iJOIN

INFSO-ICT-317941

INFSO-ICT-317941 iJOIN

D5.2

Final Definition of iJOIN Requirements and Scenarios

Editors:	Jorge Ortín, Pablo Caballero, IMDEA; Peter Rost, NEC
Deliverable nature:	Public
Suggested readers:	iJOIN GA
Due date:	October 31th, 2014
Delivery date:	November 30th, 2014
Version:	0.1
Total number of pages:	83
Reviewed by:	Dario Sabella, Beppe Coffano, GA members
Keywords:	RAN, Backhaul, Cloud Computing, RANaaS, iJOIN Common Scenarios, iJOIN Candidate Technology, iJOIN Architecture, LTE
Resources consumed	28.4 PM

Abstract

This deliverable presents the final set of reference scenarios and system requirements derived by the iJOIN project, focusing in dense small cell deployments as a way to respond to the increasing data rate demand, but always with realistic backhaul limitation in mind. Relying on the progress in cloud computing, iJOIN introduces the concept of "Radio Access Network as a Service" (RANaaS) to deploy functionalities, which are usually processed within a small cell, partially or fully in a cloud platform. This allows to benefit not only in computing power but also in centralisation coordination gains.

In particular, this deliverable presents an overview of the activities carried out by the Work Package 5 (WP5) during the second year of the project. The report gives an overview of the current status of iJOIN activities, definitions and system concepts, while specific aspects are contained in deliverables D2.2, D3.2 and D4.2 coming from the respective technical work packages. This report provides the final set of reference scenarios and system requirements considered in iJOIN based on the output of WP2, WP3, WP4 and proof-of-concept work in WP6. For each scenario, a RAN/Backhaul and a RANaaS deployment is proposed which is suitable for the specific scenario characteristics. Furthermore, the different hardware limitations that will affect the RANaaS implementation are analysed. These limitations together with the requirements imposed by 3GPP LTE and with the proposed candidate technologies are used to define the preferred functional split options. A similar analysis is performed for the joint radio access and backhaul network optimisation. Finally, metrics are defined that are used to evaluate candidate technologies and to analyse operating points of the iJOIN system.

List of authors

Company	Author
CEA	Antonio De Domenico (antonio.de-domenico@cea.fr)
HP	Marco Di Girolamo (marco.digirolamo@hp.com)
	Marco Consonni (marco_consonni@hp.com)
IMC	Umer Salim (umer.salim@intel.com)
IMDEA	Jorge Ortín (jortin@unizar.es)
	Albert Banchs (albert.banchs@imdea.org)
	Pablo Caballero (pablo.caballero@imdea.org)
NEC	Andreas Maeder (andreas.maeder@neclab.eu)
	Peter Rost (peter.rost@neclab.eu)
SCBB	Massinissa Lalam (massinissa.lalam@sagemcom.com)
TI	Dario Sabella (dario.sabella@telecomitalia.it)
	Marco Caretti (marco.caretti@telecomitalia.it)
TID	Ignacio Berberana (ibfm@tid.es)
TUD	Jens Bartelt (jens.bartelt@ifn.et.tu-dresden.de)
UC3M	Carlos Jesús Bernardos Cano (cjbc@it.uc3m.es)
UNIS	Atta Quddus (a.quddus@surrey.ac.uk)
	Yingli Sheng (y.sheng@surrey.ac.uk)
UoB	Dirk Wübben (wuebben@ant.uni-bremen.de) Henning Paul (paul@ant.uni-bremen.de)

History

Modified by	Date	Version	Comments
Jorge Ortín	30.11.2014	1.0	Final Version

Table of Contents

List of authors ... 2
History ... 3
Table of Contents ... 4
List of Figures ... 7
List of Tables ... 9
Abbreviations ... 10
1 Executive Summary ... 13
2 Introduction and key contributions ... 14
2.1 Background and Scope ... 14
2.2 Key Contributions ... 14
3 Definitions ... 16
4 iJOIN Common Scenarios ... 19
4.1 Common Scenario 1: Stadium ... 19
4.1.1 Physical Description ... 19
4.1.2 Traffic Demand ... 19
4.1.3 Proposed RAN/Backhaul and RANaaS deployment ... 20
4.1.4 Performance evaluation parameterization ... 20
4.2 Common Scenario 2: Square ... 21
4.2.1 Physical Description ... 21
4.2.2 Traffic Demand ... 22
4.2.3 Proposed RAN/Backhaul and RANaaS deployment ... 22
4.2.4 Performance evaluation parameterization ... 23
4.3 Common Scenario 3: Wide-area continuous coverage ... 24
4.3.1 Physical Description ... 24
4.3.2 Traffic Demand ... 24
4.3.3 Proposed RAN/Backhaul and RANaaS deployment ... 24
4.3.4 Performance evaluation parameterization ... 25
4.4 Common Scenario 4: Shopping Mall / Airport ... 26
4.4.1 Physical Description ... 26
4.4.2 Traffic Demand ... 27
4.4.3 Proposed RAN/Backhaul and RANaaS deployment ... 27
4.4.4 Performance evaluation parameterization ... 28
4.5 Network Sharing Enablers ... 29
4.5.1 Network Sharing in the context of 3GPP ... 29
4.5.2 Benefits of Network Sharing ... 31
4.6 Proof-of-concept assumptions ... 32
5 Functional Split Implementation Requirements of RANaaS ... 35

5.1 Implementation aspects of RANaaS hardware 35
5.1.1 Implementation Options 35
5.1.2 Virtualization infrastructure....37
5.1.3 Computational Outage 41
5.1.4 Load balancing....42
5.1.5 Migration of Virtual eNodeBs 43
5.1.6 Implementation requirements 43
5.2 Implementation constraints of 3GPP LTE 44
5.3 Preferred functional splits 46
5.4 Flexible functional split assignment 48
5.5 Preliminary Results....49
5.5.1 Opportunistic HARQ 49
5.5.2 Computational Outage 51
5.5.3 Computational Diversity....52
6 Joint Radio Access and Backhaul Network Requirements 55
6.1 Required interfaces and interaction....55
6.2 Limitations in 3GPP LTE 56
6.2.1 3GPP interfaces and requirements....56
6.2.2 Recommendations....57
6.3 Backhaul Technologies....57
6.4 Preliminary results 59
6.4.1 Joint RAN/BH Coding....59
6.4.2 Distributed IP Anchoring and Mobility Management 60
6.4.3 Network Wide Energy Optimisation 62
7 iJOIN System Performance Evaluation....65
7.1 Relevant metrics....65
7.1.1 Area Throughput....65
7.1.2 Energy Efficiency 65
7.1.3 Utilisation Efficiency....66
7.1.4 Cost Efficiency 68
7.2 iJOIN Concept Evaluation 69
7.2.1 iJOIN System Design 69
7.2.2 iJOIN System Evaluation 70
7.2.3 iJOIN System Evaluation Scenarios....70
7.2.4 Intra-Layer Interaction of CTs....72
8 Summary and Conclusions....75
Acknowledgements and Disclaimer 76
References 77
Annex A Power consumption models of iJOIN architectural entities 80

A.1 iSC Power Consumption....80
A.2 RANaaS Platform Power Consumption....81
A.3 Backhaul Power Consumption....82

List of Figures

Figure 3-1: “Generic” Mobile Network Architecture 16
Figure 3-2: Mobile Radio Network Definition 17
Figure 4-1: RANaaS and virtual eNodeB configuration options 19
Figure 4-2: Stadium – iSCs and macro cell positions (left) and details on iSCs antenna tilt (right) 20
Figure 4-3: Stadium – Physical deployment example 20
Figure 4-4: Stadium Layout in high load scenarios 21
Figure 4-5: Square Use Case 22
Figure 4-6: Square - Physical Deployment example 22
Figure 4-7: Small cell deployment in the square. 23
Figure 4-8: Wide Area - Physical Deployment example 25
Figure 4-9: Small cell deployment for the Wide Area Coverage scenario 26
Figure 4-10: Airport Use Case 26
Figure 4-11: Shopping Mall Use Case 27
Figure 4-12: Shopping Mall / Airport: Physical deployment example 28
Figure 4-13: Small cell deployment in the Shopping Mall / Airport: Sparse (left) and dense (right) deployment. 28
Figure 4-14: 3GPP support of network sharing 30
Figure 4-15: Degrees of integration in network sharing solutions 30
Figure 4-16: iJOIN Testbed Coverage 33
Figure 5-1: Implementation choice applied to iJOIN system 35
Figure 5-2: Example of splitting of digital signal processing across GPP, DSP, and FPGA 35
Figure 5-3: Server Virtualization 37
Figure 5-4: Virtual Machine Cluster 38
Figure 5-5: Hybrid Programming Model on an IaaS VM Cluster 40
Figure 5-6: Raw throughput and computational effort for rate-maximizing and computationally aware scheduler 41
Figure 5-7: Functional split options for the PHY layer [34] 45
Figure 5-8: Implementation options of 3GPP LTE RAN functionality 47
Figure 5-9: Preferred functional splits considered in iJOIN 48
Figure 5-10: Achievable outage rate depending on the SNR for an outage probability of 0.1% 50
Figure 5-11: Considered network deployment 51
Figure 5-12: Results for single-cell under a computational complexity constraint 51
Figure 5-13: Results for multi-cell network for different computational complexity constraints 52
Figure 5-14: Numerical and analytical complexity model for 3GPP LTE uplink 52
Figure 5-15: Expected and outage computational complexity for one cell 53
Figure 5-16: Scaling of computational complexity as a function of number of users/cells 53
Figure 5-17: Relative outage complexity as a function of the outage probability 54
Figure 5-18: Complexity-rate-tradeoff as a function of the number of centralized cells 54

Figure 6-1: Ordering of backhaul technologies and preferred functional splits 58
Figure 6-2: Code rate adaption and channel quality measurements required for joint RAN/BH en-/decoding 59
Figure 6-3: Throughput when using encoded BH as compared to an uncoded BH and employing a SISODQ or a ERD 60
Figure 6-4: Partial functional architecture implemented on SDN-Testbed 61
Figure 6-5: Handover time CDF 61
Figure 6-6: Total handover processing time 62
Figure 6-7: Achievable Energy Savings vs. Thresholds 63
Figure 6-8: Achievable Energy Savings vs. Switching-off period 64
Figure 7-1: Utilization gains in different network domains 66
Figure 7-2: Computational utilisation efficiency as a function of number of cells and outage probability 67
Figure 7-3: Network model for cost-efficiency analysis 68
Figure 7-4: Expected capital expenditures as a function of data center density 69
Figure A-1: Power consumption of different network elements 80

List of Tables

Table 4-1: Wide Area Coverage settings....20
Table 4-2: Stadium settings....21
Table 4-3: Square settings....23
Table 4-4: Wide Area Coverage settings....25
Table 4-5: Shopping Mall / Airport settings....29
Table 4-6: Performance metrics addressed in each testbed platform....34
Table 5-1: 3GPP timing requirements [35]....46
Table 6-1: Overview of interfaces for CTs investigated in WP2 and WP3....55
Table 6-2: 3GPP standardized QCI values....57
Table 6-3: IEEE 802.1Q Priority Code Point recommendations [22]....57
Table 6-4: Backhaul Classification [48]....58
Table 7-1: Description of iJOIN system evaluation scenario based on stadium scenario....71
Table 7-2: Description of iJOIN system evaluation scenario based on wide-area scenario....71
Table 7-3: Application of iJOIN candidate technologies to iJOIN system evaluation scenarios....72
Table 7-4: Interaction of CTs operating on PHY layer....72
Table 7-5: Interaction of CTs operating on MAC layer....73
Table 7-6: Interaction of CTs operating on network level....73
Table A-1: Power consumption model for the iSC and exemplary realistic parameter values....81
Table A-2: Power consumption model for the RANaaS and exemplary realistic parameter values....82
Table A-3: Power consumption model for the Backhauling and exemplary realistic parameter values....83

Abbreviations

3GPP	3rd Generation Partnership Project
ADSL	Asymmetric Digital Subscriber Line
AM	Acknowledged Mode
AMM	Anchoring and Mobility Management
API	Application Programming Interface
ARM	Advanced RISC Machines
ARQ	Automatic Repeat Request
ASIC	Application Specific Integrated Circuit
ATM	Asynchronous Transfer Mode
BB	Base Band
BH	Backhaul
CAPEX	Capital Expenditures
CAS	Computational Aware Scheduler
CDF	Cumulative Distribution Function
CoMP	Cooperative Multi-Point
CPRI	Common Public Radio Interface
CPU	Central Processing Unit
C-RAN	Centralized RAN
CS	Common Scenario
CSI	Channel State Information
CT	Candidate Technology
DCI	Downlink Control Information
DFT	Discrete Fourier Transform
DHCP	Dynamic Host Configuration Protocol
DL	Downlink
DSP	Digital Signal Processor
EE	Energy Efficiency
eNB	Evolved Node B
EPC	Evolved Packet Core
ERD	Error Resilient Decoder
E-UTRAN	Evolved Universal Terrestrial Radio Access Network
FDD	Frequency Division Duplexing
FEC	Forward Error Correction
FER	Frame Error Rate
FFT	Fast Fourier Transform
FPGA	Field Programmable Gate Array
GOPS	Giga Operations per Second
GFLOPS	Giga Floating Point Operations per Second
GPON	Gigabit Passive Optical Networking
GPP	General Purpose Processor
GWCN	Gateway Core Network
HARQ	Hybrid Automatic Repeat Request
IaaS	Infrastructure as a Service
iLGW	iJOIN Local Gateway
iNC	iJOIN Network Controller
IP	Internet Protocol
ISA	Instruction Set Architecture
iSC	iJOIN Small Cell
ISS	Industry Standard Server
IT	Information Technology
iTN	iJOIN Transport Node
ITU	International Telecommunications Unit
iveC	iJOIN virtual eNodeB Controler
KPI	Key Performance Indicator
KSR	Kendall Square Research

KVM	Kernel-based Virtual Machine
LOS	Line-of-Sight
LTE	Long Term Evolution
MAC	Medium Access Control
MCS	Modulation and Coding Scheme
MIMO	Multiple Input Multiple Output
MIPS	Microprocessor without Interlocked Pipeline Stages
MME	Mobility Management Entity
MOCN	Multi-operator Core Network
MPI	Message Passing Interface
MPLS	Multi-Protocol Label Switching
MPTD	Multi-Point Turbo Detection
MRS	Maximum Rate Scheduler
MUD	Multi-User Detection
NEO	Network Energy Optimizer
OpenMP	Open Multi-Processing
OpenMPI	Open Message Passing Interface
OPEX	Operational Expenditures
OS	Operating System
OSS	Operations Support System
PA	Power Amplifier
PCP	Priority Code Point
PDCP	Packet Data Convergence Protocol
PDU	Packet Data Unit
PHY	Physical Layer
PLMN	Public Land Mobile Network
POSIX	Portable Operating System Interface for Unix
QCI	QoS Class Indicator
QinQ	Q-in-Q; refers to IEEE 802.11ad
QoE	Quality of Experience
QoS	Quality of Service
OSS	Operation Support System
RAM	Random Access Memory
RAN	Radio Access Network
RANaaS	Radio Access Network as a Service
RAP	Radio Access Point
RAT	Radio Access Technology
RF	Radio Frequency
RLC	Radio Link Control
RNC	Radio Network Controller
RoF	Radio over Fibre
RRC	Radio Resource Control
RRM	Radio Resource Management
RRH	Remote Radio Head
RTT	Round-Trip Time
SA1	System Architecture, WG 1
SDN	Software Defined Networking
S-GW	Serving Gateway
SISODQ	Soft-Input Soft-output Dequantizer
SMP	Symmetric Multi-Processor
SNR	Signal-to-Noise Ratio
SPTD	Single Point Turbo Detection
TCO	Total Cost of Ownership
TCP	Transport Control Protocol
UE	User Equipment
UK	United Kingdom
UL	Uplink

UM	Unacknowledged Mode
vCPU	virtual Central Processing Unit
veNB	virtual evolved Node B
VM	Virtual Machine
VPN	Virtual Private Networks
WG	Working Group
WP	Work Package

1 Executive Summary

This report presents an overview of the activities carried out by the work package 5 (WP5) during the second year of the project. The previous report D5.1 [3] summarized the activities of the first year and focused on a comprehensive analysis of the state-of-the-art and a detailed derivation of the iJOIN architecture including a logical, functional, and physical architecture. In this report, we focus rather on the two core innovations of iJOIN, i.e. functional split and joint RAN/BH operation. In addition, we detail how iJOIN will perform a project-wide evaluation of novel technologies.

The first part of this report, Section 4, provides an update of the iJOIN common scenarios (Stadium, Square, Wide-area continuous coverage and Shopping Mall / Airport) with a detailed description of physical deployment options for RAN/BH as well as RANaaS, traffic demand assumptions, and performance evaluation parameterization. A summary of main scenarios and their physical architecture is provided; this is important in order to determine how the logical interfaces map to physical constraints and requirements. This allows for characterizing the individual logical interfaces.

Furthermore Section 4 gives a detailed overview of how network sharing impacts the iJOIN architecture and how iJOIN facilitates network sharing, particularly in a small-cell environment. Network sharing is discussed as a possible option for operators, starting from a description of use cases and support in 3GPP, and with a subsequent discussion on potential benefits for operators. Finally, proof-of-concept assumptions are presented, together with a set of metrics relevant for the assessment of testbed platforms.

Section 5 details different aspects of the functional split. First, it provides a comprehensive overview of implementation aspects and how different hardware options impact the implementation of RAN functionality. We further discuss a virtualized infrastructure which may have a significant impact on how algorithms are implemented, how they interact with each other, and how they can be scaled with the RAN. Virtualized environments hide resource constraints efficiently through virtualized interfaces. However, these constraints still need to be considered and shall be exploited in a centralized RAN environment. Load balancing is an option to exploit large-scale computing resources more efficiently. We further detail how well known concepts from cloud-computing will affect the RAN operation, e.g. how migration of virtual machines and therefore virtual eNodeBs can be implemented. Constraints originating from computing platforms are linked with constraints originating from the RAN, e.g. latency and throughput requirements. In addition, Section 5 explores preferred splits of RAN functionality and how these splits may be implemented flexibly. In particular, the flexible split of RAN functionality needs to consider practical constraints in small-cell networks.

In Section 6, this report focuses on the second main iJOIN innovation, i.e. joint RAN and BH operation. First, this part elaborates on interfaces which are required to support joint RAN/BH technologies. This is described individually for all candidate technologies and their interaction with other logical entities is explained. In addition, results for joint RAN/BH coding are provided to exemplify how joint RAN/BH operation can improve the system performance.

Finally, Section 7 defines main objectives and metrics which are applied in iJOIN, i.e. energy-efficiency, cost-efficiency, utilization-efficiency, and area throughput. In particular, the first three metrics were updated in order to provide a consistent framework for the final evaluation towards the end of the project. Moreover, this section presents the methodology used by the project for the iJOIN concept evaluation by summarizing the interaction of CTs within each layer, the interaction across layers, and by providing an overview of feasible operating points of the iJOIN system.

2 Introduction and key contributions

2.1 Background and Scope

The iJOIN project aims at providing a solution for heterogeneous small-cell based networks to incorporate partially centralized radio access network functionality. This centralization will improve the radio access performance through advanced processing such as joint transmission and reception. It will further improve the energy-efficiency through pooling gains at the central processor. In addition, the usage of a central processor based on commodity hardware will improve the cost-efficiency. Exploiting multi-user, traffic, and computational diversity further allows for improved utilization efficiency.

In order to implement the iJOIN vision, two main innovations need to be further developed, i.e. flexible functional split and joint RAN/BH operation and design. The previous deliverable D5.1 [3] provided a basic understanding for the requirements of both technologies and how the iJOIN architecture must be designed in order to allow for an efficient evolution towards the iJOIN system. In this report, the actual implementation of both innovations and resulting requirements within iJOIN's main scenarios is put in focus. The main challenges for an implementation of the flexible functional split are the question for the right usage of different hardware options, how RAN functionality can be implemented in a virtualized environment, how 3GPP LTE RAN constraints impact the implementation of the functional split, and which functional splits should be preferred. In the case of joint RAN/BH operation, the main challenge is the interfaces definition and the clarification of how the individual components in RAN and backhaul will interact.

The iJOIN project further introduces a set of novel technologies which improve the following individual performance objectives: energy-efficiency, cost-efficiency, utilization-efficiency, and area throughput. While each technology on its own may improve the performance, it may also impact other technologies and deteriorate or emphasize their improvements. Hence, it is important to understand how these technologies interact, whether they are complementary and contradicting, how their gains will be combined, and how they are integrated in the two main concepts functional split and joint RAN/BH operation. This harmonization work is the main task of work package 5 and will use the output from work packages 2, 3, and 4 where novel technologies for physical layer, medium access and radio resource control layer, and for the network operation are derived, respectively. This report is the first step towards this harmonization.

At the end of the project, a comprehensive and consistent evaluation of the iJOIN architecture and system performance will be provided. In order to avoid loosely coupled results from individual candidate technologies, a harmonized simulation campaign is required. This could be achieved through different means, e.g. a joint simulation effort where all partners apply the same simulation framework, a joint calibration effort where all partners calibrate their individual simulation tools, or a joint parameter derivation where all partners apply the same set of parameters to both the novel technologies and the baseline system. In iJOIN, the last option has been chosen due to resource constraints. The first two options require substantial resources. In the case of iJOIN, relevant parameters for each main scenario are derived and all candidate technologies incorporate these parameters. In a next step, for each of these parameters a range of meaningful values has been defined and is used for the evaluation of each candidate technology. The comparison of candidate technologies is done based on a relative basis, i.e. each technology is compared to the baseline system and then relative gains across multiple candidate technologies are compared.

2.2 Key Contributions

This report provides first the functional architecture definition covering an assessment of interaction across candidate technologies and work packages. This includes the description of how candidate technologies impact different objectives, which is essential to perform a project-wide assessment at the end of the project.

We further provide a detailed analysis of the proposed functional split concept, including implementation aspects resulting from different hardware options, impact of virtualized infrastructure, and how data processing complexity can be measured. Analytical results show how data processing complexity in a 3GPP LTE RAN system scales, how the centralization gain is reflected, and how the central processor can be dimensioned.

Furthermore, a comprehensive overview of functional splits and practical constraints of a flexible implementation are detailed. Constraints from the cloud-computing platform and 3GPP LTE RAN are related through the required data processing capabilities for a given quality of service. Analytical results for

3GPP LTE RAN show that latency constraints at the interface of physical and medium access layer can be efficiently mitigated without performance loss. Also, results for joint RAN/BH coding show how coding across both domains, distributed IP anchoring, and network-wide energy optimization can improve the system-performance. Moreover, the descriptions of evaluation metrics are updated.

Finally, this report provides a detailed description of the simulation campaign, the individual objectives and how they are going to be measured. Furthermore, first results for the utilization efficiency of cloud-RAN with higher layer split are provided.

3 Definitions

This section presents the concepts and definitions used within iJOIN to guarantee a common understanding of all partners and external readers about the glossary used across the document.

Standard Terms

Radio Access (RA): Wireless radio frequency (RF) link between the User Equipment (UE) and the Radio Access Network (RAN).

Radio Access Network (RAN): Network elements and functions required to support the Radio Access operation (E-UTRAN).

Small Cell (SC): Low power base station with intelligence, part of the Radio Access Network. A Small Cell

- is an operator-controlled equipment;
- supports the same functionalities as an evolved Node B (eNB);
- can be deployed indoors or outdoors;
- can be within or outside the coverage of a macro-cell.

Radio Remote Head (RRH): Radio frequency processing unit without intelligence (e.g. optical to radio conversion), i.e. Radio-over-Fibre.

Backhaul (BH): Links connecting the Radio Access Network (E-UTRAN) and the Core Network (EPC).

Backhaul Network: Network elements and functions required to support the backhaul operation.

Fronthaul (FH): Link within the Radio Access Network allowing a distributed implementation of the RF layer of a base station, e.g. optical link between baseband processing units and RRHs. As a convention in iJOIN, all links within the Radio Access Network allowing for a distributed implementation of the upper OSI layer(s) (L1/L2/L3) will be referred as backhaul.

Figure 3-1 illustrates the mapping of the previous generic definitions on the 3GPP LTE architecture, and Figure 3-2 describes the backhaul network. They do not represent the iJOIN architecture, but the existing architecture on top of which iJOIN will provide its evolutionary path.

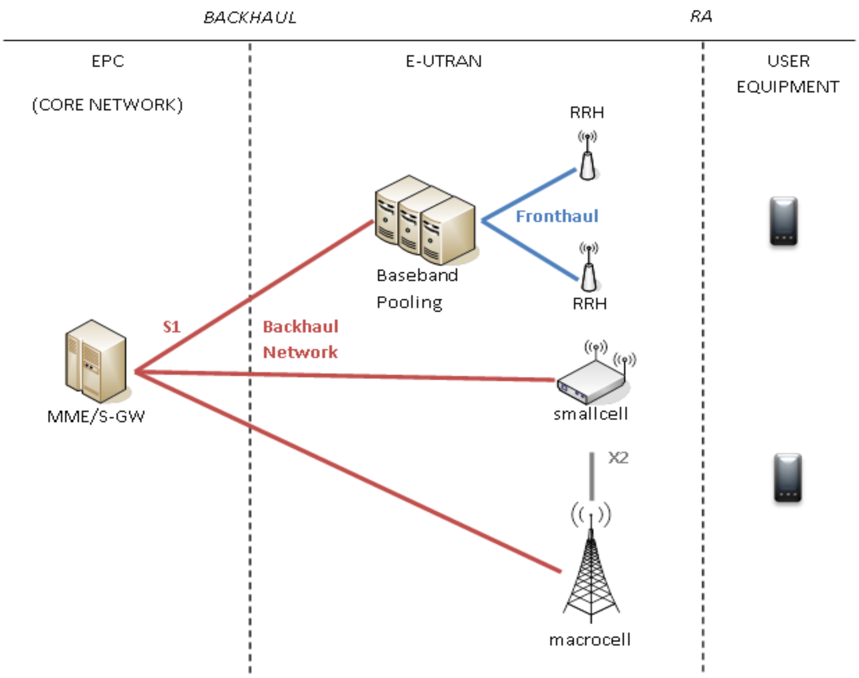

Figure 3-1: "Generic" Mobile Network Architecture

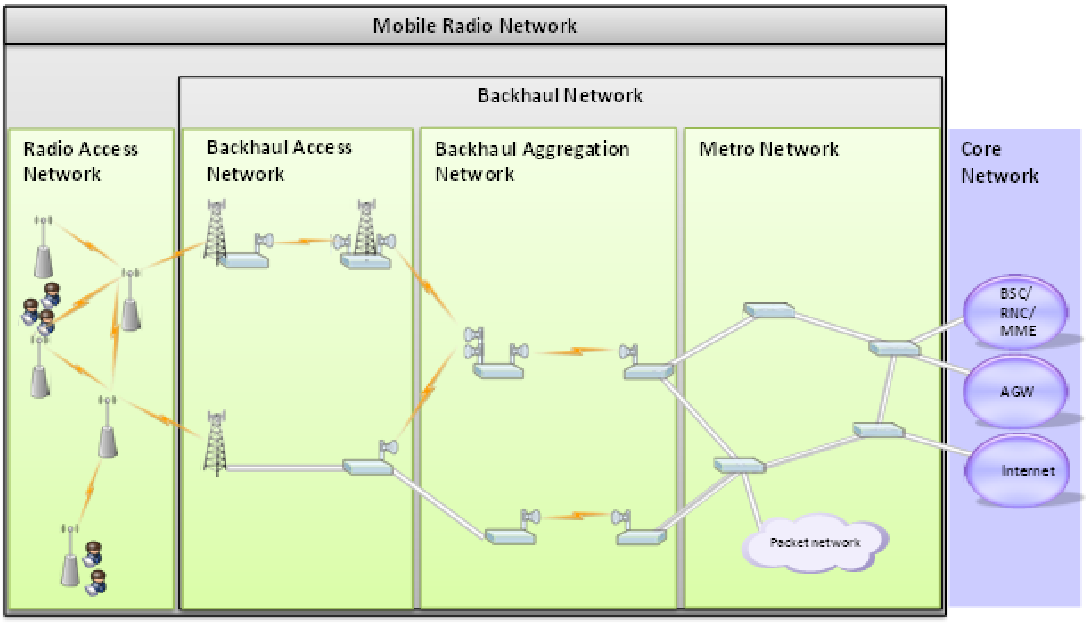

Figure 3-2: Mobile Radio Network Definition

iJOIN-specific Terms

RAN as a Service (RANaaS): concept introduced by iJOIN. Indicates a set of computing and storage infrastructure resources (typically a Cloud computing IaaS platform, but potentially any industry standard, general purpose computing infrastructure) where, according to the technology developed inside iJOIN, part of the processing functions of the lowest OSI layer(s) (L1/L2), normally distributed across a number of base stations (eNodeB), are moved and centralized.

RANaaS platform: technology baseline implementing one or more specific instances of RANaaS. For example, in the case of cloud IaaS based RANaaS, the RANaaS platform encompasses the physical infrastructure resources (e.g., servers, storage) and the cloud management software.

RANaaS instance: a specific implementation of a RANaaS platform, serving a set of iJOIN small cells with which it makes up a virtual eNB (see below).

RANaaS Point of Presence (RANaaS-PoP or iPOP): physical location where one or more RANaaS implementations are deployed and executed. Examples of RANaaS-POP locations might include fully owned enterprise datacenters, co-locate datacenters, and virtual private clouds. Different RANaaS instances can potentially be placed in a common RANaaS-PoP, assuming that there are mechanisms ensuring a full separation among them.

iJOIN Small Cell (iSC): logical entity introduced by iJOIN. Low power flexible radio access point implementing fully or partially the lower OSI layer(s) (RF/L1/L2/L3) of a base station, along with the RANaaS platform. Apart from the functional split concept introduced by iJOIN, an iSC has all the same properties of a standard small cell. An iSC is connected to the RANaaS platform through the logical J1 interface, and to another iSC through the logical J2 interface.

virtual eNB (veNB): logical entity enclosing the set of functions and interfaces which, in the iJOIN architecture, correspond to the implementation of an eNodeB according to 3GPP specifications. A veNB is composed by a RANaaS instance and one or more iSCs. Functions and interfaces are not necessarily executed or placed in the same physical or logical network entity.

iJOIN veNB Controller (iveC): logical sub-entity of a veNB, located in the RANaaS platform, responsible for functional distribution across the veNB, consistent execution of the distributed functionalities, management and configuration of the different veNB components.

iJOIN Network Controller (iNC): logical entity in charge of controlling the joint RAN/BH operation. To minimise the impacts for the operator in terms of deployment cost and complexity, the iNC may be physically co-located with the RANaaS instance.

iJOIN Local Gateway (iLGW): logical entity implementing a subset of the functions of a P-GW. It is logically connected with a (v)eNB, but can be physically located anywhere in the RAN.

iJOIN Transport Node (iTN): physical entity located between iSC and RANaaS, or between RAN and core network. Each iTN is essentially a transport node operating at a different protocol stack layer depending on the particular functional split; a set of iTNs is forming a backhaul network whose forwarding plane can be configured by an iNC.

