Document Properties

Document Number:	D 2.3
Document Title:	**Final assessment on MAC enhancements and cooperation techniques**
Document Editor:	Vincenzo Mancuso (IMDEA)
Authors:	Vincenzo Mancuso (IMDEA), Christian Vitale (IMDEA), Arash Asadi (IMDEA), Vincenzo Sciancalepore (IMDEA), Arianna Morelli (INCS), Eva Pierattelli (INCS), Erick Bizouarn (ALBLF), Laurent Roullet (ALBLF), Pablo Serrano (UC3M)
Target Dissemination Level:	PU
Status of the Document:	Submitted to EC
Version:	1.0

Production Properties:

Reviewers:	Sébastien Auroux (UPB) Vincent Kotzsch (SIG)

Document History:

Revision	Date	Issued by	Description
1.0	2015-04-30	IMDEA	First EC submission

Disclaimer:

This document has been produced in the context of the CROWD Project. The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7) under grant agreement n° 318115.

Abstract:

This document reports on the final assessment on MAC enhancements and cooperation techniques designed and studied in the CROWD project and in particular in WP2. The document contains results and algorithms for the enhancement of LTE networks, 802.11 networks and HetNets in extremely dense deployments. A key result of the study presented in this deliverable consists in the assessment of a CPS-based modelling technique to evaluate the capacity region of wireless networks, which can be adapted to analyse either LTE or IEEE 802.11, and all techniques that use LTE and IEEE 802.11, for instance D2D schemes based on such wireless technologies. As concerns LTE networks, the document presents the assessment of innovative user association strategies for energy efficiency, interference control techniques, and D2D-based offloading schemes. As concerns IEEE 802.11 networks, the document reports on new distributed MAC control techniques and D2D ad hoc schemes. As concerns HetNets, the document includes a study on the optimisation of spectrum utilization in heterogeneous LTE/LTE-A networks and the analysis of D2D mode selection for offloading in heterogeneous LTE/LTE-A environments. The key factor behind the success of the proposed MAC enhancements and cooperative schemes is the deployment of a multi-tier and multi-time-scale control architecture like the one proposed in the CROWD project.

Keywords:

802.11, LTE, HetNet, MAC enhancements

Contents

List of Figures

List of Tables

List of Project Partners

Name	Acronym	Country
Intecs S.p.A. (*coordinator*)	INCS	Italy
Alcatel-Lucent Bell Labs France	ALBLF	France
Avea Iletism Hizmetleri A.S.	AVEA	Turkey
Fundacion IMDEA Networks	IMDEA	Spain
National Instruments	SIG	Germany
Universidad Carlos III de Madrid	UC3M	Spain
Universitaet Paderborn	UPB	Germany

List of Acronyms

3GPP Third Generation Partnership Project

ABSF Almost Blank Sub-Frame

ACM Association for Computing Machinery

ADOS Adaptive algorithm for Distributed Opportunistic Scheduling

AIMD Additive Increase Multiplicative Decrease

AP Access Point

ARP Address Resolution Protocol

BER Bit Error Rate

CA Collision Avoidance

CIO Cell Individual Offset

CLC CROWD Local Controller

CPS Coupled Processor System

C-RAN Cloud Radio Access Network

CRE Cell Range Extension

CQI Channel Quality Indication

CROWD Connectivity management for eneRgy Optimised Wireless Dense networks

CSI Channel State Indication

CSMA Carrier Sense Multiple Access

CSMA/CA Carrier Sense Multiple Access/Collision Avoidance

D2D Device-to-Device

eICIC Enhanced Inter-Cell Interference Coordination

eNB Evolved NodeB

GPS Generalized Processor System

H_2(IC)$_2$ Hybrid ICIC

HetNet Heterogeneous Network

ICI Inter-Cell Interference

ICIC Inter-Cell Interference Coordination

IEEE Institute of Electrical and Electronics Engineers

ILP Integer Linear Programming

IP Internet Protocol

JFI Jain's fairness index

KPI Key Performance Indicator

LTE Long Term Evolution

LTE-A Long Term Evolution Advanced

MAC Medium Access Control

MCS Modulation and Coding Scheme

NC Network Calculus

NDOS Non-cooperative algorithm for Distributed Opportunistic Scheduling

NoA Notice of Absence

OFDMA Orthogonal Frequency Division Multiple Access

PF Proportional Fairness

PHY Physical

PRB Physical Resource Block

PSNR Peak Signal to Noise Ratio

QoS Quality of Service

RB Resource Block

REFIM REFerence based Interference Management

RF Radio Frequency

ROI Region of Interest

RSRP Reference Symbols Received Power

SC-FDMA Single Carrier-FDMA

SDN Software Defined Network

SINR Signal to Interference plus Noise Ratio

SIR Signal to Interference Ratio

SNR Signal to Noise Ratio

SOLOR Self-Optimizing wireless local area networks with Legacy-compatible Opportunistic Relays

SON Self Optimising Network

SSID Service Set ID

TDOS Team-game algorithm for Distributed Opportunistic Scheduling

TTI Transmission Time Interval

UBPC Utility-Based Power Control

UE User Equipment

V-RAN Virtualised Radio Access Network

WiFi Wireless Fidelity

WLAN Wireless Local Area Network

Executive summary

The main purpose of this document is to present the final assessment of techniques studied and designed in the frame of WP2 of the CROWD project. Such techniques propose enhancements for MAC protocols and cooperation mechanisms in LTE/LTE-A, IEEE 802.11-based WLANs, and HetNets consisting of devices using LTE and IEEE 802.11 technologies. Such specifications aim at the improvement of LTE and IEEE802.11 standard features to augment the capabilities of standard MAC protocols in extremely dense and heterogeneous wireless networks. In our proposals, all such networks operate under the control of CROWD controllers, according to the CROWD architecture designed in the project and described in Deliverables D1.1 [1] and D1.3 [2], which adopts an SDN-based approach. Such choice allows to implement multi-tier and multi-time-scale control functions, and here we focus on mechanisms that benefit from the control operated on a short time-scale and with a local scope, i.e., in a *district* of interfering or geographically close devices with overlapping coverage of multiple cells and WLANs. The specifications for the enhancements described in this document have been reported in Deliverables D2.1 [3] and D2.2 [4], while here we report on the final assessment of such techniques and on the further refinements of the techniques previously designed in WP2. Such refinements have been based on the feedback from other work packages, and in particular from early implementation and integration experiments carried out in WP5.

The document shows that the performance of LTE/LTE-A networks can be substantially improved by means of controlling user association, intercell interference, and inband relay schemes (using LTE D2D). Similarly, WLANs can leverage control theory-based optimisation of MAC parameters, cooperative relaying schemes—including D2D-base relay—and cooperative and distributed channel access schemes to dramatically improve spectral efficiency. The document also shows how to optimise resource and user allocation in cellular networks with cells of heterogeneous size, by jointly leveraging base station cooperation at both resource allocation and user association levels. Last but not least, our results on network cooperation and relay prove to be suitable for the optimisation of HetNets. However, to be effective, the control functions for LTE, IEEE 802.11 and HetNets have to operate at different geographical scales and at different time scales. In particular, this document focuses on mechanisms that work at a local time scale (i.e., in a CROWD district [2] and on packet/frame time scale, to cope with interference, shared resource access, distributed and centralized access control, cooperative resource utilization, opportunistic packet relay and energy efficiency.

The scientific achievements of WP2 described in this document include design and assessment of mechanisms for the enhancement of LTE networks, IEEE 802.11-based networks, and HetNets consisting in LTE cells of various sizes, legacy IEEE 802.11 devices, WiFi Direct links and D2D relay links using licensed and unlicensed frequencies with inband and outband resources. Specifically, the key mechanisms assessed in this document are: (*i*) energy efficient user association strategies for cellular networks, based on the minimisation of the number of active base stations in LTE/LTE-A networks; (*ii*) distributed and semi-distributed algorithms for deciding ABSF patterns based on interference conditions; (*iii*) innovative analytical methods to predict network capacity and performance figures for wireless networks (LTE, WLANs and D2D) based on CPS modelling methods using network calculus and stochastic queueing theory results; (*iv*) control theory-based mechanisms for the optimisation of distributed channel access in WLANs; (*v*) novel IEEE 802.11 relay techniques using Wireless Fidelity (WiFi) Direct D2D communications; (*vi*) algorithms for time-frequency resource allocation and user allocation to optimise eICIC in LTE/LTE-A networks

with cells of heterogeneous sizes. (*vii*) hybrid LTE/802.11 networking cooperation schemes with opportunistic multi-mode relay protocols for HetNets.

Key contributions

The main technical contributions of this deliverable are as follows:

- The assessment of user association strategies for saving energy without wasting network capacity in very dense networks, especially for the case of femtocells.
- The assessment of novel game theory-based semi-distributed mechanisms for Enhanced Inter-Cell Interference Coordination (eICIC) in LTE/LTE-A, using the ABSF mechanism.
- The assessment of D2D inband techniques for LTE/LTE-A networks, based on the analysis of the capacity region for cellular and D2D users.
- The assessment of ADOS, a distributed channel access scheme in IEEE 802.11 networks running under the guidance of control theory-based distributed schemes.
- The assessment of SOLOR, an inband D2D relaying scheme for IEEE 802.11 networks using WiFi Direct.
- The assessment of optimisation of resource allocations in IEEE 802.11 networks, based on the characterisation of stability region and performance figures of D2D inband schemes achieved by means of CPSs models.
- The assessment of techniques to optimise eICIC parameters, namely for CIO and ABSF to achieve proportion fairness under strict power constrains in LTE/LTE-A networks with macro and small cells.
- The assessment of D2D schemes and D2D mode selection mechanisms when heterogeneous technologies are available in HetNets.

The key contributions to the research and standardization communities resulting from the work performed within the aforementioned activities are the following:

- A journal article on experimental 802.11 research works has been published in IEEE Communications Surveys and Tutorials [5].
- A journal article on optimised D2D relay modes in HetNets has been published in ACM SIGMETRICS Performance Evaluation Review [6].
- A journal article focusing on the control theoretical design of an adaptive distributed opportunistic scheduling mechanism for IEEE 802.11-based WLANs has been published in IEEE/ACM Transactions on Wireless Communications [7].
- A conference paper focusing on game theory for semi-distributed control of ABSF has been published in the proceedings of IEEE SECON 2015 [8].
- A conference paper focusing on CPS techniques for analysing and optimising inband D2D transmissions in LTE/LTE-A networks has been published in the proceedings of IEEE COMSNETS 2015 [9].

- A conference paper focusing on CPS techniques for analysing and optimising inband D2D transmissions in 802.11/WiFi networks has been published in the proceedings of IEEE ICC 2015 [10].
- A conference paper focusing on optimisation of D2D mode selection in HetNets has been published in the proceedings of IEEE WoWMoM 2015 [11].

1 Introduction

As discussed in Deliverables D2.1 [3] and D2.2 [4], wireless and heterogeneous dense deployments exacerbate the issues related to radio interference and medium access control. Indeed, network performance suffers due to interference in terms of degradation of spectral efficiency, while the emergence of mobile Internet has burdened cellular networks with new types of traffic which are definitely more resource consuming than traditional voice services.

Existing solutions such as ABSF and similar mechanisms for eICIC, the introduction of cells of different sizes and technologies for improving coverage and bandwidth density, or the adoption of D2D schemes with inband and outband characteristis, need to be controlled to achieve substantial network performance gains in terms of Key Performance Indicators (KPIs) like throughput, fairness and energy efficiency. CROWD offers such a control architecture, which is able to decouple local and regional control and to differentiate between different time scales for the needed control actions, as described in Deliverables D1.1 [1] and D1.3 [2]. Therefore we adopt the CROWD architecture and study controllable mechanisms. Specifically, in this document we focus on short time scale control mechanisms and local controlling functions, i.e., affecting the behavior of nearby and interfering network devices.

In this document we assess those enhanced Medium Access Control (MAC) and cooperative mechanisms proposed in CROWD to counteract the effects of interference and congestion in dense wireless networks, which includes (*i*) the optimisation of radio resource allocation schemes, (*ii*) the quest for energy efficiency, for sustainability reasons, and (*ii*) the possibility to allow for *network controlled* collaborative schemes using, e.g., the D2D paradigm or game theory based approaches to enhance spectral efficiency and offload the cellular network infrastructure. Specifically, the document contains results and algorithms for the enhancement of LTE networks, IEEE 802.11 networks and HetNets in extremely dense deployments. As concerns LTE networks, the document presents the assessment of innovative interference control techniques and D2D-based offloading schemes. As concerns IEEE 802.11 and Wireless Local Area Network (WLAN) networks in general, the document reports on new distributed techniques and D2D ad hoc schemes. As concerns HetNets, the document includes a study on the optimisation of spectrum utilization in heterogeneous LTE/LTE-A networks and the analysis of D2D mode selection for offloading in heterogeneous LTE/LTE-A environments. With respect to the characterisation and optimisation of wireless networks, we also propose a novel technique based on the CPS modelling paradigm, for which we show how to achieve computationally acceptable approximations in LTE, WLANs, and HetNets.

1.1 Document structure

The remainder of this document is structured as follows:

- In Chapter 2 we present the results of our research on MAC enhancement and cooperative mechanisms. Specifically, Section 2.1 presents the performance assessment of mechanisms for cellular networks like LTE and LTE-A, Section 2.2 presents the assessment for cooperative and control-based mechanisms for IEEE 802.11 networks, and Section 2.3 focuses on the assessment of mechanisms proposed in CROWD for the control of interference and D2D operations in HetNets.

- Concluding remarks are presented in Chapter 3.
- At the bottom of the document, we report in the appendix many additional details on the techniques, the analysis, and the experiments pointed out in Chapter 2.

2 Research activities

2.1 LTE enhancements

We start with the description of the assessment of new mechanisms to boost the efficiency of cellular, LTE-like networks, in presence of very dense deployments. Specifically we showcase the potentials of (*i*) energy efficiency yet throughput efficient user association policies in cellular networks, (*ii*) controlled but distributed eICIC schemes, and (*iii*) D2D relay mechanisms aware of network capacity and interferences generated to and from D2D transmission using inband resources (i.e., sharing cellular resources with LTE users).

2.1.1 User association strategies

Dense heterogeneous networks constitute the paradigm for the future wireless networks. For this reason operators and standardisation bodies are particularly eager to solve density issues. Here we focus on extremely dense networks of small cells, e.g., femtocells, that could be found in crowded public places, offices or alike environments. In such deployments, energy consumption must be constrained to have sustainable operational costs. To achieve this goal, we have introduced new mechanisms in Connectivity management for eneRgy Optimised Wireless Dense networks (CROWD) for the association of the users to base stations, aiming at minimising the energy consumption of the LTE access network. In this document we assess the trade-off between capacity and energy consumption that can be achieved thanks to the CROWD control architecture and a suitable load balancing of users among eNBs. Specifically, in CROWD, we have studied new user association algorithms for dense LTE networks. Here we assess the performance of the algorithm proposed in our previous deliverable, and in particular in Deliverable D2.2 [4].

2.1.1.1 System model

We consider a dense network of femtocells, with U UEs and B eNBs physically connected to a hierarchy of switches with limited capacity. As in [12], we assume that UEs and eNBs are distributed over a finite flat surface using the MT19937 algorithm for random number generation [13]. An example network is illustrated in the reference diagram depicted in Figure 2.1.

In order to keep the complexity of the presented dense access network low, we propose to set the transmission power P_{TX} to the same value for all the eNBs. Specifically, P_{TX} is a constant value set to 15 dbm. Regarding the power received from the UEs, the following result holds:

$$P_{RX} = P_{TX} - (L(k,i) + L_{NS}(k,i)), \tag{2.1}$$

where $L(k,i)$ and $L_{NS}(k,i)$ are respectively the path loss for femto-cells, computed as in [14]—i.e., $L(k,i) = 37 + 30\ \log(d(k,i))$, where $d(k,i)$ is the distance UE i - eNB k—and the log-normal shadowing expressed in decibel units. We call *association* the process by which every UE is assigned to at most one eNB, based on the static positions of all UEs and eNBs. When the association phase ends, we assume that the maximum Modulation and Coding Scheme (MCS) supported, based on its Signal to Interference plus Noise Ratio (SINR), is assigned to each UE i associated to eNB j.

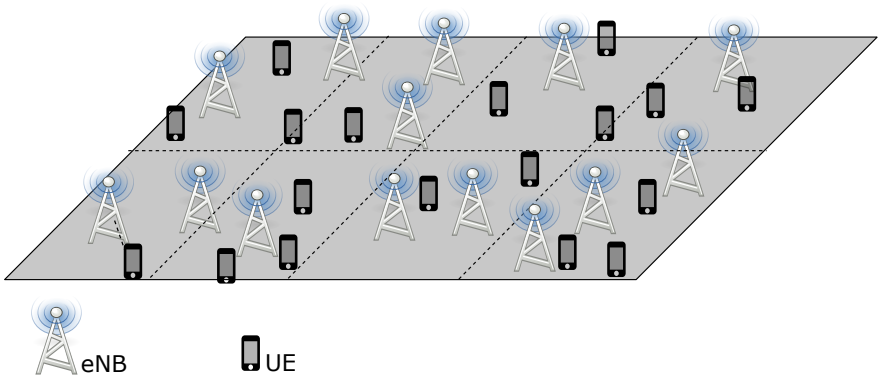

Figure 2.1: Example of LTE dense network, with UEs and eNBs.

