$

(i.e., $\Pr(C_l \le C_{out,l}) \le \varepsilon$). Here, taking into account the principle of multiuser diversity at each femtocell, the resulting sum rate at each femtocell can be given by

$$C_l = \log_2(1+\rho_l) = \log_2(1+\bar{\rho}_l x_l P_d) \text{ for } l \in \{1,\ldots,L\}, \quad (6.6\text{-}1)$$

where $x_l = \max_i x_{li}$ and $\bar{\rho}_l$ denotes the average normalized SINR, i.e., $\bar{\rho}_l = \Theta/(I^{mf} + \sigma^2)$.
We account for optimal selection cost (i.e., optimal number) of active sub–channels per femtocell user equipment (FUE) and their maximum power allocation levels. In addition, the sum energy usage (E_l) on the total available sub–channels per femtocell is limited to its maximum level E_o (i.e., $E_l = E_o$) so that the resulting interference (I_{fm}) by the use of E_l from femtocell *l*, $\forall l$ at the nearest incumbent receiver remains equal to the given threshold. Therefore, the problem at each femtocell can be posed for all *l* as

$$\begin{aligned} \max \quad & C_{out,l} \\ \text{s.t.} \quad & \Pr(C_l \le C_{out,l}) \le \varepsilon, \\ & E_l \le E_o, \\ & I_{fm} \le I_o. \end{aligned} \quad (6.6\text{-}2)$$

To resolve the above problem, we notice the fact that there exists a tradeoff between the energy usage by femtocells and the resulting interference rise at the nearest macrocell victims. That is, the higher the energy usage at femtocells, the stronger the interference rise at the victim macrocell UEs. Therefore, our main idea is to restrict the maximum power allocation level at the femtocells such that this aggregate power (or energy) usage is traded to guarantee an acceptable level of the interference rise (i.e., $I_{fm} \le I_o$).

In addition, based on the properly chosen amount E_o of the maximum energy, the next goal is to develop a method for optimally balancing the energy usage between the channel feedbacks and the data at each femtocell. For this, we take into account an optimal selection of the number of active sub–channels per FUE that achieves the maximization of $C_{out,l}$. This energy usage balance is performed in a distributed fashion among femtocells.

Here, when restricting the amount E_l of the energy usage per femtocell, the calculation of the maximum amount of energy usage is taken into account subject to the coexistence constraint $I_{fm} \le I_o$. To obtain this maximum energy amount, it is assumed that nearby femtocells coordinate to simply refer to static context information such as the number of available FUEs at each femtocell. Based on this coordination, all the nearby femtocells can determine I_{fm} and they uniformly restrict the maximum energy amount per subchannel, due to the fairness. Therefore, a femtocell consisting of more numbers of subchannels will turn out to use higher energy.

6.6.2 Optimal Solution

While satisfying the requirements in the above optimization problem, it turns out that the outage capacity expression can be seen as a function of the number of active subchannels per FUE for given system parameters. As per [D2.4], the expression can be given by

$$C_{out,l} = \log_2\left(1 + \bar{\rho} P_c (N(\bar{n} - n_l) + \eta^{-1})(a - b\log(-\log\varepsilon))\right) \quad (6.6\text{-}3)$$

where $N, \bar{n}$ and n_l are the numbers of FUEs per femtocell, available subchannels per FUE and activated subchannels per FUE, respectively. Here, P_c stands for the power level for scheduling and its ratio to the power level P_d for data transmission is fixed to η, *a* and *b* are parameters related to the inverse of the CDF of the x_l and they will be shown later in this chapter.
Therefore, the above optimization problem can be rewritten as

$$\max_{n_l \in \{1,\ldots,\bar{n}\}} \log_2\left(1+\bar{\rho}P_c(N(\bar{n}-n_l)+\eta^{-1})(a-b\log\log(-\varepsilon))\right). \tag{6.6-4}$$

This problem is a concave optimization problem since the second derivative of the outage capacity in terms of n_l is not positive, i.e., $\partial^2 C_{out,l} / \partial^2 n_l \leq 0$. To mathematically solve this problem, a Lagrange multiplier method can be applied.
Via this method, it turns out that the optimal solution in our case can be simplified to find the value minimizing the first derivation of the outage capacity in terms of n_l. Particularly, let us assume the case when using a single antenna and then, the optimum value of n_l can be given by

$$n_l^{opt} = \arg\min_{n_l}\left|\log\frac{N}{-\log\varepsilon}+1+\left(\log n_l - n_l^{-1}(\bar{n}+(\eta N)^{-1})\right)\right|. \tag{6.6-5}$$

Notice that when using the single antenna, this optimum value can also be achievable by using the inverse of the simplified CDF, $F_x(x)$ (i.e., $F_x(x)=1-\exp(-x)$) without referring to some of the well-known results on the asymptotic behaviour of the maximum among multiple independent random variables. In such a case, we can define *a* and *b* as $a=\log n_l$ and b=1, respectively. The corresponding expression for n^{opt} is still equivalent to n^{opt} in the above expression.
It can be observed that the optimum n^{opt} depends on ε for given *N*, η and $\bar{n}$. Particularly, note the fact from the above expression for n^{opt} that for a given ε, the argument of $|(\cdot)|$ can be shown to be a monotonically increasing function of n_l since $\partial()/\partial n_l > 0$ while $(\cdot)$ being strictly concave (i.e., $\partial^2()/\partial^2 n_l < 0$). Then, for given values of ε producing the term ($\log N - \log(-\log\varepsilon)$) negative, it turns out that the optimum n^{opt}, equating the argument $(\cdot)$ closest to the zero, is an intermediate value between 1 and $\bar{n}$. Specifically, this means that as far as $\varepsilon \in (0, \mathrm{e}^{(-N.\mathrm{e})}]$ (e.g., $\varepsilon \leq 0.07$ when $N = 1$), the value of the optimum n^{opt} between 1 and $\bar{n}$ can be found.
In addition, it can be found from the above expression that for given *N* and $\bar{n}$, less ε (while $\varepsilon \geq \mathrm{e}^{-N.\mathrm{e}}$), closer the term log(−log ε) is to *N*. Therefore, this reveals that when $\varepsilon \geq \mathrm{e}^{-N.\mathrm{e}}$, the smallest among candidate n_l's is the optimum, accordingly.