4 iJOIN Common Scenarios

The following subsections describe the four physical architectures identified by respective Common Scenarios defined in iJOIN (see also D5.1 [3]). In general we can notice that every physical architecture differs in terms of deployment scale, number of nodes and particular placement of physical interfaces (realizing logical connections in different ways). Moreover, in all scenarios a RANaaS instance is coordinating iSCs and its implementation on a cloud platform may consist of many Virtual Machines (VMs) each representing the baseband processing units of the coordinated iSCs.

According to the definition of the veNB, a set of cells (each one with corresponding ID cell) belongs to the same veNB. Figure 4-1 exemplarily shows two scenarios for the implementation of veNBs.

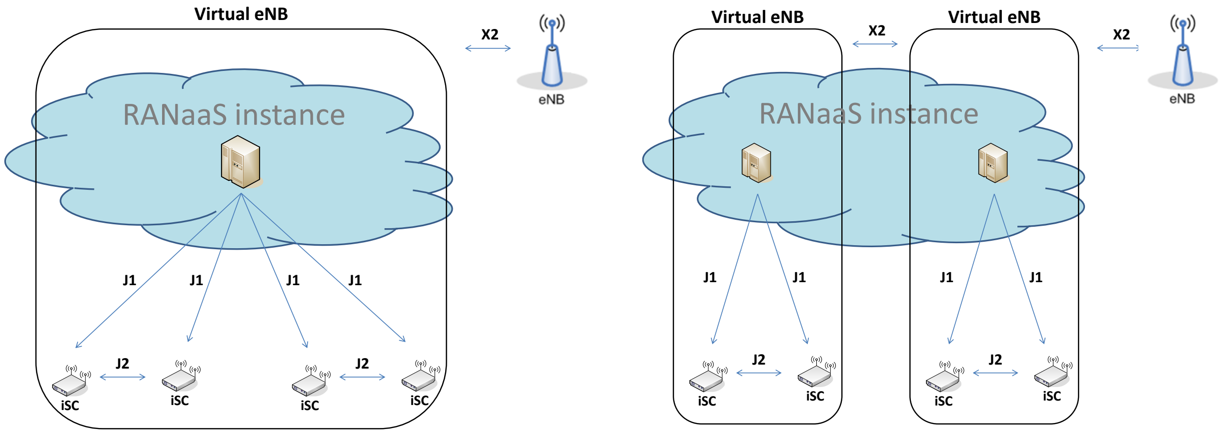

(a) A single virtual eNB in the RANaaS instance (b) Several virtual eNBs in parallel in the RANaaS instance

Figure 4-1: RANaaS and virtual eNodeB configuration options

If more than one virtual eNB (Figure 4-1(b)) is executed at the same RANaaS instance, it may need to involve the X2 interface in order to realize coordination. Each veNB is seen from the core network as an eNB and can communicate with other (v)eNBs through the 3GPP LTE X2 interface. In this case, in the view of realizing fast coordination among the cells, also X2 protocol limitations should be taken into account.

4.1 Common Scenario 1: Stadium

4.1.1 Physical Description

This scenario considers the coverage of a stadium with small cells during a sport event such as a football match. A typical stadium covers an area in the order of 50.000m^2, and can contain several thousands of users. The average number of users that can be taken into account is 40.000. In this case the communication network should comprise several small cells to support the high traffic demand (mainly media content). Users are characterized by reduced mobility and due to the outdoor nature of the deployment macro coverage can cause interference in the case of co-channel deployment between macro and small cells.

4.1.2 Traffic Demand

The average number of spectators that can be taken into account is 40.000. By assuming to guarantee an average DL/UL throughput of 10Mbps during peak of traffic with 5% of active UEs, the network should be able to manage a DL/UL throughput of 20Gbps. The general traffic parameters are reported in the following table.

Table 4-1: Wide Area Coverage settings

General Parameters	**Value**
UE density	~ 1 UE/m^2
Number of small cells	50 -300
UL/DL Throughput	1.5 – 10 Mbps/UE
Traffic density	0.1 – 0.4 Mbps/m^2

4.1.3 Proposed RAN/Backhaul and RANaaS deployment

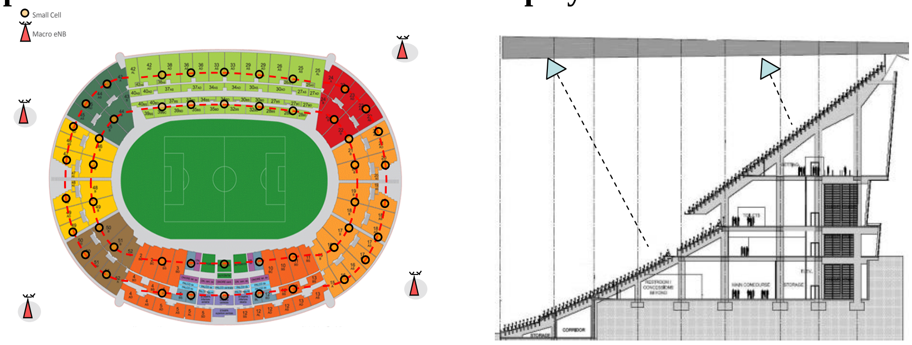

Figure 4-2: Stadium – iSCs and macro cell positions (left) and details on iSCs antenna tilt (right)

The key characteristics of this scenario are:

- Multiple rings of iSCs providing coverage in the stadium (two rings are envisaged in Figure 4-2).
- Multiple macro cells can be present to provide sufficient overage also outside the stadium, as depicted in Figure 4-2.
- All iSCs and the macro eNBs are coordinated by one iNC node controlled by the same RANaaS data centre.
- A tight coordination is envisaged among the iSCs. A loose coordination between macro and iSC layers can be considered under the control of iNC node.

Based on these characteristics, Figure 4-3 illustrates a possible physical deployment.

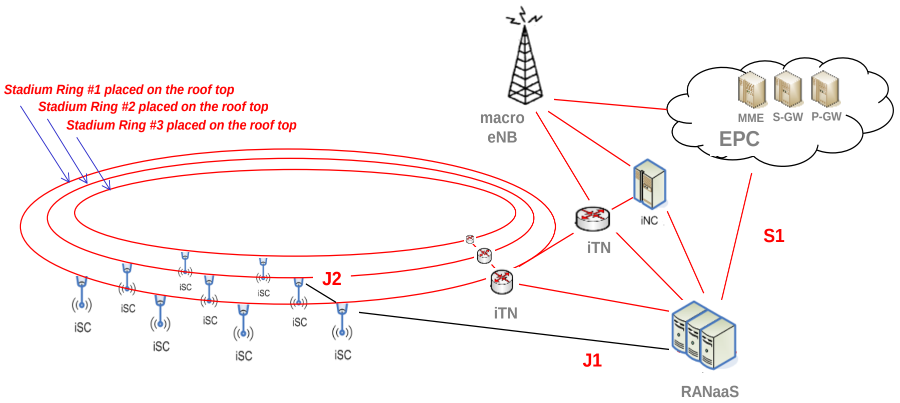

Figure 4-3: Stadium – Physical deployment example

4.1.4 Performance evaluation parameterization

In this section we present the model to assess the iJOIN CTs in CS 1 (the stadium). Performance evaluation in the whole stadium is not feasible. Hence, we focus on a limited area of the stadium. The considered stadium layout is illustrated in Figure 4-4.

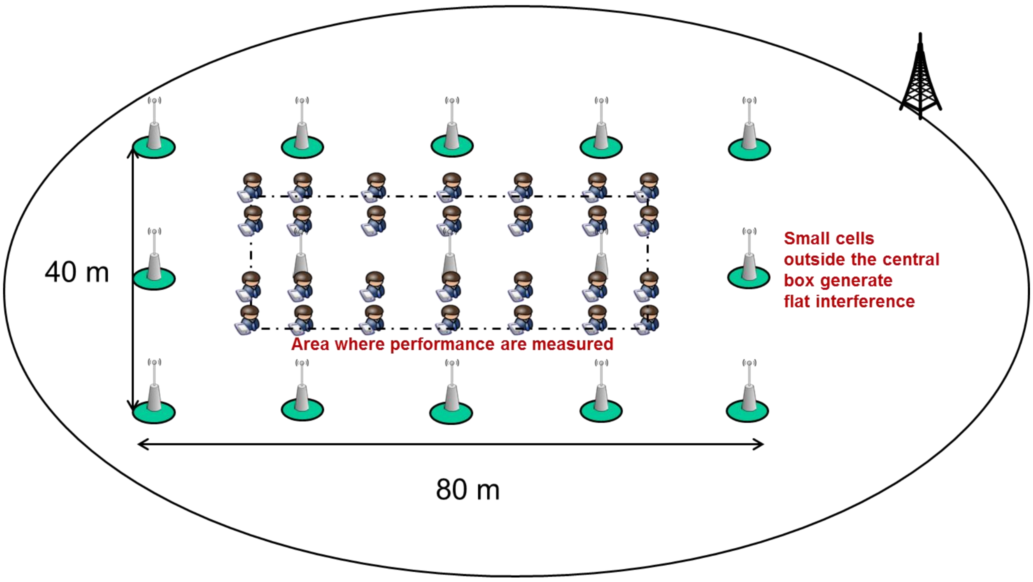

Figure 4-4: Stadium Layout in high load scenarios

Furthermore, the corresponding evaluation parameters which applied across all work packages are detailed in Table 4-2.

Table 4-2: Stadium settings

Parameters	Stadium
Number of small cells per cluster	15 but focus on central 3
Number of UEs	320 (high load, 5% active UEs) 64 (medium load, 1% active UEs)
Covered area	40 m x 80 m Uniform dropping
Minimum distance 3GPP TR 36.872 [1]	iSC-iSC 20 m UE-iSC 5 m Macro eNB-iSC cluster center 105 m
Backhaul Capacity / Latency	100 Mbps / 1-10 ms 200 Mbps / 1-10 ms 10 Gbps / 5 µs

4.2 Common Scenario 2: Square

4.2.1 Physical Description

This scenario is illustrated in Figure 4-5 and considers a square area which encloses an outdoor environment, surrounded by cafes, shops and recreation parks. Typical squares are busy almost all the day with thousands of people traversing them or visiting them in order to relax and meet with other people.

Figure 4-5: Square Use Case

In this densely populated environment, user mobility and requirements imposed by multimedia broadband services must receive particular attention in order to provide a uniform quality of experience. To provide these services, small cells will be densely deployed in an unplanned nature within or surrounding the square in non-traditional locations (i.e. lamp-posts, utility poles). Also, macro-cell coverage is expected, which may require extra coordination between small cells and macro cells for interference and mobility management.

4.2.2 Traffic Demand

In this scenario, we can observe people with different traffic demands. There are people who want to use their smartphones, tablets, and laptops to watch online video, read email, or interact in social networks with low or no mobility. We can also observe people traversing the square who mainly use their terminals for voice services and messaging, e.g. Email or text messages, while walking.

The traffic demand in a town square strongly depends on the size of the square which might vary from 20.000 m^2 to 100.000 m^2 in typical scenarios. Consider an exemplary typical square of 50.000 m^2. In this case, the number of users may range from 1000 (off-peak hours) to 5000 (busy hours). Assuming about 5% active UEs with 1.5-10Mbps total throughput demand, the overall throughput for one square will range from 1.5Gbps to 50Gbps. In order to meet this demand, a multitude of small cells must be deployed to enhance the spatial re-use.

4.2.3 Proposed RAN/Backhaul and RANaaS deployment

Figure 4-6 illustrates an example of square deployment. The key characteristics of this are:

- The RAN deployment for the square is based on a dense random deployment of iSCs.
- All deployed iSCs within one area are connected to the same RANaaS datacenter. Hence, all iSCs from the square are processed at the same RANaaS datacenter.

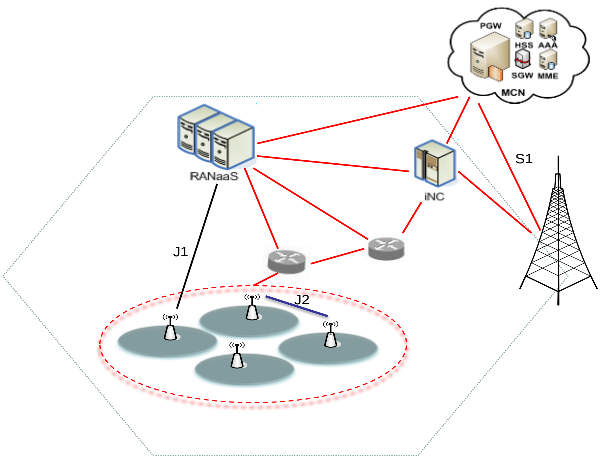

Figure 4-6: Square - Physical Deployment example

4.2.4 Performance evaluation parameterization

In this section we present the model to assess the iJOIN CTs in the CS 2 (the square). The square layout is based on the small cell deployment described by 3GPP in TR 36.872 (A1.1 and A1.2) [1].

The main characteristics of the square hotspot are:

- Random small cell deployment
- Random user deployment
- Static/Nomadic user
- Heterogeneous backhaul

The layout for the evaluation of the square scenario is shown in Figure 4-7 and further details are listed in Table 4-3.

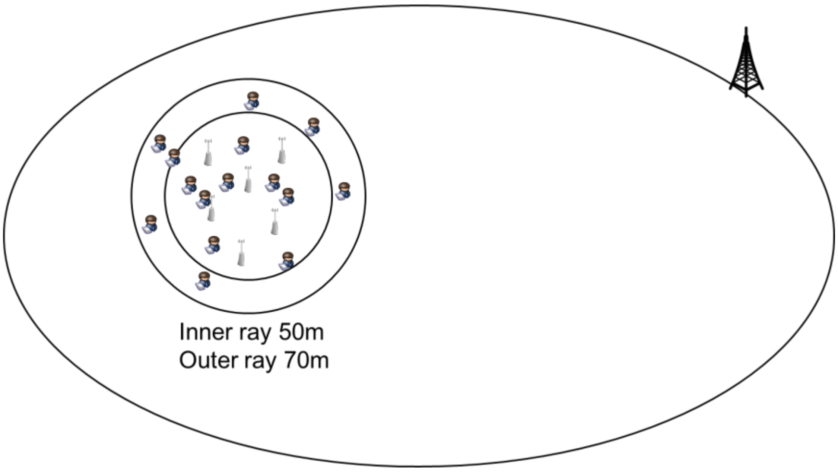

Figure 4-7: Small cell deployment in the square.

Table 4-3: Square settings

Parameters	Square
Number of small-cells	4-10 (sparse to dense deployment)
Number of UEs	15-30 (lightly to highly loaded scenarios)
Radius for small cell dropping in a cluster	50 m (3GPP TR 36.872 [1]) Random Dropping
Radius for UE dropping in a cluster	70 m (3GPP TR 36.872 [1]) Random Dropping
Minimum distance **3GPP TR 36.872 [1]**	iSC-iSC 20 m UE-iSC 5 m Macro eNB-iSC cluster center 105 m
Backhaul Capacity / Latency	~50 - 100 Mbps, 1-10ms 10 Gbps, 5µs

4.3 Common Scenario 3: Wide-area continuous coverage

4.3.1 Physical Description

Most commonly, the deployment of small cells is associated to the existence of hotspots and high demand areas, on one hand, or coverage holes of the macrocell layer, on the other hand. This scenario is intended to explore a different context where the small cell layer is not specifically associated to high demand points but is expected to provide continuous coverage over a wide area, representing an alternative or a complement to the macrocell layer. The small cells are expected to support both indoors and outdoors users although in the latter case only for a limited mobility level.

The scenario would be characterized by the regular distribution of the small cells over a wide urban area. There are several deployment options that can ensure the continuous coverage design objective:

- Use of urban furniture, like lampposts, bus stop marquees, telephone booths.
- Use of infrastructure provided by local businesses with wall mounted cells.
- Reuse of the FTTH/xDSL fixed access infrastructure, like cabinets or wall mounted optical splitters.

Associated with these deployment options there are also different options for the backhaul that will be used, either fibre based or wireless.

4.3.2 Traffic Demand

Traffic demand in this scenario is expected to have similar characteristics to the one that is attended by the macrocell cell layer in an urban scenario, with a mixture of indoor and outdoor users, high and low mobility levels, and different kinds of services, including voice, data services, messaging services, etc. It is expected, however, that small cells will provide a higher capacity per cell than macrocells. Hence, throughputs in the order of 10-15 Mbit/s in the busy hour are considered reasonable for the scenario proposed, assuming 10 MHz bandwidth.

4.3.3 Proposed RAN/Backhaul and RANaaS deployment

In this scenario, iSCs are used to provide continuous coverage over a wide area, up to several square kilometres, preferably in an urban environment. This layer can be used to provide mobile broadband services or as an additional layer to the macro-cell layer in order to offload traffic. The key characteristics of the scenario are:

- iSCs are expected to be deployed taking into account the topography and morphology of the area to be covered (e.g., short distance LoS propagation between iSCs should be avoided as far as possible in order to reduce potential interference issues).
- Different backhaul supporting technologies may be employed, e.g. taking advantage of deployed fibre infrastructure (either for FTTH services or for other purposes) or potential line-of-site (LoS) wireless links (most likely in the mmW frequency bands) with macrocells or other aggregation points. Wireless inter-iSCs links may be considered as well.
- Connecting the iSCs with the RANaaS data centre may require an aggregation network, which imposes limitations regarding the supported functional split.
- Aggregation points can be used also to host the RANaaS infrastructure.

In terms of topology of the backhaul network, different options may be feasible: rings, stars, or even meshes. However, a hybrid topology is the most likely option, with sets of rings or meshes connected to higher hierarchical nodes through a star topology. In addition, the iNC may actively reconfigure the network topology depending on different objectives such as load balancing, energy efficiency, or congestion control. Further details on the backhaul topology are provided in report D4.1 [31].

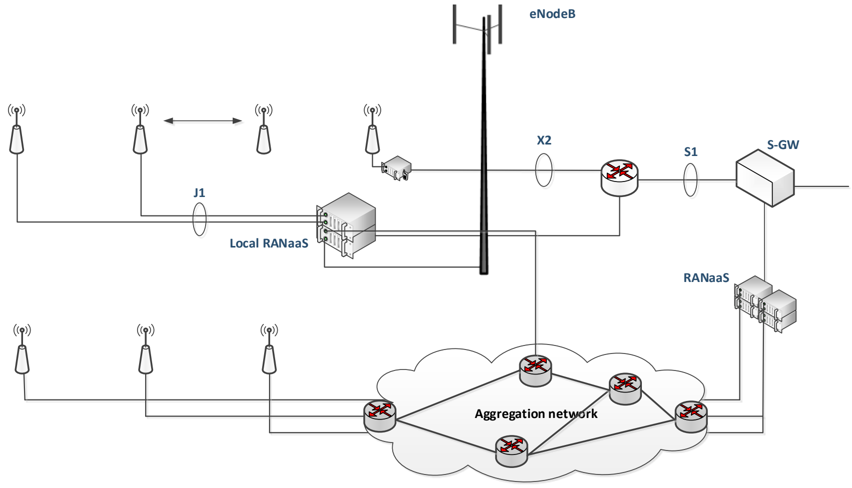

Figure 4-8: Wide Area - Physical Deployment example

4.3.4 Performance evaluation parameterization

In this section we present the model to assess the iJOIN CTs in the CS 3 (Wide Area Coverage). The Wide Area Coverage layout is based on regular small cell deployment in a hexagonal grid, which covers 1 Km^2.

The main characteristics of the Wide Area Coverage are:

- Regular small cell deployment
- Random user deployment
- Slow/High mobility
- Heterogeneous backhaul

Further details are presented in Table 4-4.

Table 4-4: Wide Area Coverage settings

Parameters	Wide Area Coverage
Number of small cells	19
Number of UEs	1-3 UE per small cell (avg.)
Small cell dropping	Regular on Hexagonal Grid
ISD	50 m
UE dropping in a cluster	Random dropping
Minimum distance	UE-iSC 5 m
Backhaul Capacity / Latency	~~50 - 100 Mbps, 1-10ms 10 Gbps, 5µs

The underlying layout for the evaluation of the Wide Area Coverage scenario is shown in Figure 4-9.

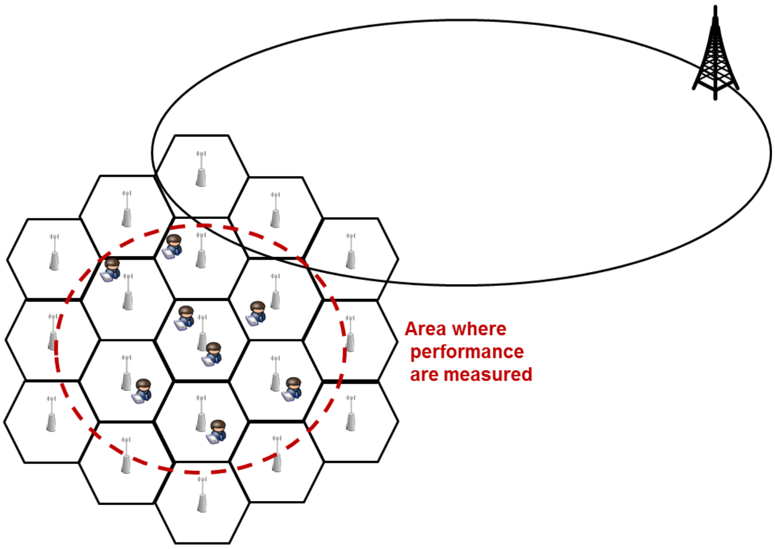

Figure 4-9: Small cell deployment for the Wide Area Coverage scenario

4.4 Common Scenario 4: Shopping Mall / Airport

4.4.1 Physical Description

This scenario considers hotpots deployed in dense indoor environment such as an airport or a shopping centre. In modern airports as shown in Figure 4-10, the waiting rooms are often aligned in a big open space per terminal leading to a natural dense small cell deployment. Usually, only one floor is present for travellers.

Figure 4-10: Airport Use Case

Similar deployments of high user and small cell density can also be found in shopping mall environment as shown in Figure 4-11, where the small cells can be deployed either by the shopping mall owner in a planned manner or by each shop in an unplanned manner to provide additional services to the customers. A shopping mall may usually have more than one floor.

Figure 4-11: Shopping Mall Use Case

The native mobility support and seamless authentication procedure offered by the cellular network (no user-driven authentication procedure to start) make cellular network technology the preferred choice for these scenarios. In both cases, the communication network should comprise several small cells to support the dense concentration of users in “small” areas (waiting room or shops). By contrast to the previous scenarios, macro coverage is not expected to play a major role due to the natural isolation (concrete or glass walls).

4.4.2 Traffic Demand

Similar to the Square scenario, this scenario is characterized by a large amount of users and high communication activity. Therefore, we assume a range of 200 to 500 active users in one hall. Given an expected throughput demand of 1 – 10Mbps, we consider a system throughput of 200Mbps up to 5Gbps.

4.4.3 Proposed RAN/Backhaul and RANaaS deployment

Figure 4-12 shows one feasible deployment of the iJOIN architecture for shopping malls or airports. The main characteristics are:

- All deployed iSCs within one premise are likely to be connected to the same RANaaS PoP hosting the RANaaS functionality.
- One central “gateway” will most likely serve as a physical concentrator of all backhaul links toward the RANaaS PoP. This gateway could host several “lines” (fibre or xDSL).
- All iSCs within one premise are connected through a heterogeneous backhaul including Ethernet, wireless, or GPON to the EPC.
- It is likely that the backhaul will be made of three main components:
    1. From iSC to a central “gateway”: This link will likely use the existing airport or shopping mall infrastructure which could be wired (fibre/GPON or Ethernet-based). Low latency and high bandwidth links are expected between the iSCs within the premises.
    2. From the central “gateway” to the RANaaS PoP: This link could use either fibre or xDSL lines with bounding to support the aggregated capacity needed. It is likely to be the physical link which will impose the overall latency and throughput constraints if wired links are used within the premises (see previous point).
    3. From the RANaaS PoP to the EPC: This link will likely be fibre-based and should not be a source of limitations. Of course, a smart dimensioning should be taken into account by the operator in order to cope for the traffic toward the EPC.
- Line-of-sight between iSC and user terminal (and iTN and iSC in case of wireless backhaul) may be feasible.

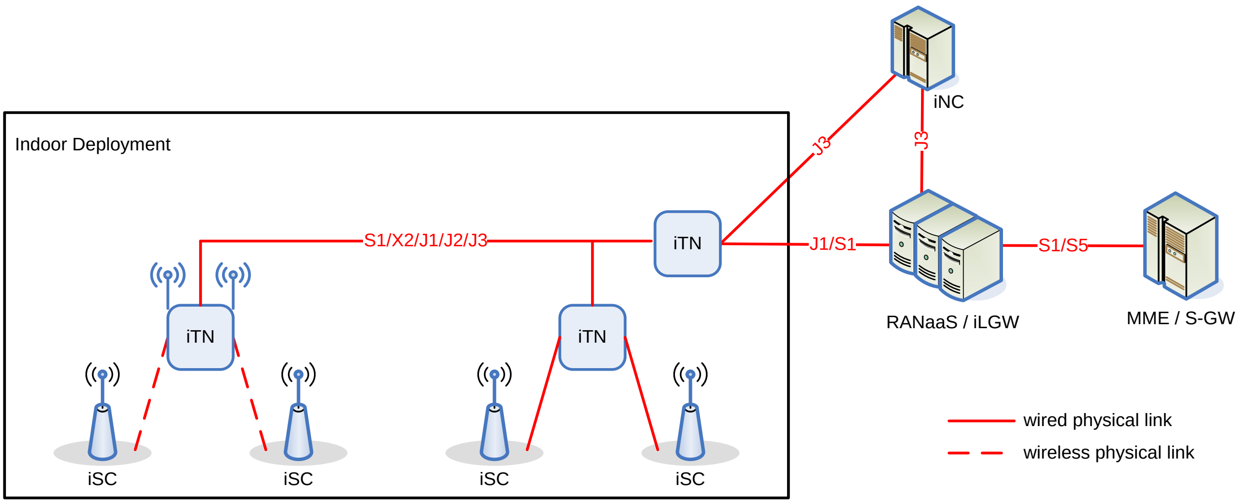

Figure 4-12: Shopping Mall / Airport: Physical deployment example

4.4.4 Performance evaluation parameterization

In this section we present the evaluation assumptions applied to assess the iJOIN CTs in this scenario. This layout is based on the ITU indoor small cell deployment described in [2] (annexes A1.5 and A1.6) [1]. Two layouts are considered, with sparse and dense small cell density, respectively. The main characteristics are:

- Regular small cell deployment
- Random user deployment
- Nomadic user
- Wireline backhaul (optical fibre and ADSL)

The layout which is considered for the evaluation of this scenario is shown in Figure 4-13.

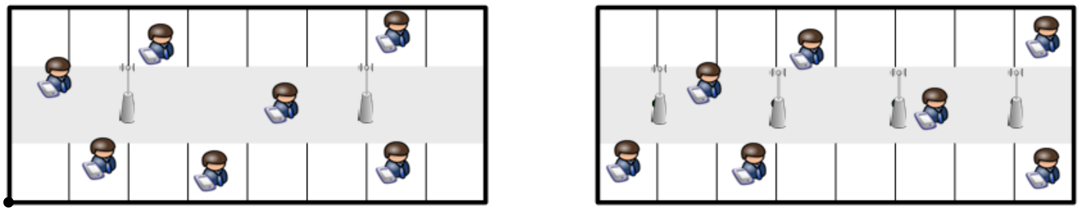

Figure 4-13: Small cell deployment in the Shopping Mall / Airport: Sparse (left) and dense (right) deployment.

Furthermore, Table 4-5 provides details on the evaluation assumptions.

Table 4-5: Shopping Mall / Airport settings

Parameters	Shopping Mall / Airport
Number of areas per floor	16 Areas;1 or 2 Floors
Floor height	6 m
Area size	15 m X 15 m
Hall size	120 m X 20 m
Number of small cells	2 (sparse) / 4 (dense) per floor
Small cell dropping	Regular
Number of UEs	10 per small cell (sparse) 5/10 per small cell (dense)
UE dropping	Random
ISD	60 m (sparse) / 30 m (dense)
Minimum distance	UE-iSC 3 m
Backhaul Capacity / Latency	>100Mbps per iSC-EPC link <1, 10, and 50 ms

4.5 Network Sharing Enablers

As the traditional model of single ownership of all network layers and elements is being challenged, network sharing is emerging as a mechanism for operators to substantially improve network costs and to efficiently utilize network capacity. More and more operators are adopting network sharing as a means of cutting Capital Expenditure (CAPEX) and Operating Expenditure (OPEX) costs involved in the initial roll-out and operation of mobile networks. The main motivations for operators to adopt network sharing schemes are:

- Increased rollout speed
- Quickly expand coverage to meet customer demand for wider coverage
- Sharing low-traffic areas leads to long-term cost advantages
- Sharing high-license obligations
- Cost efficiency (CAPEX and OPEX)
- Joint effort to offer availability of services more cost-efficiently

4.5.1 Network Sharing in the context of 3GPP

3GPP has been working on providing standardised solutions for different alternatives of RAN sharing. The main milestones are collected in Figure 4-14.