We compute such SINR ($SINR_i(j)$) as follows:

$$SINR_i(j) = \frac{P_{RX}(j,i)}{\mathcal{N} + \sum_{k=1,k\neq j,k\notin\mathcal{I}}^{B} P_{RX}(k,i)}, \tag{2.2}$$

where $P_{RX}(j,i)$ is the power received by UE i from eNB k, considering both propagation and fading, while $\mathcal{I}$ is the set of eNB to which no UE is assigned to. Finally, we assume that the throughput of UE i can be computed as follows. Once all the associations are performed and the actual SINR_i is computed, (2.2), the efficiency of user i, η_i, is obtained through:

$$\eta_i = \log_2\left(1 + \frac{\text{SINR}_i}{\Gamma}\right) \tag{2.3}$$

where Γ is a parameter that depends on the Bit Error Rate (BER) target. Assuming that UE i is the only user associated to eNB j, the throughput of i can be obtained from Table I in [15], based on the related η_i. When instead UE i shares the access through eNB j with other UEs, we assume it receives a portion of the total capacity that is proportional to its maximum throughput (or its η). For instance, if $\eta_1/\eta_2 = 2$ then the physical resources assigned to UE1 will be twice as much as those assigned to UE2.

When computing the power consumption of the elements of the LTE access network, we assume that each eNB with no associated UE enters an idle mode from which it can be easily brought back to full functionality [16]. Then, according to the model in [17], the power consumption of each eNB $P_{in_{eNB_j}}$ is assumed to be:

$$P_{in_{eNB}} = P_{BB} + P_{RF} + P_{PA} + P_{OH}, \tag{2.4}$$

where P_{BB} is the power spent in base band processes, P_{RF} is the maximum Radio Frequency (RF) output power per transmit antenna (i.e., 1 W), P_{PA} is power spent by the pre-amplifier (650 mW to 8 W) and P_{OH} is the power overhead (power systems and cooling, accounting for ~ 3.72 W).

Now we introduce the association algorithm that tries to maximise the energy saving of the LTE access network, namely the *Min-Energy* association algorithm. Basically, the performed algorithm aims to minimise the number of active eNBs, using the following assumptions: (*i*) an eNB can serve up to four UEs, (*ii*) one or more eNBs are available for each UE, (*iii*) each UE is attached to any of the eNBs still *available*, the one which guarantees the best channel condition. The ultimate goal of this approach is to keep active as few eNBs as possible, while keeping their utilization to the highest level in terms of the number of served UEs.

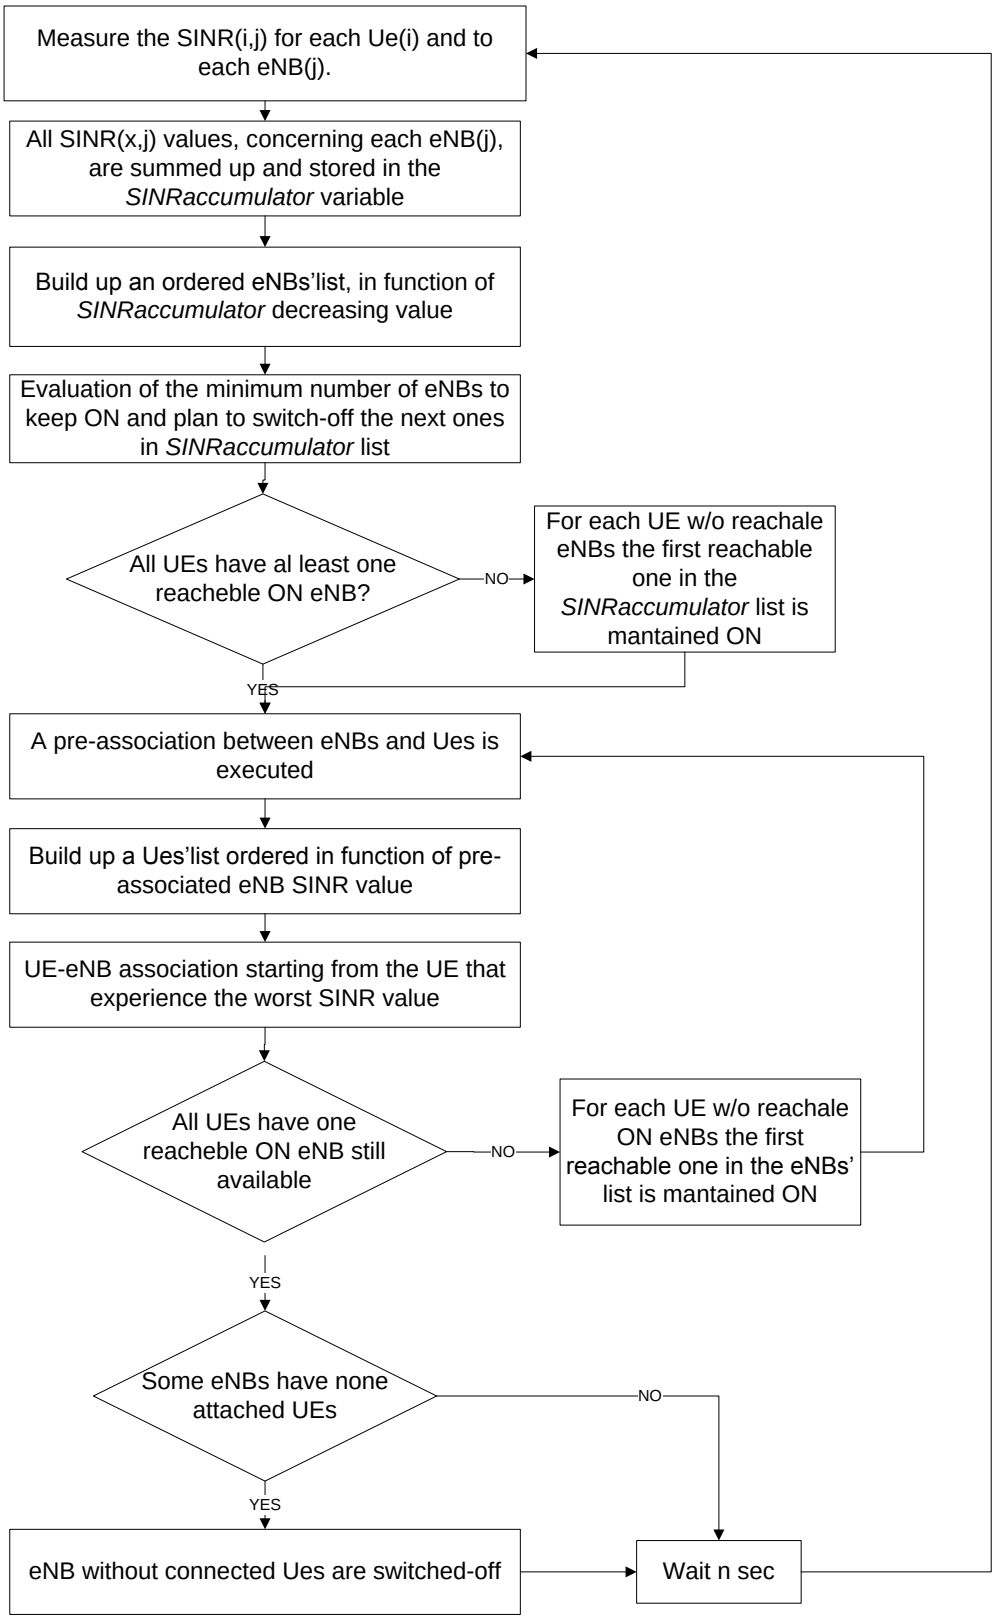

Figure 2.2: Min-Energy Algorithm.

A detailed description of the proposed *Min-Energy* algorithm is given in Figure 2.2. It consists of 3 main phases contributing to switch-off the maximum number of eNBs, taking into account the Signal to Noise Ratio (SNR) and channel interference to preserve network Quality of Service (QoS).

In the proposed algorithm, During the starting phase (*1*), we measure, for each UE, the SINR to all eNBs. Such values are summed up and stored in a variable called *SINRaccumulator* concerning each eNB; at the end of this procedure the list of eNBs is sorted in decreasing *SINRaccumulator* order.

In the second phase (*2*) we *pre-select* the minimum number of eNBs to keep active, according to the number of UEs in the network (starting with the ones with greater *SINRaccumulator*). A maximum number of 4 UEs could be attached to each eNB, while the remaining eNBs are scheduled for switch-off. In case one UE could be attached only to one eNB that should be switched-off, we decide to keep that eNB active.

Each UE owns a list of eNBs ordered in function of SINR value and the best one is chosen for *pre-association*. Now we have a list of UEs ordered in function of the pre-associated eNBs SINR value. Starting from the one that experiences the worst SINR value, all UEs are attached to the network. In this last phase (*3*) if an UE has not any more available eNBs among the ones maintained active, then the best switched-off eNB, according to the *SINRaccumulator* indicator, is switched-on again. After this phase, all the pre-associations are reevaluated again, i.e., the algorithm is repeated starting from phase (*2*). At the end of this procedure, it is possible to have eNBs without associated UEs, and such eNBs are actually switched-off.

2.1.1.2 Performance assessment of the user association algorithm

Now we analyse the performance of the proposed association algorithm in several network conditions. Furthermore, we benchmark the throughput and power consumption achieved by means of the *Min-Energy* association policy against a simple algorithm that associates each UE i to the eNB j whose received the best SNR. We name this association policy *Best-SNR*. In order to have a fair comparison, both the P_{TX} setting and the power consumption computation are performed as for the the *Min-Energy* association.

In the following, the numerical analysis assumptions are reported first. Then we discuss the obtained results. The eNBs are randomly distributed in a 400m×400m area. The UEs are uniformly randomly distributed and move according to a random walk mobility model with random velocity uniformly distributed in the interval [0.6 m/s, 1.5 m/s]. They are attached to the eNBs according to the specific algorithm. Once all the UEs are attached, the throughput is calculated as previous described executing a single run, 60 minutes long. Throughput and power consumption results repaired in the following figures are average values compute over the simulation time.

Basically the tests are focused on the analysis of capacity in terms of UEs average throughput (in Mb/s) and on the eNBs network energy consumption (in KWatt/h). The scope of the analysis is to highlight the gains introduced by a wise association of the UEs to the eNBs, while facing a limited reduction of the throughput if compared to traditional association techniques. For completeness, the simulation parameters for the scenario are reported in Table 2.1.

Once we have fixed the number of eNBs, we vary the number of UEs from a chosen minimum value up to the maximum number of UEs allowed, which is based on the total amount of eNBs in the scenario and on the maximum number of allowed UEs per eNB (L_{TOT}).

In Figure 2.3 the average throughput achieved with the *Best-SNR* association algorithm and the throughput obtained using the *Min-Energy* association are compared. Against expectation, when the *Min-Energy* algorithm is used, there is a gain of throughput and this gain is higher when the UEs are half of those that could be served. In this conditions, the *Min-Energy* associates as many UEs as possible to the same set of eNBs, while the *Best-SNR* does not concentrate the UEs in the same geographical area. Therefore, the reason of the difference in throughput achieved is

Table 2.1: User association assessment: Simulation parameters and values

$P_{TX_{eNB}}$	15 [dbm]
N_{TX}	1
P_0	4.8 [W]
Δp	8.0
P_{idle}	2.9 [W]
L_{NS}	$\sigma_S = 8dB$
$\mathcal{N}$, $\sigma^2 = K \cdot T \cdot B$	$\sigma^2 = -101dBm$
BER	10^{-5}
Gamma	$-\log_n(5 * BER)/1.5$
$SINR$	$(2^{eta} - 1) * Gamma$
Area	400m×400m
Max UEs per eNB	4
eNBs number	15
UEs number	5-50

mainly due to the different interference condition they experience. When the number of switched on eNBs increases, the effect of interference increases too. As a consequence, when the number of UEs reaches the maximum number that the network could serve, almost all the eNBs are turned on and the throughput values offered by the two compared approaches converge. Similarly, with few UEs the throughput is similar due to the low network load.

We can conclude that the *per-user* perception of the service received is positively affected by the *Min-Energy* association policy in medium load conditions, while the compared algorithms perform equally well in extreme operational conditions.

In Figure 2.4 we finally show the load dependent component of the energy consumption for both the association techniques analysed. Such choice allows us to focus on the ability of the *Min-Energy* association technique to handle the real load of the radio network. However, we leave out of the scope of the document the discussion on how to reliably turn off idle elements of the network. There is a considerable gain in terms of energy saving when using the *Min-Energy* algorithm. The maximum energy saving achieved is close to the 55% when the number of UEs and eNBs is small, whereas it obviously decreases when the access network approaches the saturation.

2.1.1.3 Conclusions

In this report we have defined a model for LTE dense networks of small cells (e.g., femtocells) that includes capacity constraints due to a limited LTE network capacity. We proposed an algorithm that uses a users association policy taking into account the minimisation of the energy consumption of the network elements in the association procedure of the users to the base stations. We analysed the obtained results and compared them with the result obtained considering the traditional user association policy used in cellular networks: thus, we have compared two very different user association policies. Results have shown that there is a huge potential of energy saving that can be enabled by smartly allocate mobiles to base stations. In practice this can be achieved by means of Self Optimising Network (SON) procedures continuously optimising network configuration based on change of conditions, which in dense networks can be hardly (if at all) predicted and optimised off-line.

These results can be improved with additional study. Currently we are working on the actualisation of the users mobility model for the used scenario, which also determines the scheduling of LTE eNBs exploiting the ABSF paradigm (see next subsection for details on ABSF and its possible

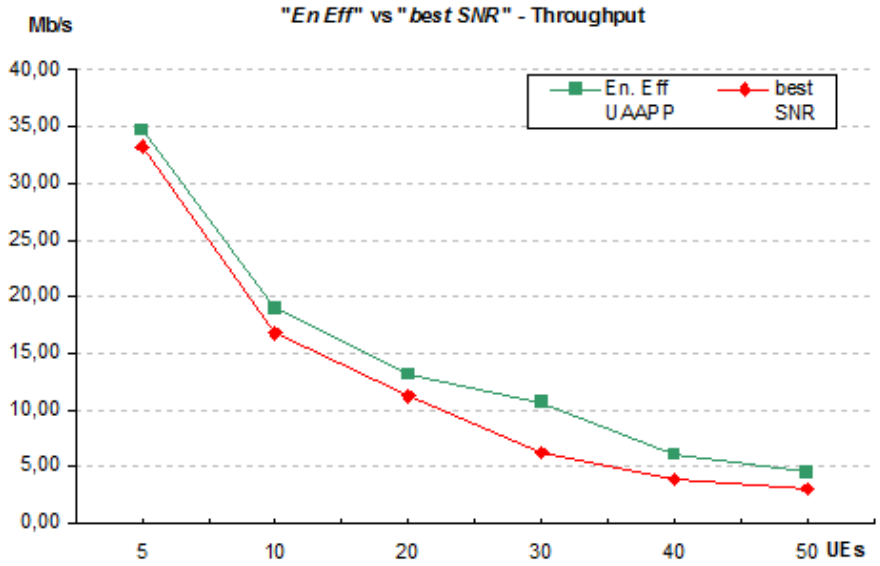

Figure 2.3: UEs average throughput, eNBs = 15.

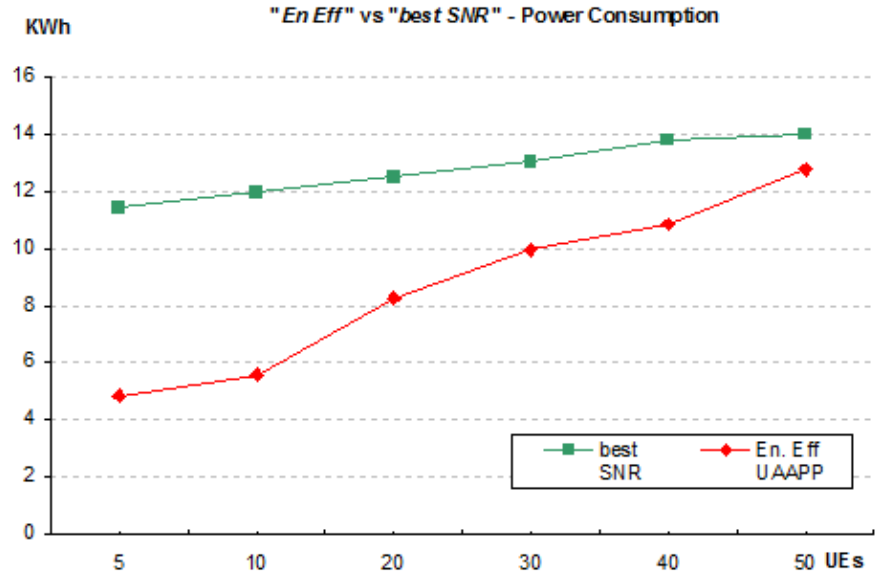

Figure 2.4: eNB power consumption, eNBs = 15.

utilization in dense cellular networks).