6.6.3 Simulation Results for stable, dense indoor Femtocells environment

For simulation scenarios, we consider a femtocells deployment in a 5x5 grid layout of geographical environment such as, enterprise environments. Here, both penetration and propagation losses are in line with 3GPP deployment parameters [3GPP ts36.211]. On this layout, the co–channel deployment of 8 femtocells is considered, where each is randomly deployed and intends to access the radio spectrum licensed to the macrocell.

Let $\bar{n}$ inactive sub-channels be given to each FUE. Here, the ratio of the energy usage between channel feedbacks and data at each sub–channel is fixed to $\eta = 3/4$ that is referred to the case when transmitting 3 and 4 OFDM symbols for the control and the data, respectively, at every frame in 3GPP LTE Femtocells. Due to the downlink communications between 8 femtocells and their FUEs, the incumbent macrocell receiver deployed near the femtocells should experience the interference from the femtocells and this should be no greater than $I_o = -30\text{dB}$. For all simulations, we use that $\varepsilon = 10^{-2}$, 4 antennas, 8 femtocells, each having N = 6 FUEs as well as $\bar{\rho} = 0\text{dB}$.

In both **Error! Reference source not found.** and Figure 6-17, we now illustrate the cases of using the optimum size of a subset of active sub–channels with respect to the number of available sub–channels. For comparison, a sub–optimum case has also been concerned and is depicted in both the figures. Particularly, for given values of $\bar{n}$, the optimum case takes into account all the possible values of $\bar{n}$ to optimize, i.e., $n_l \in \{1,...,\bar{n}\}$, while the sub–optimum case does only a subset of all the values of n_l, i.e., $n_l \in \{1,2,4,6\}$. Here are two main observations to highlight. As per Figure 6-16, the optimum case always outperforms the sub–optimum case in terms of the outage capacity. However, as seen in Figure 6-17, the sub–optimum case benefits from the less size of the set of discrete regions achieved, as compared to the optimum case.

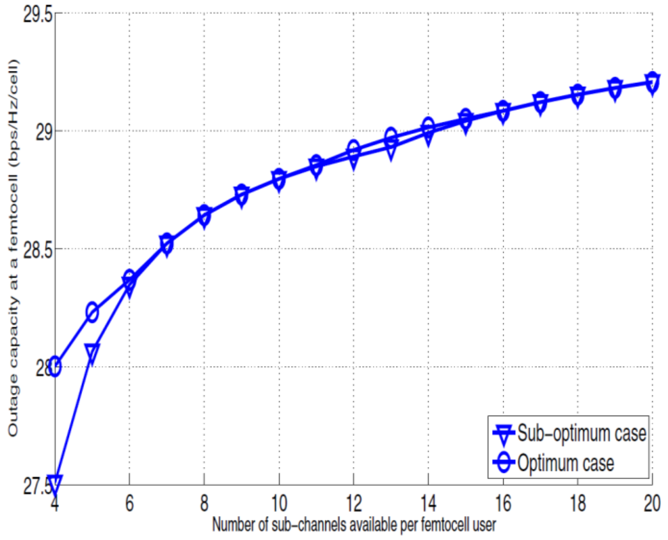

Figure 6-16: Outage capacity exploiting either the optimum or the sub–optimum sizes of active sub–channels has been depicted with respect to the number of sub–channels available.

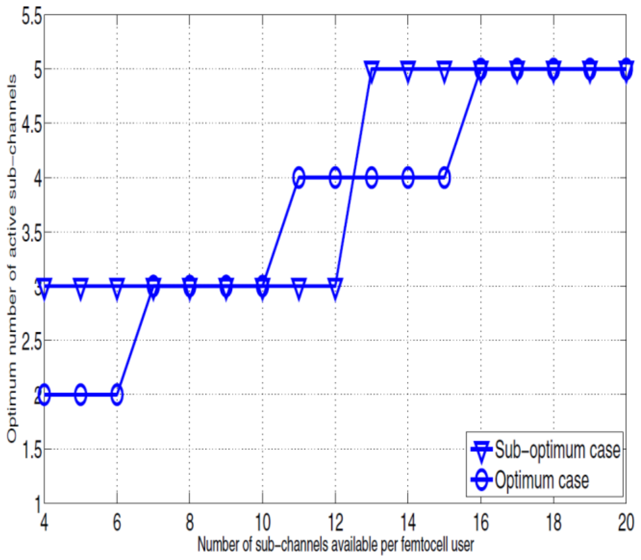

Figure 6-17: Comparison of optimum and sub–optimum sets of discrete regions, each set achieving optimum and sub–optimum sizes of active sub–channels, with respect to the number of sub–channels available.