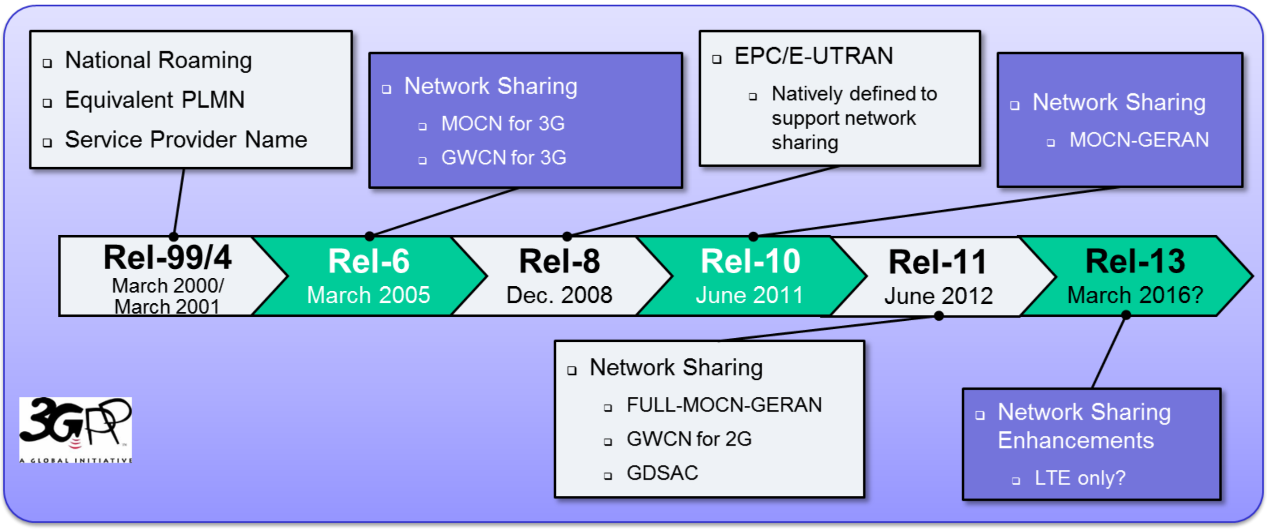

Figure 4-14: 3GPP support of network sharing

In general, the solutions supported differ in terms of the level of infrastructure integration between operators, from roaming agreements to complete network (both access and core) sharing. Most of the benefits are usually associated to RAN sharing (where iJOIN is focused), which is responsible for most of CAPEX and OPEX. In Figure 4-15, different possible degrees of integration of network sharing solutions are shown. They range from roaming agreements between network operators to full RAN and core network sharing.

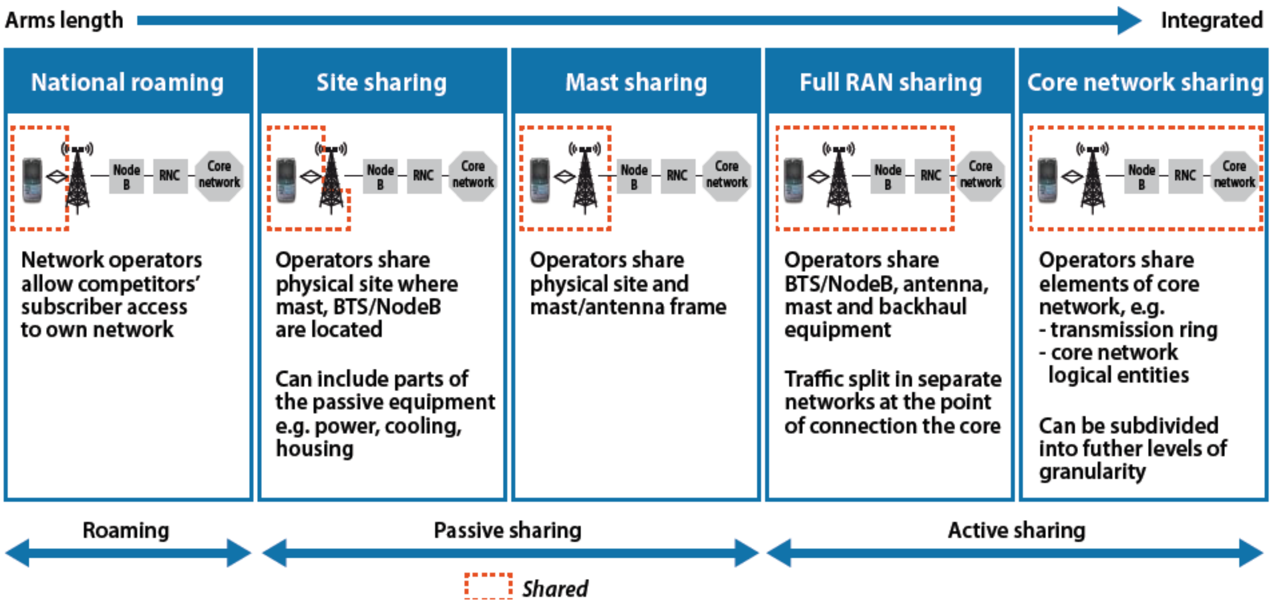

Figure 4-15: Degrees of integration in network sharing solutions

In the case of RAN sharing the 3GPP System Architecture WG SA1 (Services) specifies in [14] five main use cases for RAN sharing:

- **Sharing a common RAN**: but not the radio frequencies (Release 99). In this case the operators connect directly to their own dedicated carrier layer in the shared radio network controller (RNC) in the shared RAN.
- **Operator collaboration to enhance coverage**: where two or more operators with individual frequency licenses cover different parts, e.g. of a country, but together provide coverage of the entire country.
- **Sharing coverage on specific regions**: where one operator provides coverage in a specific geographical area with other operators allowed using this coverage for their subscribers. Outside such RAN sharing area, coverage is provided by each of the operators independently.

- **Common spectrum sharing**: considering the following two variants: one operator has a frequency license and shares the allocated spectrum with other operators, and a number of operators decide to pool and share their allocated spectral resources.
- **Multiple RANs share a common core network**: where the multiple RANs can belong to different PLMNs (public land mobile networks) and network operators. Due to operators' deployment choices, different nodes or part of the common core network can be shared.

Active RAN sharing enables partitioning or pooling of radio resources enhancing the overall RAN utilization. At the same time, investments for installing new infrastructure may be reduced as well. In 3GPP, WG SA1 conducted a study on RAN sharing which analyses a set of use cases and derives business requirements [15]. This study aims to outline ways for sharing RAN resources, maintaining and sharing policies, and providing flexibility in RAN resource sharing on-demand within shorter time periods. The architecture and operations that enable different mobile operators with a separate core network to share the RAN are specified by the 3GPP System Architecture WG SA2 in [16], detailing the following two approaches:

- **Multi-Operator Core Network (MOCN)**, where each operator has its own EPC providing a strict separation among the core network and RAN. This enables certain benefits regarding service differentiation and interworking with legacy networks. Shared eNBs are connected to core network elements of each operator, i.e. Mobility Management Entity (MME) and Serving-Gateway (S-GW), using a separate S1 interface and allowing load balancing policies to be provided within each operator's core network.
- **Gateway Core Network (GWCN)**, where operators share additionally the MME. This approach enables further cost savings compared to MOCN, but at the price of reduced flexibility, i.e. no mobility for inter-Radio Access Technology (RAT) scenarios and no Circuit Switching (SC) fall-back for voice traffic.

In general, MOCN is more expensive but more flexible, addressing conventional operators' needs. In both cases, the UE can distinguish up to six different operators that share the RAN infrastructure based on broadcast information, i.e. PLMN-ID, and can signal to obtain connectivity or perform a handover irrespective of the underlying RAN sharing arrangement. Specifically, the S1 interface supports the exchange of PLMN-IDs between eNBs and MMEs in order to assist the selection of the corresponding core network [17]. The X2 interface supports a similar PLMN-ID exchange among neighbouring eNBs for handover purposes [18]. Considering broadcasting, the Uu interface supports the PLMN-IDs enabling the UEs to perform the network selection [19].

4.5.2 Benefits of Network Sharing

In the framework of the iJOIN project, it is worth understanding which can be the benefits in considering network sharing in conjunction with the main enablers defined by the project, RANaaS implementation and joint access/backhauling design. The following underlying assumptions must be made:

- The main objective of sharing is to reduce costs, both capital and operational. If sharing does not result in a cost efficient solution, it probably should not be pursued. In other words, if, as expected, iJOIN technological solutions result in a reduced cost for operators, they may become inhibitors for RAN sharing.
- The possibility of sharing the spectrum is usually precluded by regulators. This may preclude the realization of some potential advantages by adopting the iJOIN architecture.

On the other hand, it is clear that the iJOIN architecture may open up new possibilities to overcome some of the issues associated to network sharing, such as the reduced flexibility for operators to differentiate from a technical viewpoint. In this sense, operators may be able to contract different network services from the RANaaS and backhaul (iNCs and iTNs) elements, e.g. supporting different functional splits and associated network services, different transport services. It may also be possible for operators to implement their own processing procedures on top of the RANaaS, even if it is shared with other operators.

In order to identify whether there are new technical requirements for iJOIN enablers for supporting network sharing or not, it is proposed to analyse a number of network sharing scenarios with different levels of integration between cooperating operators. The following conclusions can be drawn based on a preliminary analysis:

- Sharing of the RANaaS infrastructure should not be problematic because cloud infrastructure and technologies are specifically designed to allow for sharing the processing tasks to be carried out. However, it is not clear which may be the advantage of sharing it.
- Supporting flows from different operators when reusing the same transport (backhaul) infrastructure should be taken into account in the protocol design.
- The iJOIN architecture may allow operators to contract different services from a backhaul provider as well as different control capabilities. The SDN based iJOIN architecture may allow this to happen but in this case the iNC should provide an open northbound interface such that operators can configure the backhaul services they want to be provided, e.g., implementing different security mechanisms or using different local breakout points. But it must not be feasible for an operator to enhance its own performance at the expense of another operator's performance, i.e. congestion control corrective procedures should be under the control of the backhaul operator.

Common Scenarios studied by iJOIN project could be suitable for network sharing purposes, but in each cases specific considerations should be made:

- CS1 (Stadium): in this case, it may happen that backhauling infrastructure is mostly provided by the same company that owns the stadium, or at least deployed initially (e.g. with fibre optics cables); also the multiple iSCs could be already pre-installed to provide the needed coverage and capacity in the area, thus all RAN infrastructure in these cases is already physically available and ready for network sharing purposes; nevertheless typically the presence of common iSCs imply the difficulty for operators to customize their needs in terms of coverage, capacity or improve performances according to a particular preferred functional split;
- CS2 (Square): in this case small cells can be densely deployed in an unplanned nature within or surrounding the square in non-traditional locations (i.e. lamp-posts, utility poles); for practical reasons (e.g. acquisition of permits from public administration to install small cells in lamp-posts) it may happen that operators need to share some physical infrastructure, but it may also happen that for costs reason this scenario perhaps could be not always suitable for network sharing;
- CS3 (Wide-area continuous coverage): in this scenario network sharing use case could be more facilitated by the use of infrastructure provided by local businesses with wall mounted cells, and the reuse of the FTTH/xDSL fixed access infrastructure, like cabinets or wall mounted optical splitters. Nevertheless in some cases of fixed-mobile operators with presence of fibre deployments, for costs and competitivity reasons they don't need necessarily to share their own infrastructure with other operators, unless obliged by the regulator; on the other hand new installations in urban environments could be more suitable for network sharing use cases.
- CS4 (Shopping Mall / Airport): this typical case of dense indoor environment is also often characterized by small cells deployed either by the shopping mall owner in a planned manner or by each shop to provide additional services to the customers in an unplanned manner. In any case, these scenarios perhaps are not always suitable for network sharing; on the other hand in some cases of airports with planned infrastructure installation it may happen that the presence of BH network owned by the airport company may be a facilitator to network sharing.

4.6 Proof-of-concept assumptions

WP6 deals with the proof of concept for the proposed optimized network design in iJOIN which encompasses radio access, backhaul and core network optimization. For demonstration activities, three different testbeds have been implemented and described in detail in D6.1 [49], each one covering aspects of the above mentioned areas:

1. The RANaaS cloud platform demonstrates how the proposed technological approaches may benefit from the flexible processing shift to the cloud centres. Two RANaaS instances are implemented: one, at University of Bremen (UoB), is used to demonstrate the WP2 physical layer proposals while the instance at Telecom Italia (TI) in Turin is used to demonstrate Media Access Control (MAC) and Radio Resource Management (RRM) proposals from WP3.

2. The 60 GHz platform, developed by Technische Universität Dresden (TUD), implements a 60 GHz backhaul following the principle of hardware-in-the-loop. This platform can be harnessed to demonstrate joint access-backhaul algorithms as envisioned in WP2.
3. The third testbed, provided by Universidad Carlos III de Madrid (UC3M), is the SDN based platform. This is exclusively used to demonstrate the novel network layer algorithms and proposals, proposed in WP4.

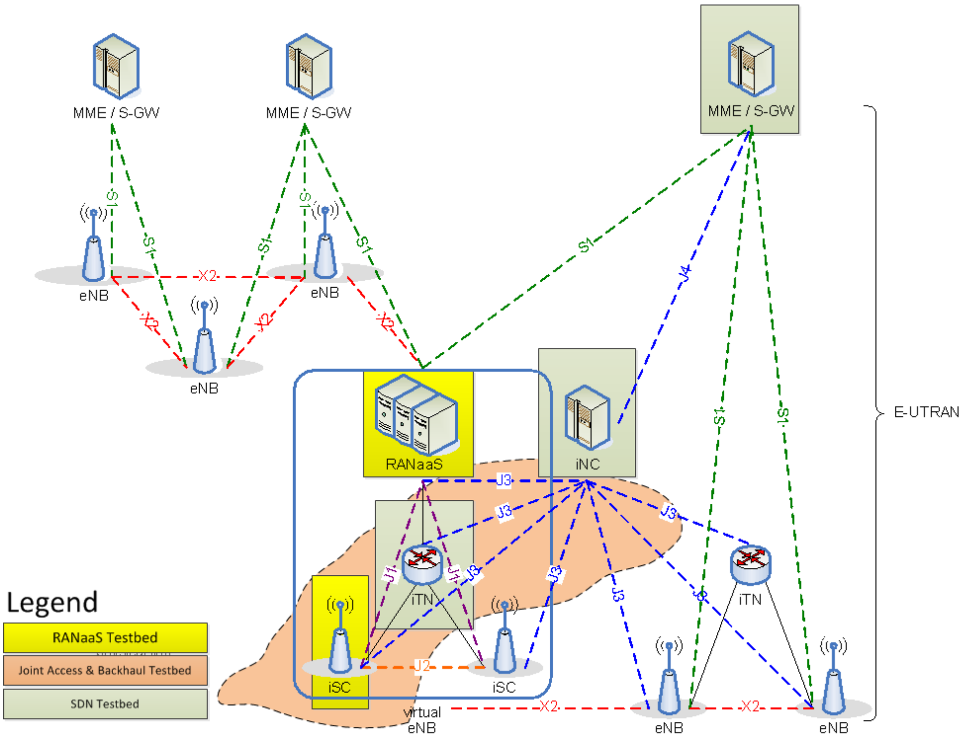

Figure 4-16: iJOIN Testbed Coverage

The three testbed platforms are thus covering the aspects of the following areas: radio access, backhaul and core network optimization. In fact, the RANaaS testbed is mainly focused on emulating veNB elements, a set of iJOIN Small Cells (iSC) and a RANaaS instance running on an Infrastructure-as-a-Service (IaaS) cloud platform. The joint access and backhaul testbed is used to investigate and validate the technology for implementing the physical layer of the communication links between the architecture elements. Finally, the objective of the SDN testbed is the implementation of algorithms for orchestrating and configuring the network entities in order to optimize the network communications.

These testbeds can be applied to the common scenarios defined in Section 4 using the candidate technologies proposed in technology work-packages WP2-4. These experiments provide quantitative results and proof of concept for the ideas proposed in iJOIN. For each testbed, multiple technology algorithms have been identified which are used for demonstration. Due to the limited size and capability of these testbeds, not all metrics and evaluation criteria as in WP2-4 can be applied. Hence, other suitable metrics for different demonstrated candidate technologies have been adopted. Table 4-6 lists the metrics which are used to evaluate the candidate technologies demonstrated on the iJOIN testbeds.

The evaluation metrics for the testbeds are relevant from the proof-of-concept perspective. For instance, the RANaaS testbed shows the applicability of provisioning RAN functions as a service. For this concept to be viable, it must satisfy the latency constraints imposed by different applications. Hence, it is highly important to analyse the processing time and backhaul delay of the RANaaS platform. Both are studied by the RANaaS testbed and joint RAN/BH testbed. The iJOIN report D6.1 [49] provides a detailed description of the evaluation metrics and preliminary experimental results.

Table 4-6: Performance metrics addressed in each testbed platform

Testbed platform	**Metrics**
RANaaS Testbed (General RAN PHY Processing on a Cloud Platform)	Area throughput[1], Processing time, CPU load usage
RANaaS Testbed (Multi-Layer Scheduling and Robust Link Adaptation)	Area throughput, Rate fairness among users
Joint Access and Backhaul Testbed	Energy efficiency[1], Area throughput, Backhaul latency, Backhaul reliability (BER/FER)
Software Defined Network Testbed	Latencies for mobility mechanisms (handover latency, anchor selection latency), signalling load, Optimal RANaaS placement

[1] Strictly speaking, testbeds use the metrics of throughput and energy consumption. As there is one-to-one mapping from throughput to area throughput and energy consumption to energy efficiency, we have taken this liberty to mention these metrics.

5 Functional Split Implementation Requirements of RANaaS

5.1 Implementation aspects of RANaaS hardware

5.1.1 Implementation Options

In general, there are three main options to implement RANaaS on iSCs and RANaaS platform. Functionality could be implemented either on dedicated hardware such as Application Specific Integrated Circuit (ASIC), Field Programmable Gate Arrays (FPGAs), or Digital Signal Processors (DSPs), or on General Purpose Processors (GPPs). Furthermore, hybrid approaches are possible where a software implementation on GPPs is complemented with dedicated hardware. Figure 5-1 illustrates the main options and how iJOIN's baseline implementation can be categorized.

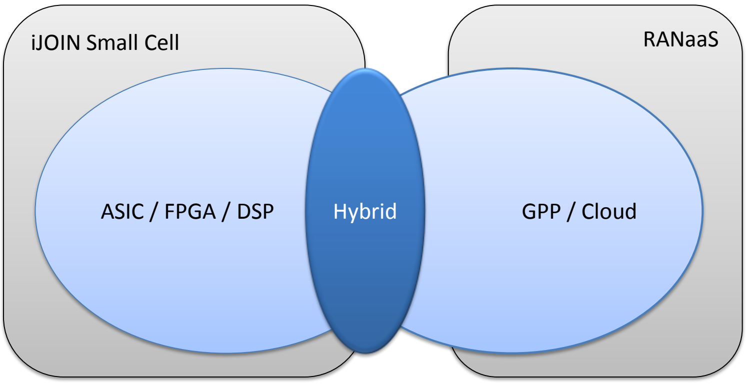

Figure 5-1: Implementation choice applied to iJOIN system

Currently, state-of-the-art virtualized baseband processing in C-RAN is based on hybrid solutions consisting of GPPs with hardware accelerators or co-processors that implement specialized digital signal processing functionalities [42]. The former can be implemented by means of DSPs, FPGAs, ASICs or a combination of them. GPPs can be based on ARM, MIPS or x86 ISAs (Instruction Set Architectures). Co-processors communicate with the CPU using a standard interface such as PCI Express. The approach is illustrated in Figure 5-2. Algorithmic bottlenecks that prevent a pure-software implementation running on GPP can be eliminated by the use of custom hardware accelerators that offload data processing from the CPU. The same approach has been followed to support in an integrated way graphic processing capabilities (combination of CPU and GPU) or packet processing capabilities. The next logical step is to define a programming model for the co-processor that is also GPU-reminiscent, akin to DirectX or OpenGL's abstraction of a computer's graphics subsystem.

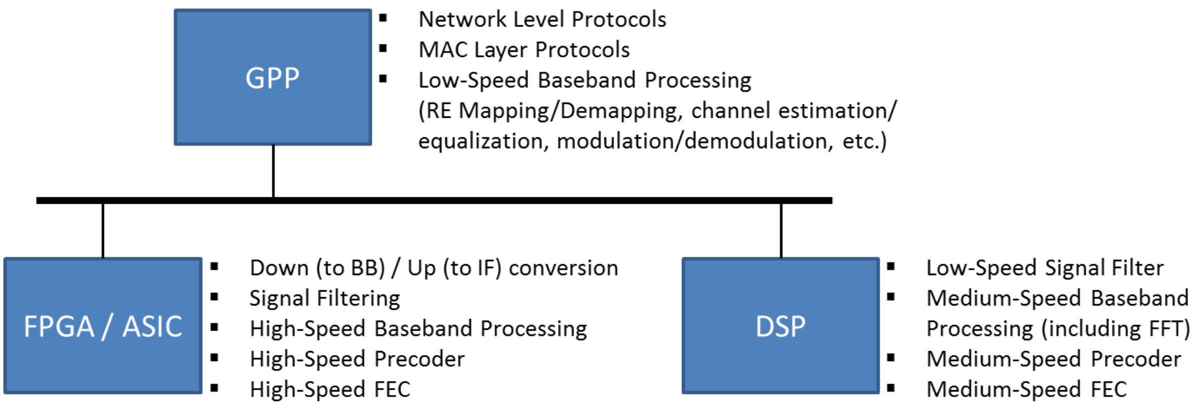

Figure 5-2: Example of splitting of digital signal processing across GPP, DSP, and FPGA

The following functions are proposed in [41] to be implemented by the co-processor:

- **FEC (decoder)**: Accelerator for decoding of turbo and convolution codes.
- **Demapper**: Extracts soft bit values from QAM signals, (LLR values), slicer decisions and slicing error values.

- **Arithmetic**: Performs all the intrinsic arithmetic functions required by the co-processor, e.g., matrix and scalar multiplications, windowing, and frequency correction.
- **DFT/FFT**: Supports orthogonal transforms (FFT and DFT). Frequency correction can be done on the input to the FFT/DFT unit.
- **Logic operations**: Specialized for scrambling, pseudo random bit stream generation, encryption solutions, and various encodings, e.g., convolutional and turbo coding. It may operate on hard or soft bits and can also perform interleaving and arithmetic operations.
- **Data rearranging**: Its main purpose is to support interleaving and data manipulation. The unit can move and interleave large volumes of data as well as handle IR (incremental redundancy), puncturing and simple decoding at high rates.

It can be noticed that functionalities identified as "software," e.g. channel estimation or MIMO processing, are not supported by the co-processor. It can be argued that the same kind of approach is already followed by the solutions provided for baseband processing in commercial base stations. However, two main differences between base station solutions and virtual RAN solutions should be noticed:

- In most baseband processing units for base stations, GPP responsibilities are limited to scheduling and coordination of DSP functionalities carried by specialized hardware, while in the virtualization solution, a significant part of the processing is carried out by GPPs.
- Solutions for the virtualized architectures should support resource virtualization as understood in the Information Technologies (IT) realm while base station solutions are dedicated to single access points.

The main reason for using a hybrid solution with co-processors is the fact that the full implementation of radio interface baseband processing by means of GPPs may be suboptimal in terms of required investment, energy consumption and other performance parameters. On the other hand, centralization of conventional baseband processing units does not allow for an easy virtualization of the resources and the reuse of IT solutions.

The support of the RANaaS concept and flexible functional split introduce a new level of complexity as it may require the solution to support different levels of processing without penalizing the network TCO (Total Cost of Ownership). It should be noted that the solutions previously described are expected to be deployed within the RANaaS instance while the distributed elements are iSCs, which only support a limited set of baseband processing functionalities depending on the functional split.

In the context of the iJOIN architecture, the iSC may implement different levels of baseband processing. Depending on the functional split, this may range from only RRH (radio remote head) functionalities as in CRAN to full support of the whole radio interface protocol stack as in a conventional distributed implementation. The same operating scenarios are applicable to the RANaaS infrastructure. The requirements for an ideal solution would be the following:

- The same solution should be reusable for both iSCs and RANaaS infrastructure in such a way that processing elements may be moved from the iSC to the RANaaS and vice versa.
- It should be possible to switch off those processing units (CPU cores, DSP co-processors) that are not required for the selected functional split.
- It should be possible to virtualize the capabilities of the processing elements in such a way that the functionalities they implement may be decoupled from their locations. For instance, the processing elements of an iSC may be used for processing connections of other iSCs if the backhaul infrastructure provides the necessary connectivity.
- It should be possible to reuse the same solution for the virtualization of other network elements, not necessarily of the mobile network, e.g., virtualization of CPEs or implementation of virtual switches.
- Utilized and provisioned data processing complexity should scale with the actual data traffic load that is processed within the RANaaS platform.

In iJOIN, the virtualisation of RAN processing receives significant attention. The main argument to deploy hybrid solutions within datacenters is the required computational complexity of RAN processing and the energy and cost-efficiency of IT hardware. While hardware-solutions offer a performance advantage over GPP based solutions, they do not offer the same scalability and flexibility as GPP based solutions. In particular, hardware-solutions are so far not able to scale the consumed data processing complexity with the actual data traffic demand. This implies a potential overprovisioning by solutions which are not based on GPPs even in the case of centralized processing. The impact of computational diversity for both required computational resources and cost-efficiency is explored in Section 5.1.3 and 7.1.4, respectively.

The assignment of functionality to dedicated hardware or GPP allows for an additional degree of freedom. In order to avoid adding this additional complexity to the scope of iJOIN, we apply the assignment as indicated in Figure 5-1, i.e. all functionality executed at the RANaaS platform is executed on GPPs and all functionality executed at iSCs is either implemented in dedicated hardware or GPPs for providing a fall-back solution at the iSC.

5.1.2 Virtualization infrastructure

In iJOIN, the initial target chosen for the RANaaS implementation is a cloud computing platform deploying general purpose computational resources. The platform implements an Infrastructure as a Service (IaaS) model where resources are provided on a "as a Service" paradigm meaning that resources are allocated and de-allocated on demand. The resources provided by an IaaS platform can be classified as follows:

- compute resources: Virtual Machines (VMs), running an operating system and application software;
- storage resources: Virtual Volumes, storage elements that can be attached to VMs;
- networking resources: Virtual Networks objects such as Virtual Level 2 (L2) trunks, subnets, or DHCP services.

IaaS platforms usually provide virtualized resources. Virtualization aims to simulate the existence of a piece of hardware which is "materialized" by a software layer running on top of the physical device. The idea is that the actual hardware is hidden to the applications and partially or temporarily used for "impersonating" the role of a virtual piece of similar hardware. The actual computation happens at the physical level but physical resources and applications are not tightly bound to each other. This makes it easier to reuse the physical infrastructure for several purposes, usually at different times.

As described in Section 5.1.1, common solution designs for supporting virtualized baseband processing use a combination of general purpose processors with hardware accelerators or co-processors for implementing specialized digital signal processing functionalities. Mapping this kind of architecture into an IaaS platform raises the problem of how to distribute the related workload on a virtualized platform. Specific attention must be paid for parallel computation which is traditionally obtained using specialized hardware devices, e.g., DSPs or FPGAs.

IaaS platforms implement server virtualization where more than one virtual server runs on top of a single physical computer. This is implemented by using a hypervisor which runs on the physical hardware and which takes care of running several virtual servers. Figure 5-3 summarizes the concept.

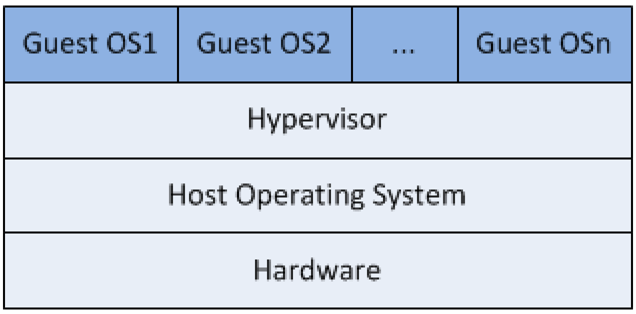

Figure 5-3: Server Virtualization

At the bottom of the stack, the physical hardware provides the actual computational resources, e.g., CPUs and RAM. An operating system is installed on the bare metal and it is integrated with the hypervisor. Each virtual server appears as an autonomous computer having its own (virtual) hardware. Users access virtual servers via network connections. Similar techniques are available for implementing storage and network virtualization.

Hypervisors are designed to minimize the processing overhead and to allow for almost the same performance as non-virtualized environments. In addition, in the case a VM is assigned a number N of virtual CPUs (vCPUs), it can execute up to N processes and threads in parallel. It's important to say that, when a VM is instantiated, it is possible to define the number of vCPUs that the VM will use. This number defines a virtual parallelism that becomes real parallelism only when the number of virtual CPUs assigned to the VM corresponds to the number of real cores dedicated to it. This aspect is regulated by the overbooking factor.

Overbooking can be defined as the ability of running a number of virtual servers requiring more hardware resources than available. For example, a physical server with 2 quad-core CPUs can run a number of virtual machines allocating a total of 10 vCPUs. This is possible because usually not all virtual machines are running at the same time. Therefore, when a VM is waiting for a “slow event,” i.e., an interrupt, the real CPUs are used for running concurrent VMs.

Similarly, a physical server with N GB RAM can accommodate a number of virtual machines requiring a total amount of memory higher than the physical RAM (values suggested by best practices can vary depending upon the hypervisor technology, anyway a value of 1.2-1.5 x N can give a realistic indication) . This is obtained with different memory over commitment techniques implemented at the hypervisor level. Anyway high overbooking can have strong impact on the VM performance; this is very often due to a VM oversizing, following the rather inefficient rule to dimension a VM as it was running on a physical server. Specially too many virtual CPUs per VM, if not needed by an application specifically written for multiprocessing, cause a huge overhead to the hypervisor in a continuous scheduling over the physical core.

Avoiding this over-provisioning allows to have more resources available for VMs that run real-time applications and, specifically, parallel algorithms: in this case an actual parallelism is obtained by allocating sufficient vCPUs for the VM hosting the algorithm. The cloud computing platform will put the VM on a physical server that has sufficient resources to satisfy the VM requirements.

Through virtualization, IaaS platforms provide the low-level building blocks for implementing systems for hosting parallel programming. In fact, each VM can utilize multiple cores of the hosting physical machine for elaborating the assigned workload. In addition, several VMs can be activated and can collaborate for expanding the computing power dedicated to the workload at hand. In this manner, parallelization can be implemented either within a single VM or across several collaborating VMs. Figure 5-4 summarizes the concept.