2.1.2 Distributed control of ABSF/eICIC

The rush towards network densification with wideband systems takes the intra-system interference problem to very critical levels. Given the high number of base stations and mobile nodes involved, to be practically implementable in a dense network, these schemes need to have low complexity and overhead. 3GPP has proposed the ABSF scheme to implement efficient inter-cell interference coordination. ABSF is used to prevent the base station transmissions when the interference is high, leading to a more efficient use of the radio channel. This paradigm is suitable to control inter-cell interference in different scenarios, including when macro cells and small cells co-exist in the same area. Indeed, ABSF allows guaranteeing a minimal quality of service within each cell. Moreover, ABSF can also manage the inter-cell interference amongst several macro-cells to boost network performance by increasing spectral efficiency. However, the choice of ABSF patterns is not a 3GPP directive, and therefore an open challenge to implement inter-cell interference coordination is the design of efficient algorithms to compute optimal ABSF patterns, able to exploit network and traffic characteristics within the constraints of system architecture and protocols.

2.1.2.1 Problem formulation and system model

To reduce the complexity of the algorithms devised to properly issue ABSF patterns while offloading this burden from the central controller, a suitable alternative consists in letting the base stations self-organise their activity patterns in a distributed fashion. Such a distributed approach shows an important advantage: since each base station has to continuously operate user scheduling, it can take ABSF and user scheduling decisions jointly, which allows for further improving performance. One critical aspect in the design of this scheme is to limit the amount of information exchanged between base stations, which can be achieved by exchanging only the activity patterns of neighbouring base stations.

A general scheme is envisaged by Figure 2.5, where each base station locally takes decision on its own activity pattern (namely, ABSF_i pattern). Then, the activity pattern is sent to the next base station, which, in turn, given the activity patterns of the other base stations, makes its own decision. We mathematically formulate the problem of joint scheduling and ABSF pattern selection and design a solution to solve this problem. We analyse the above approach as a game played by the base stations to optimise the aforementioned problem in a distributed way. We consider each base station as a player whose moves consist in selecting ABSF patterns and announcing them to the neighbours. We cast the problem into a particular class of congestion games, called weighted player-specific matroid congestion game. By means of game theory notions we prove that this particular class of congestion game may not admit a Nash equilibrium. Although uncommon network deployments brings the system to unstable state, some time-constraints are introduced to limit the stability issue impact and leads the system to an efficient working point. We call the resulting algorithm $\text{H}_2(\text{IC})_2$. The details on the algorithm can be found in Appendix A.1.2.

Problem formulation principles. The goal of ICIC is to improve system spectral efficiency. To this end, ICIC optimally orchestrates base station activities and performs user scheduling on a time-slot basis, i.e., per Transmission Time Interval (TTI). Here, we cast the ICIC problem into an LTE-Advanced network that implements the ABSF mechanism. With this mechanism, each base station uses an ABSF pattern, which is a bitmap that specifies which TTIs must be blanked by the base station.

In what follows, we first formulate the ICIC problem from a *centralised* scheduling perspective, which is practically unfeasible due to computational and signaling overhead. Then we show how to abate and distribute the computational load of the ICIC problem over the base stations. However, as we show in the final part of this section, introducing such a *fully distributed* approach requires some game theory tools, and does not always guarantee that base stations' decisions *converge*. To

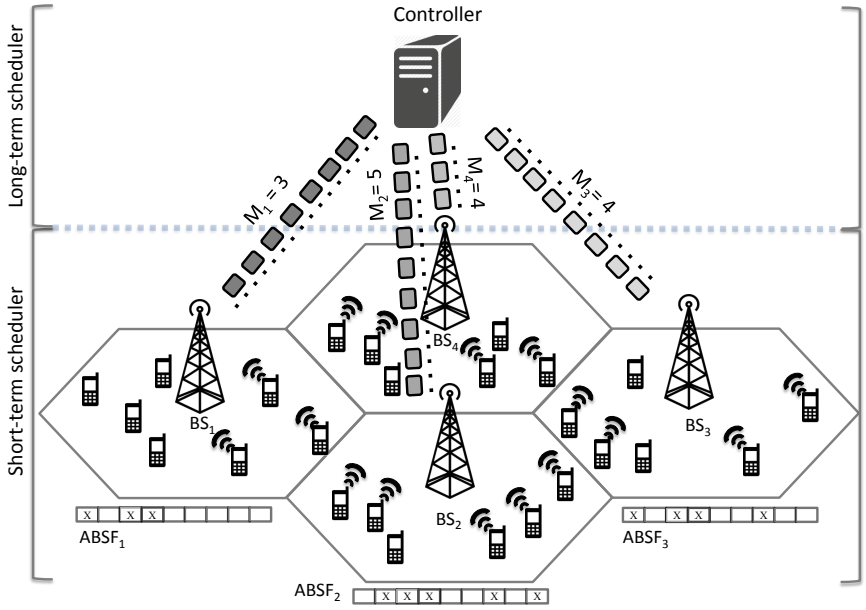

Figure 2.5: $H_2(IC)_2$: Hybrid two-level mechanism for intercell interference coordination. In the short-term level (bottom side of the figure), the distributed approach is played amongst the base stations, while in the long-term level (top side) the controller decides the number of available TTIs per base station.

solve these issues, in Appendix A.1.2 we propose our two-tier semi-distributed mechanism.

For the sake of simplicity, problem formulations presented in this section consider downlink traffic only; however, very similar techniques could be used for uplink traffic. Additionally, we focus on elastic traffic (i.e., traffic for which there are no stringent requirements in terms of latency and bandwidth) since it represents the most common traffic type in mobile data networks. Following current cellular deployments, we consider that base stations transmit at fixed power; therefore, it is sufficient to know which base stations are active to determine the level of interference suffered by a transmission.

A centralised problem. The main assumption behind the centralised problem is that users' Channel State Indication (CSI) is perfectly known. Such information is gathered and updated by a centralised controller—the CROWD Local Controller (CLC) defined in [1]—which uses it to compute the optimal scheduling. Specifically, the centralised controller maps each user u onto any available TTI t and issues the resulting user scheduling information to every base station i. In this way, transmissions during each frame are entirely controlled by the centralised controller. The time horizon of the optimisation consists in a set of TTIs $\mathcal{T} = 1..T$, in which base stations' activities are coordinated. While this scheme is clearly unpractical, the optimal solution to this problem provides us with the benchmark corresponding to the best possible performance of any implementable algorithm.

The objective function to be maximised by the centralised problem, $\widehat{\eta}$, is the sum of the utilities of the individual base stations. Following the widely accepted *max-min* fairness criterion, we define the utility of base station i as the minimum rate of all the users in the base station.[1]

[1]Note that the selected objective function provides a trade-off between maximising the spectral efficiency and guaranteeing a minimum level of service quality, as pointed out, e.g., in [18]. Nevertheless, different objective functions can be considered as well, without substantially changing the proposed approach and the following analysis.

With this, we can formulate the centralised optimisation problem with the following Integer Linear Programming (ILP) model.

Problem `CENTRAL`:

$$\begin{aligned} \text{maximise } & \widehat{\eta} = \sum_{i \in \mathcal{N}} \left(\min_{(u,t) \in \mathcal{U}_i \times \mathcal{T}} R^r \cdot x_u^{r,t} \right), \\ \text{s.t. } & \sum_{u \in \mathcal{U}_i, r \in \mathcal{R}} x_u^{r,t} \leq y_{i,t}, \quad \forall i \in \mathcal{N}, t \in \mathcal{T}, \\ & \frac{P\, G_{u,i}}{N_0 + \sum_{k \in \mathcal{N}: k \neq i} P\, G_{u,k} \cdot y_{k,t}} \geq \gamma^r \cdot x_u^{r,t}, \\ & \qquad \forall i \in \mathcal{N}, u \in \mathcal{U}_i, r \in \mathcal{R}, t \in \mathcal{T}, \\ & y_{i,t}, x_u^{r,t} \in \{0;1\}, \ \forall i \in \mathcal{N}, u \in \mathcal{U}_i, r \in \mathcal{R}, t \in \mathcal{T}; \end{aligned} \tag{2.5}$$

where $\mathcal{N}$ is the set of base stations, $\mathcal{U}_i$ is the set of users associated to base station i, $G_{u,i}$ is the channel gain between user u and the base station i, P is the transmitting power of any base station and N_0 is the background noise. Binary decision variables $x_u^{r,t}$ indicate whether user u is scheduled or not into TTI t, and binary variables $y_{i,t}$ take value 1 if base station i is active at TTI t, and 0 if the TTI is blanked. When scheduled, u's transmission rate is r (the available rates each user can adopt are listed in the vector R^r corresponding to the Modulation and Coding Schemes (MCSs) of the 3GPP standard [19], where set $\mathcal{R}$ is the index set of vector R^r).

The first set of constraints of Problem `CENTRAL` impose that at most one user per base station is accommodated in a single TTI t (note that this assumption can be easily relaxed in order to address different user schedulers). The second set of constraints ensures that, when setting a rate r for a user in a TTI, the perceived SINR is not lower than the activation threshold γ^r for this rate. Note that, while these constraints are not linear, they can be straightforwardly linearized in order to solve the problem with state-of-the-art solvers.

Problem `CENTRAL` can be reduced to a bin-packing problem in which the sum of interferences cannot exceed a threshold. Therefore, this problem is NP-hard [20]. Moreover, it involves a very high overhead to deliver CSI information to the centralised controller, which needs this information to select the ABSF patterns and compute the user scheduling. Thus, while the centralised approach can be an attractive option for small networks, a less complex and more distributed approach is required to deal with the case of very dense wireless networks consisting of hundreds of base stations and thousands of wireless nodes.

A distributed problem. We next present a distributed formulation of Problem `CENTRAL`, whose implementation distributes the computational burden of the original problem over the base stations present in the network. Specifically, to reduce complexity, in the distributed problem each base station only optimises the scheduling of its own users and considers that other base stations use fixed ABSF patterns. However, this approach needs an iterative mechanism to find the optimal ABSF pattern of all base stations. Note that, with the distributed approach, *the complexity of the problem to solve is dramatically reduced, while the number of iterations required to converge will be shown to grow at most quadratically with the network size.*

To formulate the distributed approach, the original problem is split into several smaller instances, which are solved locally by each base station. To solve a problem instance, the base station is provided with the activity pattern declared by other base stations. This is given by *ABSF patterns*, $ABSF_{i,t}$, which are exchanged among base stations ($ABSF_{i,t} = 0$ if base station i blanks TTI t). Such information is needed by each base station to estimate the interference in each TTI suffered by any possible candidate scheduled user. With the above information, and without explicitly forcing any additional constraint, each base station i would schedule users selfishly in the entire set of T TTIs, in order to optimise the local utility. Therefore, to avoid that base stations use all available TTIs, in the distributed problem formulation, we grant a single base station i access to up to M_i

TTIs over T available TTIs; such M_i value plays a key role in the distributed mechanism, as it will be clarified in Appendix A.1.2.

The above description corresponds to the following instance of the local problem for base station i, which can be formulated as an ILP model as follows:

Problem `LOCAL`:

$$\begin{aligned}
\text{maximise} \quad & \widehat{\eta}_i = \min_{(u,t)\in\mathcal{U}_i\times\mathcal{T}} R^r \cdot x_u^{r,t} \cdot a_{u,i}, \\
\text{s.t.} \quad & \sum_{u,r} x_u^{r,t} \cdot a_{u,i} \leq 1, \qquad \forall t \in \mathcal{T}, \\
& \frac{PG_{u,i}}{N_0 + \sum_{k\in\mathcal{N}:k\neq i} PG_{u,k}\cdot ABSF_{k,t}} \geq \gamma^r \cdot x_u^{r,t}, \\
& \qquad\qquad \forall u \in \mathcal{U}_i, r \in \mathcal{R}, t \in \mathcal{T}, \\
& \sum_{u,t,r} x_u^{r,t} \leq M_i, \\
& x_u^{r,t} \in \{0;1\}, \quad \forall u \in \mathcal{U}_i, r \in \mathcal{R}, t \in \mathcal{T};
\end{aligned} \tag{2.6}$$

where all parameters and constraints have the same meaning as in Problem `CENTRAL`, except for the third constraint, which limits the number of usable TTIs to M_i. Note that a feasible solution of Problem `LOCAL` can be computed by using any available max-min scheduling heuristic.

As stated above, each base station i is in charge of solving Problem `LOCAL`, by computing the optimal user scheduling into available TTIs. Note that the solution of this problem depends on the solutions computed by the other base stations, since the SINR of each user is given by the interference generated by the other base stations in the system when they are active. Therefore, in the distributed approach formulation, each base station simply schedules local users in order to maximise the objective function defined in Problem `LOCAL`. However, the schedule defines the activity of the base station, and the interference generated towards other base stations, which, in turn, can react readjusting their scheduling in order to adapt to changed interference conditions. A new scheduling may cause new interference levels, therefore each base station must iteratively solve Problem `LOCAL`, until the system converges to a stable solution.

In addition, as explained in details in Appendix A.1, to avoid the player selfishness, we introduce an upper bound M_i on the number of time-slot each base station can use for scheduling its users. Therefore, the CLC needs to dynamically assess the network performance and decide, on a long-term timescale (in the order of seconds), the M_i value of each base station which directly affects how fast the system reacts to traffic changes. Strictly specking, this CLC-aided distributed mechanisms make our approach *semi-distributed*.

2.1.2.2 Performance assessment of semi-distributed ABSF control

A toy-scenario is provided in Figure 2.6, where 7 base stations are placed in a 450m×550m area serving 70 users. Some users are placed in a particular position to show how the inter-cell interference problem affects the system performance and the robustness of the proposed distributed scheme. The parameters we used for our simulations are summarised in Table 2.2.

Table 2.2: List of parameters for the LTE-A wireless scenarios used in ABSF control experiments

$\lvert\mathcal{N}\rvert$	Number of Base Stations	7
$\lvert\mathcal{U}_i\rvert$	Number of UEs per Base Station	10
T	ABSF Pattern Length	70 TTIs
BW	Spectrum Bandwidth	20 MHz
P	Transmitting Power	1 W
ISD	Inter-Site Distance	200 m
N_0	Background Noise	1.085×10^{-14}

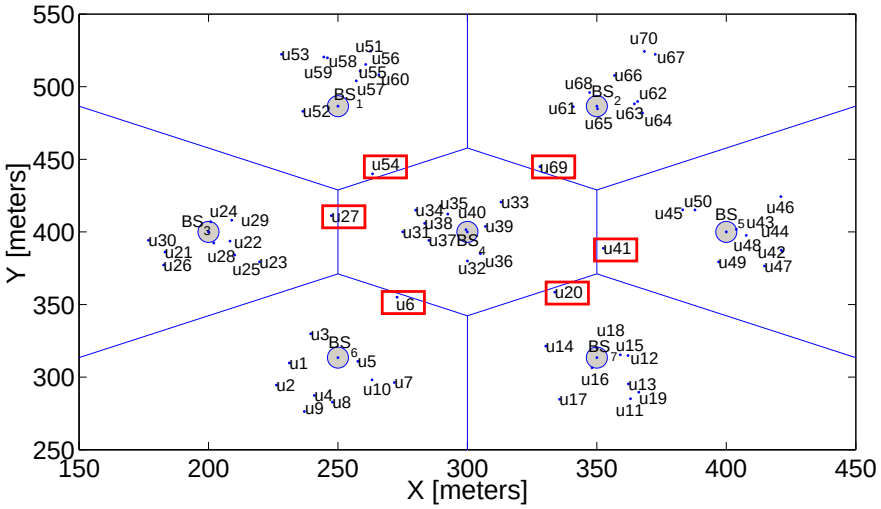

Figure 2.6: Toy-scenario representing 7 base stations and 70 users randomly dropped into the network area.

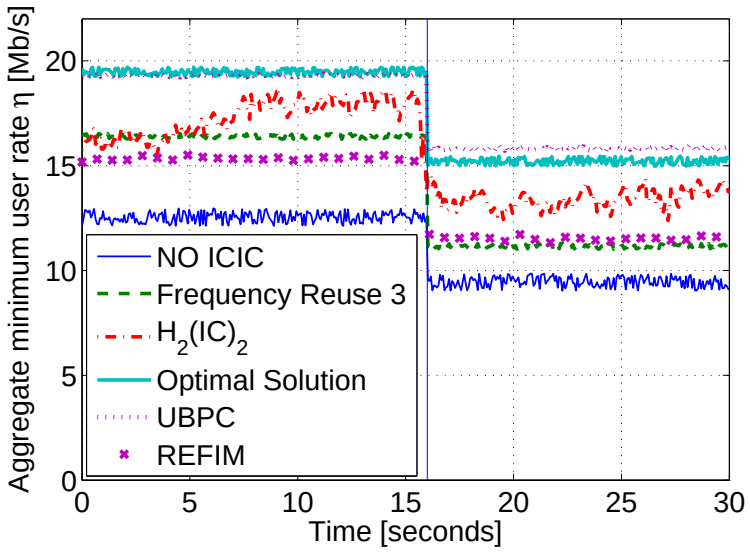

Figure 2.7: Dynamic behaviour of $H_2(IC)_2$ applied to a changing scenario. On the left side, the scenario has $|\mathcal{N}| = 7$ base stations, $|U_i| = 10$ users and $T = 70$ TTIs. On the right side, the number of users is increased up to $|U_i| = 20$ users. $H_2(IC)_2$ quickly adapts to a network change while keeping high the accuracy of the solution (w.r.t. Optimal solution and no ICIC).