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.

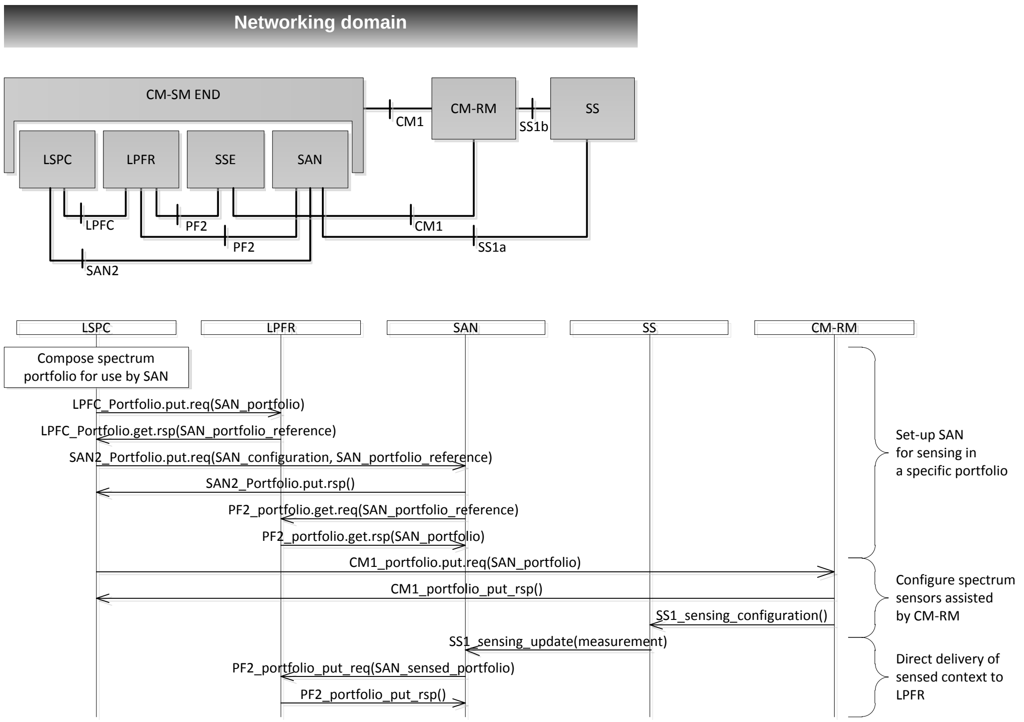

Figure 7-1: Accessing a SAN and sample MSC (SAN configured to create a portfolio from raw sensor data for storing to LPFR)

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 than 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.

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

7.4 Evaluation of Spectrum analyser and selector

7.4.1 Evaluation of Spectrum Selection Functions based on Incumbent User Statistics from Spectrum Sensing

Four different SSE functions are proposed as described below. These do not replace the spectrum management functionality specified in the IEEE 802.22 standard, but will be complementary and coexist to enhance performance by considering statistics calculated over longer time periods. To evaluate the performance of these SSE functions, the IEEE 802.22 standard has been implemented in NS-2 with mobile opportunistic users (OUs).

SSE Functions

SSE-Power

The *SSE-Power* function is a basic algorithm that will be used to benchmark the other three SSE algorithms. SSE-Power selects the channel where the spectrum sensor has detected lowest signal from WMs.

The SAN function receives sensing results $r_{i,j}$ from OU$_i$ on channel j and sends this information to the SSE. The *SSE-Power* function then selects the optimal channel ch_{Power} based on the following criteria:

$$ch_{Power} = \min_j \left\{ \max_i \left(r_{i,j} \right) \right\}, \tag{7.4-1}$$

$$subject\ to{:}\, r_{i,j} < \delta, j \in [0, M], i \in [0, N] \tag{7.4-2}$$

where N is the total number of OUs connected to the BS, M the total number of channels to be sensed and δ the sensing detection threshold (-107dBm over 200kHz in the simulator).

SSE-Distance

The second spectrum selection algorithm is one that enhances QoS in presence of mobile cognitive radio terminals. It uses sensing to predict the distance to WMs and then to proactively select the channel with farthest distance from the closest OU to WM on that channel (e.g. each 5th second, which is used in the simulator).

Upon receiving sensing results from the IEEE 802.22 system, the SAN finds the distance $d_{i,j}$ for each OU$_i$ to the closets WM on each channel j. Note that the SAN function is able to predict the location of the WM based on received sensing results from the cognitive radio devices. This information is forwarded to the SSE-Distance function which uses it to select the optimal channel $ch_{Distance}$ based on the following criteria:

$$ch_{Distance} = \max_j \left\{ \min_i \left(d_{i,j} \right) \right\}, \tag{7.4-3}$$

$$subject\ to{:}\, r_{i,j} < \delta, j \in [0, M], i \in [0, N] \tag{7.4-4}$$

where N is the total number of OUs connected to the BS, M the total number of channels to be sensed and δ the sensing detection threshold.

SSE-Distance and *SSE-Power* will give similar results in many cases. The difference is that the *SSE-Distance* uses historical sensing data from several OUs to predict the location of the WMs, hence it will give more reliable prediction of the WM location.