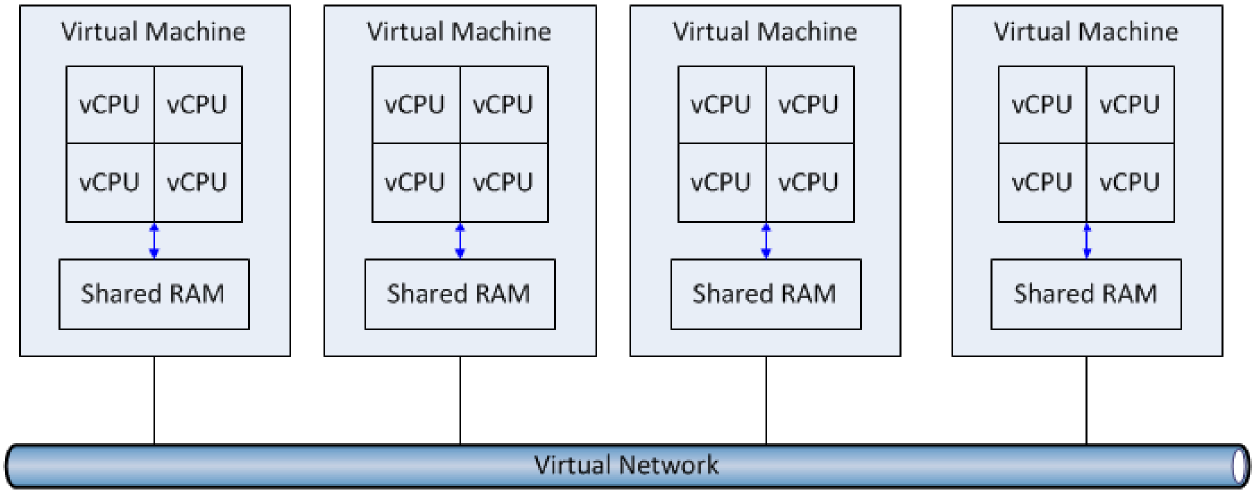

Figure 5-4: Virtual Machine Cluster

Every single virtual machine works as a Symmetric Multi-Processor machine (SMP), a computer with multiple processors which all share a single address space. SMP units, in turn, can be connected through a virtual network giving the possibility of creating parallel computer clusters for further parallelization of algorithms. It is important to mention that every single VM runs inside the boundaries of a physical node. On the other hand, two different VMs can be hosted in a single node. In the latter case, the communication of two VMs co-located on a single physical node is obtained through a virtual network.

Through virtualization and IaaS platforms, it is possible to implement various cluster topologies for supporting parallel computing. Specifically, parallelism can be alternatively supported within every single node or by distributing the algorithm across several VMs. Therefore, it is important to analyse the advantages and disadvantages of implementing parallelism within a single element of the cluster (i.e., a VM) versus the advantages and disadvantages of implementing parallelism activating several collaborating VMs. These indications provide the guidance for finding the best balance that fits the problem at hand.

In a single VM, all vCPUs access the memory as global address space. Multiple processors can operate independently but share the same memory resources and changes in a memory location caused by one processor are visible to all other processors. From a programming point of view, the global address space facilitates data sharing between parallel tasks and access to data is both fast and uniform. On the other hand, single VMs lacks of hot internal scalability in term of vCPUs; anyway the number of vCPUs can be increased powering off the VM. Adding more vCPUs can geometrically increase traffic on the shared memory-CPU path and for cache coherent systems, geometrically increased traffic associated with cache and memory management. From a programming perspective, synchronization for ensuring the correct access to global memory must be explicitly indicated by the programmer through constructs such as semaphores, barriers, and queues. On this side, standard software libraries such as POSIX threads (Portable Operating System Interface for Unix threads or Pthreads, for short [36]) or OpenMP (Open Multi-Processing [32]) can significantly support the implementation task.

Considering all VMs in a cluster, it should be noted that each VM has its own memory address that does not map to other VMs. Each VM operates independently and changes to its local memory have no effect on the memory of other VMs, even in the case they run on the same physical node. When a VM needs access to data in another VM, it is usually up to the programmer to explicitly define how and when data is communicated. Elasticity is the main advantage of this approach because resources, e.g., processors and RAM, can be easily adapted to the workload changes. Whenever the workload changes, VMs are provisioned or de-provisioned following a so-called scale-out paradigm (or horizontal scalability) that spreads the workload over several virtual machines, possibly running on different physical hosts. The flip side is that the programmer is responsible for the details associated with data communication between tasks running on different VMs. Depending on the algorithms, it could be difficult to map the data structures used by the algorithm on a distributed memory architecture and there is a non-uniform memory access times, i.e., data residing on remote nodes takes longer to access than local data.

IaaS platforms provide tools for creating VM clusters that, from the programmer view point, can be considered and treated as clusters of ‘real’ machines. In this perspective, the programmer can take advantage of software technologies specifically designed for implementing parallel programming on top of server clusters. These technologies come as software libraries or compiler directives that permit the programmer to distribute the tasks on the servers participating to the cluster. Every technology implements a parallel programming model that, theoretically, can be deployed independently on the underlying cluster topology.

For example, it is possible to implement a shared memory model where all the tasks access a virtually shared memory even when they run on different machines of a distributed topology, e.g. Kendall Square Research (KSR) ALLCACHE [37]. In such a case, it is up to the used technology to emulate a single shared memory, hiding the underlying implementation complexity. In addition, it is possible to implement a purely distributed memory model where each task has its own private memory and communicates with other tasks through a message passing mechanisms, e.g., Open Multi-Processing Interface (OpenMPI) [33], even when both operate on the same machine.

It is important to note that there is not a unique programming model to apply because this choice strongly depends on both the problem to solve and the parallel algorithm to implement. Usually, hybrid solutions, where the two models are applied at the same time, permit a better usage of the underlying infrastructure at the price of some additional complexity. Figure 5-5 shows an example where different parallel programming technologies are used on a VM cluster created on an IaaS platform.

Figure 5-5 shows a parallel application distributed on a cluster of VMs, each VM runs a certain number of tasks. A software technology such as Open Multi-Processing (OpenMP) [32] or Pthreads is used for managing the tasks running within a single VM. It provides mechanisms for creating, coordinating and synchronizing tasks which share the memory of the VM. Another software technology, such as a Message Passing Interface (MPI) implementation, is used for managing tasks running on different VMs.

The software technologies used for implementing parallel computing constitute a layer shown as the Parallel Computing Mechanisms layer in Figure 5-5. The following paragraphs give a short description of the most relevant characteristics.

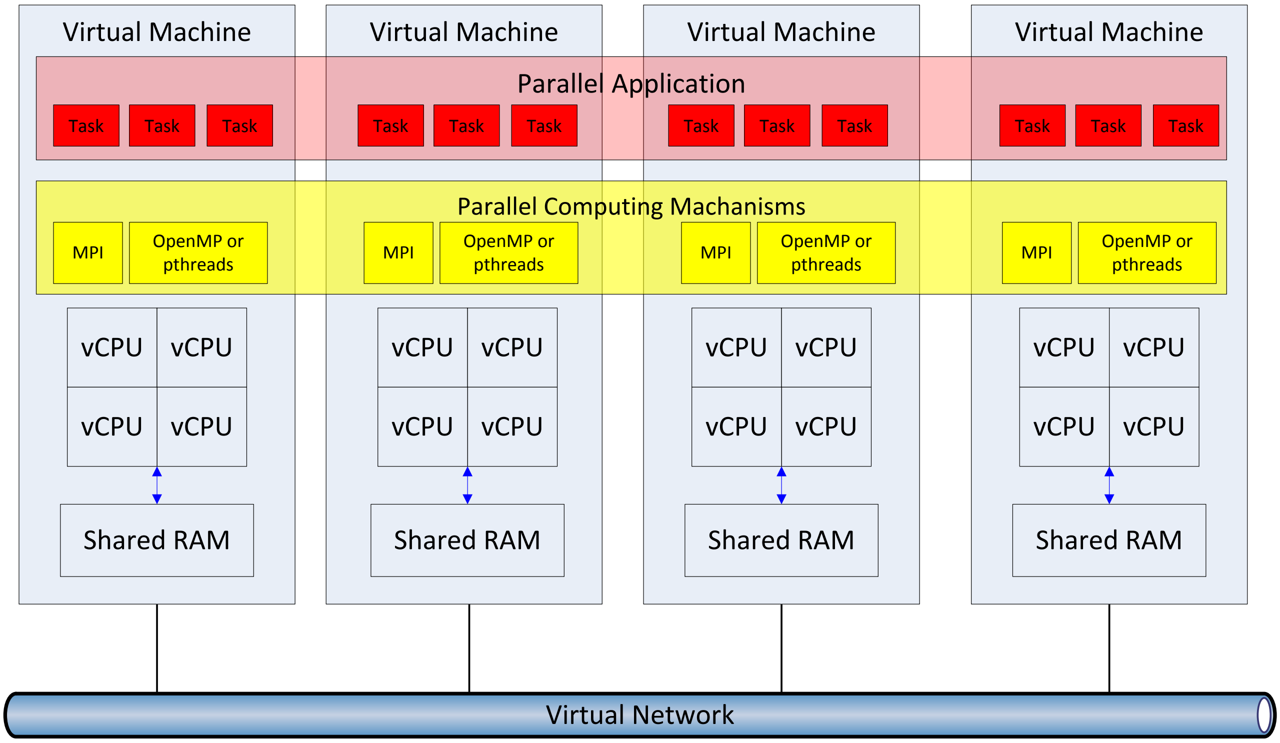

Figure 5-5: Hybrid Programming Model on an IaaS VM Cluster

OpenMP is a standard Application Program Interface (API) that allows for the implementation of portable, shared memory applications [32]. Through compiler directives or explicit library calls, a programmer implements multi-threaded, shared memory applications. The programmer defines the portions of code that are to be executed in parallel as well as synchronization points for coordinating parallel computation streams. At run-time, the compiled program runs as a process of the underlying operating system split into light-weight threads all sharing the memory address space of the 'hosting' process. OpenMP run-time allocates the processors, i.e. CPUs or cores, to thread execution and takes care of maximizing the "actual" parallelism. OpenMP is available for C/C++ and Fortran programming languages and runs under several operating systems, e.g., Solaris, AIX, HP-UX, Linux, Mac OS X and Windows.

Pthreads is an alternative mechanism for implementing multi-threaded/shared memory applications [36]. It comes as a set of library calls originally implemented for Unix operating systems and currently available under other platforms such as Linux. Using Pthreads, the programmer implements a parallel program as a single operating system process providing the same computational resources such as memory, network connectivity, and file system access. Inside a process, parallel threads can be created and managed by explicitly invoking suitable Pthreads library calls. It's up to the Pthreads run-time library to ensure that concurrent threads are actually executed on different processors in order to obtain the maximum level of parallelism. In these aspects, Pthreads and OpenMP are very similar. Pthreads is available for C/C++ binding and provides, in addition to functions for managing threads, also functions for coordinating their executions and concurrent access to memory, i.e., mutual exclusion management, condition variable management, and synchronization.

Message Passing Interface (MPI) is a standard definition for a software library that allows for the implementation of parallel programs realizing a pure distributed memory [33] based on a message passing model. It was originally designed for clusters of single CPU computers communicating through network connections but has been more recently adapted to run on clusters of multi-processor computers. The underlying programming model relies on distributed memory, meaning that each parallel task accesses its own private chunk of memory and, in case two or more tasks need to share some data, they need to explicitly exchange messages. However, some MPI implementations optimize the mechanism for exchanging messages. They use shared memory if the tasks run on the same machine and network messages if they do not run on the same physical or virtual machine. MPI provides several library calls for sending and receiving messages among tasks and, more importantly, implements a cluster 'concept' where several nodes can be considered as a unique computing platform for distributing the computational workload. The programmer writes parallel programs as monolithic entities but has complete control on the location where parallel tasks will be executed, i.e. on which node of a cluster. MPI is available for many programming languages such as C/C++, Fortran, Java, Perl, Python, and runs on several computing platforms ranging from PC to supercomputers equipped with a range of operating systems, e.g., Unix, Linux, Mac OS.

The combination of the two parallel programming models, i.e., shared-memory/multi-threaded and distributed-memory/message-passing model, opens the possibility of taking advantage of the best characteristics of both. The former model is more suitable to situations where parallel tasks need to exchange data with minimum communication overhead. The latter can be used successfully for addressing scalability issues. How these two technologies are used strongly depends on the problem at hand and on the parallel algorithm that solves the problem.

For example, in case a certain algorithm shall be executed with real-time constraints, i.e. the response must be guaranteed within strict time constraints, the best solution is to 'deploy' the related tasks on a single virtual machine of the cluster. This avoids unpredictable or unacceptable delays caused by network communications between tasks running on different machines.

5.1.3 Computational Outage

A virtualized infrastructure is able to provide an abstract interface between the network platform (RAN) and the underlying physical computational resources. However, these resources are still limited which may lead to computational outage of the RAN data processing rather than channel outage due to difficult channel conditions. For example, a computational outage would occur if the employed Forward Error Correction (FEC) decoding software would require more computational resources, such as following from a high number of decoding iterations, than can be provided by the virtualized infrastructure. Computational outage leads to a waste of spectral resources and loss of throughput, in much the same way as a channel outage. In the case of small-cell deployments where each small-cell needs to implement the full RAN protocol stack, computational outage may even dominate channel outage if only limited computational resources are available at a small-cell base-station.

Therefore, it may not be throughput-optimal if the scheduler only considers the raw throughput while ignoring the required computational resources. We refer to this as Maximum Rate Scheduling (MRS). By contrast, a Computationally Aware Scheduler (CAS) would consider both the achievable rate as well as the required computational resources. Specifically, the CAS uses a link-adaptation table which has been obtained for max two iterations compared max eight iterations in the case of MRS. Figure 5-6 (a) shows the achievable raw throughput of both schedulers [50]. Obviously, the MRS achieves a slightly higher raw throughput than the CAS. However, Figure 5-6 (b) shows the required computational resources. As can be seen, the required computational resources of the MRS may be up to 4 times higher than the required resources of the CAS. There are more complex schedulers possible which take the current computational load into account and adjust the maximum number of iterations according to the available resources. In Section 5.5.2, we provide further detailed results on the computational outage using a detailed system level uplink simulation employing an actual LTE decoder. This result is further exploited for the cost-efficiency study performed in iJOIN.

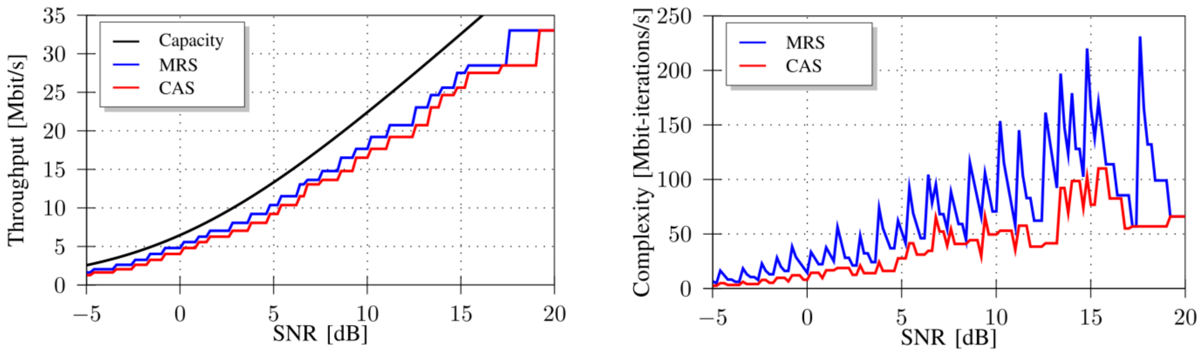

(a) Raw throughput **(b) Computational complexity**

Figure 5-6: Raw throughput and computational effort for rate-maximizing and computationally aware scheduler

Centralizing multiple base-stations and their associated computational load allows for exploiting computational diversity gains. Figure 5-6 shows the strongly varying computational complexity depending on the actual channel quality and the corresponding MCS (Modulation and Coding Scheme). These fluctuations can be well exploited in a centralized system where multiple base-stations share the same pool of resources. As a consequence, the provisioned amount of computational resources is reduced while maintaining the same performance level, in terms of computational outage or net-throughput. More formally, let $C_{outage,N}(\varepsilon)$ be the outage complexity for N centralized base stations and an outage probability ε, i.e. $C_{outage,N}(\varepsilon)$ computational resources are provided, e.g. number of CPUs; then the probability that one particular BS cannot be processed due to insufficient computational resources is given by ε, similar to the notation of outage capacity and outage probability. Computational diversity gain is now defined by the ratio:

$$c(N) = \frac{N C_{outage,1}(\varepsilon)}{C_{outage,N}(\varepsilon)}. \tag{5.1}$$

This measure gives the overprovisioning ratio in the case of a distributed implementation compared to a centralized implementation. It further gives us an indication of the utilization of the system because the centralized system is able to perform the operations with only $1/c(N)$ of the resources. In Section 5.5.3, we provide numerical results for the computational diversity gain. A detailed description and formal definition of computational outage and complexity is given in [43].

5.1.4 Load balancing

The previously explained computational diversity allows for computational load balancing. The goal of load balancing is to distribute the dynamic workload across multiple processing nodes, to achieve optimal resource utilization, and to avoid computational overload. It prevents bottlenecks of the system that may occur due to overburdened nodes and further helps in promoting equal availability of computational resources. One of the most important challenges in implementing load balancing algorithms for the RANaaS concept comes from the highly variable computational complexity of the various functionalities that have to be implemented.

Load balancing algorithms follow different classifications, according to whether the workload is distributed between the processing nodes in a static, dynamic, or adaptive manner [6]. In the static approach, the load balancing is defined when the system is implemented. The dynamic approach takes into account the current state of the system during load balancing decisions. The adaptive approach further allows dynamically changing the properties of the implemented functionality, e.g. switch from an optimal to a suboptimal algorithm, according to the state of the system when the load balancing decisions are made. The adaptive approach seems more appropriate in the RANaaS context, since the computational load can vary significantly in time due to fluctuations in the traffic load. Dynamically adapting the implemented functionalities to the computational load allows for further optimization of use of the available resources in a RANaaS instance.

Load balancing can also be used to implement failover [5], i.e. the continuation of a service after the failure of one or more of its components. The components are monitored continually, and when one becomes non-responsive the load balancer is informed and no longer sends traffic to it. This is an inhered feature from grid-based computing for cloud-based platforms.

Another issue that can be addressed by using load balancing algorithms is related to energy optimization [7]. In traditional server cluster systems, the workload is distributed in an equal fashion in order to achieve the best possible performance and scalability. However, distributing the work across many servers may result in low levels of utilization, thus yielding excessive energy consumption with respect to the amount of useful work done. The reason is that the power consumption of current systems is not proportional to how much work they are doing, with low levels of utilization incurring disproportionate amounts of energy. In [4] , it has been pointed out that it may be possible to rewrite load balancing algorithms to be more energy aware and introduce the concept of “load-skewing.” If servers were continually allocated work while they have resources remaining, then we would be able to power down unused servers and therefore save on energy consumption. Switching off or powering down components and entire systems effectively when not in use can be considered a key area of energy aware computing. However, the effect and extent of these power state transitions requires careful consideration. For example, powering down a CPU can be an effective means of saving energy. Suspending also the system cache, memory and controllers will save even more energy, but at the penalty of increased cost and time to return the system to a useful state [7]. A balance must be achieved between energy savings and system performance.

5.1.5 Migration of Virtual eNodeBs

The virtualization of physical resources allows for a virtually unlimited availability of resources and the possibility to provide resources on-demand. While the user of a virtualized environment does not want to care about the actual physical deployment and assignment of virtual machines, the operator of the cloud-platform needs to take care of it. In particular, the following events require special attention:

1. Maintenance of physical resources which requires to turn off part of the infrastructure,
2. Poor resources of the host where the VM is running on,
3. Failure of equipment.

All three events require a migration of virtual machines across physical hosts.

Alternatively, virtual machines may be grouped while maintaining a common abstract interface to the user of this group. In terms of virtual eNBs, one virtual eNodeB may be composed of multiple virtual machines which appear as one black box. In this case, it may be possible to increase or decrease the amount of consumed resources through adding or removing virtual machines. For this process, a template of new virtual machines and their state is required. However, if the cloud service is based on the Openstack framework, the better way to change the number of VMs running the same services is to exploit the autoscaling service that integrates Heat (the provisioning component) with Ceilometer (the monitoring component). Heat templates can be built so that Heat itself can provision or deprovision a VM upon an alarm coming from Ceilometer.

Finally, a failure recovery mechanism is required. This mechanism needs to ensure continued service of user terminals if a virtual machine fails. This can be done in two ways. Firstly, a stand-by copy of each virtual eNB is maintained and ready to be used at any time. Obviously, this would require significant resources and contradicts the idea of improving the utilization efficiency. Note that any hypervisor has its own HA (High Availability) feature embedded that is able to restart a VM on another host in case of failure or just only isolation too of the VM itself. This is not (yet) true in Openstack at the cloud VM scheduler level, but works are in progress to implement this feature also there. As an alternative, 3GPP mechanisms could be used to re-connect user terminals to a new virtual machine after a failure.

In the case of virtual eNBs, the following specific cases need to be considered:

1. Reassigning UEs between virtual eNodeBs at the same RANaaS instance: Since the virtual eNB determines important parameters of L1-3 processing, this kind of "migration"[2] would involve a handover process between both virtual eNBs. However, both virtual eNBs are able to exchange relevant information and to prepare the handover very efficiently. This reduces the handover time and involved signaling significantly.
2. Reassigning UEs between virtual eNodeBs at different RANaaS instances: Similar to the previous case with the difference that the communication between virtual eNBs may be subject to higher delays. In the case that it requires multiple milliseconds, it is not possible due to the tight real-time constraint of 3GPP.
3. Reassigning iSCs between different veNBs at the same RANaaS instance: This case is similar to the first one, except for the fact that in this case only part of the processing and virtual machine state is moved to another virtual machine. This still involves copying the relevant information between source and target virtual eNB. The amount of information depends significantly on the maintained state for each iSC. This determines the required downtime and whether a seamless migration is possible. If a seamless migration is not possible, 3GPP mechanisms may be applied to re-assign users.
4. Reassigning iSCs between different veNBs at different RANaaS instances: Similar to the previous case whereby the communication latency between both RANaaS instances determines whether a seamless migration can be supported. In the case that it requires multiple milliseconds, it is not possible due to the tight real-time constraint of 3GPP.

[2] In principle, according to a correct IT terminology, VM migration is related to the entire VM, while in this section this term (whereas indicated as "migration") is used in a wider sense, and applied to the base station virtualization use case (e.g. referring to the partial transfer of relevant information between two VMs).

5. Moving complete veNBs between different RANaaS instances: Again, the required downtime for copying the state of a veNB determines whether a seamless migration is possible. In the case that it requires multiple milliseconds, it is not possible due to the tight real-time constraint of 3GPP.

5.1.6 Implementation requirements

iJOIN does not envision a monolithic porting of the 3GPP LTE stack into the veNB. On the contrary, it targets a modular and even dynamic environment, where only the best suited part of the stack functions is executed within the veNB hosted in the RANaaS platform. Generally speaking, IaaS platforms provide virtual machines as a natural mechanism for implementing modularity. In this perspective, programs implementing CTs functionality can be integrated inside proper virtual machine images. Hence, whenever we need to activate a new instance of the CT, it is only a matter of creating and activating a new virtual machine configured with the right software.

The actual feasibility or effectiveness of the implementation depends on some key parameters. Such parameters are different for different candidate technologies, being tied to the characteristics of each algorithm in terms of distribution, computational intensity and timing. In addition, when considering the porting of CTs into an IaaS platform, it is fundamental to consider some limitations imposed by the target environment.

Processing power can be a critical parameter for CPU bound algorithms, since GPPs like the ones powering industry standard servers cannot reach the top processing performance rates provided by a DSP or even an ASIC or FPGA. The limitation is both in the CPU own computational power and in the fact that industry standard servers do not execute microcode but software programs whose interaction with the processor is mediated by an operating system and which are written in non-machine languages. This aspect may worsen when using virtualization, a foundation technology of cloud computing. As detailed in Section 5.1.2, virtualization implements virtual hardware resources by using the physical underlying resources, e.g., RAM, CPUs, and storage. This can raise issues when running real-time applications because a critical computation allocated to a VM could be periodically interrupted by the hypervisor for permitting other concurrent VMs to proceed. Cloud computing can partially address this issue by ‘regulating’ the actual number of VMs running on a single physical server and ensuring that the total amount of resources required for running does not exceed the actual amount of the available resources.

However, as reported by [8], hypervisors introduce significant overhead on interrupt management that results in higher latency with respect to situations where the processing is executed on ‘bare metal’ computing environment. In the case of an interrupt, the hypervisor must dispatch the event to the ‘right’ VM, which has originally been waiting for it. For example, when a network packet is received on a network interface, it must be dispatched to the correct VM. Results in [8] show that the typical latency to interrupt on a virtualized environment with Kernel-based Virtual Machine (KVM) hypervisor [38] ranges from 300 to 700 µsec against a typical latency of 20 µsec on non-virtualized environments (Intel® Romley Server with 2 Intel® Xeon® CPU E5-2697 v2 processors,70GHz 12 cores x86_64 architecture). This is not necessarily a problem as it strictly depends on the type of the applications running on the virtualized environment. The potentially more serious challenge is the non-real-time behaviour of commodity hardware while usually applied FPGA and DSP architectures are real-time systems. Furthermore, commodity hardware may require more processing time which potentially exceeds the achieved processing times on FPGAs and DSPs. The problems of higher computational latency and jitter need to be solved in order to implement a 3GPP RAN system on commodity hardware.

5.2 Implementation constraints of 3GPP LTE

The previous section has addressed the capability of a cloud platform using GPPs to perform RAN functionalities with the PHY layer being the most extensive computational task. If nothing prevents one functional split to be applied in theory with any kind of backhaul, there are still strong timing constraints to consider if the 3GPP LTE compliancy is targeted. The lower we perform the functional split in the protocol stack, the more bandwidth is required to support the forwarding of the data between the iSC and the RANaaS instance. In addition, the split point within the PHY processing chain itself (see Figure 5-7) can also significantly impact the required bandwidth as shown in [34] and [45].

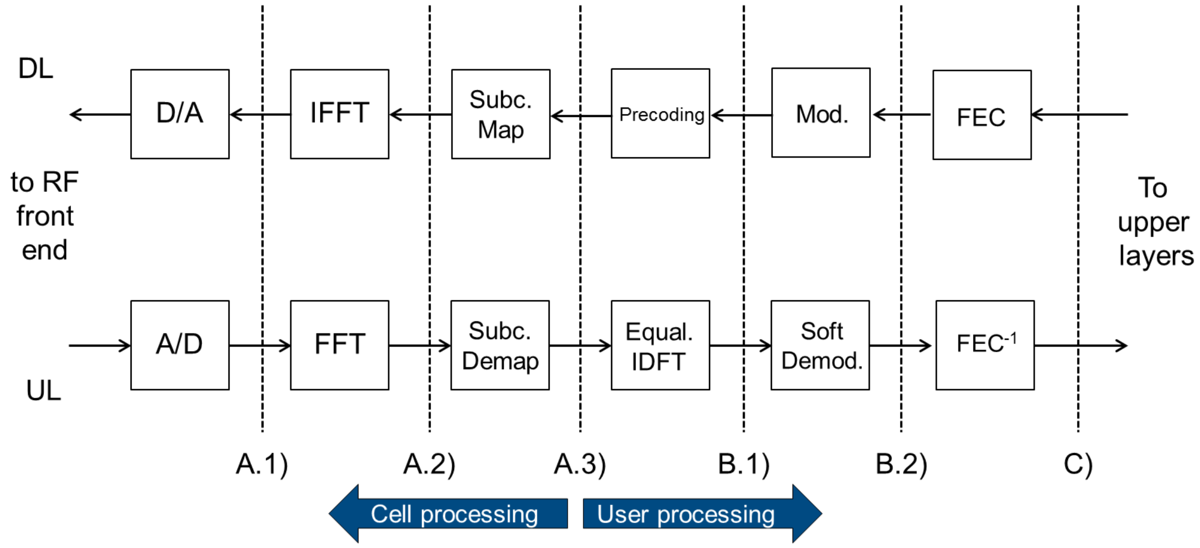

Figure 5-7: Functional split options for the PHY layer [34]

More important than the bandwidth requirements, the timing requirements must be carefully considered in the functional split decision. Indeed, 3GPP has defined many timers for each of the layer (from MAC to RRC) which dictate the behaviour of the complete LTE system. They may impose some serious constraints on the feasibility of one specific functional split within a legacy 3GPP LTE ecosystem. In [35] and [46], those timers have been all gathered and analysed. Table 5-1 only presents the identified ones that may possibly impact a functional split decision. Many of the higher layer timers are configurable within a specified range definition large enough to allow for a setting adapted to the backhaul and processing time required.

However, the heaviest timing constraint will appear as soon as the HARQ process will be performed in the RANaaS instance and not at the iSC. In the case a transmission attempt failed, the HARQ procedure repeats the same message, i.e. chase combining, or a different encoding of the same message, i.e. incremental redundancy. To avoid a communication to stall waiting for the acknowledgment of a message transmission, several HARQ processes are used in parallel, allowing one message transmission to be done while waiting for a previous message transmission to be acknowledged. In FDD, eight HARQ processes are defined which fits well the way data are acknowledged in particular in the uplink. When one UE is scheduled for transmission in subframe *n*, it received this scheduling information in subframe *n*-4. At subframe *n*+4, the UE is expecting an acknowledgement from the RAP which can either implicitly schedule a retransmission or explicitly schedule a (re)transmission at subframe *n*+8. Since the data should not be flushed until a positive acknowledgment is received, the HARQ process used at subframe *n* cannot be used before subframe *n*+8. Thus, the round-time trip (RTT) timer of one HARQ process is defined by 8ms as listed in Table 5-1.

In uplink LTE, a synchronous behavior is implemented. This implies that one HARQ process can only be used in modulo eight subframes, i.e. eight HARQ processes in parallel. This has a strong impact on the functional split implementation in the case that decoding processing is performed at the RANaaS instance. Indeed, a UE expects a positive or negative acknowledgement of each transmission within 4 subframes, i.e., 4ms. This implies that the two-way transmission from the iSC to the RANaaS PoP and the decoding itself must finish in less than 3ms such that the UE can expect to receive the acknowledgement in the next subframe (n+4). This timing requirement is the heaviest constraint with such functional split.

If a legacy UE fails to receive this acknowledgement in subframe *n*+4 for an HARQ process used at subframe *n*, it will not be able to use this HARQ process at subframe *n*+8, and a next scheduling occasion which will only be at subframe *n*+12 for subframe *n*+16. An LTE-compliant solution can be for the iSC to send a positive acknowledgement on the PHICH, even if the result of the decoding is not yet known. By not sending scheduling information on the PDCCH, the UE will not flush its HARQ process unless receiving an explicit order to do so, i.e. with the NDI indicator set to 1. This compatibility comes with the cost that the possible maximum throughput of the UE will be reduced.