Due to the simplicity of our approach and its limited control overhead, at best of our knowledge, this is the first attempt to move towards a practical, efficient, scalable and adaptive implementation of ABSF in real networks. Indeed, our numerical results, detailed in Appendix A.1.4 show that our proposal achieves near-optimal results with respect to a centralised omniscient network scheduler, and achieves performance levels similar to advanced schemes using complex power control approaches. Here we report a key result, which concerns the near-optimality of $H_2(IC)_2$. Specifically, Figure 2.7 shows η averaged over the time horizon of $T = 70$ TTIs[2] for the case of $|\mathcal{N}| = 7$ base stations and 10 users per base station. Due to the adaptive nature of the algorithm, $H_2(IC)_2$ shows a dynamic behavior that is close to the optimal one.

[2]Typical values for the ABSF pattern length are between 60 and 80. We use 70, which yields a round number for M_i^* in our simulated scenario consisting of 7 cells.

2.1.2.3 Conclusions

In CROWD we have proposed $H_2(IC)_2$, a practical eICIC scheme leveraging the ABSF paradigm in LTE-A networks. $H_2(IC)_2$ uses semi-distributed architecture, which allows base stations (eNBs) to jointly compute ABSF patterns and operate the local scheduling. The behavior of the base stations can be mapped onto the one of the players of a distributed game. The architecture is semi-distributed because it requires the presence of a central coordinator to drive the system to the best achievable network utility by controlling the number of TTIs to blank at each base station. Due to the simplicity of our approach and its limited control overhead, at best of our knowledge, this is the first attempt to move towards a practical, efficient, scalable and adaptive implementation of ABSF in real networks. Indeed, our numerical results show that $H_2(IC)_2$ achieves near-optimal results with respect to a centralised omniscient network scheduler, and achieves performance levels similar to advanced schemes using complex power control approaches (see Appendix A.1 for details on algorithm design, complexity and performance figures). The results of the research activity presented in this section have been published in [8].

2.1.3 Optimisation of LTE inband D2D

In cellular access networks, D2D communications represent a direct communication mode between two mobile users, in which the exchanged information does not traverse the eNB or the core network. Initially proposed for multi-hop relays in cellular networks, D2D communications are finding an increasing number of use cases in vehicular networks, in content distribution and cellular offloading [21]. In such applications, D2D has the potential to lower packet delays, increase energy efficiency, fairness and throughput while improving the spectral efficiency of dense cellular networks (e.g., adopting new flexible paradigms like in LTE-A).

Different types of D2D schemes have been proposed. *inband* schemes either allow D2D transmissions to occur over dedicated cellular resources (*inband overlay* schemes [22]) or over the same resources used by legacy cellular users (*inband underlay* schemes [23]). Here we focus on the latter, though our approach can be easily extended to inband overlay schemes.

2.1.3.1 Characterisation of LTE inband D2D

The analysis of the performance of inband underlay schemes is challenging, due to the complexity of D2D systems and of their interaction with cellular system operations. As D2D and cellular users all share the same resources, indeed, their performance is strongly correlated through interference, through their traffic patterns, and through the specific scheduling algorithm adopted. Available results are mainly based on simulations [24], and they assume either the system to be saturated, or that only a single D2D pair can be scheduled in the same Resource Block (RB) as a cellular user [23]. This leads to pessimistic, overly conservative results, particularly in non-saturated settings, and in general it does not enable the characterisation of the main performance trade-offs of such D2D systems essential for the design of efficient scheduling and rate allocation algorithms.

In Deliverable D2.2 [4] we proposed a queuing-theory based approach to performance characterisation of D2D schemes, which captures the dependencies between interfering transmissions and achievable rates. We report in Appendix A.3 a short summary of such analysis. The approach is based on the CPS [25] modelisation of the D2D environment. The CPS model naturally applies to D2D systems, as it explicitly accounts for the achievable transmission rates when the correlation between the service rates of multiple queues is known. The particular resource scheduling policy we consider is a variation over the one proposed in FlashLinQ [26]. FlashLinQ is a state-of-the-art Physical (PHY)-MAC architecture for D2D that allows the scheduling of different transmitters (D2D or cellular) in the same time and frequency resource, through an Orthogonal Frequency Division Multiple Access (OFDMA)-like access selection mechanism. In [4], by applying our approach, we derived new sufficient conditions for stability of transmission queues in a D2D system, and we showed how to evaluate the effects of D2D transmissions over cellular user performance.

2.1.3.2 Optimisation of D2D activity

By exploiting the knowledge of the conservative estimate of the stability region of the system, herein we show instead how to achieve proportional fairness among D2D transmissions. Given a CPS with D queues, each representing one of the D D2D pairs in the scenario, we first formalize a proportional fairness throughput optimisation problem that exploits the knowledge of the stability region, and then we introduce an heuristic that significantly reduces the complexity.

The goal of the optimisation is to introduce filters at the transmitters which allows just part to the demand to be queued at the transmitters. The allowed rates ρ_i of the leaky bucket rates are chosen in order to (i) ensure stability at the transmission queues of the D2D users, (ii) maximise a weighted sum of the logarithms of the rates, thus achieving *proportional fairness* of user's throughputs. The solution of the optimisation problem and the distribution of the long term rates of the shapers to

the D2D transmitters can be easily performed by the eNB where the D2D transmissions are taking place. Formally, we express the optimisation problem as follows:

$$\underset{\rho^*}{\text{maximise}} \sum_{i=1}^{D} w_{d_i} \log \rho^*_{d_i}, \tag{2.7}$$

where ρ^* values represent the long term rates of the shapers we want to introduce in the network and where w_{d_i} is the weight assigned to the transmitter d_i in order to achieve proportional fairness.

In the CPS model for D2D communications, each node affects the others due to interference, since multiple transmitters are scheduled on the same resources. To make the problem tractable, we proposed a method to break circular dependencies in the analysis. In practice, we create an ordered list out of the D2D transmitters in the system, and we evaluate the impact of the interference following such list. If we refer to the k-th transmitter in the list, we evaluate what is the effect of interference just due to the transmitters $\{1, ..., k-1\}$. With this, we model inter-queue dependencies in one direction (top-down in the list). To model the dependencies in the other direction (bottom-up in the list) without incurring in complex calculations, we consider that a queue in position j in the list is considered as always active by all the queues listed in position $k < j$. For further details, please refer to Appendix A.2. This is clearly a worst case approach, which will lead to the identification of performance bounds. Nevertheless, this represents, for the first time, a fully analytic approach for the study of performances of the D2D system, and enable the optimisation of system recources we mentioned above.

Each of the ordering of the D2D transmitters leads to a valid bound for the stability region. Therefore, the optimisation described by (2.7) should be performed on the union of the bounds achieved from each of the possible ordering of the transmitters. In other words, we should evaluate what is the optimum for (2.7) in each of the orderings of the D2D transmitters, and then select the one which is the maximum among the achieved working points for the system. In theory, this means that we should solve a number of optimisation problems that scales factorially with the number of transmitters in the system. Nevertheless, we present in Appendix A.3 a simple heuristic that allows to select, among all the possible orderings, just the ones presenting with high probability the higher values for the optimisation problem described in (2.7).

2.1.3.3 Performance assessment of D2D resource allocation based on CPS analysis of the capacity region

We evaluated through simulations both the achieved gains in fairness after applying the optimisation expressed in (2.7), when compared to the system that does not exploit the knowledge of the stability system, both the complexity of the heuristic proposed, i.e., the number of orderings analysed so to achieve such gains. As it is easy to see from Table 2.3, even if we achieve just a conservative estimate of the whole stability region, the shapers improve significantly the value of the achieved utility. In particular, when the scenarios get larger, the utility improves up to 35.2% on average and 131.53% in the best case.

In order to evaluate the complexity of the heuristic proposed, Table A.6 in Appendix A.3 shows the number of orderings evaluated by the heuristic. In particular, the table shows that the complexity of heuristic proposed remains computationally feasible, even though the complexity of an approach that evaluates all the orderings grows fast. For further details on the simulations performed, and for richer numerical results, please refer to Appendix A.3.

2.1.3.4 Conclusions

In this work, we have introduced a novel analytical approach to D2D systems based on a CPS model. We have shown how to use such characterisation to optimise a D2D system under any given

Table 2.3: Optimisation vs. Saturation: Utility

D2D transmitters	Mean Sat.	95% Conf. Interval	Mean Opt.	95% Conf. Interval
3	2.3289	2.11-2.54	2.5303	2.37-2.68
4	2.2870	2.06-2.50	2.5750	2.44-2.70
5	1.9881	1.80-2.17	2.3424	2.21-2.47
6	1.7460	1.39-2.09	2.2153	2.01-2.41
8	1.5185	1.25-1.77	2.0537	1.91-2.19

Table 2.4: Heuristic Complexity

D2D transmitters	Mean # of Orderings (Heuristic)	95% Confidence Interval	Networks Available
3	3.44	3.36-3.52	6
4	4.60	4.36-4.84	24
5	9.56	8.95-10.18	120
6	11.70	10.49-12.92	720
8	22.70	20.52-24.88	40320

operational condition, i.e., also when saturation does not hold. Specifically, our method accurately estimates the stability region of D2D users and the saturation throughput of cellular users. Such knowledge is exploited to achieve, with a relatively low complexity, a proportionally fair allocation of throughput among D2D users. One of the main issues left open by the present work is the inclusion in the model of the coupling between multiple eNBs. The results reported here and in Appendices A.2 and A.3 have been published in [9, 27].

2.2 IEEE 802.11 enhancements

We now move to the description of the assessment of new mechanisms to boost the efficiency of 802.11 networks, in presence of very dense deployments. Specifically we showcase the potential improvements achievable by (*i*) using a control theoretic design of distributed channel access mechanism, (*ii*) adopting SOLOR, a practical inband relaying scheme for 802.11 networks using WiFi Direct, and (*iii*) optimising resource allocations based on the characterisation of stability region and performance figures of D2D inband schemes in general.

2.2.1 Control theoretic design of distributed channel access

In Deliverable D2.1 [3] we presented a model for distributed channel access and, based on a given set of power consumption parameters, derived the configuration that results in the most energy-efficient performance. However, one of the key weaknesses of the proposed solution is that it lacks a mechanism to smoothly converge to that point of operation when conditions of the WLAN change, e.g., more stations join the network. In the following, we present a mechanism based on control theory that tunes the operation of a distributed channel access scheme—for simplicity, we do not consider power consumption, although it could be extended to maximise energy efficiency along the lines of Deliverable D2.1 [3].

2.2.1.1 System model and optimal configuration

Similarly to the model presented in Deliverable D2.1 [3], we model our system as a single-hop contention-based wireless network with N stations where time is divided into mini slots of fixed duration τ. At the beginning of each slot, station i contends for channel access with a given channel access probability, p_i. A slot can be empty if none of the stations attempt to access the channel. If $N > 1$ stations access the channel in the same slot, a collision occurs and the channel is freed for the next slot. There is a successful contention if only one station accesses the channel, which then probes the channel. After this channel probing (which we assume takes only one slot), the station has perfect knowledge of the instantaneous link conditions which can be mapped into a reliable transmission bit rate $R_i(\theta)$ at time θ. If the available rate is below a given threshold $\bar{R}_i$, station i gives up its transmission opportunity and frees up the channel for re-contention. Otherwise, the station transmits data for a fixed duration of time $\mathcal{T}$.

Under these conditions, we have that the optimal transmission policy is a threshold policy [28]: given a threshold $\bar{R}_i$, station i only transmits after a successful contention if $R_i(\theta) \geq \bar{R}_i$. With the above, stations throughputs are a function of the access probabilities, $\mathbf{p} = \{p_1, \ldots, p_N\}$, and the transmission rate thresholds, $\bar{\mathbf{R}} = \{\bar{R}_1, \ldots, \bar{R}_N\}$.

In Appendix A.4 we show that optimal configurations for p_i and $\bar{R}_i$ exist and how they can be computed. With such results we can design a controller that aims to drive the system to the computed optimal value for p_i and $\bar{R}_i$. We design such a mechanisms to be adaptive, and we call it Adaptive algorithm for Distributed Opportunistic Scheduling (ADOS).

2.2.1.2 Performance assessment

In this section, we present the fundamental assessment result for the performance of ADOS. Such results have been achieved by means of simulations. Further simulation results are available in Appendix A.4.4.

Unless otherwise stated, we assume that different observations of the channel conditions are independent and that the available transmission rate for a given SNR is given by the Shannon channel capacity:

$$R(h) = B \log_2(1 + \rho|h|^2) \quad \text{bits/s} \tag{2.8}$$

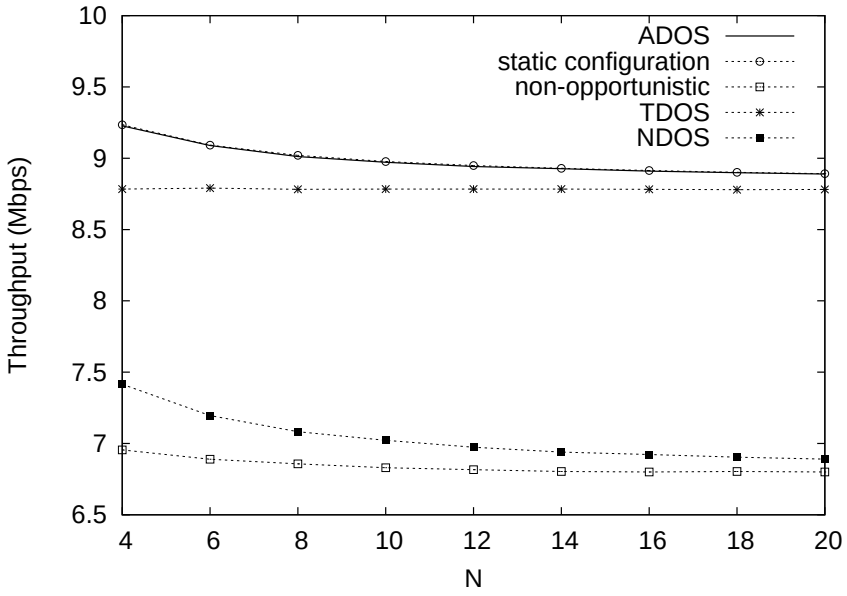

Figure 2.8: Homogeneous scenario with saturated stations.

where B is the channel bandwidth in Hz, ρ is the normalized average SNR and h is the random gain of Rayleigh fading.

We implemented the ADOS mechanism in OMNET++.[3] In the simulations, we set $B = 10^7$ and $\mathcal{T}/\tau = 10$. For all results, 95% confidence intervals are below 1%.

Homogeneous scenario with saturated stations. We start by considering a homogeneous scenario where all stations are saturated and have the same normalized average SNR ($\rho_i = 1\ \forall i$). We compare the performance of ADOS to the following approaches:

(i) The static optimal configuration obtained from performing an exhaustive search over the $\{p_i, \bar{R}_i\}$ space and choosing the best configuration ('*static configuration*').

(ii) An approach that does not perform opportunistic scheduling but always transmits after successful contention ('*non-opportunistic*').

(iii) The Team-game algorithm for Distributed Opportunistic Scheduling (TDOS) proposed in [28] ('*TDOS*'). This approach requires that each station knows the channel state of all the stations in the network, and hence incurs substantial signaling overhead.

(iv) The Non-cooperative algorithm for Distributed Opportunistic Scheduling (NDOS) approach proposed in [28] ('*NDOS*'). This approach, like ours, only requires information that can be observed locally, and hence does not involve any signaling.[4]

Figure 2.8 shows the total throughput as a function of the number of stations in the network. The figure confirms that ADOS is effective in driving the system to the optimal point of operation, providing the same throughput as the benchmark given by the *static configuration*. The *TDOS* and *NDOS* approaches provide lower throughput as they only optimise the transmission rate

[3] `http://www.omnetpp.org/`

[4] Since [28] only optimises the transmission rate thresholds but not the access probabilities, for the '*TDOS*' and '*NDOS*' approaches we take the configuration of access probabilities that are used in the simulation results of [28]. For the '*non-opportunistic*' approach, we choose the access probabilities that maximise the performance.

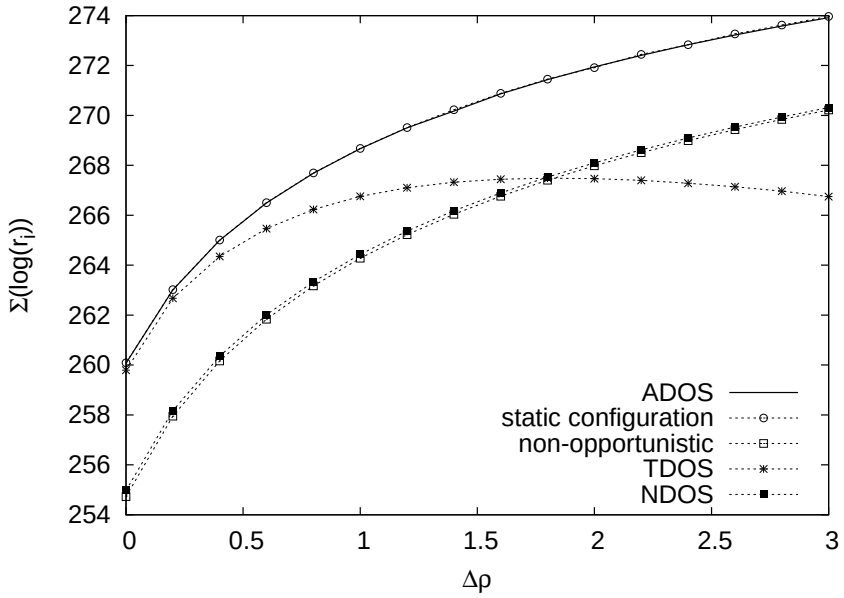

Figure 2.9: $\sum_i \log(r_i)$ as a function of $\Delta\rho$.

thresholds; among them, *NDOS* performs substantially worse as it has less information. Finally, the *non-opportunistic* approach provides the lowest throughput due to the lack of opportunistic scheduling. In conclusion, the proposed ADOS mechanism provides optimal throughput performance, outperforming the other approaches.