Furthermore, the *SSE-Distance* function can use information about WM locations to optimize the sensing scheme used. If this is used, it will be referred to with "+"-symbol as *SSE-Distance+*. We focus on the two-stage sensing function in IEEE 802.22, where at the coarse sensing stage (first stage) an energy detector is used for frequent and short sensing duration T_c=1ms. If coarse sensing detects a

WM signal it switches to the fine sensing stage (second stage) that uses a more detailed WM detection process for a longer duration T_s=30ms. The coarse sensing period, i.e. the interval between two coarse sensing stages, is given as $T_P = n_c \cdot T_F$, where T_F=10ms is the OFDMA frame length in IEEE 802.22. The channel detection time (CDT) which is set to 2 seconds by FCC gives an upper bound on n_c, hence we have $1 \leq n_c \leq CDT$.

As n_c and hence T_P increases, the sensing overhead will reduce. Note that coarse sensing senses at the end of the OFDMA frame, which will cause a reduction in data transmitted uplink. However, since the number of coarse sensing stages is reduced, the number of false alarms will also decrease reducing the number of fine sensing stages. This will reduce overhead.

We want to set T_P to protect the WM and the IEEE 802.22 nodes from interfering with each other. Hence, when distance $d_{i,j}$ between the OU i and WM j on channel j is low, T_P should be set to a low value such that WMs can be detected faster. This will reduce interference. If the $d_{i,j}$ is high such that the OU is outside detection range, T_P should be set high to increase throughput in the IEEE 802.22 network. To find the WM detection range for the OU, we set up the equation for received transmit power RX_{OU} for the OU:

$$RX_{OU} = TX_{WM} - PL + G_{TX} + G_{RX} \tag{7.4-5}$$

where $TX_{WM} = 17$dBm (50mW) is the WM transmit power, PL is the path loss with the Okumura-Hata model, G_{TX} and G_{RX} are antenna gains for the CR and WM respectively ($G_{TX} = G_{RX} = 0$ dBi in the simulator). We then set $RX_{OU} = -107$ dBm for the detection threshold as given by FCC and used in the IEEE 802.22 standard, such that:

$$-107dBm > 17dBm - PL\left(d_{i,j}\right)dB \tag{7.4-6}$$

$$PL(d_{i,j}) > 124dB \tag{7.4-7}$$

$$d_{i,j} > 10^{\left(\frac{124-69.55-26.16\cdot log_{10}(f)+13.82\cdot log_{10}(h_b)+C_H}{44.9-6.55\cdot 13.82\cdot log_{10}(h_b)}\right)} \tag{7.4-8}$$

where f is the frequency in MHz, h_b is the base station height in meters and C_H is a correction factor for user terminal height. Next, we find $d_{i,j}$ for the parameters used in the simulation scenario and find that; for the BS with parameters (f=600, h_b=15, C_H =0), $d_{bs,j} > 0.885$ km, and for the OU with parameters (f=600, h_b=1.5, C_H =0), $d_{i,j} > 0.436$ km.

For *SSE-Distance+*, the value n_c for the coarse sensing period T_P is selected based on the following criteria for the BS with distance $d_{bs,j}$ to the WM *j* on channel *j*:

$$n_c = \begin{cases} 10, & if\, d_{bs,j} \leq 0.9km \\ 100, & if\, 0.9km < d_{bs,j} \leq 1km \\ 200, & if\, d_{bs,j} > 1km \end{cases}, \tag{7.4-9}$$

and based on the following criteria for the OU:

$$n_c = \begin{cases} 10, & if\ d_{i,j} \leq 0.5km \\ 100, & if\ 0.5km < d_{i,j} \leq 0.6km \\ 200, & if\ d_{i,j} > 0.6km \end{cases}. \tag{7.4-10}$$

which gives T_P=0.1, 1 and 2 seconds respectively. A higher margin is considered for the OU compared with the BS since the OUs are mobile. Note that the protection ratio for a WM as set by FCC when using the database approach only is considered to be 1 km.

SSE-OnOff

The *SSE-OnOff* function aims to enhance the QoS and performance of the mobile OUs in scenarios where the density of the incumbent WMs is high. The function uses sensing to predict the probability that a channel will not be occupied by a WM. To do this, the SAN function uses sensing results from the cognitive radio devices to calculate the mean values for the ON (channel busy) and OFF (channel idle) periods for the WMs on each channel j, denoted $T_{\mathrm{j,ON}}$ and $T_{\mathrm{j,OFF}}$ respectively. The SAN function then sends this information to the *SSE-OnOff* function which selects the optimal channel ch_{OnOff} based on the following criteria:

$$ch_{OnOff} = \max_{j} \frac{T_{j,OFF}}{T_{j,OFF} + T_{j,ON}} \quad (7.4\text{-}11)$$

$$subject\ to: r_{i,j} < \delta, j \in [0, M] \quad (7.4\text{-}12)$$

where N is the total number of OUs connected to the BS, M the total number of channels to be sensed and δ the sensing detection threshold.

SSE-Hybrid

This *SSE-Hybrid* function combines *SSE-Distance* and *SSE-OnOff*. The aim of this function is to enhance QoS when the OUs are mobile and when the density of WMs and/or OUs is high.

We want to use *SSE-Distance* when the distance $d_{i,j}$ between the OU i and WM j on channel j is high such that the probability of a WM appearing within the detection range reduces. On the other hand, we want to use *SSE-OnOff* if the $d_{i,j}$ is low such that the WM might appear within detection range. By using the same detection ranges as calculated for the *SSE-Distance+*, the *SSE-Hybrid* function can be configured to select the optimal channel ch_{hybrid} based on the following criteria:

$$ch_{hybrid} = \begin{cases} ch_{OnOff}, & d_{i,j} \leq 0.436km, d_{bs,j} \leq 0.885km \\ ch_{Distance}, & d_{i,j} > 0.436km, d_{bs,j} > 0.885km \end{cases} \quad (7.4\text{-}13)$$

$$subject\ to: r_{i,j} < \delta, j \in [0, M], i \in [0, N] \quad (7.4\text{-}14)$$

Summary of Performance Evaluation

Simulation Scenario

The simulator implementation of IEEE 802.22 and parameters used are described in detail in [D6.6].