Table 5-1: 3GPP timing requirements [35]

	Timer	Short description	Max Value
PHY	Subframe	Physical subframe length	1 ms (fix)
	Frame	Physical frame length	10 ms (fix)
MAC	HARQ RTT Timer	When an HARQ process is available	8 ms (fix)
RLC	t-PollRetransmit	For AM RLC, poll for retransmission @TX side	500 ms
	t-Reordering	For UM/AM RLC, RLC PDU loss detection @RX side	200 ms
	t-StatusProhibit	Prohibit generation of a status report @RX side	500 ms
PDCP	discardTimer	Discard PDCP SDU / PDU if expiration or successful transmission	Infinite
RRC	TimeToTrigger	Time to trigger of a measurement report	5.12 s
	T300	RRCConnectionRequest	2 s
	T301	RRCConnectionReestablishmentRequest	2 s
	T304	RRCConnectionReconfiguration	2 s or 8 s
	T310	Detection of physical problem (successive out-of-sync from lower layers)	2 s
	T311	RRC connection reestablishment (E-UTRA or another RAT).	30 s

5.3 Preferred functional splits

The objective of this section is to provide an initial assessment of the different implementation options in order to support the flexible functional split proposed by iJOIN. For these purposes, a three step process is required.

The first step in this process is to identify those functionalities that are best implemented on specialised hardware such as ASICs, FPGAs or DSPs, from those that may benefit or are not significantly impaired from an implementation on GPPs. Whether a functionality falls into either category depends on whether it is composed of repetitive tasks that can be accelerated by specific hardware solutions or whether it is of larger algorithmic complexity that is provided by a software based solution. A corresponding overview of preferred implementation options is provided in Figure 5-8 while Section 5.1.1 illustrates the implementation options considered in iJOIN.

In the following, an identification of the LTE radio interface functionalities in terms of the optimal implementation option is carried out. In general, it can be assumed that splits that are user dependent will potentially provide statistical multiplexing processing gains when the traffic generation is not homogeneous. Furthermore, latency, throughput, and execution jitter need to be considered to decide upon the functional split. In general, 3GPP LTE is a real-time system and therefore hard deadlines must not be violated. However, the RAN protocol stack may be decomposed into a time-critical and a less time-critical part. The time-critical should preferably be implemented on hardware while less time-critical parts may be implemented in software.

Finally, a third step that needs to be performed is the evaluation of the time resilience of the solutions adopted, i.e. how well they can be adapted to the support of new functionalities to be incorporated in the evolution of the networks. Examples may be the support of massive MIMO solutions or full duplex communications as well as the incorporation of new operational frequency bands.

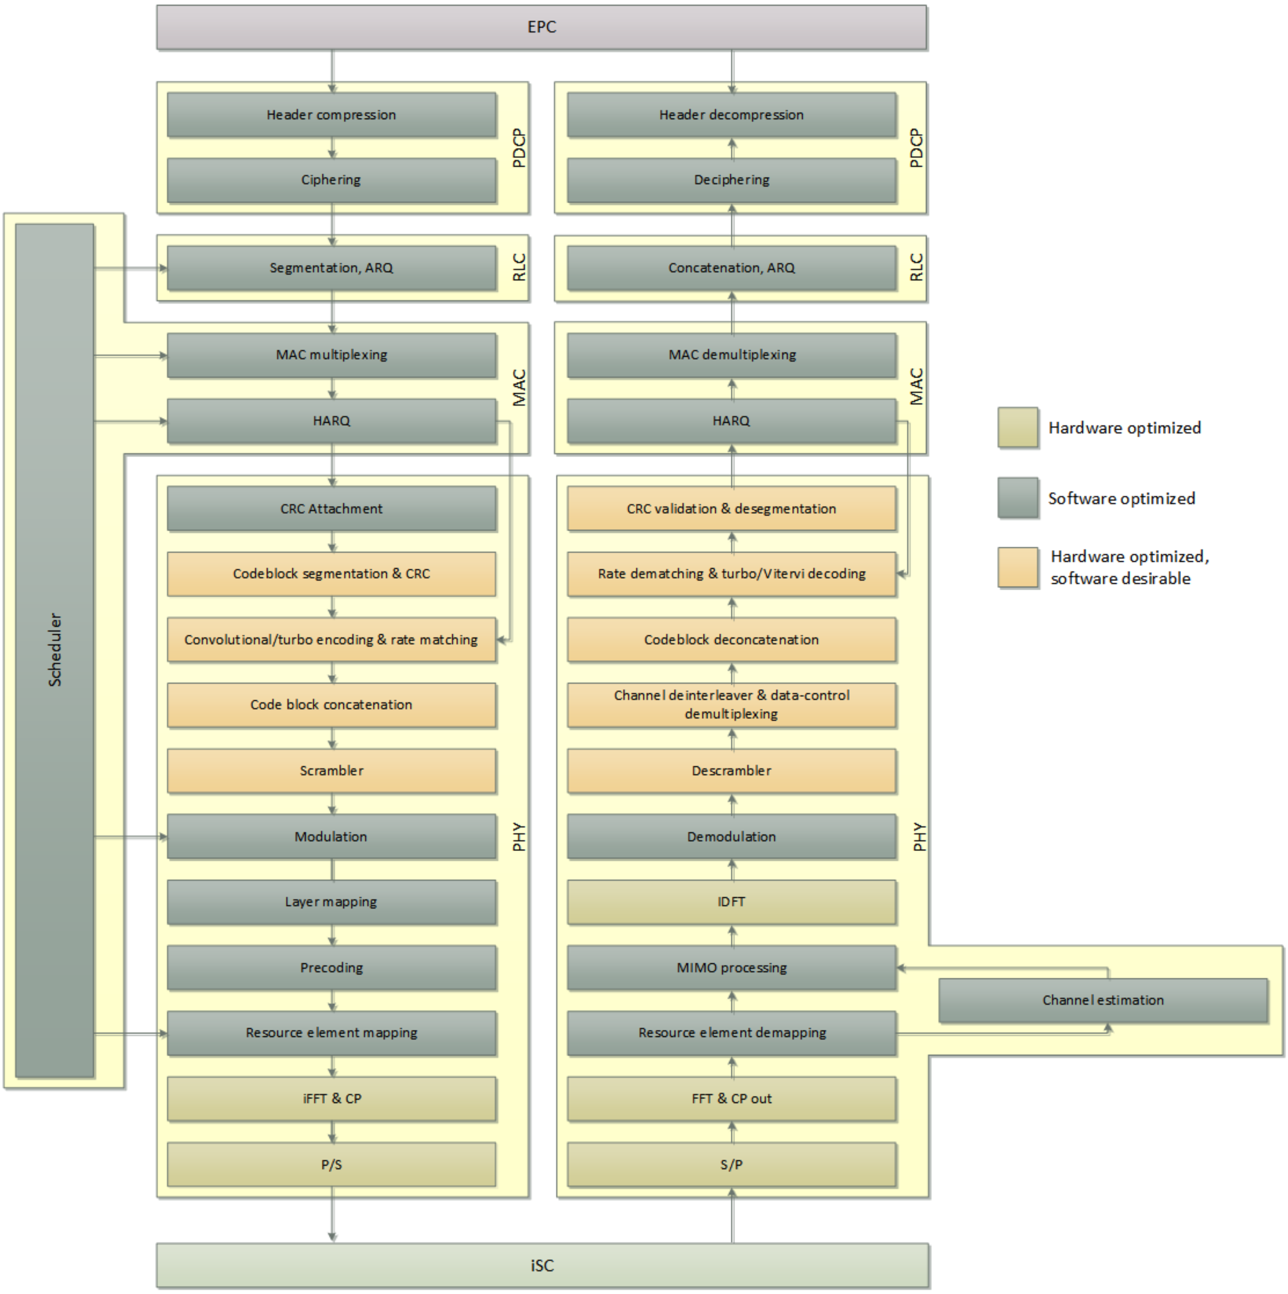

Figure 5-8: Implementation options of 3GPP LTE RAN functionality

Based on the previous discussion and the results in D2.2 [45] and D3.2 [46], Figure 5-9 shows the four functional splits which were identified as preferred functional splits. They are further detailed in the corresponding reports D2.2 (Splits A and B) and D3.2 (Splits C and D). The four preferred functional splits are:

A. Similar to CPRI (Common Public Radio Interface), most of digital processing is centralized. In this case, a very low latency high-capacity backhaul is required. Furthermore, this option does not allow for exploiting multiplexing gains in the backhaul. However, it offers most centralization gains through joint processing.

B. In this case, user-based functionality is centralized including forward error correction while cell-specific processing such as FFT remains at the iSCs. This allows for exploiting multiplexing gains in the backhaul, computational diversity gains at the central processor, and centralization gains through multi-point algorithms.

C. The third option is to only centralize MAC functionality and part of the scheduler. In this case, time-critical processing at the central processor is avoided while advanced coordination algorithms can be executed.

D. Finally, Split D only centralizes the PDCP and RAN control-plane in order to avoid time-critical or computationally intense task at the central processor, but it still allows for gaining from coordination gains across multiple cells. This option is 3GPP LTE compliant and would allow for a split of control- and data-plane.

The above splits are further detailed and divided into "sub-splits" in D2.2 [45] and D3.2 [46].

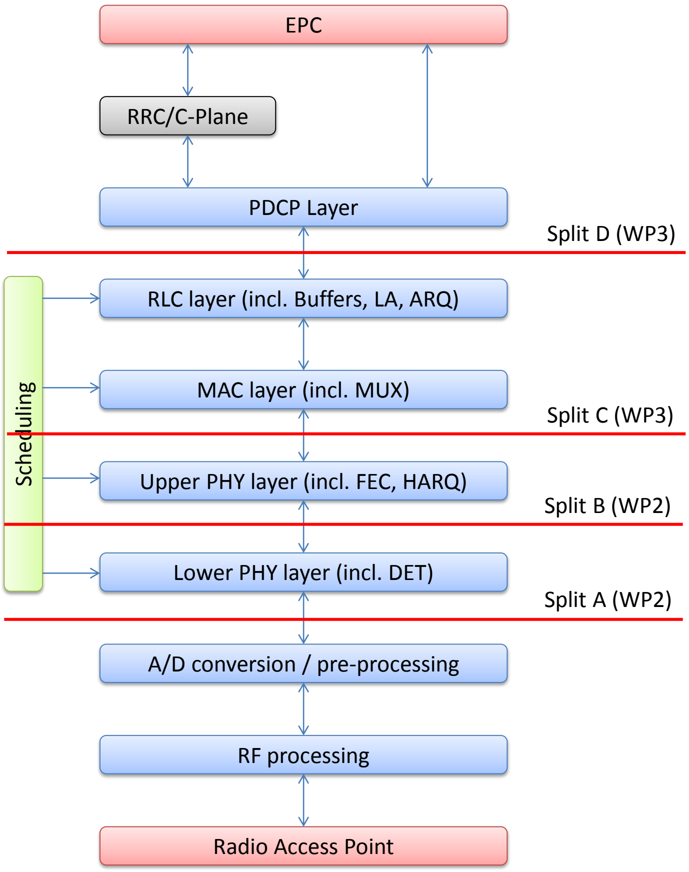

Figure 5-9: Preferred functional splits considered in iJOIN

5.4 Flexible functional split assignment

The objective of this section is to analyse, from a practical viewpoint, the actual flexibility that can be achieved in terms of functional split. Ideally, a fully virtualized environment where each function could be moved at any place would be feasible. In practice, there will be different small-cell implementations which may or may not support different functional splits, or different co-processors which can be turned on/off. In addition, the assignment of the RANaaS PoP cannot be assumed to be changed on the fly.

In practice this means that the set of supported functional splits will be a reduced set. If only one functional split would be supported, it will be the one that provides the best ratio of potential benefits associated to the centralization degree with respect to potential cost variations. The benefits that can be obtained from centralization are mainly associated to increased spectral and energy efficiencies. This ratio depends not only on technical factors but also on other factors such as the traffic demand, user distribution, and the possibility of reusing deployed infrastructure. For example, the centralization of the baseband processing may allow for an implementation of cooperation mechanisms that may help to improve the spectral efficiency, reducing the need for new deployments in high traffic demand areas and making it a sensible option from a techno-economic viewpoint. But if demand is relatively low or other solutions such as carrier-aggregation are available, then it may be that the opposite is reached. On top of this, different technical criteria should also be taken into account.

- Cell based vs. user based processing: cell based processing should be distributed as far as possible, because it reduces the transport requirements and does not exhibit potential processing multiplexing gains. On top of this, this processing is better implemented using hardware solutions.

- Software based processing vs. hardware based processing: As has been indicated in previous sections, some functionality is more efficiently implemented by means of hardware based solutions, while others benefit from a software based implementation (see Figure 5-8).
- Latency requirements: Some processing functionalities are very sensible to latency. Obviously, this factor determines whether a potential centralization in the RANaaS infrastructure is feasible or not.

Based on this description and the previous section, it may be feasible that the implementation at the iSCs split into two parts: a hardware implementation and a software implementation. Then, based on the actual functional split, individual modules would be turned on and off. Each module may be composed of the functionality shown in Figure 5-9. Some of these modules may be implemented in hardware and some in software based on the recommendation in Figure 5-8. In a practical setup, an iSC may support only two, at most three, functional splits:

- A preferred functional split where functionality at the iSC is executed on hardware and all remaining functionality is executed in software at the RANaaS entity. The individual modules at the iSC may be implemented on different co-processors in order to allow for flexible functional split configurations based on a single platform.
- A fall-back solution which is applied in the case that the RANaaS platform is not used and the iSC connects directly to the core network. In this case, all upper layers are executed in software at the iSC in addition to the functionality that is executed in the case of the preferred functional split.

In addition, an iSC may choose an intermediate functional split between preferred functional split and full decentralization. This could be useful in order to adapt to the backhaul network as well as the current computational load of the RANaaS platform.

5.5 Preliminary Results

5.5.1 Opportunistic HARQ

The preferred functional split B would centralize the decoding process while Radio Frequency (RF) processing is still performed at the iSC. This split of RAN functionality allows for implementing advanced decoding algorithms at the central processor and would centralize a major part of computational complexity. However, in this case stringent latency requirements must be fulfilled such as the discussed HARQ timing constraints (see Section 5.2), i.e., HARQ requires that all uplink processing must be finished within 3 ms after receiving a subframe (in 3GPP LTE frequency division duplexing (FDD)). These 3 ms include both the round-trip delay between RANaaS PoP and iSC, as well as the actual decoding operation. Non-ideal backhaul may imply significantly higher latencies. In addition, deploying general purpose hardware at the RANaaS PoP will lead to computational jitter violating real-time constraints.

Hence, we need a solution which is applicable to non-ideal backhaul, which allows for centralizing the computationally intense part of the PHY layer, which meets the stringent timing requirements imposed by HARQ, and which does not impose significant performance penalties. If we were applying currently available technology, with most of the current backhaul solutions only the backhaul round-trip delay may already exceed the HARQ timing requirements and therefore lead to RAN protocol errors. The solution must be standards compliant as any change particularly to mobile terminals should be avoided.

In the following, we provide numerical results for an opportunistic HARQ approach where the iSC estimates the probability of decoding success based on the received SNR. Using this estimate, the iSC sends HARQ feedback to the mobile terminal and forwards the received packets as well as information on the HARQ feedback to the RANaaS instance. If this approach is applied, the RANaaS instance could then combine the received packets, taking into account the HARQ feedback provided by the iSC. The iSC needs not to decode any packet and only deploys one mapping curve using an effective SNR based on the channel state information from all transmissions. The effective SNR incorporates all retransmissions and therefore a one-dimensional mapping can be applied, i.e. only the effective SNR depends on the number of HARQ retransmissions but not the mapping of outage rate to effective SNR.

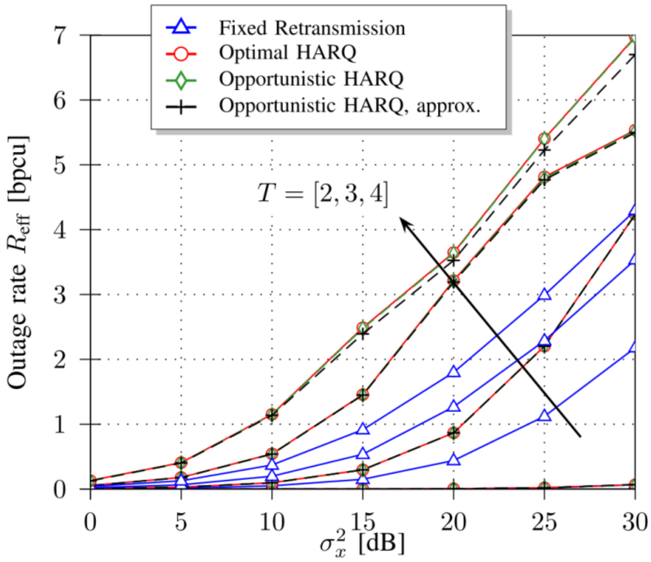

Figure 5-10: Achievable outage rate depending on the SNR for an outage probability of 0.1%

Figure 5-10 shows numerical results for the considered opportunistic HARQ approach which has been described in more detail in [44]. It shows the results for

- Fixed retransmission: For each codeword, T transmissions are independently encoded, transmitted and combined,
- Optimal HARQ: HARQ feedback is provided based on the actual decoding result,
- Opportunistic HARQ: HARQ feedback is provided based on the exact outage probability expression,
- Opportunistic HARQ, approx.: HARQ feedback is provided based on a single mapping curve which employs an effective SNR computed over all transmitted codewords. The effective SNR is given by

$$\gamma_{eff} = \max\left(\sum_{t=0}^{T-1}\gamma_t, (e/4)^T \prod_{t=0}^{T-1}\gamma_t\right). \tag{5.2}$$

The results shown were acquired for identical and independent block Rayleigh fading, 360 information bits, effective outage probability $\varepsilon = 10^{-4}$, and T=[1, 2, 3, 4] transmission rounds. The figure shows the effective spectral efficiency under the given outage probability constraint and considering the actual number of transmissions.

For T = 1, the effective rate of all four approaches coincides below 0.1bpcu and is therefore only recognizable at the bottom of the figure. The results show that opportunistic HARQ is able to maintain the benefits of HARQ and offers the same diversity gain. The benefits compared to a fixed number of transmissions would increase with decreasing outage probability. We can further observe that the effective SNR in (5.2) implies only a minor performance loss for T = 3 and T = 4 compared to optimal HARQ and opportunistic HARQ.

In currently deployed centralized RAN, a very high capacity and very low latency connection between RRH and central processor is required in order to provide HARQ feedback within the required time, i.e. 3 ms in the case of 3GPP LTE FDD. Our approach divides the HARQ process into a time-critical part and computationally intense part. The time-critical part, i.e., determining HARQ feedback, is implemented at the iSC based on the channel state information and without the need to decode the received codeword. The computationally intense part, i.e., decoding the received codeword, is moved towards the RANaaS instance where advanced and computationally intense algorithms may be implemented. In addition, this implies that non-real time general purpose hardware may be deployed at the RANaaS instance, which would imply computational jitter.

The time-critical part has been removed from the RANaaS instance, at least the part of particular relevance to PHY and MAC. Hence, it is possible to relax real-time constraints and deploy general purpose hardware. It further allows for using non-ideal backhaul in a centralized RAN architecture which is critical in areas of high deployment costs, e.g., small-cell deployments where trenching optical fibre would constitute a major part of the capital expenditures.

5.5.2 Computational Outage

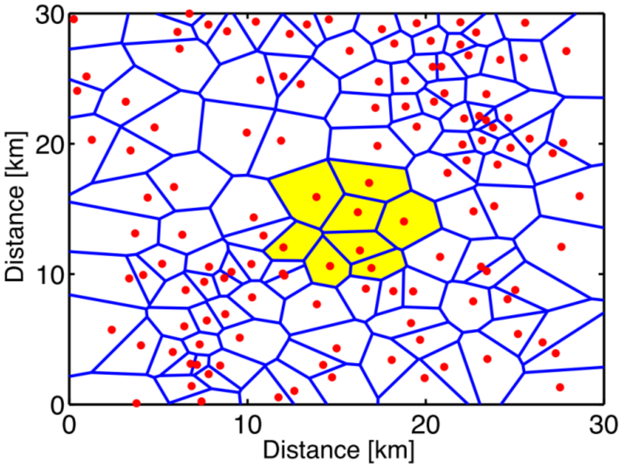

Figure 5-11: Considered network deployment

In the following, we provide numerical results for the impact of computational outage on the system. This analysis uses the network as illustrated in Figure 5-11, originating from an actual deployment of a UK telecommunications operator. We further assume block Rayleigh fading, a simplified distance-dependent path-loss model, and that each iSC serves exactly one user on the whole bandwidth within one sub-frame. In future reports, we will extend this network model to be aligned with the evaluation assumptions of iJOIN. Based on this network model, Figure 5-12 shows performance results for a single link ignoring any inter-cell interference. For these results, we limited the complexity to 50 Mbit-iterations/s which corresponds to about 6-12 processor cores. The left side shows the outage probability depending on the average channel SNR (corresponding with large-scale fading). We can observe a sudden increase of outage between 10-20dB. The reason for this increase is that the likelihood of higher modulation schemes increases where more computational power is required (see Figure 5-6). In the case of CAS, the outage is much lower than in the case of MRS. The right hand-side of Figure 5-12 shows the effective throughput of both approaches. We can see that in the case of limited computational resources, CAS provides higher effective throughput than MRS although the latter uses more spectral efficient modulation and coding schemes.

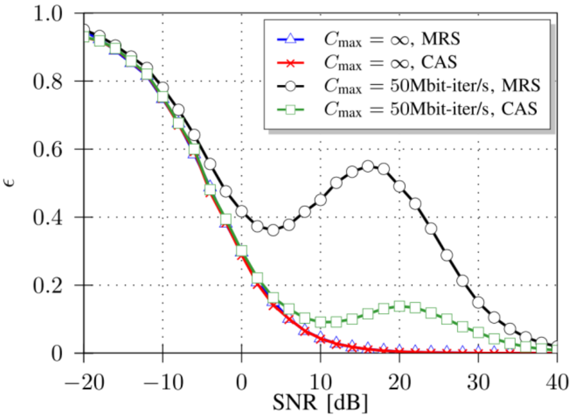

(a) Overall outage probability

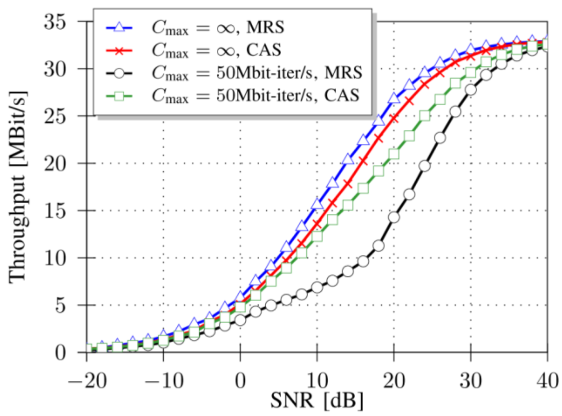

(b) Effective throughput

Figure 5-12: Results for single-cell under a computational complexity constraint

In Figure 5-13, we show the expected sum-throughput as a function of the maximum normalized per-iSC computational complexity and for eight centralized iSCs. We can observe that in the case of strong computational limitations, the centralized CAS (CAS, CP) outperforms all other approaches, e.g. if local processing and MRS is applied, the maximum performance is only achieved with more than 100 Mbit-iter/s while CAS with local processing as well as MRS with central processing achieve their maximum performance with 60 Mbit-iter/s, and CAS with central processing requires only 25-30 Mbit-iter/s.

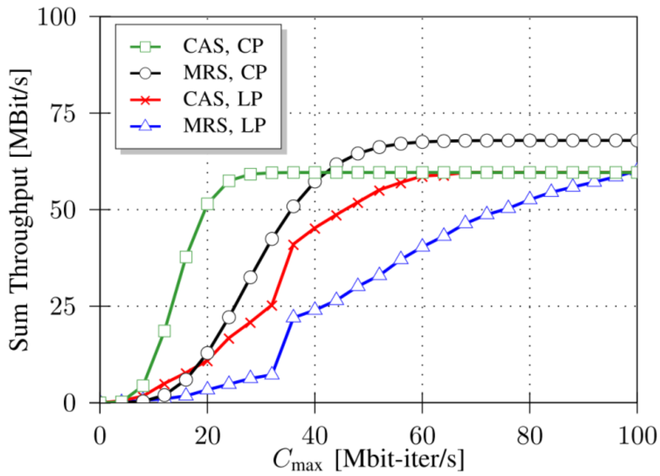

Figure 5-13: Results for multi-cell network for different computational complexity constraints

These results show that a RAN implementation on cloud-platforms requires an intelligent design of channel scheduler and resource scheduling in the virtualized infrastructure. It is necessary that both are aware of each other in order to optimize the throughput performance and resource usage. In the next subsection, we further elaborate on the resource usage by showing the scaling behaviour of computational diversity gain.

5.5.3 Computational Diversity

In Section 5.1.3, we introduced the idea of computational outage and computational diversity $c(N)$. In order to evaluate both quantitatively, we applied and extended the complexity model which was introduced in [12]. In Figure 5-14, we show both the measured curves using system level evaluations (blue line) as well as the theoretical complexity (black line). We use as a measure of complexity again bit-iterations, here, normalized to a channel use.

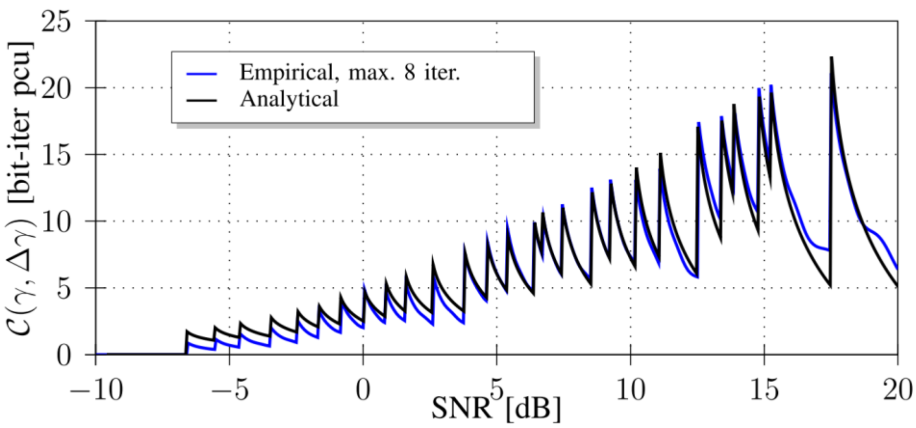

Figure 5-14: Numerical and analytical complexity model for 3GPP LTE uplink

Based on this model, we derived in [43] the expected computational complexity at a single base-station as a function of the decoder quality and the number of MCS schemes. The decoder quality is represented by a constant SNR offset between theoretical Shannon AWGN capacity, $C = \log(1+\gamma)$, and the actual link-adaptation curve. Therefore, smaller rate-offset values will be closer to Shannon capacity but also require more computational complexity. Furthermore, the number of MCS schemes impacts both complexity and achievable rate, i.e., the more MCS schemes are employed, the closer we operate to Shannon's capacity which drives complexity but also improves spectral efficiency. In 3GPP LTE, we would apply a maximum rate scheduler (MRS) which always tries to achieve maximum rate. By contrast, we compare it with a computational aware scheduler (CAS) applying a 0.4dB and 0.9dB link-adaptation offset compared to MRS.

Figure 5-15 shows the complexity scaling in the number of modulation and coding schemes as well as comparing MRS and CAS. Apparently, as we apply more MCS schemes, also the complexity increases. However, the more significant complexity increase is observed when the link-adaptation offset is reduced. In this case, the complexity increases significantly. Furthermore, the results compare numerical results (markers) and results of an analytical framework (solid line) [43].

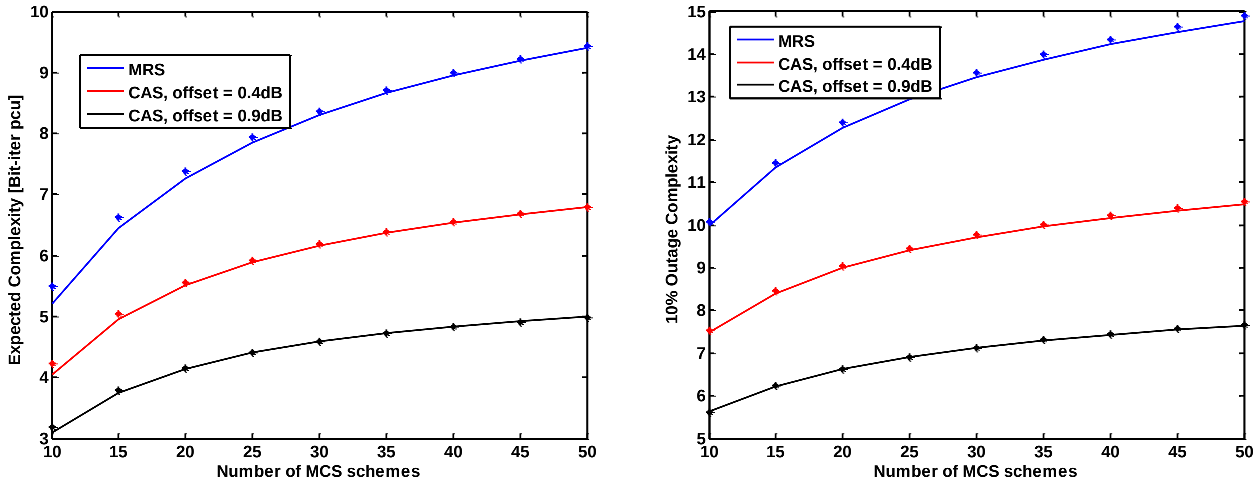

(a) Expected complexity **(b) 10% Outage complexity**

Figure 5-15: Expected and outage computational complexity for one cell

This model is further extended to a network of cells. We assume that each cell serves exactly one user, each user experiences block Rayleigh fading and path-loss dependent fading with path-loss exponent of 2. Furthermore, we normalize the per-cell computational outage as described in Section 5.1.3. In Figure 5-16(a), we can see the absolute normalized complexity $C_{outage,N}(\varepsilon)/N$ for varying decoder quality and different number of cells. We can see that already for a small number of cells, the full computational complexity is exploited, i.e. for $N > 10$ the computational complexity decreases slowly with the number of cells while for $N < 10$ the complexity is reduced significantly. This is again illustrated in Figure 5-16(b) which shows $C_{outage,N}(\varepsilon)/NC_{outage,1}(\varepsilon)$. By contrast to the absolute complexity, we can see that the relative gain is rather independent of the decoder quality. Again a saturation is seen at N = 10 to 20 where about 40% of the computational resources of a distributed system is required.