Heterogeneous scenario with saturated stations. In the case of heterogeneous channel conditions, performance does not only depend on the total throughput but also on the way this throughput is shared among the stations. To analyse performance in this scenario, we consider $N = 20$ saturated stations divided into four groups according to their channel conditions. The normalized SNR of the stations from group i is given by $\rho_i = 1 + (i-1)\Delta\rho$, with $i \in \{1, 2, 3, 4\}$. Figure 2.9 shows $\sum_i \log(r_i)$, the figure of merit for proportional fairness, as a function of $\Delta\rho$. We observe that ADOS performs at the same level as the benchmark given by the *static configuration*, while the other approaches provide a substantially lower performance. *TDOS* exhibits an increasing degree of unfairness as $\Delta\rho$ grows that harms its performance in terms of proportional fairness. *NDOS*, in contrast to *TDOS*, does not show this behaviour: with *NDOS*, each station sets its threshold based on its local radio conditions and therefore the fact that other stations have better radio conditions does not impact fairness. The price that *NDOS* pays for this non-cooperative behaviour, however, is that the overall throughput performance is substantially degraded for all $\Delta\rho$ values. The *non-opportunistic* approach also provides a poor throughput performance, similar to *NDOS*.

In order to gain additional insight into the throughput distribution with heterogeneous radio conditions, Figure 2.10 depicts the throughput obtained by a station of each group with the different approaches, along with the Jain's fairness index (JFI) of each distribution. The results confirm that *TDOS* suffers from high unfairness with heterogeneous radio conditions, since with this approach the stations with worst radio conditions (r_1) are almost starved while the stations with best radio conditions (r_4) obtain a very large throughput. In contrast, the *TDOS* and *non-opportunistic* approaches do not suffer from unfairness but provide significantly smaller throughputs than ADOS. We conclude that ADOS substantially outperforms all other approaches with heterogeneous radio conditions. For the evaluation of stability and for the behaviour of ADOS under changing radio conditions and number of users, we refer the reader to Section A.4.4 in the appendix of this

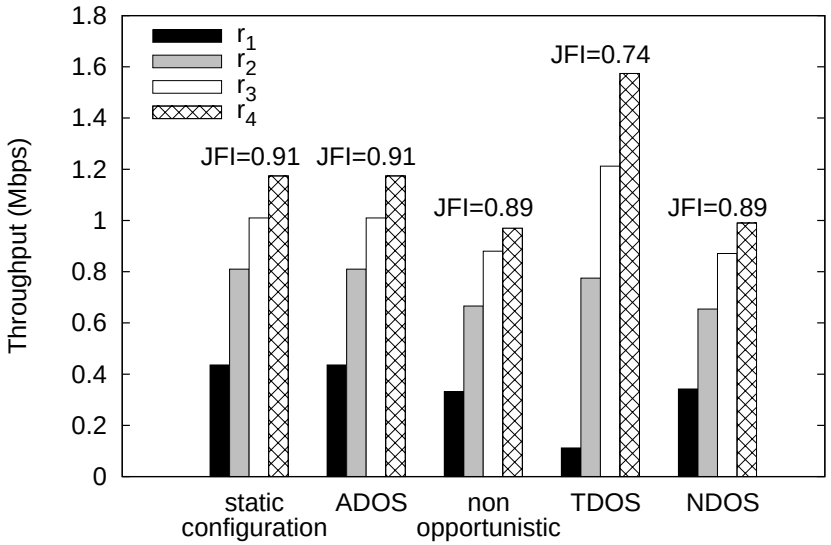

Figure 2.10: Throughput achieved by a station of each group for $\Delta\rho = 2$.

document.

2.2.1.3 Conclusions

ADOS, and distributed opportunistic scheduling techniques in general, provide throughput gains in wireless networks without requiring a centralized scheduler. One of the challenges of these techniques is the design of an adaptive algorithm to adapt the scheduling parameters to their optimal value. With ADOS. we have shown that it is possible to (*i*) jointly optimise both the access probabilities and the rate transmission thresholds of distributed scheduling mechanisms; (*ii*) provide a good tradeoff between total throughput and fairness, and (*iii*) guarantee convergence and stability. The performance of ADOS has been extensively evaluated via simulations, and results confirm that ADOS provides significantly better performance than previous proposals.

2.2.2 Practical relaying with WiFi Direct

In Deliverable D2.2 [4] we presented the theoretical analysis and performance of Self-Optimizing wireless local area networks with Legacy-compatible Opportunistic Relays (SOLOR) [29], a novel framework for scenarios in which some nodes judiciously decide to opportunistically relay for others, in order to maximise a performance metric that takes into account both throughput performance and energy consumption. Here, we detail the operation of SOLOR as a first step towards its implementation, describing how it is designed so it results legacy compatible, and then we report the implementation experiences and results. The operational details of SOLOR are reported in Appendix A.5.

2.2.2.1 Experimental Evaluation of SOLOR

Here we describe the results from the implementation of the SOLOR framework. Our 802.11g testbed, represented in Figure 2.11, is comprised of seven nodes, all using Ubuntu 11.10 with kernel 3.00. There are four legacy nodes, one of which is the AP, and three relay-enabled nodes. The legacy nodes are standard laptops equipped with WLAN cards based on the Atheros AR5413 chipset, using the `ath5k/mac80211` wireless subsystem, while the relay-capable nodes are desktop machines, each equipped with *two* WLAN cards based on the Atheros AR922X chipset and using the `ath9k/mac80211` subsystem. We decided, for simplicity, to use two NICs (Network Interface Cards) to emulate a single NIC with the ability to serve as Access Point (AP) on one channel and to connect to an AP on a different channel, as existing open-source drivers do not support this feature yet. On the other hand, our implementation will not require any modification once this feature becomes available. Note that, throughout our experiments, we take great care in confirming that only one of the two NICs is active at any point in time.

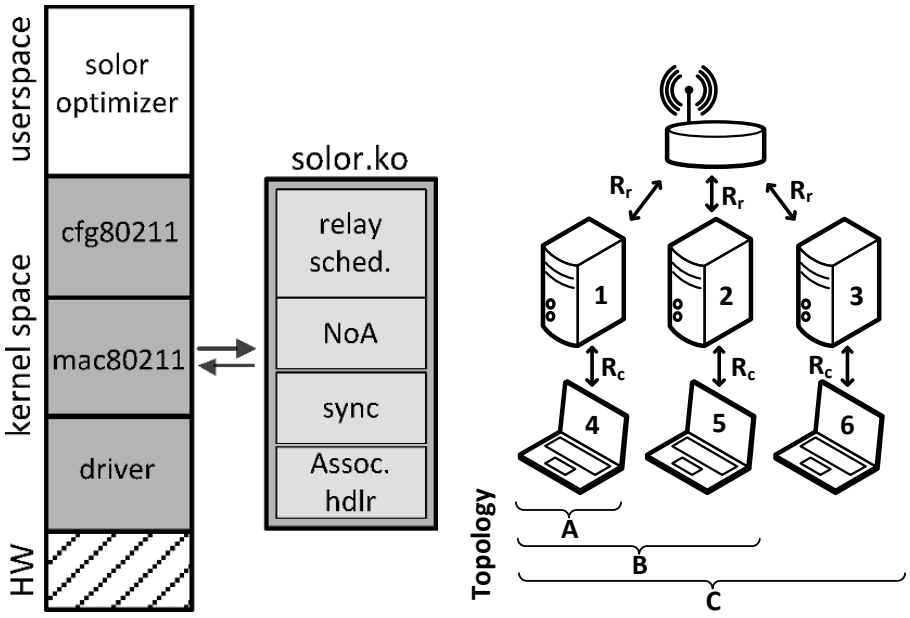

(a) Software modules. (b) Testbed deployed

Figure 2.11: Implementation architecture.

Implementing SOLOR: In order to implement SOLOR, three main functionalities are required: a) to analyse the WLAN deployment and compute the optimal configuration; b) to implement the

resulting relay schedules; c) to force legacy nodes to connect to the proper relay and to sleep when needed. This is achieved by the software architecture depicted in Figure 2.11a, consisting of a user-space application that computes the optimal configuration, and a kernel module (`solor.ko`) to interact with the Linux wireless subsystem.

The optimal configuration of the network is independently computed by the *SOLOR optimiser* of each SOLOR node; given the policies described before, using MAC addresses as node IDs, this will result in all relays computing the same joint schedule with fractions $\vec{F}$. Unless otherwise stated, the individual preference parameters α_i's are set to 1, and the timers are set to T_1 =500 ms and T_2=1 s.

To implement the schedule, the `solor.ko` module builds on the *synchronisation* provided by beacon frames sent by each parent, and triggers the corresponding notifications to the *relay scheduler*. This one reacts upon a notification and apply the required context change in the driver through mac80211 (i.e., transmit buffered data, received and buffer data, or sleep). The setup of the links computed by the new topology is handled by the *Association handler* which, as explained in Appendix A.5, forges a disassociation message and announces to the network as an AP (which will have better SNR with the target clients). Finally, `solor.ko` implements the *Notice of Absence* protocol to advertise the sleeping policies to the relay's clients.

Performance Evaluation: Finally, we analyse the performance of SOLOR in a real scenario under both static and dynamic scenarios.

Static conditions: We start our experimental evaluation by measuring the throughput performance of different static settings with a fixed topology, in order to validate the results from the previous sections. To this end, we consider the three topologies depicted in Figure 2.11b and different settings of the transmission rate between the laptops and the relays (denoted as R_c), and the rates between the relays and the AP (denoted by R_r), and compare the per-node throughput figures X_n obtained in the testbed with the analytical ones both for uni-directional and bi-directional flows. The results are depicted in Table 2.5, showing that in all cases the experimental figures match remarkably well the results from the analytical model (described in Deliverable D2.2 [4]), which are provided in parenthesis.

Table 2.5: Per-node throughput (in Mbps) for the topologies in Figure 2.11b.

Topo.		R_c, R_r (Mbps)	X_1, X_2, X_3 (X_{model})	X_4, X_5, X_6 (X_{model})
Uni-directional	A	48, 48	14.60, -, - (14.62)	7.31, -, - (7.31)
		48, 24	14.22, -, - (14.62)	5.51, -, - (5.57)
		24, 24	8.7, -, - (9.00)	4.65, -, - (4.5)
	B	48, 48	7.46, 7.42, - (7.31)	6.98, 7.12, - (7.31)
		48, 24	7.64, 7.63, - (7.31)	7.16, 7.23, - (7.31)
		24, 24	4.11, 4.92, - (4.50)	4.32, 4.10, - (4.50)
	C	48, 48	5.30, 4.21, 4.42 (4.87)	3.80, 4.12, 3.84 (4.87)
		48, 24	4.53, 4.98, 4.41 (4.87)	4.30, 4.56, 4.52 (4.87)
		24, 24	2.92, 3.22, 3.15 (3.00)	2.63, 2.52, 2.77 (3.00)
Bi-idirectional	A	48, 48	Up: 6.59, -, - (6.08)	Up: 2.49, -, - (3.04)
			Dwn: 6.01, -, - (6.08))	Dwn: 2.96, -, - (3.04)
	B	48, 48	Up: 3.21, 2.94, - (3.04)	Up: 2.81, 3.05, - (3.04)
			Dwn: 3.09, 3.34, - (3.04)	Dwn: 2.71, 2.98, - (3.04)
	C	48,48	Up: 1.75, 2.11, 1.87 (2.03)	Up: 2.22, 1.89, 2.08 (2.03)
			Dwn: 2.12, 1.99, 2.23 (2.03)	Dwn: 1.87, 2.04, 1.91 (2.03)

Dynamic conditions: We next assess the performance of SOLOR in a dynamic scenario, in which nodes activate the relaying functionality in real-time and thus the topology changes over time. Nodes 1–3, which do not have the relay functionality activated at the beginning of the experiment, can transmit to the AP at 48 Mbps, while nodes 4–6 transmit to the AP at 6 Mbps, and could transmit to nodes 1–3 at 48 Mbps. Our experiment is divided in stages of approximately 20 seconds each. During the first stage, all nodes are transmitting to the AP, this being the "default" scenario; during the second stage, node 1 enables the SOLOR functionality and as a consequence starts relaying traffic for nodes 4–6; in the third stage, node 2 also enables the SOLOR functionality and relays the traffic from node 6, while node 1 keeps relaying for nodes 4 and 5; finally, in the last stage, node 3 is also enabled as a SOLOR node and, as a consequence, each relay-enabled node serves one client, i.e., the topology C depicted in Figure 2.11b.

We display the evolution of the per-node throughput figures over time in Figure 2.12 (top), in which the transient caused by the re-association periods can be easily identified. The corresponding overall utility of the WLAN is depicted in the bottom subplot, along with the theoretical values. We conclude from this experiment that enabling the relay functionality supports increasing the utility of the network, with a good match between experimental and analytical results, and that the SOLOR framework is easily implementable using commercial, off-the-shelf hardware.

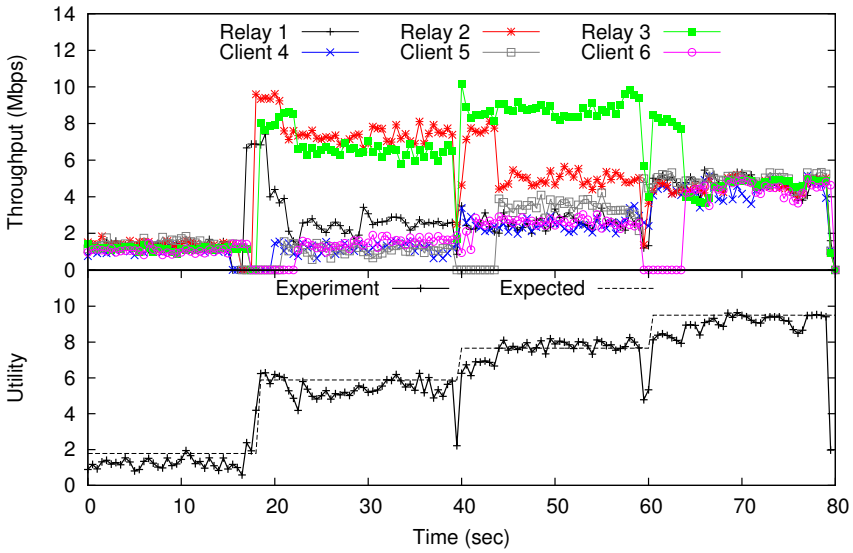

Figure 2.12: Dynamic experiment

Energy performance and per-node preferences: We now evaluate our prototype with dynamic individual preferences and show the results in Figure 2.13. For the sake of readability, we only use Relay 1-2 and Client 4-5 and initialise a static topology with 1 and 2 serving 4 and 5, respectively.

We start off by selecting $\alpha_i = 1,\ \forall i$, just like we did in our previous evaluations and we vary each node preferences sequentially every 10 seconds, illustrating that the larger the α, the more emphasis is given to throughput performance. We conclude from this experiment that SOLOR succeeds at tuning the per-node preferences in the throughput vs. power consumption trade-off.

2.2.2.2 Conclusions

We have assessed the performance of SOLOR, a novel self-optimising, legacy-compatible opportunistic relaying framework which addresses the rate anomaly problem by taking into account three major considerations to achieve an efficient deployment in real-world systems: (i) relaying

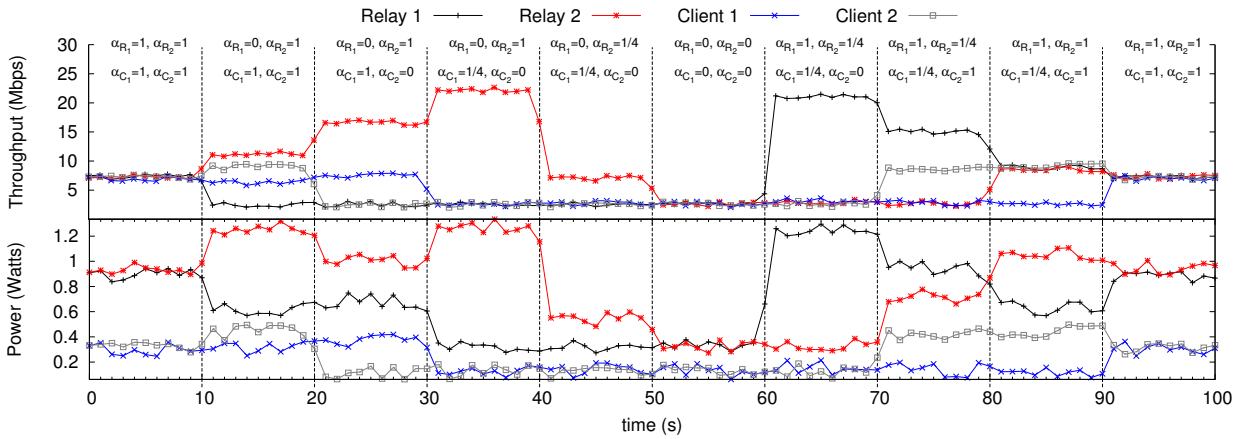

Figure 2.13: Per-node preferences

could imply increased power consumption, and nodes might be heterogeneous, both in power source (e.g., battery-powered vs. socket-powered) and power consumption profile; (*ii*) similarly, nodes in the network are expected to have heterogeneous throughput needs and preferences in terms of the throughput vs. energy consumption trade-off; and (*iii*) any proposed solution should be backward-compatible, given the large number of legacy 802.11 devices already present in existing networks. Our results show that SOLOR greatly improves network throughput performance and energy consumption even in systems comprised mostly of vanilla nodes and legacy access points.