The QoSMOS scenario considered for the simulation scenario is cellular extension in white spaces. Furthermore, the scenario can be described by a light urban environment with one BS (height 15m, EIRP 36dBm) of cell radius 1.2 km and mobile OUs (height 1.5m, EIRP 20dBm) that move following a random waypoint model with a random speed between 1 and 20 m/s. Their initial location is randomly selected within the BS radius. Constant bit rate traffic will be transmitted in the downlink to each OU with load 200kbit/s.

It is assumed that 4 channels of 6MHz bandwidth are available for access after consulting the geo-location database. This means that these 4 channels neither are allocated by TV broadcasters nor registered wireless microphones. However, unregistered WMs (height 1.5m, EIRP 50mW, bandwidth 200kHz) will appear in these channels in the simulation scenario following the ON/OFF pattern with inter-arrival rate and departure rates with average 20 and 5 seconds respectively are selected randomly with ±10 and ±2.5 seconds respectively. During simulation, both the WM inter-arrival and departure rates will vary according to the negative exponential distribution for the initial selected inter-arrival and departure rates. There are totally 4 WMs in the scenario, each one appearing separately on one of the 4 channels. Their location is randomly selected within the area of 1.4 km radius from the BS. The Cost Hata model is used for propagation modelling.

The simulations are computationally demanding since all levels in the protocol stack are simulated and packets transmitted on the wireless medium are received by all nodes in the simulator. We make a note that more simulations are required to provide more accurate results than have been done for the results presented here.

Performance Evaluation

The average throughput for all OUs for the *SSE-Power*, *SSE-Distance*, *SSE-OnOff* and *SSE-Hybrid* functions when compared with the optimal case for *SSE-Distance* without presence of WMs Figure 7-2.

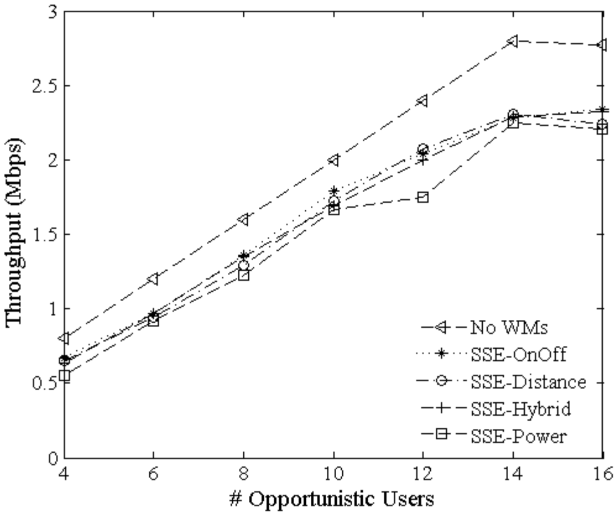

Figure 7-2: Performance evaluation of the SSE-Power, SSE-Distance+, SSE-OnOff and SSE-Hybrid functions compared with the near optimal performance for “No WMs” where WMs are located at 3km distance from the cell edge.

It can be seen that *SSE-OnOff* achieves highest throughput for most number of OUs. This is because when the WM activity level is quite high as in the considered scenario, *SSE-OnOff* will most often select the channel that stays idle for the longest period. Hence, the number of channel switches is reduced and harmful interference to the OUs is reduced resulting in a more stable network. It can also be seen that *SSE-Distance* achieves slightly higher and similar throughput for 12 and 14 OUs. This comes at the cost of higher interference to the WMs. The higher throughput and increase in interference to the WM here can be explained by the fact that the OUs often are located outside of WM detection range while still not experiencing harmful interference from the WM. *SSE-Hybrid* does not achieve maximum throughput as desired. Hence the *SSE-Hybrid* function might not always select the optimal SSE function in the considered scenario. This could potentially be different if other values or another optimization were used for the SSE function selection.

It is observed that *SSE-Distance* achieves better performance than *SSE-Power*. The two will mostly select the same channel in the case that WMs are present on all channels since the distance is derived from the sensor measurements. However, in the case that no WM is present on one or more channels during sensing, *SSE-Power* will in many cases select a channel without knowledge about where the WM might appear. Therefore, the WM might appear close to the OU such that both the OU and WM experience harmful interference. It was observed especially for *SSE-Power* that this resulted in the OU losing synchronization with the BS which reduced throughput dramatically. This was also observed for the other SSE functions, but less frequently.

Not shown in Figure 7-2 is a simulation presented in [D6.6] that evaluates performance for different WM activity levels. When WM activity is higher (e.g. for lower WM inter-arrival values), SSE-OnOff performs better than the other SSE functions. When WM activity is lower (e.g. for higher WM inter-arrival values), *SSE-Distance* performs better than the other SSE functions. When the WM activity level is very low, the SSE functions perform quite similar.

7.5 Interfaces

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 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 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 QoSMOS reference model.