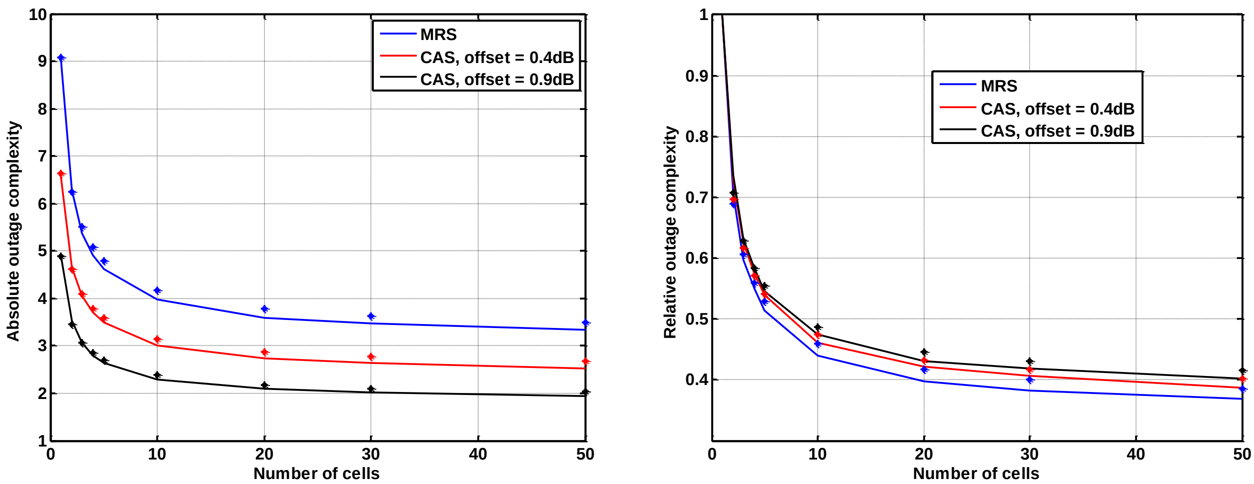

(a) Absolute normalized complexity in bit-iterations per channel use **(b) Relative normalized per-cell complexity**

Figure 5-16: Scaling of computational complexity as a function of number of users/cells

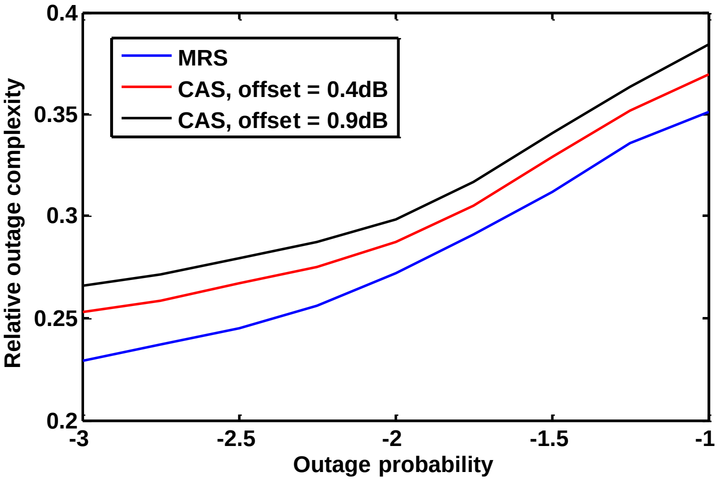

Figure 5-17: Relative outage complexity as a function of the outage probability

Similar to outage capacity in channel coding, also the outage complexity depends on the outage probability that is supposed to be achieved. As the outage probability decreases, more resources need to be provisioned in order to handle also peak demands. This implies, on the other hand, that the system is highly underutilized most of the time and therefore centralization will allow for higher diversity gains. This is shown Figure 5-17 where the relative outage complexity is drawn as a function of the computational outage probability (given in logarithmic domain). Smaller outage probabilities imply a lower relative outage complexity, i.e. for a computational outage probability of 0,1% the RANaaS instance would need to provide only 25% of the computational resources needed in a conventional distributed LTE deployment. The gain is slightly higher for MRS because MRS operates closer to capacity and therefore requires higher over-provisioning.

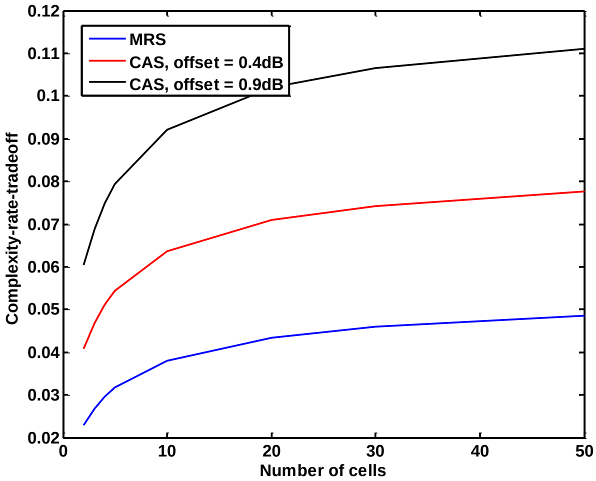

Figure 5-18: Complexity-rate-tradeoff as a function of the number of centralized cells

Finally, Figure 5-18 shows the complexity-rate-tradeoff which measures the number of additional bits per channel use (spectral efficiency) that can be obtained if one additional bit-iteration per channel use (computational complexity) is invested. It therefore measures how complexity and spectral efficiency can be traded off against each other. If we apply a MRS then the system already operates very close to capacity and therefore any further gain in spectral efficiency requires significant additional computational resources. This is reflected by a rather low complexity-rate-tradeoff. By contrast, a system which operates farther away from capacity, e.g. CAS with 0.9 dB offset, gains more spectral efficiency if additional computational complexity is invested.

6 Joint Radio Access and Backhaul Network Requirements

6.1 Required interfaces and interaction

In this section, we describe the interfaces required for a joint RAN/BH operation. In particular, we focus on the required interaction between iTN, iSC, iNC and RANaaS instance. We will pay special attention for the timescale on which information is exchanged. The timescale is a key parameter for the integration of the candidate technologies. Based on the information in reports D2.2 [45], D3.2 [46], and D4.2 [48], Table 6-1 summarizes the interaction of candidate technologies with network elements.

Table 6-1: Overview of interfaces for CTs investigated in WP2 and WP3

CT	J1	J2	iTN	iveC	iNC	EMS/NMS	MME
2.1	X	X		X	X	X	
2.2	X	X		X	X	X	
2.3	X	X		X	X		
2.4	X			X	X	X	
2.5	X	X		X	X	X	
2.6				X	X	X	
2.7				X	X	X	
3.1			X		X		
3.2			X			X	
3.3			X				X
3.4	X	X			X		X
3.5		X	X		X		X
3.7			X		X		X
3.8					X		
3.9					X		
4.1					X		X
4.2	X		X	X	X		
4.3	X		X	X	X		
4.4			X		X		
4.5	X		X		X		

In general, CTs developed in WP2 require low-latency backhaul. In the case of CT 2.2 and 2.4, latencies in the order of microseconds is assumed, and in the case of CT 2.1, 2.3, and 2.5 latencies in the order of few 100µs up to 1ms is assumed. CT 2.6 and 2.7 play a special role as they introduce novel backhaul technologies and therefore rather serve as a transport medium for J1 and J2 traffic. By contrast, CTs developed in WP3 cope with higher latencies. In particular, CTs 3.1, 3.4, 3.5, 3.6, and 3.9 can cope with latencies in the order of milliseconds, and CT 3.1 and 3.5 can cope with latency higher than 10ms (operating on functional D). Finally, CTs developed in WP4 operate on much larger timescale. These CTs regard the control and management of the RAN and backhaul network. The temporal granularity of this operation determines the frequencies with which information must be exchanged with the iTNs, iNC, and veNBs.

6.2 Limitations in 3GPP LTE

6.2.1 3GPP interfaces and requirements

3GPP considers the transport network underlying the mobile network as out of scope of its standardization focus. Consequently, 3GPP specifications are in general agnostic to transport network technologies and, in particular, 3GPP assumes that underlying transport networks are not contended. They are therefore assumed to satisfy the requirements for network operation.

In the mobile backhaul, the following traffic types based on 3GPP interface definitions can be differentiated:

- S1-U traffic destined for the S-GW; note that S1-U traffic can be further differentiated according to the assigned QCI value;
- S1-C traffic destined for the MME;
- X2-U and X2-C traffic destined for other eNodeBs;
- OSS (operations support system) traffic destined for core applications that provide fault, configuration, and performance management;
- Network synchronisation traffic.

All these traffic types have different requirements regarding quality of service (QoS). It can be generally stated that control plane traffic, e.g. S1-C, X2-C, and synchronization traffic, have higher requirements in terms of latency and reliability but have lower bandwidth requirements compared to user-plane traffic, e.g. S1-U and X2-U.

In today's networks, traffic differentiation for 3GPP traffic types is implemented via traffic type, e.g. control plane or user plane, and traffic class, e.g. based on QCI. Both are mapped on transport network traffic differentiation techniques, which depend on the employed transport network technology. For example, legacy ATM defines four different traffic classes which describe bandwidth requirement characteristics such as constant bit rate or variable bit rate. However, no delay requirements are specified. In LTE-Advanced, all-IP networks with layer 3 routing and VPN technologies, e.g. MPLS, or QoS and IP-aware layer 2 switching technologies, e.g. based on 802.1q/p, are expected to play a larger role due to the availability of Ethernet-capable eNodeBs in the access network and corresponding cost benefits.

While 3GPP defines a set of standardized QCI values [21], there is no standardized guideline available on how mobile network traffic is mapped to service classes on the transport layer. The problem is amplified by differences in the implementation between different vendors.

To illustrate the challenge, Table 6-2 shows the standardized QCI values in 3GPP for different service classes. The quantitative parameters include the packet delay budget and the packet error loss rate, both referring to the overall connection from access to core or vice versa. In Table 6-3, IEEE 802.1Q Priority Code Point (PCP) recommendations are shown (note that there are no standardized parameter sets), which are often applied to Ethernet or similar link technologies in the backhaul network. The challenge is now to map QCI traffic to according PCP values, which additionally need to be parameterized appropriately.

It can be concluded that neither 3GPP nor other standardization bodies offer a standardized methodology on how to map interface and protocol requirements of the mobile network to the backhaul network. Configuration is thus a case-by-case issue which needs fine-tuning for each deployment and equipment scenario.

Table 6-2: 3GPP standardized QCI values

QCI	Resource Type	Priority	Packet Delay Budget	Packet Loss Rate	Example Services
1	GBR	2	100 ms	10^{-2}	Conversational Voice
2		4	150 ms	10^{-3}	Conversational Video (Live Streaming)
3		3	50 ms	10^{-3}	Real Time Gaming
4		5	300 ms	10^{-6}	Non-Conversational Video (Buffered Streaming)
5	Non-GBR	1	100 ms	10^{-6}	IMS Signalling
6		6	300 ms	10^{-6}	Video (Buffered Streaming) TCP-based (e.g., www, e-mail, chat, ftp, p2p file sharing, progressive video, etc.)
7		7	100 ms	10^{-3}	Voice, Video (Live Streaming) Interactive Gaming
8		8	300 ms	10^{-6}	Video (Buffered Streaming) TCP-based (e.g., www, e-mail, chat, ftp, p2p file sharing, progressive video, etc.)
9		9			

Table 6-3: IEEE 802.1Q Priority Code Point recommendations [22]

PCP	Priority	Acronym	Traffic Types
1	0 (lowest)	BK	Background
0	1	BE	Best Effort
2	2	EE	Excellent Effort
3	3	CA	Critical Applications
4	4	VI	Video, < 100 ms latency and jitter
5	5	VO	Voice, < 10 ms latency and jitter
6	6	IC	Internetwork Control
7	7 (highest)	NC	Network Control

6.2.2 Recommendations

It can be concluded that for 3GPP, the ongoing discussion on the impact of virtualization both in core network and RAN is an opportunity to also discuss the necessity of standardized interfaces or mechanisms on how to deal with different backhaul characteristics.

Within iJOIN, a coordination of backhaul characteristics and CT functions is foreseen in the interplay between the SDN-based iNC and the iveC, which is potentially deployed in the RANaaS PoP. The coordination has to take into account the functional split, the CT requirements and the backhaul characteristics.

6.3 Backhaul Technologies

Regarding the different technologies that can be employed in the backhaul, a thorough analysis is performed in D4.2. Table 6-4 [48] summarizes different backhaul technologies that are considered within the scope of iJOIN. The decision on the functional split and the implementation feasibility of the different candidate technologies will depend mainly on the latency and the throughput imposed by the technology available for the backhaul. In this regard, for wireless backhaul, the latency per hop and the achievable throughput depend on the range of frequencies employed, on availability of Line of Sight transmission and on the topology (PtP or PtmP). On the other hand, for fibre transmission, the main factor is the topology and the multiplexing technology.

Table 6-4: Backhaul Classification [1][48]

Number	BH technology		Latency (per hop, RTT)	Throughput	Topology	Duplexing	Multiplexing Technology
1a	Millimeter wave	60GHz Unlicensed	≤5 ms	≤800 Mbit/s	PtP (LOS)	TDD	
1b			≤200 µsec	≤1Gbps	PtP (LOS)	FDD	
1c		70-80GHz Light licensed	≤200 µsec	≤2.5 Gbit/s	PtP (LOS)	FDD	
2a	Microwave (28-42 GHz) Licensed		≤200 µsec	≤1Gbps	PtP (LOS)	FDD	
2b			≤10 ms	≤1Gbps	PtmP (LOS)	TDD	TDMA
3a	Sub-6 GHz Unlicensed or licensed		≤5 ms	≤500Mbps	PtP (NLoS)	TDD	
3b			≤10 ms	≤500Mbps (shared)	PtmP (NLoS)	TDD	TDMA
3c			≤5 ms	≤1 Gbit/s (per client)	PtmP (NLoS)	TDD	SDMA
4a	Dark Fibre		5 µs/km × 2	≤10 Gbps	PtP		
4b	CWDM		5 µs/km × 2	≤10·N Gbps (with N≤8)	Ring		WDM
4c	Metro Optical Network		250 µs	≤1 Gbps	Mesh/Ring		Statistical Packet Multiplexing
4d	PON (Passive Optical Networks)		≤1 ms	100M – 2.5Gbps	PtmP		TDM (DL) TDMA (UL)
5	xDSL		5-35 ms	10M – 100Mbps	PtP		

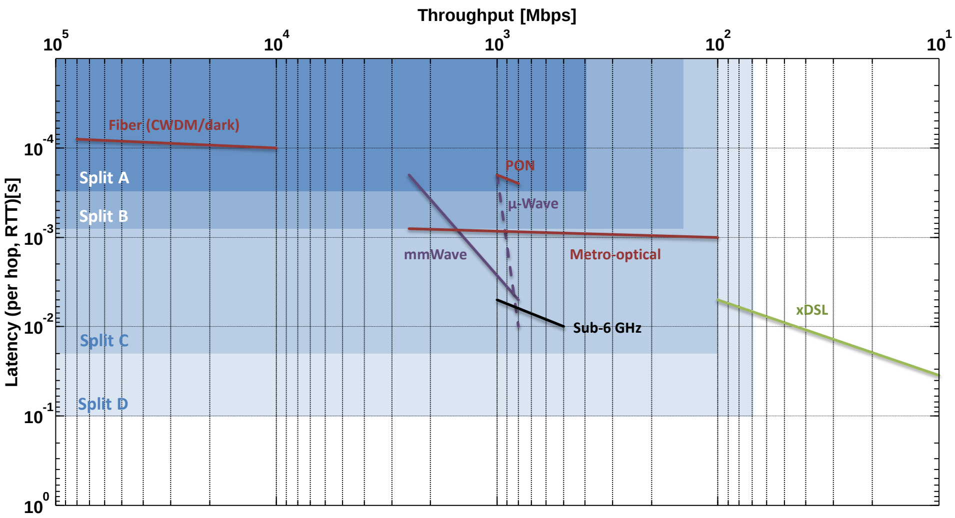

Figure 6-1: Ordering of backhaul technologies and preferred functional splits

Based on the backhaul technologies listed in Table 6-4 and the preferred functional splits introduced in Section 5.3, we can group backhaul technologies according to their ability to support the individual functional splits. This ordering is illustrated in Figure 6-1 and will applied for the iJOIN concept evaluation discussed in Section 7.2.

6.4 Preliminary results

6.4.1 Joint RAN/BH Coding

In current mobile networks, RAN and BH links are often perceived as separate links in terms of coding. The RAN link is encoded to fit the channel quality of the radio access channel and the BH link is encoded according to the BH channel quality. While a BH channel code is not required in the fibre-based CPRI backhaul, it is important when using outdoor E-band links as they face varying channel conditions due to environmental changes. In CT2.7, iJOIN investigates the possibility of jointly encoding and decoding both links in the uplink. If the decoding of the RAN link is offloaded to the RANaaS instance, then the data on the BH link is already protected by the RAN channel code. By adapting the code rate of the RAN channel code not only to the RAN channel's quality but to that of the BH link as well, a second en-/decoding is unnecessary and an uncoded BH transmission can be used. Further details on this can be found in D2.2 [45]. This not only reduces the hardware requirements in the iSCs but also reduces the latency. Every additional processing performed on the BH link increases the overall latency between the UE and the RANaaS instance. If the decoding is offloaded to the RANaaS instance, all timing constraints discussed in Section 5.2, especially the tight constraint for the HARQ acknowledgement, have to be met by the combined RAN/BH transmission.

The joint encoding integrates very simply into the current standard. The RAN code rate is decided by the veNB by taking the current frame error rate into account and communicating the decision in form of DCI (Downlink Control Information) during the UL grant. A low-quality BH link increases the frame error rate, which will be noticed by the veNB and adjusts the code rate accordingly, which is depicted in Figure 6-2.

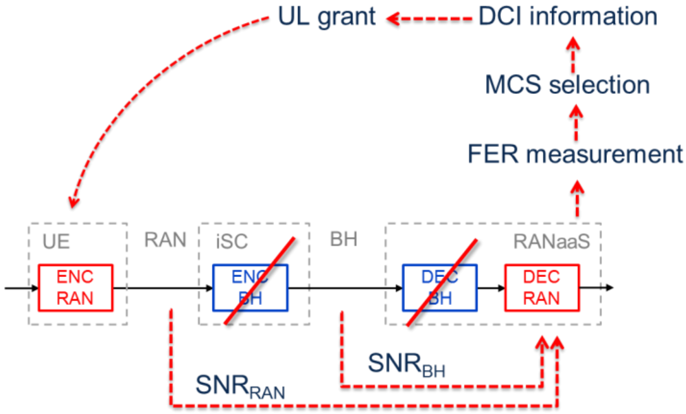

Figure 6-2: Code rate adaption and channel quality measurements required for joint RAN/BH en-/decoding

The joint encoding/uncoded BH scheme slightly decreases the throughput compared to using a separate BH code [9]. However, the additional BH transmission should not decrease the end-to-end throughput of the system. To mitigate the lower performance of joint decoding, two new schemes can be deployed. The first is a soft-input/soft-output dequantizer (SISODQ), which enables forwarding soft information between the demodulation modules of the RAN and the BH link [20]. The second is an error resilient decoder (ERD) that takes an erroneous BH channel into account when calculating log-likelihood ratios before decoding in the RANaaS instance [47]. Both schemes increase the throughput even beyond the value of a separately coded BH, which can be seen from Figure 6-3. The SISODQ performs better than the ERD at the cost of a higher complexity. The increased throughput can also be traded off for energy efficiency by using lower transmitting power on the BH. As investigated in [20], the SISODQ is faster than an additional BH en-/decoder under certain circumstances, thereby reducing the latency as required. However, the exact latency depends very much on the actual implementation.

In conclusion, CT2.7's approach on joint RAN/BH coding integrates easily into the RANaaS concept, enables the low latency required for centralized processing, and reduces the complexity of the iSCs, while at the same time increasing throughput or energy efficiency of the network.

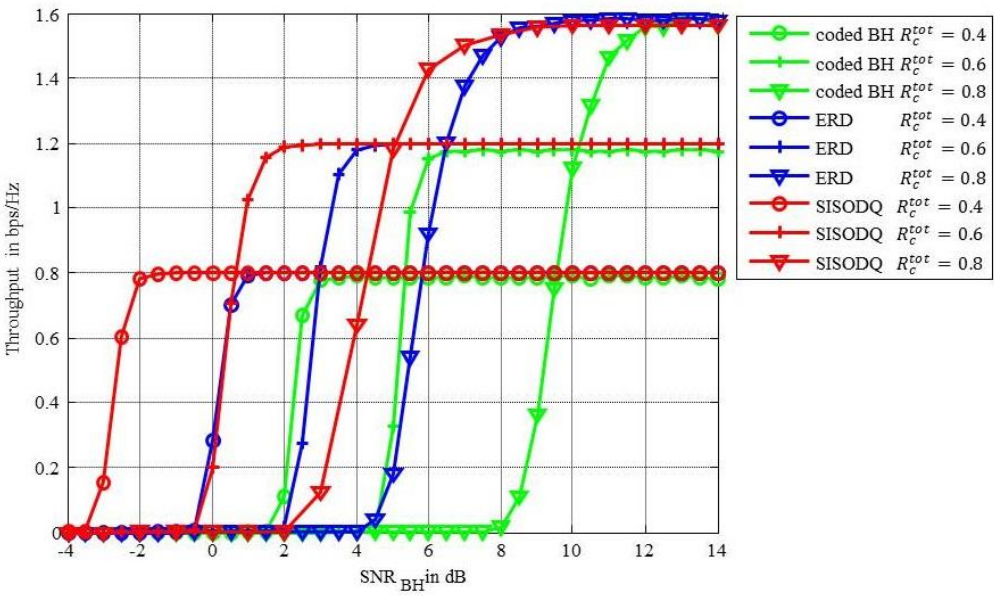

Figure 6-3: Throughput when using encoded BH as compared to an uncoded BH and employing a SISODQ or a ERD

6.4.2 Distributed IP Anchoring and Mobility Management

In CT4.1, iJOIN investigates an SDN-based distributed IP anchoring and mobility management. As reported in D4.2 [48], the AMM module (Anchoring and Mobility Management) is in charge of managing the mobility of the UEs attached to the network. When a UE moves from a point of attachment to another, AMM runs different algorithms to select the optimal anchor and configure the new path in the network. The AMM module follows the SDN-paradigm and is executed within the iNC using information provided by iTNs and iSCs.

At the time of writing this document, an initial implementation of the AMM module is available. The AMM runs on the iNC of the SDN Testbed as Ryu application. Ryu is the iOpenFlow [10] controller running on the iNC. By now, the following functionality has been implemented:

- UE attachment detection: veNBs detect the attachment and inform directly the AMM module. At the moment no communication occurs with the MME.
- Upon the UE attachment, the AMM selects the anchor statically. This means that the anchor selection algorithm is not implemented yet.
- Once an anchor is selected, the AMM configures properly the anchor rules for both new and old anchors.

In order to evaluate the AMM module, a basic Traffic Engineering Enforcement Module (TEEM) version has been implemented. The TEEM module provides the functionalities required by AMM. The AMM module requests the TEEM module to compute the best path and set up the new best path for UE or traffic flow. Figure 6-4 shows the software architecture implemented on the SDN-testbed running on the iNC. The communication between the modules occurs as internal Ryu event following an event-driven communication paradigm. Each module exports the required events which make subscription to other modules available.

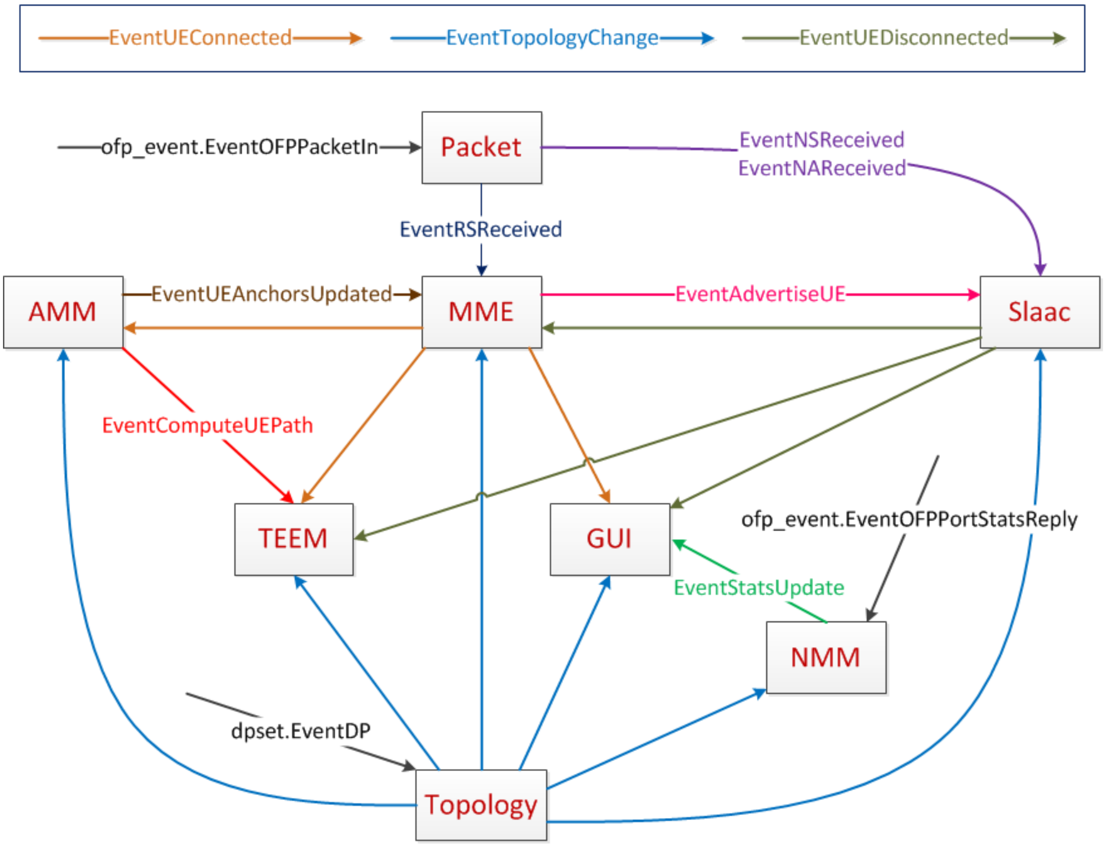

Figure 6-4: Partial functional architecture implemented on SDN-Testbed

We obtained preliminary results for the inter-anchor mobility scenario using the following measurement methodology:

- One node external to the SDN Testbed starts pinging the UE by sending ICMPv6 echo request packets every 2 ms.
- During this pinging procedure we detach the UE from the current iSC and attach it to a new iSC. The attachment triggers the AMM procedure.
- By measuring the ICMPv6 sequence number gap we can roughly estimate the overall handover time with a granularity of 2 ms.

Figure 6-5 shows the Cumulative Distribution Functions (CDFs) for the total handover time considering separately layer 2, layer 3 and ping disconnectivity, in the case of three anchors assigned to a single UE.

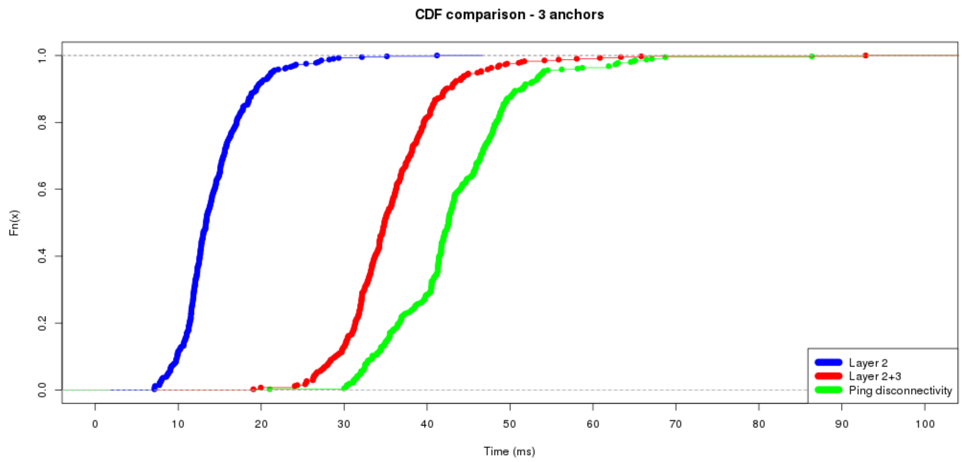

Figure 6-5: Handover time CDF

Using this methodology, the total handover time in terms of 95% percentile is:

- 21 ms, for layer 2 handover;

- 46 ms, for layer 3 handover;
- 54 ms, for ping disconnectivity.

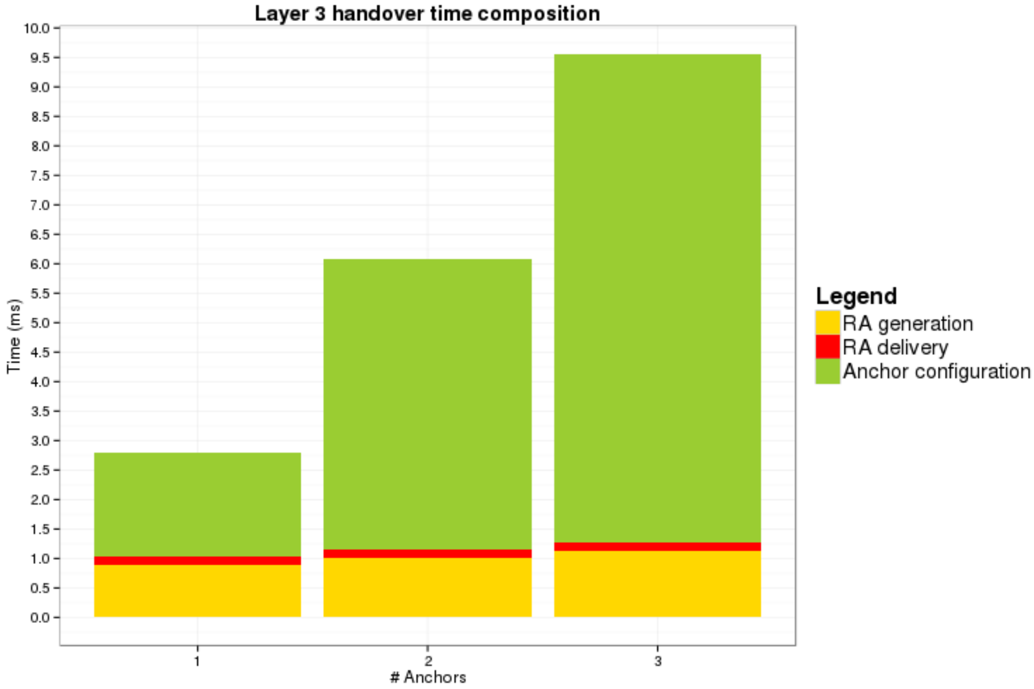

Figure 6-6: Total handover processing time

The selection of the anchor is actually performed statically. Therefore, we believe that the total handover time will be slightly higher. Currently the main contribution to the total handover time is given by the configuration of OpenFlow rules on the anchors as depicted in Figure 6-6. In the SDN-Testbed, TEEM takes 1 ms for sending one OpenFlow configuration packet to the anchors. This time is also highly dependent on the distance between the controller and the anchors. Further measurements will be made in order to evaluate this impact. The second main contribution is given by the creation and the delivery of Router Advertisement packets by AMM to iSCs. In our SDN Testbed this procedure takes 1 ms.