2.2.3 Optimisation of WiFi inband D2D

Wireless mesh networks based on IEEE 802.11 are nowadays inexpensive, well widespread solution to easily, effectively and wirelessly connect entire cities. Thanks to such pervasiveness, they are poised to play a central role in many "Internet of Thing" application scenarios, with very diverse QoS requirements. Such wide deployment makes it crucial to develop models for analytical performance study of such networks, whose users will use D2D to communicate, as empirical studies in such a complex environment hardly give clear indications on general properties of such systems.

2.2.3.1 CPS-based analysis of 802.11 capacity

Performance analysis of the IEEE 802.11 CSMA/CA mechanism has traditionally focused on saturated traffic assumptions [30]. Many of the available results for non-saturated conditions do not capture the effects of traffic dynamics on system performance. For example, the authors of [31] assume that the probability of transmission and the probability of success of the stations in non-saturated conditions is always the same during time, despite traffic dynamics do alter significantly such quantities. Furthermore, the totality of the approaches in the state of the art are based on a Poissonian traffic assumption, and they depend heavily on the assumption of traffic stationarity. [32] analyses non saturation in heterogeneous traffic conditions, but its results still requires a complete stochastic traffic characterisation to be parametrized. This has been done despite the fact that traffic in real networks is well far from being Poissonian (see [33] and related literature). In particular, traffic from live audio/video streaming exhibits a periodic behavior which substantially departs from the Poisson model, and which is characteristic of several known examples of instability [34]. This leaves open the issue of how to derive valid performance guarantees in ad hoc networks in realistic settings, when little is known about traffic statistics.

In the CROWD project, we tackled the problem through a different approach, which is based on a CPS characterisation of the network that can be used for transmissions based on the presence of AP, ad hoc configurations or D2D schemes as well. First, we assume traffic to be constrained by *leaky bucket arrival curves* [35], which limit just the maximum amount of bits which can arrive in a given time interval. The leaky bucket arrival curve is characterised by two parameters, i.e., σ (burstiness) and ρ (rate). That is, if $A(t)$ are the cumulative arrivals, then $\forall t \geq 0$, $\forall t' \leq t$, $A(t) - A(t') \leq \sigma + \rho t$. Such approach widen the applicability of the presented results, since traffic patterns of different nature can be described through leaky bucket curves. Second, as anticipated above, our analysis is based on CPS description of the system, i.e., a *coupled processors* model [25] of the ad hoc system, which allows capturing the dependencies between user achievable rates due to sharing of the wireless transmission medium (mediated by the CSMA/CA mechanisms) and traffic dynamics, which such coupling entails. A CPS is a set of parallel queues (i.e., queues which do not exchange traffic among them) served by work conserving schedulers, and whose service rates at any time t is completely determined by the set of active queues at that time. Similarly, in an ad hoc system that applies IEEE 802.11 MAC layer the probability of successful transmission at time t, and consequently the throughput achieved by stations, is univocally determined by the set of active transmitters at t. We assessed the quality of the CPS model of the ad hoc system through simulations and, for the 99.73% of the evaluations performed, the difference among the model and the simulations was less than the 10% of the real throughput achieved by transmitters. We report the details of this validation in Appendix A.6 Third, in order to study the performance of the underlying CPS model of the ad hoc system, we used the analytical solution proposed within the CROWD project and presented in [27], that we summarize in Appendix A.2.

Through the analysis proposed, we were able to determine conditions on the arrivals such that the queues at the transmitters are guaranteed not to explode, i.e., sufficient conditions for the stability of the system. In particular, sufficient conditions for the stability of the system are represented by

bounds on the long term rate of the arrivals ρ at the transmitters.

We denote as $\underline{\rho} = \{\rho_1, ..., \rho_j, ..., \rho_N\}$ the vector of the long term rates of the arrivals at the ad hoc transmitters, when N transmitters are present in the scenario. Exploiting the additional knowledge provided by our analysis, we present a computationally feasible method to choose, among the set of stable long term rates $\underline{\rho}$, the rates that optimise a utility function, e.g., to guarantee that some form of fairness is maximised. In particular, in the present work, the utility function that we choose to optimise is a weighted fairness function, which is one possible way of balancing some notion of fairness among users with, for instance, different classes of service. Its expression is as follows:

$$U = \sum_{i=1}^{N} w_i \log\left(\frac{\rho_i}{\rho_0}\right), \tag{2.9}$$

where ρ_0 is the minimum bit rate for an acceptable performance for the application. The details of the presented optimisation are reported in Appendix A.6.

2.2.3.2 Performance assessment of proportional fairness optimisation of 802.11 networks based on a CPS characterisation of the capacity region for D2D transmissions.

We assessed the results achieved through the optimisation presented, evaluating a large set of scenarios. In the following, we just report an example of results. For more full details and results, please refer to Appendix A.6. Here is enough to say that the weights w_i we use in the utility function, defined in Equation (2.9), are uniformly distributed in $[0, 1]$, and that, overall, for each value of D2D transmitters N, we considered a total number of instances of our setting (i.e., set of weights for U and positioning of the transmitters in the scenario) sufficient to get a 95% confidence interval. In any case, we never used less than 100 instances.

In Figure 2.14 we compare the average log-utility, together with the 95% confidence interval, from our optimisation and the one obtainable when no knowledge of the stability region is present, i.e., when all transmitters are considered saturated. Figure 2.14 also contains the median of the utility U in the same cases.

We can see how in all cases the average log-utility derived by optimising (through our heuristic) over the set of operating points which are stable according to our method is always at least 18.43% larger than the one derived by assuming the system in saturation. Moreover, we see that the relative improvement brought over by our heuristic over the utility achieved under saturation assumption grows with the size of the scenario. The larger is the number of the stations in the system, indeed, the higher is the rate of contentions and, consequently, the inefficiency of the MAC under saturation assumptions.

In order to have a better idea of the difference between the operating points resulting from the optimisation and from the saturation assumption, we have compared them on the basis of the total average throughput, weighted in order to take into account the relative contribution of each host to the utility of the system. That is, the weights in these sums are the same as those adopted in the utility function. The results are shown in Figure 2.15, where the case under analysis is exactly the same used in Figure 2.14. We see how our heuristic brings the system to an operating point for which the total average throughput is at least 114.43% higher than the total average throughput achieved under saturation assumptions.

2.2.3.3 Conclusions

To summarize, in the present research work we have proposed a new analytical method for the analysis of ad hoc networks, valid for any number of nodes. Our method does not require the traffic to be Poissonian, nor to be stationary, but only to be constrained by a deterministic arrival

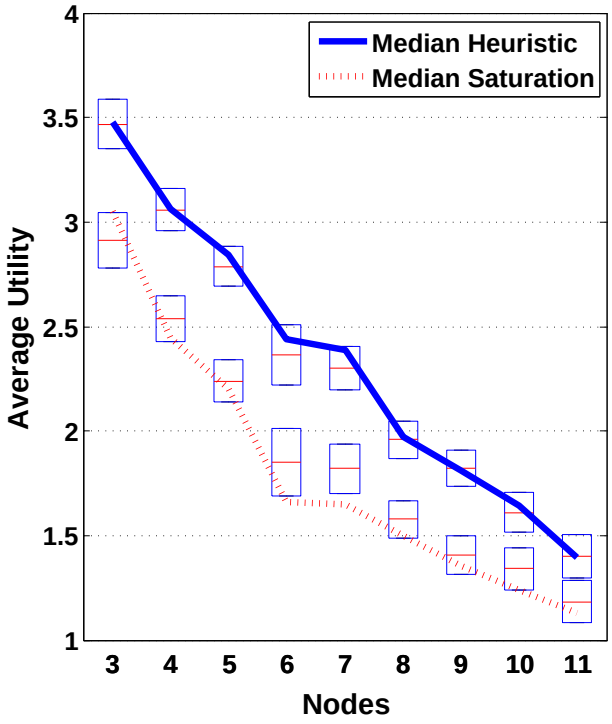

Figure 2.14: Average and Median Log-utility. Optimisation vs. Saturation approx.

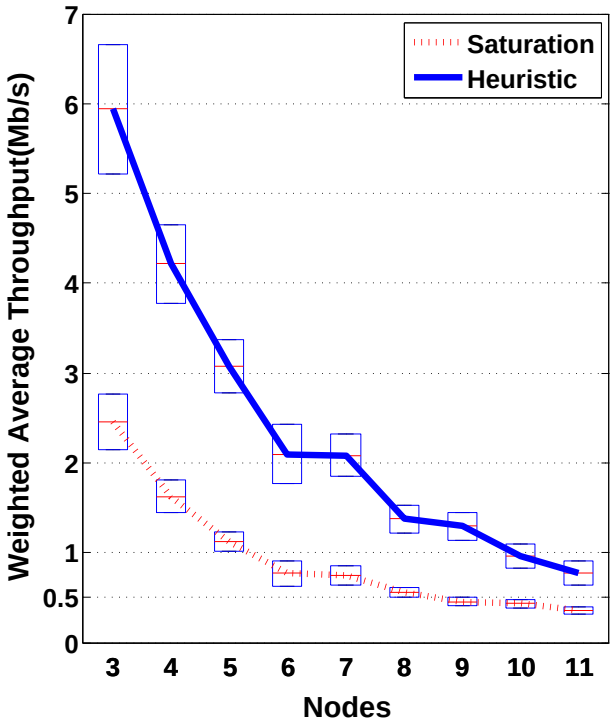

Figure 2.15: Weighted Av. Throughput. Heuristic vs. Saturation Condition.

curve. We proposed an optimisation for the derivation of a stable and proportionally fair allocation of resources and we showed how such optimisation derives system operating points that, besides maximising the utility function, bring to a much more efficient utilization of network resources. The results shown in this section and in Appendix A.6 have been published in [27, 10].

2.3 HetNet enhancements

We finally address the description of the assessment of new mechanisms to boost the efficiency of HetNets, which includes the cases of heterogenous cell sizes and heterogeneous technologies densely deployed in the same coverage area. Specifically we showcase the potential improvements achievable by (*i*) optimising cell association and inter-cell interference in presence of small cells which complements the coverage of macro cells, and (*ii*) optimising LTE D2D operation by selecting the D2D mode of D2D transmissions according to network conditions, assuming that D2D operations are under the control of the network.

2.3.1 HetNet Optimisation

To cope with limited radio spectrum and exponential traffic growth, a promising approach has been proposed by 3GPP LTE/LTE-A, called HetNet. It consists in the deployment of small cells complementary to macro cells in order to increase the network capacity via higher spatial reuse of spectrum [36]. There are two main challenges faced by HetNet technology: Cell association and Inter-Cell Interference (ICI) management. To address these issues, eICIC has been introduced in 3GPP Release 10 [37]. Besides, Cell Range Extension (CRE) has been proposed to increase small cell coverage footprint by adding Cell Individual Offset (CIO) to user's Reference Symbols Received Power (RSRP) measurements. Traditionally a UE would attach to an eNB that provides the strongest received signal. As macro eNBs are transmitting with higher power than small cells (e.g., 40 W vs. 1 W), only few users would attach to these low power nodes [38]. Therefore, these small cells may be underutilized while macro cells are overloaded. By CIO, a UE will be biased to attach to a different eNB.

2.3.1.1 Problem formulation

A first question is what CIO values should be used for each eNB. Secondly, since a UE in the CRE of a small cell may experience strong interference from the macro cell, there is a need of enhanced inter-cell interference management. One can see that the user association and interference mitigation problems are closely related [39]. Techniques to address these issues can be divided into three major categories: time domain, frequency domain, and power domain methods. KPIs include system and cell-edge user throughputs and also service fairness, energy efficiency, second-order statistics (e.g., for mobility robustness optimisation), etc. Here, we establish a generic framework and then provide solutions of optimising eICIC parameters such as ABSF and CIO and user association [40, 41]. Analytical and simulation results have shown their effectiveness. In particular, we consider a HetNet which consists of macro and small cells, and over which Physical Resource Blocks (PRBs) are synchronized. More details on the analysis of the problem and the formal problem statement are reported in Appendix A.7.

In our optimisation, each eNB is authorized to use a PRB only if such decision optimises a utility function. The problem to solve can be expressed as follows. Given a network state $\mathbf{s}$, which has three elements (the CIO, time and frequency pattern, and transmit power of each eNB, in the set of eNBs K), we look for the optimal values of those parameters so as to maximise a cost utility function $Ut(\mathbf{s})$:

$$Ut(\mathbf{s}) = \sum_{k \in K} Ut_k(\mathbf{s}), \tag{2.10}$$

where Ut_k is the utility of the eNB k given by a proportional fair approach on the acheved user throughput, and the optimisation is subject to power constraints. More details on the utility functions and the optimisation problem are given in Appendix A.7.1.

Based on the formulate problem, we propose a generic framework using game theory to determine the optimal configuration using an iterative approach and enabling a fully distributed scheme as in [36] for 3G networks and in [42] for 4G networks. Each UE reports to its serving eNB only long-term statistics, which are then explicitly shared between the own serving eNB and its neighbours so that each eNB can make, iteratively, decision that are not purely selfish but also considers the impact over neighbours. The details of the game theory-based framework are described in Appendix A.7.2.

2.3.1.2 Performance assessment of HetNet optimisation algorithms

We use a MATLAB-based LTE-compliant simulator developed by the TU Wien's Institute of Telecommunications [43]. This tool allows both link level and system level simulations. Then, it is possible to analyse link-level related issues and also the entire LTE system where the physical layer is abstracted by simplified models that capture its essential characteristics. Globally, the simulator is structured in two main building blocks or layers: link measurement model and link performance model (see Appendix A.7, and in particular Figure A.20 and Table A.12 for more details on the simulations). The simulation runs using a Region of Interest (ROI) in which the eNodeBs and UEs are positioned and total simulation duration is expressed in TTIs. As output, the simulator provides traces containing the main KPIs such as throughput and bit error rates. Generally, this simulator allows the study of cell planning, scheduling, interference coordination, etc. The throughput results are presented in Table 2.6 for various configurations:

- Config 1: A virtual case where cells are considered as non-interfering; this situation gives us a maximum theoretical capacity.
- Config 2: Interfered system and no optimisation is performed.
- Config 3: Only CIO is optimised, some UEs get offloaded from macro cells to small cells but they remain highly interfered.
- Config 4: Joint optimisation of CIO and ABSF is performed.

Table 2.6: Throughput performance results for HetNet optimisation algorithms.

	Config1	Config2	Config3	Config4
Total throughput (Mbps)	420	157	170	187.5
Average UE throughput (kbps)	—	303	329	362
Average cell edge UE throughput (kbps)	—	218	285	673

When performing only CIO optimisation (Config 3), the average throughput of the network increases from 157 to 170 Mbps. This is due to the change in CIO values, hence associating some UEs to small cells offering them higher bandwidth. However, these cell-edge users remain highly interfered by macro cells which continue to schedule at all sub-frames as there is no ABSF optimisation in this case. The best performance is given by the joint optimisation of CIO and ABSF. We can see that the gain obtained for the entire network's throughput is about 20% while the cell-edge users' throughput is enhanced by about 200%.

Figures 2.16 and 2.17 show the distribution of throughputs of all user and cell-edge users to compare performance obtained with and without eICIC optimisation. We observe a significant improvement: 50% of cell-edge users get at least 0.5 Mbps by the eICIC optimisation, while only 10% of them get the same throughput when the joint optimisation is de-activated.