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

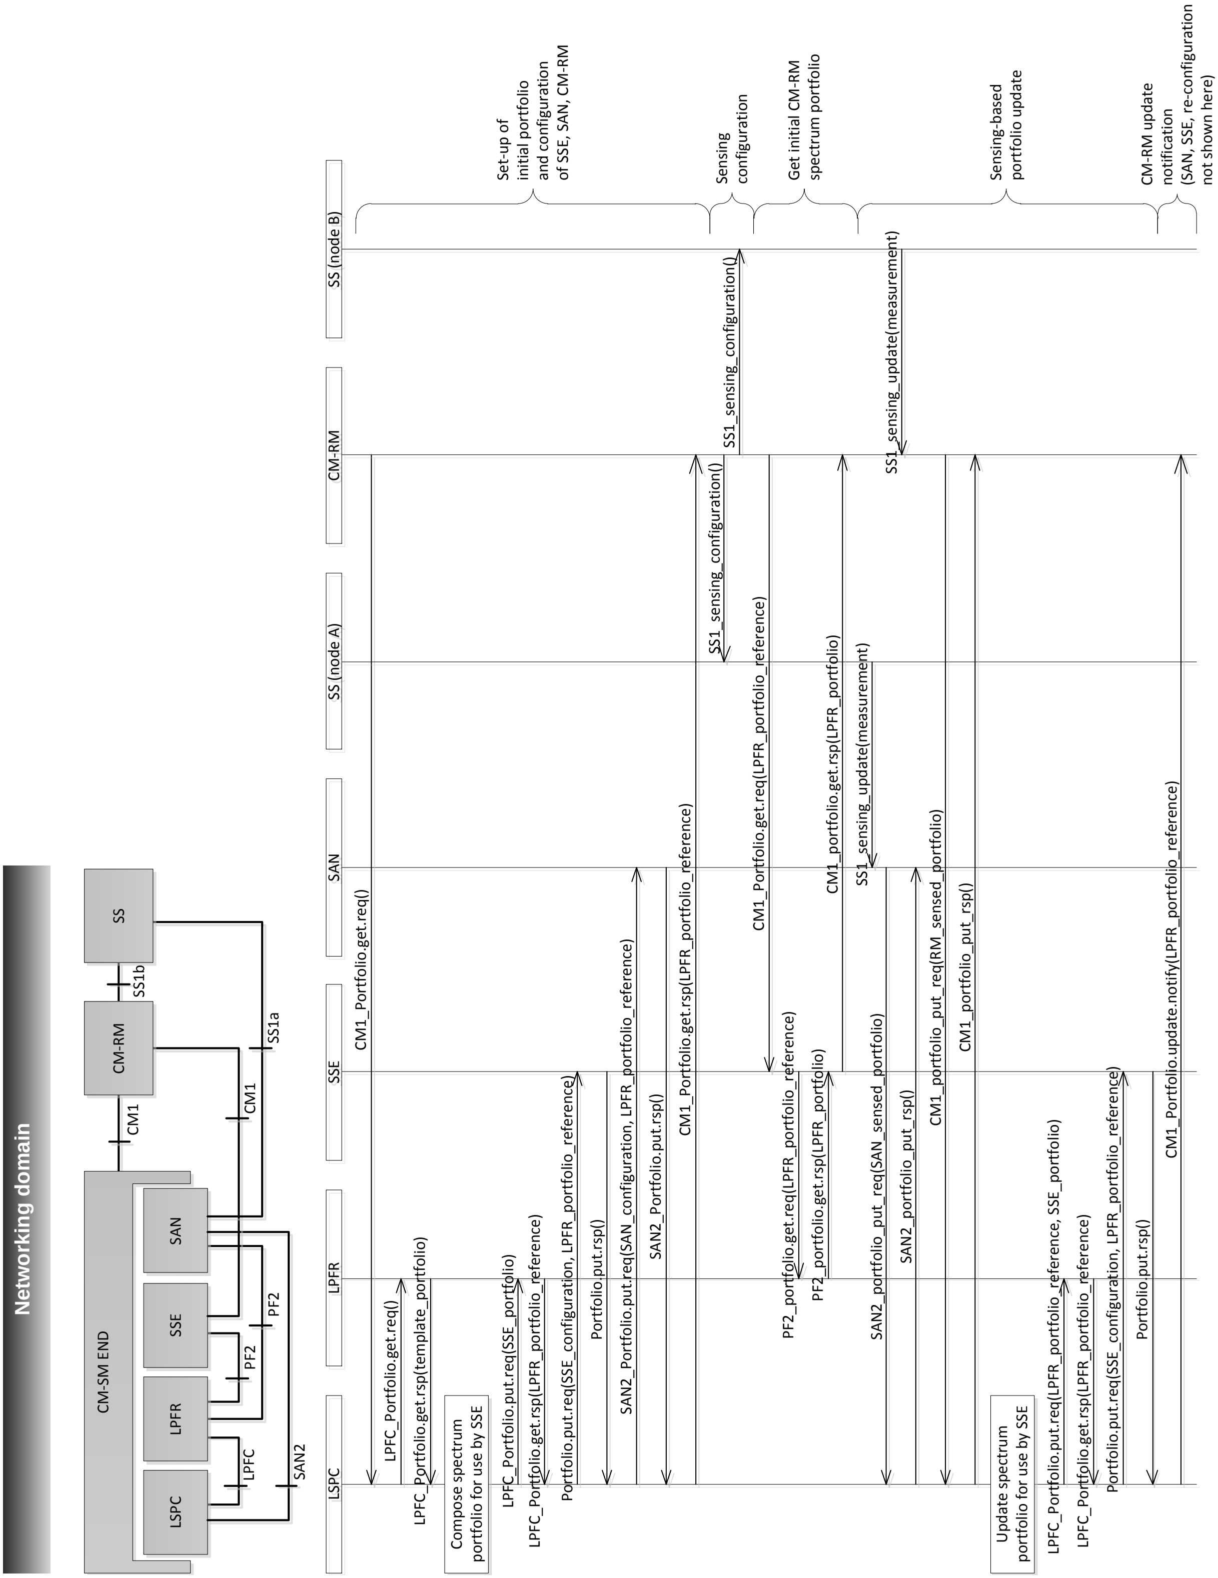

Figure 7-3: Accessing an SSE and sample MSC (initialization and update of an SSE-deployed spectrum portfolio involving SS and CM-RM context updates)