6.4.3 Network Wide Energy Optimisation

CT4.2 investigates an SDN-based network wide Energy optimisation. As reported in D4.2 [48], the Network Energy Optimizer (NEO) module is in charge of managing the energy savings that the cellular network can achieve. NEO runs an algorithm that tries to decrease the energy consumption of the cellular network under a given side constraint on the throughput performance.

In general, significantly more opportunities arise for switching off iSCs in smaller timescale than in larger timescales due to (a) coverage overlaps stemming from heterogeneous deployment of cells, (b) larger spatio-temporal load variations, and (c) power-proportional and load-dependent iSCs. Thus, NEO not only guarantees the user Quality of Experience (QoE) while switching-off an iSC but also considers the achievable energy savings even for short time-scales.

At the time of writing this document, an initial implementation of the NEO module is available. Thus, we have already considered the iSCs (access network) and present initial results about feasible energy savings. While NEO switches off a cellular node, it must guarantee desired levels for the user QoE [1][39]. Specifically, NEO considers:

- **Network coverage**, i.e. the probability that a random user experiences poor signal quality when it needs to use the network is defined as failure probability. While switching-off an iSC, then some users are going to be attached to other iSCs, so the average failure probability of a random user increases. We denote as $p_{failure}$ the maximum allowed failure probability, which would usually be defined by the operator.
- **Admission control and blocking probabilities** *(for flows that require a "dedicated" amount of bandwidth)*; these probabilities are not only related to user admission but also admission of flows that require a certain amount of dedicated bandwidth. While switching-off an iSC, some users are

going to be attached to other iSCs.Hence, some iSCs will have to deal with more flows that require a certain amount of bandwidth, thus the blocking probability of such a flow due to lack of resources increases. We denote as p_{block} the maximum allowed blocking probability, again defined by the operator.

- **Service delay** *(for "best-effort" flows);* the probability that the delay exceeds a desired upper bound. We denote as D_{max} the maximum delay threshold that is defined from the operator.

NEO decides to switch-off an iSC fulfilling the above criteria. The larger the QoE thresholds ($p_{failure}$, p_{block}, D_{max}) defined by the operator, the less strict we are with the switching-off criteria, such that more iSCs can be switched off and more energy can be saved.

Figure 6-7 illustrates the achievable energy savings for different values of the above parameters and for the user QoE in a scenario of 120 iSCs and 2 macrocells. For example, in the up-right picture, the top curve corresponds to the portion of energy saved when we consider only the first constraint to be active, if the switching-off duration is supposed to last 10 min. On the x-axis we increase the constraint threshold and plot the respective energy savings. As can be seen there, increasing the threshold, i.e. making the constraint less strict, increases savings as it allows for more iSCs to be switched off. We can save up to 68% of the total energy consumption of our cellular network for $p_{failure} = 0.4$. The bottom curve also shows the energy savings but now with the other two constraints active as well: the blocking threshold is fixed at 10^{-3} and the delay threshold at D_{max}=50 ms. As can be seen, savings increase again but less sharply as the other two constraints can become the bottleneck for a switch-off decision, especially as $p_{failure}$ increases. For example, with $p_{failure}$=0.4 and the other two thresholds fixed, the portion of energy saving can be up to 30%. Similar behaviour is observed in the other two pictures of the figure.

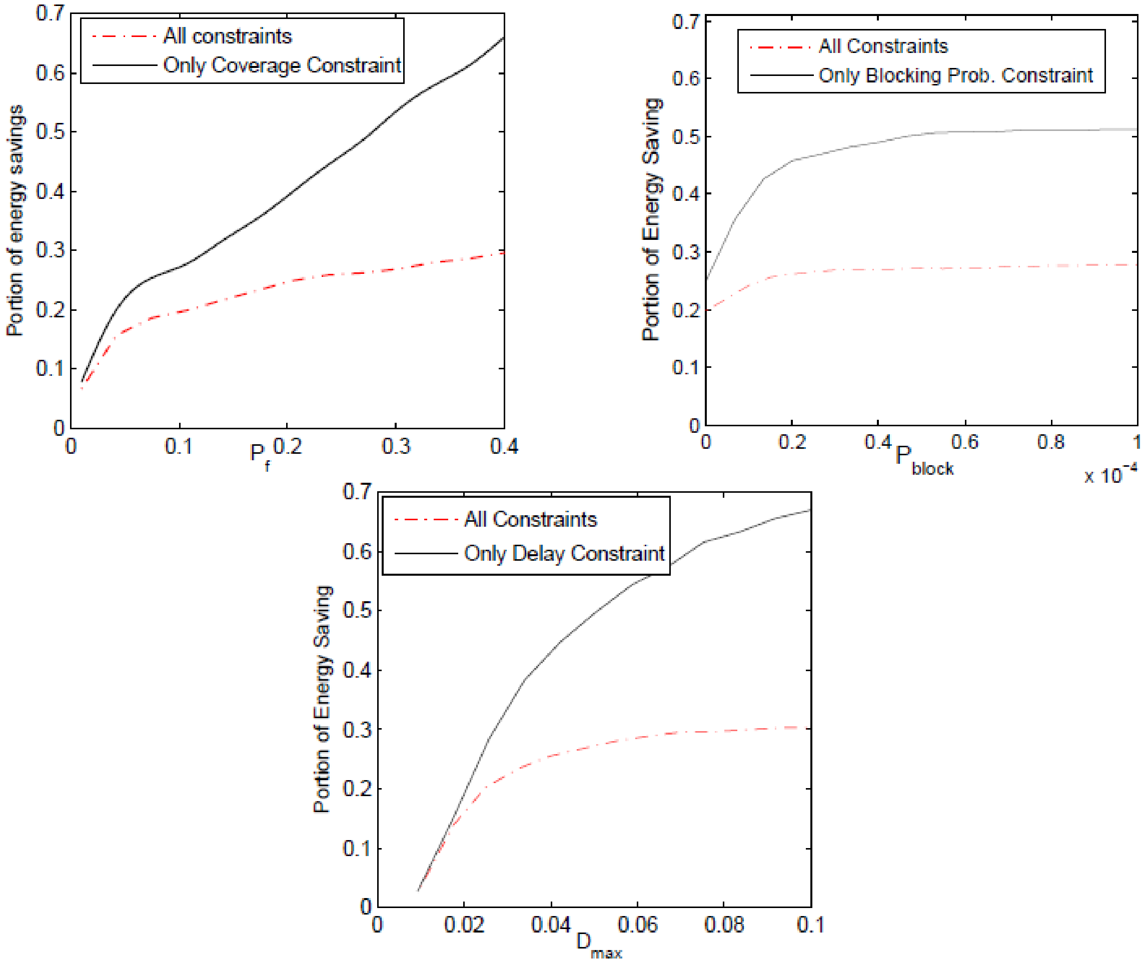

Figure 6-7: Achievable Energy Savings vs. Thresholds

Figure 6-8 depicts the portion of energy saved for different values of the switching-off period (X). As can be seen, energy savings are maximized when X is relatively small but start decreasing and eventually flatten out as X increases. The reason is that, for small X, it needs to be considered only the impact of active users when evaluating the constraint and the impact of hand overs to neighbouring iSCs. However, as X increases, there is a higher chance that connected and potential disconnected users will add to the total transferred load and thus there might be a bigger impact on existing and remote users, which might prevent us from switching off an iSC. Finally, the plot for each respective constraint is not always linear, as some additional phenomena, such as convergence to stationarity for the stochastic systems we use in constraints 2 and 3, also affect systems' behaviour.

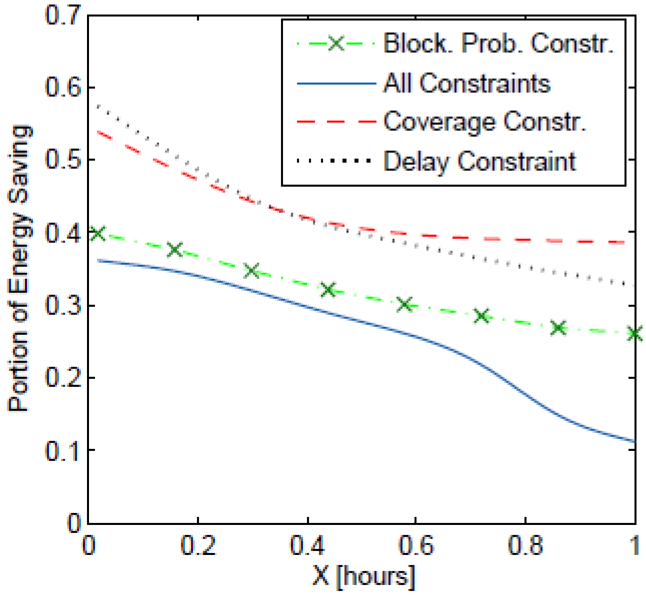

Figure 6-8: Achievable Energy Savings vs. Switching-off period

7 iJOIN System Performance Evaluation

7.1 Relevant metrics

7.1.1 Area Throughput

The past few years have seen a tremendous growth in internet data carried by mobile cellular networks, i.e. the mobile traffic demand in year 2020 will be at least 1000 times more than the capacity of current cellular networks [11]. In this context, iJOIN targets to increase the system throughput within the same spectrum by a factor of 50-100 as a result of:

- High density of small cells, re-use of spectrum, and PHY / RRM improvements enabled by RANaaS to adequately address interference (≥10x)
- Shorter distances and increased LOS probability (5-10x)

Area Throughput is expressed in terms of *bits/area*. It measures the utilization of the radio spectrum over a given geographic area and also represents the capacity that a mobile operator offers to its subscribers. Observing the network over time period T, one can measure the traffic flowing through the network. Denoting by $r_i(t)$ the rate by which bits are correctly delivered at (from) the UE i, the total information (number of bits) delivered, within the time period T, in a network comprising N UEs is calculated as:

$$I = \sum_{i=1}^{N} \int_{0}^{T} r_i(t)dt \qquad \text{[bit]} \tag{7.1}$$

The average rate R in the network is then given by I/T. Often, it is helpful to normalise the rate R by either the number of cells or the network area. In order to allow for a normalised measure, we choose here to work with rate per area unit expressed in square kilometres, i.e.:

$$R_A = \frac{R}{A} = \frac{I}{A \cdot T} = \frac{1}{A \cdot T} \cdot \sum_{i=1}^{N} \int_{0}^{T} r_i(t)dt \qquad [\text{bps/km}^2] \tag{7.2}$$

A shortcoming of this definition is that it provides only an average value and does not reflect the distribution of throughput in a given area. Capacity distribution in a given area greatly impacts Quality of Service (QoS) to a mobile user, for example, in terms of session dropping, session blocking, and data rate requirements. Simulations are often used to produce not only average values but also the CDF of the cell or user throughput in order to give more complete information about the system behaviour. In particular, an important metric that must be taken into account is the cell edge user throughput: a good system design should take into account also this statistic, so that also minimum radio performance is guaranteed in the covered area.

7.1.2 Energy Efficiency

Energy efficiency (EE) evaluation of the iJOIN system and each CT is strictly related to the proposed logical architecture. The power consumption of the veNB should take into account the iSCs, RANaaS instance and the backhaul network including iTNs. Energy consumption of the different network entities in iJOIN depends on the particular functional split adopted by each single CT as well as on the particular topology of the backhaul network considered for system level evaluations.

At a first glance, it is not intuitively clear whether the iJOIN system implies a reduced energy consumption than today's networks, and whether the consumed per delivered bit reduces. The iJOIN architecture is enabling advanced and convenient RAN sharing scenarios that can significantly improve energy efficiency and long term sustainability. At network level, energy efficiency evaluations are traditionally performed [3], [13] by considering two kind of metrics: energy per information bit expressed by [J/bit] or equivalently [W/bps], and power per area unit expressed by [W/m^2]. Thus, given a specific evaluation scenario, it is possible to compare a certain EE metric of a classical flat architecture with the iJOIN architecture, considering both RAN and backhaul parameters.

In iJOIN, the main metric for energy efficiency is the consumed energy per information bit (see deliverable D5.1 [3] for further details). The traditional architecture considers a distributed implementation composed of overlaid macro and small cells. By contrast, the iJOIN architecture considers a partially centralized implementation composed of iSCs and RANaaS PoPs. The energy efficiency analysis must consider the sum of all contributions in the network. The system energy consumption is directly related to the power usage of each network element over a given period of time. Considering the system architecture as introduced in Section 4 a holistic power model for a RANaaS instance comprising N_{iSC} iSCs can be given by:

$$P_{\text{Total}} = P_{\text{RANaaS}} + P_{\text{Bh}} + \sum_{n=1}^{N_{\text{iSC}}} P_{\text{iSC-}n} \quad (7.3)$$

where P_{RANaaS}, P_{Bh} and $P_{\text{iSC-n}}$ represent the consumed power at the RANaaS instance, backhaul network, and each iSC n, respectively. The power consumption at iSCs and RANaaS instance depends, among others, also on the particular functional split. Furthermore, in the case of a flexible functional split, the EE analysis must consider time-variant power consumption according to the functional split switching. The power consumption of the backhaul network is mainly due to the presence of iTNs. Further details of the power model are given in [24] and Annex A.

Energy Efficiency is measured as an Energy Consumption Index (ECI) [3]:

$$EE = \frac{E_{Total}}{I} = \int_0^T \frac{P_{Total}(t)}{\sum_{i=1}^{N} r_i(t)} = \int_0^T \frac{P_{RANaaS}(t) + P_{BH}(t) + \sum_{i=1}^{N} P_i(t)}{\sum_{i=1}^{N} r_i(t)} dt \quad (7.4)$$

Note that energy saving enablers adapt the system characteristics to the load variations. Hence, full buffer traffic should not be applied to assess the EE. By contrast, a given user demand should be provided, i.e. $\int_0^T \sum_{i=1}^{N} r_i(t)$, which must be maintained while reducing the energy consumption of the system.

7.1.3 Utilisation Efficiency

Utilisation efficiency is defined as a metric which expresses how well the available resources are utilised for a given performance metric. Therefore, high utilization efficiency means the following:

- The system such as a network is highly utilized, and therefore not over-provisioned.
- The system is capable to exploit utilized resources efficiently to provide the desired output, such as cell throughput or other metrics.

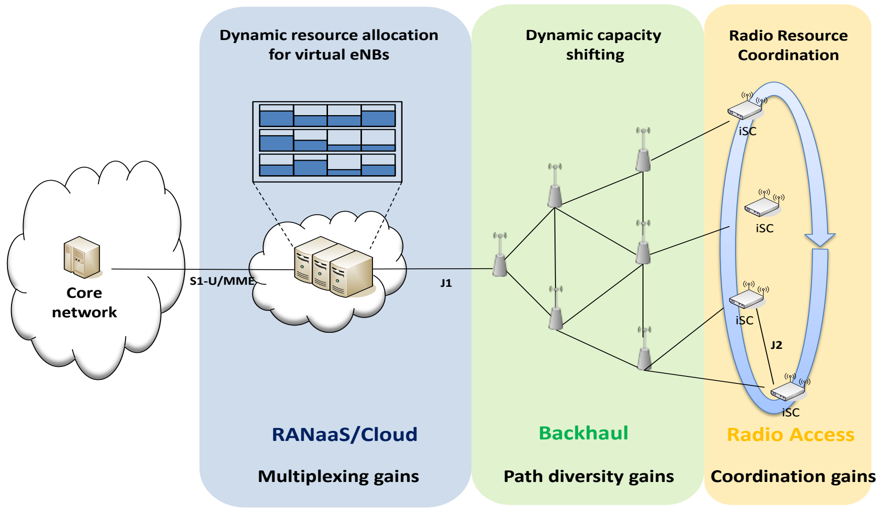

Figure 7-1: Utilization gains in different network domains

Figure 7-1 shows an example of how different resource allocation techniques in different iJOIN network domains can lead to different types of gains, e.g. multiplexing, diversity and coordination gains. It also illustrates a fundamental problem of defining a network-wide metric for utilization efficiency: different network domains, i.e. RANaaS instance, backhaul, and radio access, utilize different types of resources, e.g. CPU cycles, link bandwidth, radio spectrum. Hence, a simple summation of domain-specific metrics is in general not possible. We define the total utilization efficiency of a system as following:

$$\eta_U = \frac{\sum_{d \in D} \alpha_d u_d}{|D|} \quad (7.5)$$

where α_d is a scaling factor s.t. $\sum \alpha_d = 1$, and u_d is the *domain utilization* for the considered domain, with D being the set of network domains, e.g. RANaaS instance, backhaul network, and RAN.

The definition of the domain utilization u_d depends on the resource of interest. As described in [23], different network domains have in many cases different resources. However on a more abstract level, resource normalization can be applied across network domains. We identified the following resource classes which will be investigated in more detail:

- Bandwidth/capacity resources. The domain utilization is defined as

$$u_d^B(X) = \frac{B_{mean,d}(X)}{B_{cap,d}(X)}, \quad (7.6)$$

  where $B_{mean,d}(X)$ is the average measured data rate and $B_{cap,d}(X)$ is the corresponding outage or theoretical maximum capacity of the system. The parameter X depends on the investigated network scenario and can be the number of cells or user arrival rate.

- Computational resources. Here, the domain utilization is defined by

$$u_d^C(X) = \frac{C_{mean,d}(X)}{C_{outage,d}(X)}, \quad (7.7)$$

  where $u_d^C(X)$ is the ratio of expected computational demand and provided computational resources, depending on the number of cells in the scenario, X. The latter is the outage complexity which is defined as the amount of computational resources to make sure that a per-cell computational outage probability ε is not exceeded. Both are defined through an analytical framework which has been described partly in Section 5.1.3. This framework resembles the characteristics of computational load of a 3GPP LTE uplink decoder.

<u>Preliminary evaluation of computation utilisation efficiency:</u>

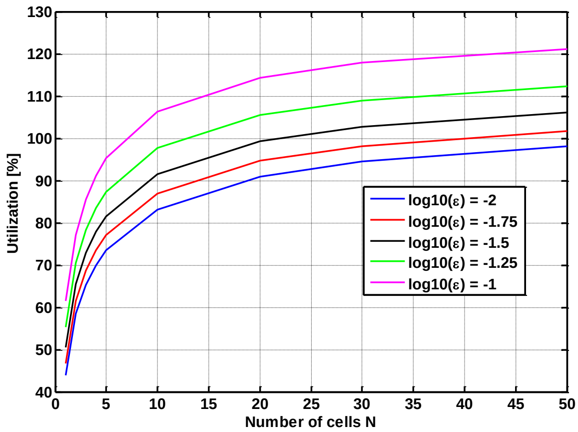

Figure 7-2: Computational utilisation efficiency as a function of number of cells and outage probability

Based on this framework, the expected utilisation of a centralised processor for different number of cells and depending on the outage probability criterion is shown in Figure 7-2. For these results typical LTE parameters including actual SNR link-adaptation thresholds have been used. Furthermore, a Rayleigh fading process is assumed with SNR of 10dB. In the next report, this investigation is extended to more complete fading processes including path-loss and power control. However, the results will differ quantitatively but not qualitatively.

From Figure 7-2 we can see that for a large number of centralized base stations, an expected utilization of more than 100% is achieved. This implies that less computational resources than the expected overall computational demand are provided. This is due to the fact that the system is optimized such that a per-cell outage probability is not exceeded. We can observe that this effect depends strongly on the chosen outage probability, e.g. for a computational outage of 10% already 7 centralized base stations would exceed the provided resources while for a computational outage of 1% more than 50 base stations need to be centralized. This utilization performance curve will be helpful to dimension the centralized resources accordingly and to design the resource scheduler. Based on the actual communication resource demand (throughput) also the computational resource demand (processing) can be scheduled, and vice versa.

In the next report, these results will be further detailed to include more practical constraints and characteristics, e.g. multiplexing gain and more complex channel models.

7.1.4 Cost Efficiency

The cost efficiency of iJOIN will be investigated by combining the large-scale analytical results based on stochastic geometry obtained in D4.2 [48] with the analysis of computational complexity and diversity as illustrated in Section 5.5.3. Users, base stations, backhaul nodes, and datacenters are modelled using independent homogeneous Poisson point processes as illustrated in Figure 7-3.

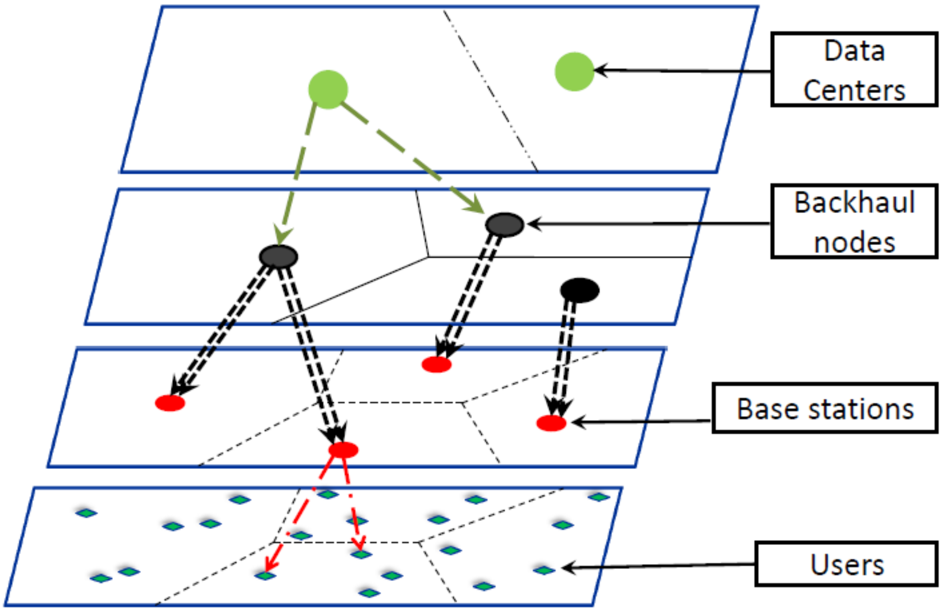

Figure 7-3: Network model for cost-efficiency analysis

A particular equipment cost for each device is assumed. Capacity and infrastructure cost are assumed to have given “base” cost for connecting two different network components and this base cost is assumed to increase with the distance between the network components. Utilising this method of modelling, we can obtain an expression for the average cost of deploying a backhaul node. From which, we obtain the total cost of the network.

Furthermore, we use the results of the computational complexity analysis to derive the expected computational diversity gain, i.e. a linear function which defines the required computational resources depending on the number of centralized users and as multiple of the required computational resources for a single base-station. The required computational resources depend on the service quality, i.e. if the LTE system operates at its maximum achievable rate more computational resources are required while at slightly reduced achievable rate fewer resources are necessary. These dependencies are taken into account by scaling the required computational resources with the offered achievable rate per base station while increasing the base-station density accordingly. We further take into account that a maximum computational outage must not be exceeded. The model will incorporate expected costs of datacenters as a function of the datacenter size.

Using this model,

- we will compare the cost performance against a traditional scenario in which there is no RAN functionality executed in the cloud;
- we will leverage on analysis of the computational resources required in the cloud vs. resources required in the eNB, size of the area served by the cloud, etc;
- we will determine the deployment cost for the iJOIN network and for a traditional one as a function of the cost of the individual parameters (such as the cost of a processing unit, the cost of a bandwidth unit in the backhaul, etc.).

Preliminary results of the cost-efficiency study are shown in Figure 7-4 which illustrates the expected capital expenditures (CAPEX) for a given datacenter density. It shows that the expected CAPEX is up to about 15% lower than in the case of DRAN. We can further observe that an implementation using MRS as well as CAS (see Section 5.1.3) has no impact on the cost-efficiency while keeping the required area throughput constant. This implies that the system can be operated at lower computational complexity at the cost of additional iSCs. However, reduced computational complexity also implies lower software latency within the RAN protocol stack which is a major obstacle for GPP based RAN implementations.

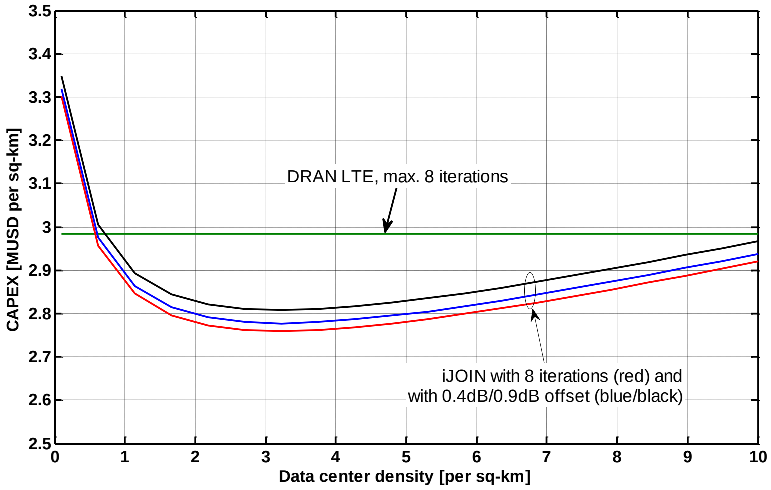

Figure 7-4: Expected capital expenditures as a function of data center density

7.2 iJOIN Concept Evaluation

In iJOIN, WPs 2, 3, and 4 investigate CTs which aim for improving the four objectives: area throughput, energy-efficiency, cost-efficiency, and utilization-efficiency. In addition, iJOIN partners introduced and analysed technologies which allow for flexible assignment of RAN functionality, for implementing RAN functionality on commodity hardware, and for controlling the RAN, BH network, and RANaaS instance jointly. In order to provide a consistent and coherent system design, iJOIN performed a system integration of CTs developed in WPs 2, 3, and 4. Part of this integration is the analysis of side-effects and impact on objectives, interaction of CTs within and across WPs, and the individual operating points of the iJOIN system. Note that a detailed analysis of CT interaction within each layer has been performed in D2.2 [45], D3.2 [46] and D4.2 [48]. This section summarises the interaction within each layer, how CTs interact across layers, and this section provides an overview of perspective operating points of the iJOIN system.

7.2.1 iJOIN System Design

iJOIN introduces a novel system concept which allows for a flexible centralisation of radio access network functionality towards a cloud-computing platform based on commoditised hardware. In addition, the iJOIN system concept allows for considering jointly RAN and BH aspects to operate and to design the mobile network. Overall, iJOIN defines four objectives which provide considerable degrees of freedom to operate the overall system at differently optimised operating points. In order to allow for exploiting these degrees of freedom, the iJOIN system investigated novel technologies for the wireless access, the management of the radio access and backhaul network, and for the operation of the centralized computing infrastructure.

These novel technologies target different objectives which may be counteracting each other leading to unpredictable behaviour and inefficient system performance. In order to avoid this, iJOIN divided its technologies into two main areas: one area addressing improvements on short timescale and mainly targeting improvements of area throughput, and secondly, another area addressing improvements on larger timescale and mainly targeting energy efficiency improvements. The first area is mostly covered by technologies in WP2 and WP3 while the second area is mostly covered by technologies in WP4. This allows for a two-level system optimisation where network management algorithms (developed in WP4) perform system optimisation towards higher energy efficiency. Decisions taken on network management level (larger timescale) are then used by technologies operating on shorter timescale and target an improvement of area throughput on shorter timescale.

For example, WP4 developed one technology which allows for turning on and off jointly RAN and BH equipment. In order to perform this task, it receives two inputs: achievable rates by the technologies applied to both RAN and BH (developed in WP2 and WP3), the data rate demands which must not be violated. Based on these two inputs, equipment will be turned on and off. Furthermore, WP4 technologies are triggered which allow for an efficient re-routing of traffic and control of the resulting network topology. In the RAN, technologies developed by WP2 and WP3 will guarantee the required data rates through innovative algorithms such as multi-cell signal processing and interference coordination. Finally, novel control technologies for RAN data processing on commodity hardware (introduced in WP3 and WP5) will make sure that the required data processing capabilities can be guaranteed.

The described approach will be applied to selected and representative iJOIN system evaluation scenarios, consisting of iJOIN common scenarios, iJOIN preferred functional splits, and iJOIN candidate technologies. The next sections provide an overview of the selected evaluation objectives, the system evaluation scenarios, and the underlying data which has been used to select the individual technologies for each system evaluation scenario.

7.2.2 iJOIN System Evaluation

The iJOIN system will be evaluated with respect to four individual objectives. In the following, we describe how the evaluation of individual objective will be addressed. First, iJOIN will perform a cost-efficiency analysis which is based on the framework presented in Section 7.1.4 as well as in D4.2 [48]. This framework allows for assessing the required capital expenditures based on the required data rates, backhaul network topology and technologies, and the applied radio access network technologies. The main purpose of this framework is to assess costs at the planning stage (CAPEX) in order to provide a tool for deciding whether the deployment of the iJOIN system would be cost-efficient or not in a particular scenario.

The objectives area throughput and energy efficiency are applied after the decision for the deployment of the iJOIN system has been made and the particular physical mobile network structure is known. Based on this knowledge, different technologies can be applied in order to ensure efficient cell clustering, multi-cell signal processing, per-link energy efficiency, or efficient network management. However, all these technologies have two objectives: 1) adjust the network topology in order to minimize the energy consumption, 2) maximize the area throughput under a given network topology in order to satisfy the data rate demands. Both objectives result inherently in a reduction of operational expenditures and improved network utilisation efficiency. Furthermore, iJOIN introduced technologies which allow for maximizing the data processing utilisation efficiency under a given area throughput.

These four objectives will be evaluated by iJOIN using the iJOIN system evaluation scenarios described next.

7.2.3 iJOIN System Evaluation Scenarios

iJOIN identified four system evaluation scenarios which will demonstrate the effectiveness of iJOIN's two main concepts, flexible functional split and joint RAN/BH operation, respectively. These four scenarios are:

1. **Stadium scenario, Split A**: In this case, we apply the scenario described in Section 4.1 as well as the preferred functional split A (see Section 5.3) which requires high-throughput and low-latency backhaul, i.e. dark-fibre technologies as in the case of CPRI.
2. **Wide-area scenario, Split A**: In this case, we apply the scenario described in Section 4.3 as well as the preferred functional split A. By contrast to the stadium scenario, this scenario covers a large area and is subject to mobility. Furthermore, Split A represents a dark-fibre solution and we will demonstrate how iJOIN technologies can improve area throughput and energy efficiency over existing technologies.