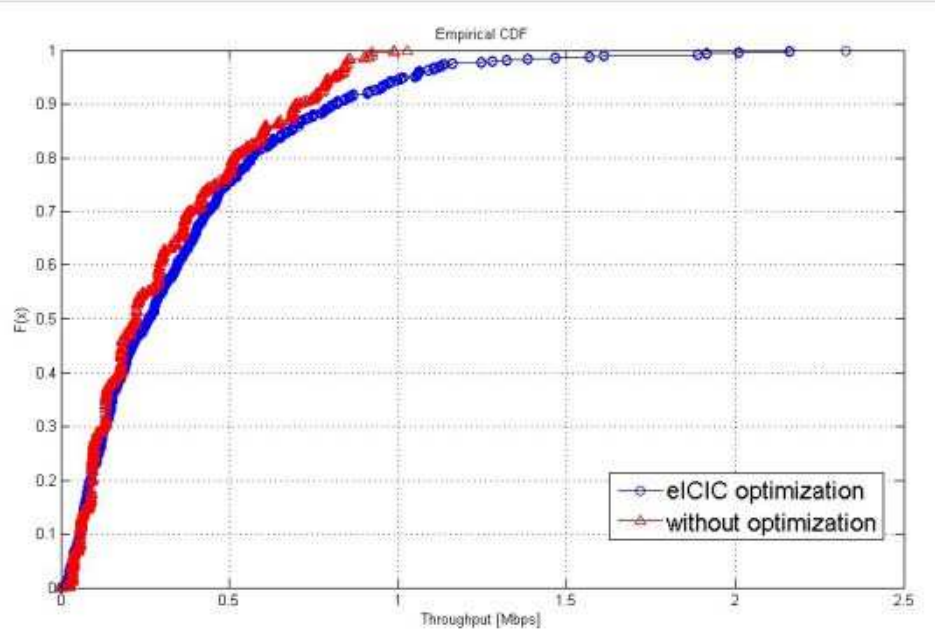

Figure 2.16: Cumulative distribution function of global throughput.

Secondly, we use the Jain index $J(x_1, x_2, \ldots, x_n) = \frac{(\sum_{i=1}^{n} x_i)^2}{n \sum_{i=1}^{n} x_i^2}$ to indicate the service fairness among all the users. The result will range from $1/n$ (worst) to 1 (best). Note that it is maximum when all users receive the same allocation. Without optimisation, the Jain index is 0.602. With eICIC optimisation, the fairness is increased and the value goes to 0.633.

Figure 2.18 shows the cell attachment of some users (in blue dots): black dotted lines refer to the UE association before optimisation and pink lines are the result after eICIC optimisation. In the initial configuration with all CIOs equal to 0, all the users were attached to macro eNBs, even those who are extremely near the small eNBs. After optimisation, some CIOs have been changed, which impacted the user attachment. We see that 39 users made handovers from the macro eNBs to small eNBs, which allows to offload traffic from macro cells and to have better cell-edge user experience.

2.3.1.3 Conclusions

Based on our mathematical tool from game theory, we have designed a flexible framework for addressing key heterogeneous network radio parameter optimisation. The proposed solution exhibits the following properties:

- It can support various utility functions and operation requirements depending on network operator's optimisation strategy;
- It can support hybrid architectures for computing the utility function ranging from a fully distributed manner to a centralized scheme;
- It can support various data collection granularity, frequency and types, according to the UE grouping performed by the local eNB;
- It can support exception sets such as race conditions (exclusion of some operation cases) or more limited optimisation such as disjoint optimisation of ABSF patterns and CIO values.

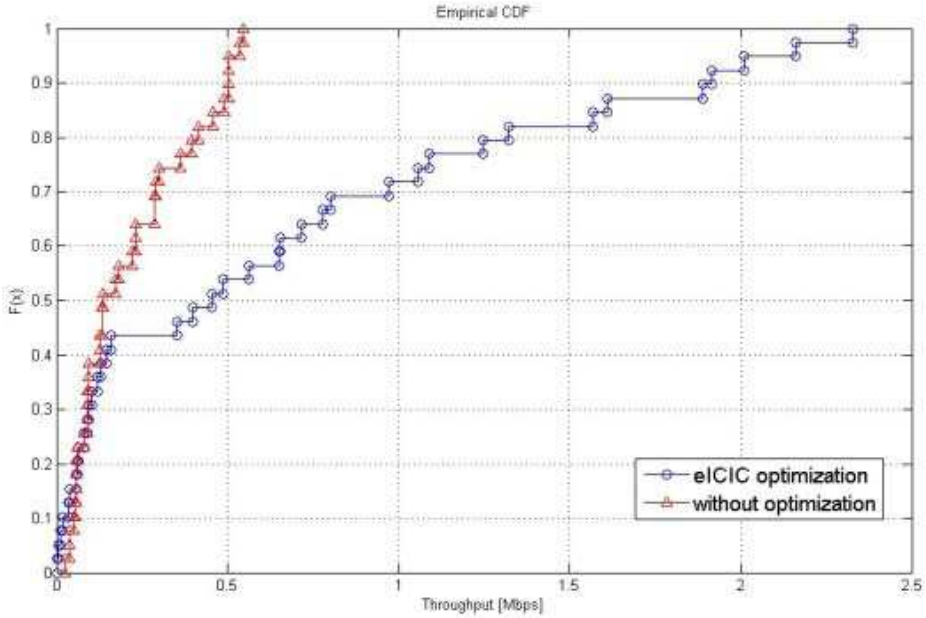

Figure 2.17: Cumulative distribution function of cell edge users' throughput.

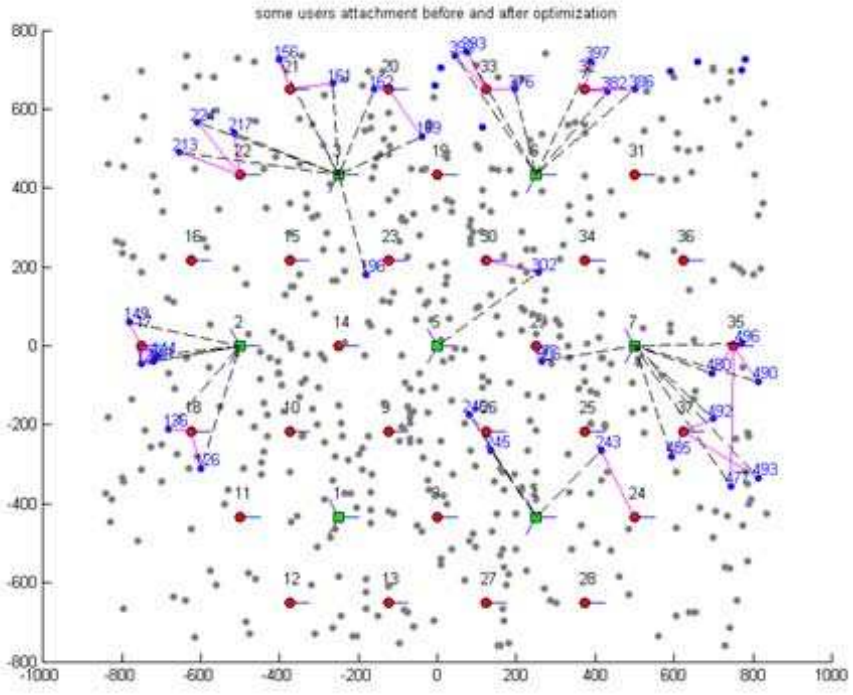

Figure 2.18: User attachment before and after optimisation.

2.3.2 D2D mode selection for offloading in LTE

Researchers have proposed a wide range of use-cases for D2D communications such as relay [44], multicasting [45], and cellular offloading [46]. Initial D2D proposals focused on D2D communication underlaying cellular network transmissions, i.e., using the same spectral resources used for cellular communications [47]. Later, other D2D techniques have been proposed, which either fall under either *inband* or *outband* D2D communication. Inband D2D communications allow D2D users to communicate over the cellular spectrum, while outband schemes demands the D2D users to access unlicensed bands for D2D transmissions [21]. Each of these D2D operational *modes* poses its own merits and disadvantages in terms of interference management, implementation complexity, achievable spectral efficiency, and therefore in terms of performance guarantees. However, the available literature proposes solutions for efficiently implementing each mode in isolation, i.e., *mode selection* has not been addressed. Nevertheless, according to the definition provided by 3GPP standards, "D2D communication is the communication between two users in proximity using a direct link between the devices in order to bypass the eNB[5] or core network" [48]. Therefore, any of these modes or perhaps all shall be used for D2D communications. Moreover, promising studies on D2D communication moved industry leaders such as Qualcomm to invest on future implementation of D2D communications, and 3GPP is considering to include generic D2D support in the next release of LTE-A standard as a public safety feature [48].

We believe that different D2D modes should not be treated as competitors but as complementary techniques. Co-existing D2D modes can immensely increase the system complexity because there should exist a mechanism to select the correct D2D mode according the overall system conditions.

Here, we describe our reference system, our proposed mode selection approach, and its practical implications, along with their feasible solution. We provide analytical insights into the mode selection problem (see Appendix A.8) resulting in an innovative multi-mode multi-band setup, which accounts for both achieved *throughput* and *energy costs*. We call such a novel approach *Floating Band D2D*, because D2D transmissions can occur on either inband or outband modes.

2.3.2.1 System Model

We consider a hexagonal multi-cell LTE-A network with a reference cell in the center and its first-tier neighbours as shown in Figure 2.19. The cell consists of N users labelled as $n \in \mathcal{N} := \{1, 2, \ldots, N\}$. Downlink and uplink channels are separated and each one has a fixed bandwidth. Users may communicate with other users in the cell or with those outside the cell. If a user wants to communicate with another user in proximity, it can use D2D communications.

Inband D2D communications use uplink cellular spectrum [48]. It is assumed that each user communicates with (at most) one user at any given time. Each connection between users n and m is referred to as (n, m), $\forall n, m \in \mathcal{N}$. For notational convenience, the eNB is addressed as user $N + 1$. In this paper, the outband D2D exploits WiFi Direct technology. With the above, we use four communication modes operating as described in the following:

- Mode 0 $\leftrightarrow$ cellular;
- Mode 1 $\leftrightarrow$ inband underlay D2D;
- Mode 2 $\leftrightarrow$ inband overlay D2D;
- Mode 3 $\leftrightarrow$ outband D2D (WiFi).

Our system operates in discrete time units and the eNB is in charge of mode selection and scheduling. The eNB makes the scheduling decisions on a per-frame basis. Each *frame* consists

[5] eNB is the 3GPP term referring to cellular base stations.

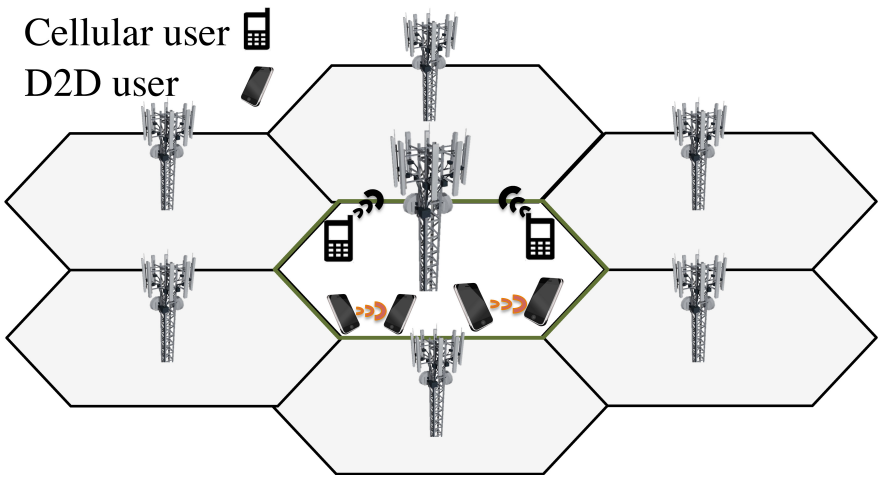

Figure 2.19: System model for the evaluation of D2D mode selection strategies. The System consists of a cell with its first-tier neighbours.

of 10 *subframes* and the length of a frame is 10 ms. In each subframe, only one cellular user is scheduled, while the number of concurrent D2D transmissions is not limited *a priori*. Therefore, there is no interference among cellular users (i.e., mode 0), but underlay users (i.e., mode 1) interfere with the cellular and other underlay users (i.e., modes 0 and 1). Overlay users only interfere with each other, while outband D2D users simply contend for the WiFi channel. A fixed portion of cellular bandwidth is dedicated to overlay D2D users. This portion is released to cellular and underlay users if there is no user in overlay mode.

2.3.2.2 Mode selection and scheduling

As mentioned, mode selection and scheduling decisions by nature require decision making schemes with a different time-scale resolution. Thus, we propose to decouple the mode selection and scheduling problems. The decoupling is mainly inspired by the fact that D2D connections last more than a few frames in a real world scenario and scheduling them on a per-frame basis is unnecessary and possibly inefficient. The inefficiency is due to the high signaling overhead, which is caused by such a high resolution mode selection. Moreover, the channel quality of D2D links is potentially less time variant in comparison to that of the cellular links due to the short-range nature of D2D communications. The decoupling also simplifies the integration of D2D communications into current cellular systems as it minimises the changes to the scheduler. Although mode selection and scheduling are decoupled, they are still highly intertwined. On one hand, the scheduling is affected by the interference, which is unknown before mode selection. On the other hand, mode selection depends on the set of cellular users scheduled along with D2D users. Hence, we choose the eNB to perform mode selection, because it is already in charge of scheduling.

We propose a mechanism in which the eNB handles these decisions in two steps: (*i*) *mode selection* and (*ii*) *scheduling*. First, in mode selection, each D2D pair is assigned a mode (modes 1 to 3), and the assignment is repeated at regular *mode intervals* of length T seconds. The eNB selects D2D modes with the assumption of a worst-case interference scenario. This approach helps to reduce the system complexity and to avoid disruptive co-channel interference. Second, in the scheduling phase, the eNB schedules users and assigns them a Modulation and Coding Scheme (MCS). Mode selection and scheduling both rely on the accuracy of CSI data gathered at the eNB, which can be challenging in terms of signaling overhead and scheduling for eNBs

We formulate the mode selection problem based on the described scenario which is an NP-Hard problem. Hence, we propose three heuristics (namely, Social, Greedy and Ranked) in order to

solve the mode selection problem in a timely manner. The details of the Heuristics and problem formulation can be found in Appendix A.8.

2.3.2.3 Performance assessment of D2D mode selection optimisation

Here, we use numerical simulations to evaluate the performance of our proposed heuristics. The evaluation scenario consists of a hexagonal multi-cell network with a reference cell in the middle and its first-tier neighbours (see Figure 2.19). The results reported in this paper pertain to the reference cell, and the neighbouring cells model the impact of inter-cell interference. Error bars in the results are the 95% confidence intervals. Although our approach can be tested with any scheduler, here we refer to the Proportional Fair (PF) scheme for scheduling cellular users, since it represents the state of the art for schedulers used in real implementations [49, 50]. In addition to our heuristics, we evaluate three benchmark schemes, namely, `Forced-LTE`, `Forced-WiFi`, and `Optimal`. In `Forced-LTE`, D2D users are forced to use legacy cellular communications (i.e., mode 0). In `Forced-WiFi`, D2D users are forced to communicate over WiFi (i.e., mode 3). `Optimal` results are based on the exact solution to the optimisation problem. The benchmarks allow to compare our proposals with the legacy cellular system, to measure the gain due to extra WiFi bandwidth, and to see how far the heuristics are from the optimum. The details of the simulation setting can be found in Appendix A.8.

Figures 2.20 to to 2.22 illustrate the impact of N on achievable system performance. We can observe the achievable throughput in Figure 2.20. The aggregate throughput has a negligible change with N under `Forced-LTE` because the distribution of channel qualities in the cell remains the same for different density of users, and therefore the average aggregated throughput. The throughput of the rest of schemes increases with N because there are probabilistically more D2D pairs in a denser cell, hence D2D throughput is higher. In `Forced-WiFi`, the throughput grows slowly due to the contention-based nature of WiFi, in which the MAC overhead increases with the number of contending users. Since some of the outband D2D pairs do not interfere with each other (i.e., they are more distant than 150 m), the aggregate throughput of `Forced-WiFi` in our experiments reaches up to 98 Mbps. More importantly, not only the simple proposed heuristics greatly outperform `Forced-LTE` and `Forced-WiFi`, but they also perform very close to `Optimal` (due to the computational complexity of such an ideal scheme, we only have the results up to 80 users).

In terms of energy cost, the aggregate cell power increases with N, as shown in Figure 2.21, mainly due to the baseline energy consumption of wireless interfaces. `Forced-WiFi` has higher energy consumption because outband users have to maintain two active wireless interfaces instead of one.

Figure 2.22 shows that the trend for system utility is similar to that of throughput because the throughput is the dominant factor with the current value of α. Our results show that, with a reasonable population, say 100 users per cell, the aggregate throughput gain over `Forced-LTE` is tenfold. This gain comes from both the frequency re-use of inband modes and additional spectrum provided by the outband mode. The significant contribution of both outband and inband modes to this gain highlights the importance of Floating Band D2D. Moreover, this gain can easily compensate for the infrequent D2D CSI feedbacks sent to the eNB (user-to-user CSI). Note that in LTE-A systems with millisecond feedback reporting, the CSI contributes to less than 20% of the total bandwidth.

To summarize, our proposed algorithms can highly improve the performance of the cellular networks with minimal complexity. We also evaluate our proposed algorithms in terms of fairness, delivery ratio and density of the network. These results can be found in Appendix A.8.