8 Summary and Conclusions

This deliverable presents the integrated final specification of QoSMOS spectrum management framework. It details the refined CM-SM reference model, its cognitive methods and opportunistic functions, internal procedures and the communication between the functional entities defined by the model.

The implementation of CM-SM operation relies on roles instantiated at QoSMOS entities. These roles were developed in details as abstract state machines. During the tasks of completion of integration each defined roles were implemented as software modules and the output was handed over to WP7 for integrating them into PoC scenarios. The documentation of these software modules are also summarized in this deliverable.

For the sake of flexibility and robustness a layered communication strategy is recommended in this work. The main benefit of such layering is in gaining flexibility regarding different communication sub-system implementations, and in increasing reusability of codes across different application entities.

The increase in flexibility is created from using a generic ‘setter/getter’ paradigm that allows choosing different implementations for the communication sub-system. CORBA, REST, XML-RPC, ETSI M2M are some of those that comply with the paradigm and have been evaluated for specific configurations and specific interfaces of the QoSMOS framework.

The collaboration of the cognitive engines CM-SM and CM-RM (Cognitive Manager – Resource Manager) is one of the key aspects of spectrum management; therefore it is presented in detail; however CM-RM specification is given in [D5.1], [D5.2] and [D5.3], where the description of the final structure of the cognitive manager for resource management is published.

Geolocation database solution is considered as the main enabler of opportunistic spectrum access and this has been developed and demonstrated. This is part of the QSMOS WP7 integrated demonstration platform.

Multi-objective portfolio optimization has been successfully used in resolving the problem of frequency selection/aggregation of cognitive radio systems as described in deliverable [D6.6]. The main conclusion was that technical and economic target functions can be conveniently reached in a cognitive radio network by means of a multi-objective port-folio optimization problem given an appropriate balance between return and risk components and certain values of economic parameters of the available frequency bands. However, the approach considered is a low level of interaction between CM-RM and CM-SM.

An overview the role of the long-term user activity in the cognitive spectrum management systems, and the algorithmic approach of the different calculations have been analysed. The main elements of this method can be found in detail in D6.4. The long-term observation of the ON/OFF activity of incumbents and opportunistic users gives the large-scale overview of a cognitive system. The activity duration statistic of different users was applied to build a model, express the distribution of the length of the activity and the activity-free periods.

One of the major challenges in cognitive radio spectrum management is to find the best suited spectrum portfolio and its power settings for each individual cognitive node. Heterogeneous cellular telecommunication scenarios in the future, requires being able to handle and resolve the large interactions and strong couplings between the different cognitive nodes, and furthermore these highly coupled parameters shall quickly be adapted and optimized to its currently given situation. In order to handle and resolve this spectrum management issue, an approach for distributed self-learning Self-Organizing-Network (SON) spectrum management has been developed as described in detail in deliverables [D6.4] and [D6.5]. This concept has been implemented in a radio system simulator for cellular wireless networks in order to validate this concept, and to evaluate its performance for cognitive radio as well as to characterize its capabilities and limitations.

This novel approach for cognitive radio spectrum management seems to perform well as the distributed SON entities in the spectrum managers configure and optimize many highly interacting configuration parameters via a self-learning prediction model which uses offline calculations without the need of direct system feedback. It is a major step towards the vision of simply put any cognitive nodes into a very diverse, heterogeneous environment, and then letting the SON entities simply adapt and optimize themselves automatically to the given current situation and to quickly changing environments while thereby resolving the major inter-cell interactions and parameter couplings. The spectrum managers adapt to both, to changing needs from the resource managers as well as to modified external constraints such as spectrum databases. This distributed self-learning approach includes that the cognitive nodes are capable to cope and optimize themselves in the best possible way to externally given non-changeable situations, such as interfering nodes from other non-cooperating networks.

Optimized cognitive spectrum utilization strategy for stable, dense indoor femtocells was developed. It has been sown that the functionality of cognitive spectrum management can opportunistically enhance reliability of the spectrum management operation from the femtocell perspective by investigating the method for a joint energy and spectrum (i.e., sub–channels) control for dense femtocell channels. It was shown also that a joint design of the energy and the spectrum to the finite, random sub-channels available at the dense femtocells can result in the increased outage capacity. Accordingly, overall energy and spectrum usage among femtocells was developed. Methods were developed for (i) formulating the aggregate amount of the energy usage by taking into account the cost of both the sub–channels at the femtocells and the energy usage for channel feedbacks and data transmissions per femtocell, and (ii) restricting the maximum power allocation level per femtocell based on the aggregate interference rise (by the nearby femtocells) at the incumbent macrocell receiver. Based on these methods, the performance of the joint design has been analysed by deriving expressions for the aggregate energy usage, the interference rise at the nearest macrocell receiver and the outage capacity with limited energy usage at the femtocells.