3. **Wide-area scenario Split B**: Again, the wide-area scenario is considered but with no particular dark-fibre technologies resulting in higher latencies and the choice of preferred functional split B. We will use this scenario to demonstrate how iJOIN can maintain a major part of the area throughput and energy gains compared to the previous scenario. In addition, this scenario will show how iJOIN allows for exploiting these gains in scenarios in which centralization was not feasible with state-of-the-art technologies and therefore these gains could not be exploited.
4. **Wide-area scenario Split C**: Similar to the previous scenario, this scenario shows how to exploit centralization gains in scenarios which so far did not support any centralized processing. Compared to the previous scenario, the backhaul requirements are further reduced as mainly legacy backhaul is considered as well as scenarios in which backhaul deployment is very difficult (e.g. NLOS backhauling).

These four scenarios will allow for showing the power of the iJOIN system concept. However, within each work package further detailed evaluation of all four common scenarios and different functional splits are performed. In order to allow for comparability of results and to facilitate the integration of results across multiple work packages, the system assumptions as listed in Table 7-1 for stadium scenario Table 7-2 for the three wide-area based scenarios will be applied.

Table 7-1: Description of iJOIN system evaluation scenario based on stadium scenario

Parameter	Value	Description
Number of small-cells	15	Placed as uniform grid
Covered area	40 m x 80 m	Representing a part of stadium standings
Number of active users (exp.)	320	UE separation of 1m along x-axis and 0.5 m along y-axis; 5% of UEs assumed to be active
UE dropping	Regular grid	Representing regular grid of seats
Minimum distance	UE-iSC = 5 m	
Backhaul capacity / latency	10 Gbps / 5 µs	Case 4a in Table 6-3

Table 7-2: Description of iJOIN system evaluation scenario based on wide-area scenario

Parameter	Value	Description
Number of small cells	19	Placed as two-tier hexagonal grid; wrap-around optional
Number of UEs per cell (exp.)	2 Randomly dropped	Reflecting high temporal and spatial traffic fluctuations
Inter-site distance	50 m	Very dense network
Minimum distance	UE-iSC = 5 m	
MIMO configuration	2 antennas at iSC, 1 antenna at UE	
Bandwidth	10 MHz	
Backhaul capacity / latency	A: 10 Gbps / 5 µs B: 2.5 Gbps / 400 µs for J2 and 600 µs for J1 C: 500 Mbps / 10 ms	- Rates per UL and DL - Considers BH technologies 4a (A), 1c (B) and 3a (C) in Table 6-4 - Assuming two hops for J2 and three hops for J1 - Neglecting switching latency

Furthermore, based on the candidate technologies summarized in the next part, the preliminary assignment of candidate technologies to iJOIN system evaluation scenarios as shown in Table 7-3 has been performed.

Table 7-3: Application of iJOIN candidate technologies to iJOIN system evaluation scenarios

Objective / WP	WA, Split A	WA, Split B	WA, Split C	ST, Split A
Area throughput / WP2	CT2.2 (MPTD) CT2.4 (CoMP) CT 2.6 (BH)	CT2.1 (INP) CT2.3 (JCNC) CT2.5 (CoMP) CT 2.7 (BH)	CT2.2 (SPTD) CT2.5 (CoMP)	CT2.2 (MPTD) CT2.4 (CoMP) CT 2.6 (BH)
Area throughput / WP3		CT 3.8	CT3.4 Optional: CT 3.6 CT 3.9	CT 3.5 Opt: CT 3.1
Energy efficiency / WP3	Optional: CT3.3	Optional: CT3.3	CT3.3	
Energy efficiency / WP4	CT4.2	CT4.2	CT4.2	N/A

7.2.4 Intra-Layer Interaction of CTs

Table 7-4, Table 7-5, and Table 7-6 provide a summary of the candidate technologies and their properties which have been considered to derive the previously introduced iJOIN system evaluation scenarios.

Table 7-4: Interaction of CTs operating on PHY layer

CT	Main objective	Side effects	Compatible CTs
CT2.1	Area Throughput	Processing shifted among nodes improves utilization. Reduced J1 rate reduces BH cost and BH energy	All DL CTs; CT 2.6, CT2.7
CT2.2	Area Throughput	Processing to enhance the detection of users scheduled on the same resources, mainly targeting the edge users. Utilization Efficiency (UE) of the RAN and BH may also be improved as the maximum throughput supported is also increased.	All DL CTs;
CT2.3	Area Throughput	Increased Utilization Efficiency (UE), as a reduced BH rate is required to provide a target system throughput. Reduction of BH rate may also reduce cost or energy.	All DL CTs;
CT2.4	Area Throughput	Increasing energy efficiency (EE) by improving area throughput (AT) with the same power consumption	All UL CTs, CT 2.6 and CT2.7; Potentially can interoperate with CT2.5
CT2.5	Area Throughput	Increased complexity of precoding at the iSCs. Hierarchical structure of the channel state information.	All UL CTs;

CT2.6	Area Throughput	Reduction of BH rate reduces cost. Utilization Efficiency (UE) of BH may also be improved by performing statistical multiplexing of the traffic generated by different iSCs over BH links	All UL and DL CTs
CT2.7	Area Throughput	Reduction of BH rate reduces cost or energy. Wireless should be cheaper than laying fibre	All UL and DL CTs

Table 7-5: Interaction of CTs operating on MAC layer

CT	**Main objective**	**Side effects**	**Compatible CTs**
CT 3.1	Area Throughput	This CTs improves the UEff thanks to the joint RAN/BH resource optimization	All
CT 3.2	Area Throughput	This CTs improves the UEff thanks to the joint RAN/BH resource optimization	All
CT 3.3	Energy Efficiency	This CTs improve the CE since OPEX is reduced through dynamic small cell on/off	All
CT 3.4	Area Throughput	This CT is enabler of RANaaS allowing for robust link-adaptation.	All except CT 3.5, CT 3.9
CT 3.5	Area Throughput	This CT improve the system spectral efficiency through ICIC, which has a positive side effect on UEff	All except CT 3.4, CT 3.9
CT 3.7	Area Throughput	Processing to enhance the detection of users scheduled on the same resources, mainly targeting the edge users. UEff may also be improved as the maximum throughput supported is also increased.	All except CT 3.8
CT 3.8	Area Throughput	Joint Processing shifted among nodes improves UEff.	All except CT 3.7
CT 3.9	Area Throughput	This CT has a positive impact on the EE due to the implemented power control scheme	All except CT 3.4, CT 3.5

Table 7-6: Interaction of CTs operating on network level

CT	**Main objective**	**Side effects**	**Compatible CTs**
CT4.1	Utilisation Efficiency		All CT4.1 and CT4.2 interact in order to avoid conflicts
CT4.2	Energy Efficiency		All CT4.2 and CT4.1 interact in order to avoid conflicts CT4.2 and C4.5 are compatible at different time scales

CT4.3	Cost Efficiency	Increased Utilisation Efficiency due optimal RANaaS placement	All
CT4.4	Utilisation Efficiency	Increased Cost Efficiency by avoiding backhaul's links congestion	All
CT4.5	Utilisation Efficiency		All, CT4.5 and C4.2 are compatible at different time scales

8 Summary and Conclusions

This report provided a comprehensive overview of the iJOIN functional architecture, i.e. how novel candidate technologies interact, which objectives they address, and which impact they have on the overall system. This will be required until the end of the project to perform a system-wide evaluation. This report further provides a detailed analysis of the split of RAN functionality. In particular, implementation aspects have been discussed which will lead to a feasible study at the end of the project. In addition, virtualised infrastructure received particular attention as it will lead to new constraints and requirements if RAN functionality is executed on top of it. 3GPP LTE RAN constraints have been identified and discussed. In this report, solutions to the most challenging constraints are discussed and results are provided. Beside the functional split analysis, the joint RAN/BH operation has been further detailed. Finally, this report discussed the evaluation campaign based on harmonized parameters for each common scenario agreed across all partners of the iJOIN project.

Based on this report, the following preliminary conclusions can be drawn

- An implementation of RAN functionality on commodity hardware appears feasible. Only a very limited set of functional splits seem to be useful, i.e. digitized received signals (similar to CPRI) if the required bandwidth and backhaul technology is available, digitized and (soft-) demodulated/modulated signals in order to perform centralized decoding, or only centralized RRC while lower-layer functionality remains with the iSC.

- Due to practical 3GPP LTE RAN constraints, an implementation of a functional split over heterogeneous backhaul network is challenging. Most importantly, latency and throughput constraints of the underlying backhaul technology determine the achievable functional split. The probably most challenging task is to mitigate the latency constraints, e.g. incurred to HARQ and radio resource control for which iJOIN introduced novel technologies which are able to cope with these constraints.

- iJOIN will perform a harmonized evaluation campaign where results will be compared on a relative performance basis, i.e. using a common set of parameters, each CT is compared to the baseline system. Based on this relative performance, CTs are compared and it is shown in which scenarios they are most efficient and how their performance scales in system parameters such as RAP density and user density.

- Basis for the comparison of CTs will be the four objectives energy-efficiency, cost-efficiency, utilization-efficiency, and area throughput. All four objectives are defined in this report.

Acknowledgements and Disclaimer

This work was partially funded by the European Commission within the 7th Framework Program in the context of the ICT project iJOIN (Grant Agreement No. 317941). The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the iJOIN project or the European Commission.

References

[1] 3GPP, "TR 25.814 v7.1.0; Physical Layer Aspects for Evolved UTRA", Sep. 2006.

[2] 3GPP TR 36.872, "Small cell enhancements for E-UTRA and E-UTRAN - Physical layer aspects (Release12)", V12.1.0, Dec. 2012.

[3] iJOIN, Deliverable D5.1 "Revised definition of requirements and preliminary definition of the iJOIN architecture," November 2013.

[4] G. Chen, W. He, J. Liu, S. Nath, L. Rigas, L. Xiao, and F. Zhao, "Energy-Aware Server Provisioning and Load Dispatching for Connection-Intensive Internet Services", *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, vol. 8, 2008, pp. 337-350.

[5] B. P. Rimal, E. Choi, Eunmi, and I. Lumb, "A taxonomy and survey of cloud computing systems", *IEEE 5th International Joint Conference on INC, IMS and IDC*, 2009, pp. 44-51.

[6] A. K. Sidhu, S. Kinger, "Analysis of Load Balancing Techniques in Cloud Computing", *International Journal of Computers & Technology*, vol. 4, no. 2, 2013, pp. 737-741.

[7] J. W. Smith, "Green Cloud: A literature review of Energy-Aware Computing", University of St Andrews, http://jamie.host.cs.st-andrews.ac.uk/files/LiteratureReview.pdf, 2010.

[8] J. Di Giglio, D. Ricci, "High Performance, Open Standard Virtualization with NFV and SDN", http://www.windriver.com/whitepapers/ovp/ovp_whitepaper.pdf

[9] iJOIN, Internal Report IR 2.2 "Preliminary definition of PHY layer approaches that are applicable to RANaaS and a holistic design of backhaul and access network," May 2014.

[10] iJOIN, Internal Report IR 4.2 "Network-layer algorithms and network operation and management: preliminary set of candidate technologies," May 2014.

[11] Nokia Siemens Networks, "2020: Beyond 4G, Radio Evolution for the Gigabit Experience", Aug. 2011

[12] P. Grover, K. A. Woyach, and A. Sahai, "Towards a communication-theoretic understanding of system-level power," IEEE Journal on Selected Areas in Communications, September 2011

[13] EARTH, Deliverable D2.4 "Most suitable efficiency metrics and utility functions," December 2011.

[14] 3GPP TR 22.951, Service Aspects and Requirements for Network Sharing, Rel.11, Sep. 2012.

[15] 3GPP TR 22.852, Study on Radio Access Network (RAN) Sharing Enhancements, Rel.12, June 2013.

[16] 3GPP TS 23.251, Network Sharing; Architecture and Functional Description, Rel.12, Dec. 2013.

[17] 3GPP TS 36.413, S1 Application Protocol (S1AP), Rel.12, Mar. 2014.

[18] 3GPP TS 36.423, X2 Application Protocol (X2AP), Rel.12, Mar. 2014.

[19] 3GPP TS 36.331, Radio Resource Control (RRC), Rel. 12, Mar. 2014.

[20] J. Bartelt, G. Fettweis, „A Soft-Input/Soft-Output Dequantizer for Cloud-Based Mobile Networks," *15th IEEE International Workshop on Signal Processing Advances in Wireless Communications (SPAWC 2014)*, Toronto, Canada, June 2014.

[21] 3GPP TS 23.203, Policy and charging control architecture, March 2014.

[22] IEEE 802.1q, Virtual LANs, October 2012.

[23] iJOIN, Deliverable D3.1 "Final report on MAC/RRM state-of-the-art, requirements, scenarios, and interfaces in the iJOIN architecture," November 2013.

[24] D. Sabella, A. De Domenico, E. Katranaras, M. Imran, M. Di Girolamo, U. Salim, M. Lalam, K. Samdanis, A. Maeder: "Energy Efficiency benefits of RAN-as-a-Service concept for a cloud-based 5G mobile network infrastructure", submitted to IEEE Networks magazine, 2014.

[25] G. Auer, V. Giannini, I. Godor, P. Skillermark, M. Olsson, M. Imran, D. Sabella, M. Gonzalez, C. Desset, O. Blume, A. Fehske, “How much energy is needed to run a wireless network?,” IEEE Communications Magazine, vol.18, no.5, pp. 40-49, October 2011.

[26] A. De Domenico, E. Calvanese Strinati, A. Capone, “Enabling Green cellular networks: A survey and outlook,” Computer Communications, Volume 37, pp. 5-24, January 2014.

[27] ICT FP7 FIT4GREEN, Project Deliverable D3.3, “Presentation of the full-featured federated energy-consumption models,” March 2012. www.fit4green.eu/sites/default/files/attachments/documents/D3.3_final.pdfwww.fit4green.eu/sites/default/files/attachments/documents/D3.3_final.pdf

[28] T. Werthmann, H, Grob-Lipski, and P. Proebster, “Multiplexing gains achieved in pools of baseband computation units in 4G cellular networks,” Proceedings of the IEEE 24th International Symposium on Personal Indoor and Mobile Radio Communications (PIMRC), September 2013.

[29] P. Monti, S. Tombaz, L. Wosinska, and J. Zander, "Mobile backhaul in heterogeneous network deployments: Technology options and power consumption," 14th International Conference on Transparent Optical Networks (ICTON), pp.1,7, 2-5 July 2012.

[30] NGMN Alliance, "Small cell backhaul requirements," White Paper, June 2012.

[31] iJOIN, Deliverable D4.1 “Report on SotA and requirements for network-layer algorithms and network operation and management,” November 2013.

[32] Open Multi-Processor (OpenMP), [online: http://openmp.org/wp/]

[33] Open Message Passing Interface (OpenMPI), [online: http://www.open-mpi.org/]

[34] D. Wübben, P. Rost, J. Bartelt, M. Lalam, V. Savin, M. Gorgoglione, A. Dekorsy, and G. Fettweis, “Benefits and impactImpact of cloud computingCloud Computing on 5g signal processing”. To appear in the 5G special issue ofSignal Processing”, IEEE Signal Processing Magazine, November 2014.

[35] A. Maeder, M. Lalam, A.D. Domenico, E. Pateromichelakis, D. Wübben, J. Bartelt, R. Fritzsche, P. Rost, “Towards a Flexible Functional Split for Cloud-RAN Networks”, EuCNC'14

[36] Blaise Barney, Lawrence Livermore - National Laboratory, “POSIX Threads Programming”, [online: https://computing.llnl.gov/tutorials/pthreads/]

[37] Kendall Square Research, [online: http://en.wikipedia.org/wiki/Kendall_Square_Research]

[38] Kernel Based Virtual Machine, [online: http://www.linux-kvm.org/page/Main_Page]

[39] N. Sapountzis, S.Sarantidis, T.Spyropoulos, N.Nikaen, U.Salim, “Reducing the Energy Consumption of Small Cell Networks subject to QoE Constraints”, Globecom 2014

[40] EARTH project, “D2.3 - Energy efficiency analysis of the reference systems, areas of improvements and target breakdown,” [online: https://www.ictearth.eu/publications/publications.html]

[41] ASOCS Ltd, [online: http://www.asocstech.com/]

[42] G. Li, S. Zhang, X. Yang, F. Liao, T. Ngai, S. Zhang, and K. Chen, “Architecture of GPP based, scalable, large-scale C-RAN BBU pool,” in IEEE GLOBECOM 2012 Workshops, International Workshop on Cloud Base-Station and Large-Scale Cooperative Communications, Anaheim, CA, USA, Dec.2012.

[43] P. Rost, S. Talarico, and M. Valenti, “The Complexity-Rate-Tradeoff of Centralized Radio Access Networks,” submitted to IEEE Transactions on Wireless Communications, October 2014.

[44] P. Rost and A. Prasad, “Opportunistic HARQ – Enabler of cloud-RAN over non-ideal backhaul,” IEEE Wireless Communications Letters, June 2014

[45] iJOIN, Deliverable D2.2 “Definition of PHY layer approaches that are applicable to RANaaS and a holistic design of backhaul and access network,” November 2014.

[46] iJOIN, Deliverable D3.2 “Definition of MAC and RRM approaches for RANaaS and a joint backhaul/access design,” November 2014.

[47] J. Bartelt and G. Fettweis, "An Improved Decoder for Cloud-Based Mobile Networks under Imperfect Fronthaul," Globecom 2014 Workshop – Wireless optical network conversion in support of cloud architectures (GC14 WS- WONC), Austin, USA, Dec. 2014.

[48] iJOIN, Deliverable D4.2 "Network-layer algorithms and network operation and management: candidate technologies specification," November 2014.

[49] iJOIN, Deliverable D6.1, "Preliminary proof-of-concept results for selected candidate algorithms," November 2014.

[50] M. C. Valenti, S. Talarico, and P. Rost, "The role of computational outage in dense Cloud-based centralized radio access networks," IEEE Global Conference on Communications, December 2014, Austin (TX), USA.

[51] D. Sabella, A. De Domenico, E. Katranaras, M. Imran, M. Di Girolamo, U. Salim, M. Lalam, K. Samdanis, A. Maeder, "Energy Efficiency benefits of RAN-as-a-Service concept for a cloud-based 5G mobile network infrastructure", IEEE Open Access 2014.

Annex A Power consumption models of iJOIN architectural entities

In the following, the power consumption of each individual network element is discussed. Furthermore, some examples of measures are provided to correlate and obtain an idea on the order of magnitude of each element's power consumption, depending on the cells' load (which is interrelated to the cells' RF output power). Figure A-1 summarizes the power models developed in iJOIN and published in [51] related to the different network elements of our proposed network architecture.

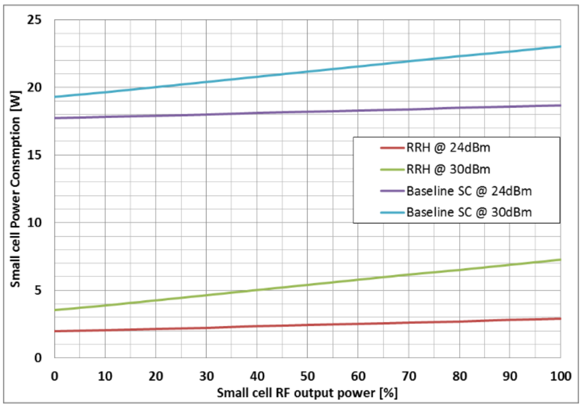

a) Complete small cell and RRH power consumption with respect to different RF output power and power constraints

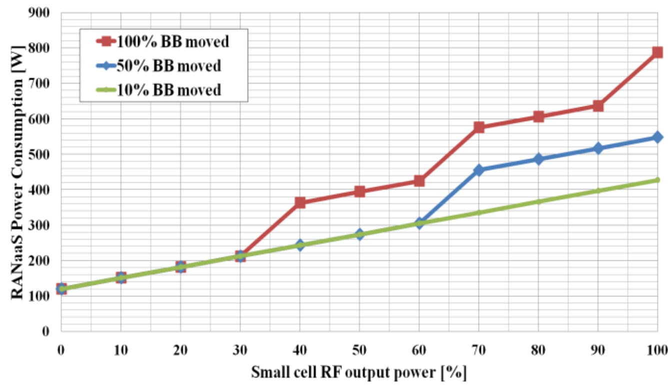

b) RANaaS power consumption with respect to the small cell RF output power for different BB shift options

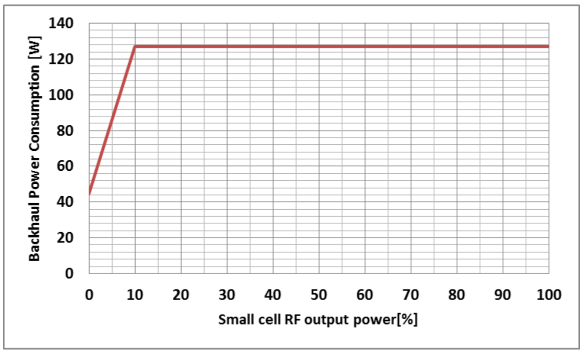

c) Backhaul Power consumption

Figure A-1: Power consumption of different network elements

In the following some detailed parameters related to the developed power models are presented and discussed, respectively for iSCs, RANaaS platform and Backhauling network.

A.1 iSC Power Consumption

The FP7 EARTH has investigated how the power consumption of distinct components of several eNBs, such as power amplifier, baseband engine, main supply, and active cooling, depends on the transmission bandwidth, the transmission power, and the number of radio chains/antennas [25]. Furthermore, it was found that a linear function of the transmission power can approximate very well the generalized model.

To adopt the aforementioned model for approximating iSC power consumption, we have taken into account the functional split. Therefore, its power consumption will be bounded by the two extreme cases: 1) RRH and 2) complete small cell, respectively (see Figure A-1: Power consumption of different network elementsa). RRHs are considered as low complexity nodes that solely perform RF operations and rely on self-backhauling (P_{BB} = 0). On the other hand, complete small cells perform all the based band (BB) operations (P_{BB} =6.8 W). Table A-1 reports the power model and the associated parameters to estimate the power consumption of iSCs [25], considering two (per-antenna) maximum transmit power, i.e., 24dBm ($P_{Tx,1}$) and 30dBm ($P_{Tx,2}$). It is worth mentioning that for the iSC power model:

1. no active cooling is considered,
2. iSCs may enter a low consumption sleep mode where only the power amplifier (PA) is turned off when no data is received or transmitted (BB engine reductions due to sleep mode are not considered here for simplicity), and
3. PA power consumption is approximated as a linear function of the PA output power (for further details see [40]).

Table A-1: Power consumption model for the iSC and exemplary realistic parameter values

iSC	$P_{\text{iSC-}n} = \dfrac{N_{\text{ant}} \dfrac{W}{10\text{ MHz}} \cdot \left(P_{\text{BB}} + P_{RF}\right) + y_n \cdot P_{\text{PA-max}}}{\left(1\text{-}\sigma_{\text{DC}}\right) \cdot \left(1\text{-}\sigma_{\text{MC}}\right)}$		
Bandwidth (W)	10 MHz	PA max consumption ($P_{\text{PA-max}}$)	0.8W if $P_{TX,1}$ 3.2W if $P_{TX,2}$
# antennas per iSC (N_{ant})	2	DC-DC conversion losses (σ_{DC})	6.4 %
BB consumption (P_{BB})	[0 ; 6.8] W	Main Supply losses (σ_{MC})	7.7 %
RF consumption (P_{RF})	0.8W if $P_{TX,1}$ 1.5W if $P_{TX,2}$	Load of cell n (y_n)	0 – 100 %

A.2 RANaaS Platform Power Consumption

To obtain an accurate estimation on the power consumption of the RANaaS platform due to BB processing moved from iSCs, we use of a model from the IT world. Fit4Green has investigated the power consumption for IT resources of datacenters [27]. In particular, results for the various computing style servers are provided using a monitoring tool and a generic power consumption prediction model. Considering the measurement results, it can be observed that a linear model approximate well the server power consumption versus its CPU workload.

Considering the RANaaS as an enclosure hosting several identical ISS Blade servers equally sharing the requested workload, the servers' processing capacity will define how many servers are required to process the system BB-related workload. Therefore, the overall power consumption due to BB processing at the RANaaS platform will be the sum of the power consumption at each of the required servers.

Furthermore, Werthman et al. have recently investigated the relation between the CPU workload and the cell load, and they have defined the resource effort required to serve an UE as a function of the number antennas, the modulation bits, the code rate, the number of spatial MIMO-layers, and the allocated frequency resources in DL [28]. Since in the iJOIN architecture some functionality can be moved towards the RANaaS, we extend this work and introduce an average sum to approximate the total average RANaaS workload required to serve all UEs. Therefore, the Giga-Operations-per-Second[3] (GOPS) required at RANaaS will depend on the number of iSCs, their load, the system bandwidth, the number of antennas per iSC, the average number of data bits per symbol per user, and the average number of MIMO layers (see Table A-2). The RANaaS power consumption with respect to the small cell RF output power for different BB shift is shown in Figure A-1b).

Table A-2: Power consumption model for the RANaaS and exemplary realistic parameter values

RANaaS	$P_{\text{RANaaS}} = \left\lceil \frac{X(y)}{X_{\text{Cap}}} \right\rceil \cdot \left(P_0^{\text{srv}} + \sum_{n=1}^{N_{\text{iSC}}} y_n \cdot \frac{\frac{W}{10\text{ MHz}} \cdot \left(30N_{\text{Tx}} + 10N_{\text{Tx}}^2 + 20\frac{e_{\text{MSC}}}{6} e_{\text{MIMO}}\right)}{X_{\text{Cap}} \left\lceil \frac{X(y)}{X_{\text{Cap}}} \right\rceil} \Delta_p^{\text{srv}} P_{\text{max}}^{\text{srv}} \right)$		
BB processing (in GOPS) moved from iSC into RANaaS	$X = \sum_{n=1}^{N_{\text{iSC}}} y_n \cdot \beta_{\text{BB}} \frac{W}{10\text{ MHz}} \left(30N_{\text{ant}} + 10N_{\text{ant}}^2 + 20\frac{e_{\text{MSC}}}{6} e_{\text{MIMO}}\right)$		
Server Capacity (X_{Cap})	324 GFLOPS	Consumption at Server Max Workload ($P_{\text{max}}^{\text{srv}}$)	215 W
Server idle consumption (P_0^{srv})	120 W	GOPS/Watt cost factor (c_{BB})	160
Linear model slope (Δ_p^{srv})	0.44	# iSCs in veNB (N_{iSC})	5 - 20
Average # of antennas used to serve a UE (N_{Tx})	2	% of BB processing moved into RANaaS from each iSC (β_{BB})	0 – 100 %
Average # of data bits per symbol per UE (e_{MCS})	4/3	Average # of spatial MIMO layers used per UE (e_{MIMO})	1.1

A.3 Backhaul Power Consumption

The last important element that we have modelled is the backhaul network. In general, centralised systems have notable backhaul load; therefore, power consumption due to data transport and switching can become a significant percentage of the total system power consumption [26].

Monti et al. provided some basic power consumption models for data transport through various backhaul technologies and topologies in small cells [29]. Considering iSCs with microwave links and omitting iTNs for simplicity, the backhaul power consumption can be estimated by modifying this model in accordance with the iJOIN architecture; backhaul power consumption shall scale with the power for transmitting and receiving the aggregate backhaul traffic at any iSC, the number of iSCs in the system, the average number of microwave antennas per iSC, and the power consumption of switches at any iSC. Note that the power consumption at any switch will depend on the aggregated traffic at the associated iSC and its maximum capacity. Moreover, the power consumption for transmitting and receiving the aggregate backhaul traffic will generally depend on the traffic conditions. In this work, we consider a two-step function (low/high capacity traffic), where the two capacity regions are distinguished by a single threshold. Our analysis shows that for

[3] It is noted that the processing capacity of the server is expressed in Giga-FLOPS (GFLOPS) **¡Error! No se encuentra el origen de la referencia.**; however, it can be converted in GOPS, and in this work we use a 1:1 ratio as a conservative estimation.

generic small cells, the backhaul always operates in low capacity region, which results in flat power consumption for medium/high cell RF output power (see Figure A-1c)).

The question that arises next is how backhaul traffic load can be translated into cell load in current LTE-based RAN. For this, we need to consider the iSC maximum bits-per-second capacity and the non-negligible overheads from X2 U- and C-plane, the transport protocol, and the IPsec [30]. Accordingly, Table A-3 presents the backhaul power model and the relevant parameters with exemplary realistic values. Note that iSC available capacity is evaluated assuming a single carrier with 10MHz bandwidth, 2x2 MIMO, 64QAM, and 28% control overhead [30]. In addition, we consider that backhaul links can enter in idle mode for energy saving.

Table A-3: Power consumption model for the Backhauling and exemplary realistic parameter values

Backhaul	$P_{\text{Bh}} = \sum_{n=1}^{N_{\text{iSC}}} P^{n}_{\text{switch}}\left(y_n\right) + N^{n}_{\text{mw-ant}} P^{n}_{\text{link}}\left(y_n\right)$		
Switch power consumption	$P^{n}_{\text{switch}} = \begin{cases} 0, & N^{n}_{\text{mw-ant}} = 1 \text{ or } y_n = 0 \\ P_s \cdot \left\lceil \dfrac{y_n \cdot [Y_{\max} \cdot f_{\text{cell-Bh}} \%]}{C_{\text{switch}}} \right\rceil, & \text{otherwise} \end{cases}$		
# microwave antennas per iSC ($N_{\text{mw-ant}}$)	2	Switch maximum capacity (C_{switch})	36 Gbps
Switch basic consumption (P_s)	53 W	Average cell capacity ($Y_{\max}$)	86.4 Mbps
% increase from cell load to backhaul traffic ($f_{\text{cell-Bh}}$)	128 %		
Backhaul link consumption	$P^{n}_{\text{link}} = \begin{cases} P_{\text{idle}}, & y_n = 0 \\ P_{\text{low-traffic}}, & 0 < y_n \le \dfrac{C_{\text{thr}}}{Y_{\max} \cdot f_{\text{cell-Bh}} \%} \\ P_{\text{high-traffic}}, & y_n \ge \dfrac{C_{\text{thr}}}{Y_{\max} \cdot f_{\text{cell-Bh}} \%} \end{cases}$		
Node power region for idle/low/high traffic conditions ($P_{\text{idle/low/high-traffic}}$)	22.2 / 37 / 92.5 W	Traffic threshold between low/high power regions (C_{thr})	500 Mbps