9 References

[D1.2] "QoSMOS consolidated scenarios", FP7-ICT-2009-4/248454 QoSMOS project deliverable D1.2, Dec. 2010.

[D2.1] "Initial description of system architecture options for the QoSMOS system", FP7-ICT-2009-4/248454 QoSMOS project deliverable D2.1, Apr. 2010.

[D2.2] "System architecture options for the QoSMOS system", FP7-ICT-2009-4/248454 QoSMOS project deliverable D2.2, Dec. 2010.

[D2.3] "System specification and evaluation criteria", FP7-ICT-2009-4/248454 QoSMOS project deliverable D2.3, Nov. 2011.

[D2.4] "System Architecture Consolidation, Evaluation and Guidelines", FP7-ICT-2009-4/248454 QoSMOS project deliverable D2.4, Dec. 2012.

[D5.1] "Initial description of framework for supporting QoS and handling mobility", FP7-ICT-2009-4/248454 QoSMOS project deliverable D5.1, Jul. 2010.

[D5.2] "Final framework description, preliminary cognitive manager structure and first mechanisms for QoS support", FP7-ICT-2009-4/248454 QoSMOS project deliverable D5.2, Jul. 2011.

[D5.3] "Final description and specification of cognitive manager and corresponding QoS support mechanisms with performance evaluation", FP7-ICT-2009-4/248454 QoSMOS project deliverable D5.3, Jul. 2012.

[D6.1] "Initial description of spectrum management framework, requirements analysis and approach selected", FP7-ICT-2009-4/248454 QoSMOS 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 QoSMOS 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 QoSMOS project deliverable D6.3, Nov. 2011.

[D6.4] "Trust and security functions of the spectrum management framework", FP7-ICT-2009-4/248454 QoSMOS project deliverable D6.4, Apr. 2012.

[D6.5] "Specification of cognitive and opportunistic functions of the spectrum management framework", FP7-ICT-2009-4/248454 QoSMOS project deliverable D6.5, Apr. 2012.

[D6.6] "Spectrum management framework integration and implementation", FP7-ICT-2009-4/248454 QoSMOS project deliverable D6.6, Dec. 2012.

[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

[1900.4] 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.

[1900.5.1] IEEE Draft Std 1900.5.1™; "IEEE Standard Policy Language for Dynamic Spectrum Access Systems"; June 2012

[1900.5] IEEE Std 1900.5™-2011;"IEEE Standard for Policy Language Requirements and System Architectures for Dynamic Spectrum Access Systems"; Jan 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

[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.

[BOYD04] S. Boyd and L. Vandenberghe, Convex optimization, Cambridge University Press, 2004.

[CelEtal2011] U. Celentano, B. Bochow, C. Lange, F. Noack, J. Herrero, B. Cendón, O. Grøndalen, V. Mérat, C. Rosik, "Flexible architecture for spectrum and resource management in the whitespace", Proc. Int. Symp. Wireless Personal Multimedia Commun. (WPMC 2011), Brest, France, 3-7 Oct 2011.

[FAB07] F.J. Fabozzi, et al., Robust Portfolio Optimization and Management, Wiley Finance, 2007.

[Fette06] Fette, B.; "Cognitive Radio Technology", Newnes, 2006.

[FRI67] B. D. Fritchman: "A binary channel characterization using partitioned Markov chains", IEEE Trans. Inform. Theory, vol. 13, Apr. 1967, pp. 221-227.

[HP09] Harris Poll, Average net user online statistics, released in 2009. http://news.cnet.com/8301-1023_3-10421016-93.html

[IEEE2003] "Traffic Models for IEEE 802.20 MBWA System Simulations", IEEE P 802.20; Draft 802.20 Permanent Document, July 8, 2003

[JOU03] W. Jouini, C. Moy, and J. Palicot, "Decision making for cognitive radio equipment: analysis of the first 10 years of exploration," EURASIP Journal on Wireless Commun. and Networking, January, 2012.

[KON05] A. Konak, D.W. Coit, and A.E. Smith, "Multi-objective optimization using genetic algorithms: A tutorial," Reliability Engineering and Safety, Elsevier, 2005, pp. 992-1007.

[LU12] L. Lu, X. Zhou, U. Onunkwo, and G.Y. Li, "Ten years of research in spectrum sensing and sharing in cognitive radio," EURASIP Journal on Wireless Commun. and Networking, January 2012, Open Access.

[MIT93] J. Mitola, "Software radios: survey, critical evaluation and future directions", IEEE Aerosp. Electron. Syst. Mag. 8(4), 25–31 (1993).

[ORACLE2008] "Public Deliverable D1.3 – ORACLE System Model and Implementation Guide", The ORACLE consortium, FP6-IST-2004 027965, www.ist-oracle.org, November 2008.

[PAL06] R. Pal, D. Idris, K. Pasari, and N. Prasad, "Characterizing Reliability in cognitive radio networks," First Int. Symposium on Applied Sciences on Biomedical and Commun. tech., 2008.

[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

[WIND112] WINNER Del. D1.1.2. Final link- and system-level channel models, http://www.ist-winner.org/.

[WIXSLS] WiMAX system level evaluation methodology, WiMAX Forum, V.0.0.1, 2006.

[XGL2004] DARPA XG Working Group, “DARPA XG Policy Language Framework, Request for Comments”, v. 1.0, prepared by BBN Technologies, Cambridge MA, USA, 2004.