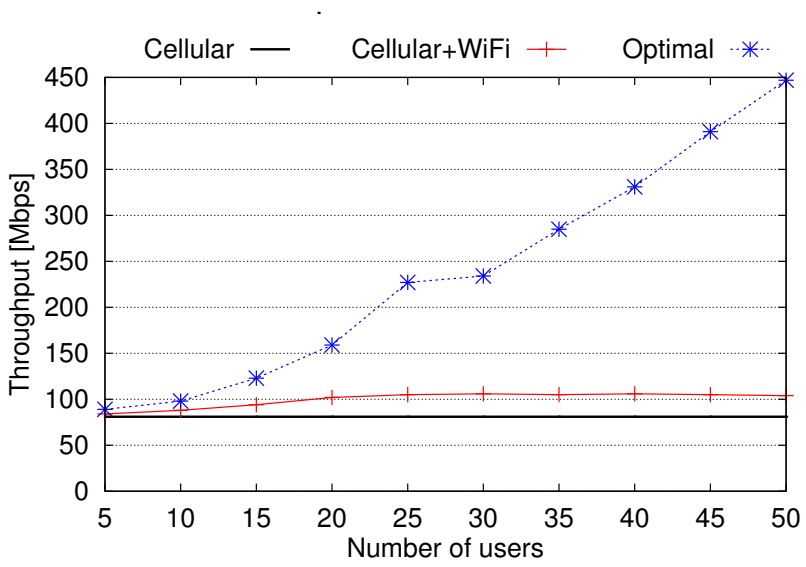

Figure 2.20: Aggregate cell throughput under different D2D mode selection schemes.

2.3.2.4 Conclusions

With our work we have shown that the performance of D2D modes is highly scenario-dependent. Thus, the most convenient mode in one scenario, say in a macro cell scenario, could be a poor choice in another, say in a micro cell scenario. To cope with this issue, we proposed the *Floating Band D2D* framework along with practical heuristics suitable for quick and adaptive mode selection in such a complex setup. Unlike existing schemes, Floating Band D2D allows D2D users to communicate over inband or outband modes, depending on network load and channel conditions. Our results demonstrate the impressive potentials of multi-band mode selection. Remarkably, our simple heuristics result in fair operation and achieve near optimal performance by dramatically ameliorating network utility, which accounts for both throughput and energy consumption. The results of this research have been published in [6, 11].

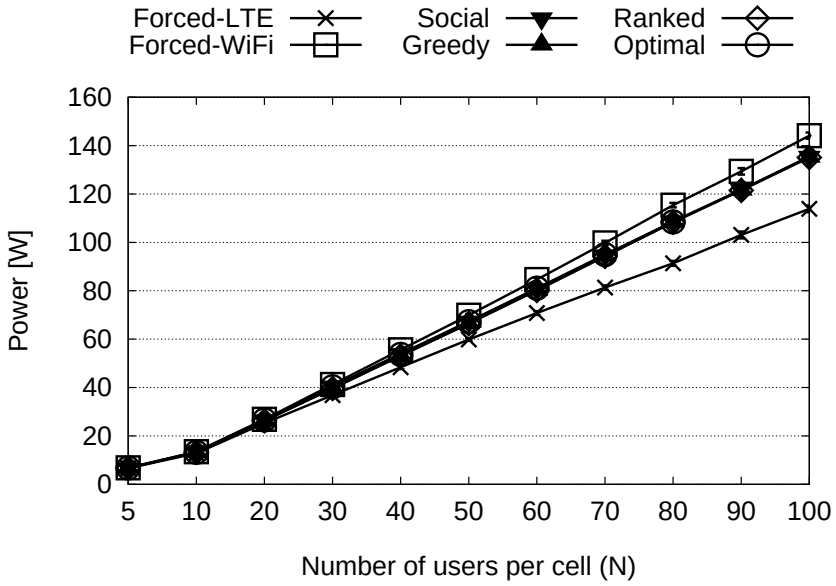

Figure 2.21: Aggregate cell power under different D2D mode selection schemes.

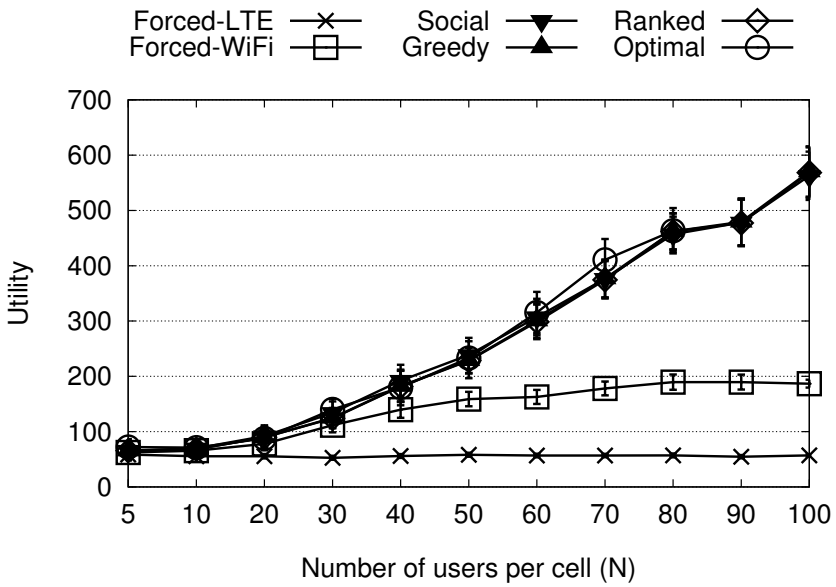

Figure 2.22: Aggregate cell utility under different D2D mode selection schemes.

3 Conclusions

In this document, we have presented the results of the research carried out in the frame of WP2 of CROWD, and we have assessed the performance of MAC and cooperative mechanisms for resource management in very dense and potentially heterogenous wireless access networks. The document has shown theoretical results and algorithms, as well as experimental results, for the characterisation and enhancement of LTE networks, IEEE802.11/WLAN networks, and HetNets with cells of heterogeneous sizes and technologies. A key result presented in the document is the assessment of a CPS modelling-based technique to evaluate the capacity region of wireless networks, which can be adapted to analyse LTE, IEEE 802.11, and D2D schemes based on such wireless technologies. Specifically, for what regards LTE networks, the document has shown that innovative user association strategies, interference control techniques and D2D-based offloading schemes are key to boost spectral efficiency and energy efficiency in the radio access part of the network. For what regards WLANs, the document has shown that distributed medium access control, based on control theory results, and D2D relay schemes leveraging the WiFi Direct paradigm, can be used to boost throughput while achieving high levels of fairness. For what concerns HetNets, the document has shown that efficiency can be optimised by means of resource allocation, user association and adaptive relay schemes using different D2D modes to offload the network traffic via either inband or outband resources. The most important factor behind the success of the proposed MAC enhancements is the availably of a control architecture that allows for short-time scale and local control decisions for the optimisation of the wireless access network. Therefore, the adoption of a multi-tier and multi-time-scale control architecture, like the one proposed in the CROWD project, is fundamental to enable the potential gains achievable through MAC enhancements. Notably, the majority of the results presented in this document have been used to produce high quality scientific publications accepted in top international venues, such as IEEE/ACM transactions and conferences.

A Appendix

A.1 Distributed ABSF mechanism for elastic traffic

In this appendix we detail the operational principles of $H_2(IC)_2$ and the game theoretical approach that leads to its design. We have formulated the ABSF/eICIC control problem in Section 2.1.2, where we have also discussed the near-optimality of $H_2(IC)_2$ as a semi-distributed heuristic to solve the aforementioned problem. Here we also present additional performance figures achieved with $H_2(IC)_2$ in a dense LTE-A network.

A.1.1 Convergence analysis of the distributed approach (Problem 2.6)

In the following, we analyse the (fully) *distributed* approach formulated above in Problem 2.6 in Section 2.1.2 from a game theoretic standpoint and show that its *convergence is not guaranteed*. Building on this result, later in Section A.1.2 we propose a *semi-distributed* approach that guarantees the convergence of the game.

Based on game theory, the distributed approach can be modeled as a game where base stations iteratively play in order to maximise their utility. Let us define this game as an *Interference Coordination Game* Γ, where each base station i acts as a player (the terms "player" and "base station" are indistinctly used in the rest of the paper). The set of strategies of each player $\mathbb{S}_i$ consists in the set of pairs (user, TTI), $(u,t) : u \in \mathcal{U}_i, t \in \mathcal{T}$, available for each base station according to constraints in Problem `LOCAL`.

In order to analyse the convergence of the above game, we rely on the concept of *Bottleneck Matroid Congestion Game* (for a detailed discussion, we refer the reader to [51]). A *Bottleneck Congestion Game* is a class of games where resources are shared among players. The utility of each player depends on the utility of the resources she chooses and the number of players choosing the same resources: the higher the congestion, the lower the utility. In particular, the individual player utility is the minimum of the utilities of the resources chosen in her strategy. In sequential improvement dynamics, players act selfishly and play a Best Response strategy (BR) $\mathcal{S}_i^* \in \mathbb{S}_i$, i.e., the strategy that maximises their individual utility function, given the strategies played by other players.

In addition to the above, regular congestion games can be generalized in *player-specific congestion games* and *weighted congestion games*. In the former, every player has her own utility function for every resource. In a weighted congestion game, every player affects the other players strategies with a different weight, namely, she causes a different level of congestion.

The following theorem shows that our game falls in the intersection between the above categories, and hence existing results on these classes of games can be applied to our problem.

Theorem 1. *The Interference Coordination Game* Γ *is a Weighted Player-specific Bottleneck Matroid Congestion Game.*

Proof. Here we provide the reader with a sketch of the proof. The Interference Coordination Game Γ is player-specific since utility is player-specific as it depends on received interference, and it is a congestion game in which congestion weights are given by the interference caused by the scheduled users in each TTI. Moreover, strategies' constraints induced by constraints in Problem `LOCAL` make the strategy space a matroid, thus Γ is a Matroid Congestion Game. □

Table A.1: Example of weighted player-specific matroid bottleneck congestion game that does not converge

Rate	Alone	With BS 1	With BS 2	With BS 3
$c_{u_1,t}$	2.0	–	1.5	**1.1**
$c_{u_2,t}$	2.0	**1.1**	–	1.5
$c_{u_3,t}$	2.0	1.5	**1.1**	–

Regular bottleneck congestion games have been proven to satisfy the *finite improvement property*, which states that an arbitrary BR sequence played by each player during the game always converges to an equilibrium in a finite number of steps [51]. However, the generalizations of player-specificity and different congestion weights introduce many degrees of freedom, which weakens the game structure and its convergence guarantees. Indeed, the following theorem shows that Weighted Player-specific Bottleneck Matroid Congestion Games do not satisfy the finite improvement property.

Theorem 2. *Weighted player-specific matroid bottleneck congestion games do not exhibit the finite improvement property in best-response improvement dynamics.*

Proof. Let us consider a scenario with $T = 2$ TTIs and 3 base stations, each of them associated with $|\mathcal{U}_i| = 1$ distinct user. For each player i, the strategy space $\mathbb{S}_i$ is defined as $\mathbb{S}_i = \{\{(u_i, t_1)\}; \{(u_i, t_2)\}; \{(u_i, t_1), (u_i, t_2)\}\}$. Let us assume an upper bound on available TTIs per base station $M_i = 1, \forall i \in N$ and a user rate $c_{u,t}$, expressed as bits/symb/TTI, according to Table A.1. Now we consider the sequence of strategies taken by each player, described by Table A.2.

Table A.2: State evolution for a weighted player-specific matroid bottleneck congestion game that does not converge (example used in the proof of Theorem 2)

$S_i^{*(s)}$, $f(S_i^*, S_{-i})$	$s = k$	$s = k+1$	$s = k+2$	$s = k+3$	$s = k+4$	$s = k+5$	$s = k+6$
BS 1	$\{u_1, t_1\}$, *(1.1)*	$\mathbf{\{u_1, t_2\}, 1.5}$	$\{u_1, t_2\}, 2$	$\{u_1, t_2\}$, *(1.1)*	$\mathbf{\{u_1, t_1\}, 1.5}$	$\{u_1, t_1\}, 2$	$\{u_1, t_1\}$, *(1.1)*
BS 2	$\{u_2, t_2\}, 2$	$\{u_2, t_2\}$, *(1.1)*	$\mathbf{\{u_2, t_1\}, 1.5}$	$\{u_2, t_1\}, 2$	$\{u_2, t_1\}$, *(1.1)*	$\mathbf{\{u_2, t_2\}, 1.5}$	$\{u_2, t_2\}, 2$
BS 3	$\mathbf{\{u_3, t_1\}, 1.5}$	$\{u_3, t_1\}, 2$	$\{u_3, t_1\}$, *(1.1)*	$\mathbf{\{u_3, t_2\}, 1.5}$	$\{u_3, t_2\}, 2$	$\{u_3, t_2\}$, *(1.1)*	$\mathbf{\{u_3, t_1\}, 1.5}$

Whenever a player i chooses a new strategy at the k^{th} step in order to maximise the utility function (bold-marked), the value of utility function calculated by the other players may decrease and they may want to change their strategy. This leads to a loop where players sequentially return on the same strategies indefinitely, such as strategies at step k and strategies at step $k + 6$. Hence, players playing arbitrary best responses do not necessarily converge to a Nash equilibrium in Weighted Player-specific Bottleneck Matroid Congestion Games, and thus, a finite improvement property does not always exist. □

The above analysis has shown that the distributed approach may not converge.[1] Moreover, it does not ensure that M_i values are selected according to a global fairness. In order to address these shortcomings, in the next section we propose a semi-distributed two-level mechanism where a central coordinator controls the behavior of the distributed game.

[1]It is worthwhile noting that the somehow pathological scheduling behavior considered in theorem's proof does not commonly exhibit in networks; indeed, according to the simulations conducted for typical realistic scenarios, the interference coordination game Γ reaches an equilibrium with very high probability. Nevertheless, we still need to design an algorithm whose convergence is guaranteed.

A.1.2 $H_2(IC)_2$: a controller-aided distributed mechanism for ABSF

In the following we present the design principles of $H_2(IC)_2$, a *semi-distributed approach* that relies on a central coordinator to implement a distributed ABSF pattern selection. The global scheme of the system mechanism is depicted in Figure 2.5. As shown in the figure, the scheme operates at two different timescales:

- On a long-term timescale (in the order of seconds), a central coordinator, the *Controller*, is in charge of adjusting the M_i value of each base station, where M_i gives the maximum number of TTIs that base station i can use to schedule its users within the time horizon T by solving Problem `LOCAL`. In addition, adapting M_i is used to react to traffic changes in the system.
- At a shorter timescale, base stations play the Interference Coordination Game Γ by sequentially exchanging their scheduling decisions in terms of ABSF patterns.[2] As described in the following, the central coordinator does not directly participate in the game, but it controls its convergence by limiting the number of iterations.

The remaining challenge for coordinator-aided approach is the design of the algorithms executed by the central coordinator to (i) ensure convergence, and (ii) adjust the values M_i. In the following we address the design of those algorithms, which aim at driving the system behavior to an optimal state in the long run.

A.1.2.1 Convergence control of game Γ

In order to guarantee the convergence of the game Γ of Theorem 1, the central coordinator imposes a deadline of Z TTIs, with $Z < T$: if the game has not finished by this deadline, it is terminated by the central coordinator.

When the game finishes before the deadline, the resulting scheduling corresponds to an equilibrium of the game, which ensures that resources are fairly shared among base stations. In contrast, when the game is terminated by the central coordinator, base stations use the scheduling that they computed in the latest iteration of the game, which does not correspond to an equilibrium. Thus, in the latter case some base stations could potentially have a better scheduling (i.e., more resources) than the others. However, as shown by our results of Appendix A.1.4, we have observed that in practice the game can be interrupted after only a very few iterations without negatively impacting fairness or performance in a significant manner.

The deadline Z has been chosen in order to have a valid scheduling before the current period T finishes: the resulting scheduling (and the corresponding ABSF pattern) will then be used for the next period. During the game, transmissions and users are scheduled according to the result of the previous period. Note that the iterations of game Γ do not need to be synchronized with the TTIs; they can be much faster, allowing for more than Z iterations within Z TTIs. Indeed, the execution of one iteration only requires passing the "current" ABSF patterns from one base station to another. As shown in Section A.1.4.2, Z can be chosen in the range $[|\mathcal{N}|, |\mathcal{N}|^2]$.

A.1.2.2 Dynamic adjustment of TTI bounds M_i

One critical aspect for the performance of the proposed mechanism is the setting of the M_i parameters, which give the maximum number of non-blank TTIs available to each base station. Indeed, if the M_i values are too small, performance is degraded because, even if base stations can be scheduled one at a time with low interference, the number of TTIs available for transmitting can be too

[2] Note that there is no need to announce which specific user will be scheduled in a specific TTI, since base stations transmit at a fixed power and thus their activity causes the same level of interference independently of the scheduled user. Therefore, it is sufficient to propagate a binary string of T bits containing the ABSF pattern.

small to accomodate all users. Conversely, if the M_i values are too large, performance is degraded as a result of too many base stations scheduled together and creating high interference. Thus, performance is maximised when the M_i parameters are optimally set to values that are neither too large nor too small. In the rest of this section, we design an adaptive algorithm that follows an Additive Increase Multiplicative Decrease (AIMD) strategy [52] to find the optimal M_i setting.

In addition to optimally setting M_i to improve the performance of the network, the adaptive algorithm also aims at dynamically adjusting the M_i configuration to follow the changes in traffic and interference. From this perspective, the adaptive algorithm is a long-term process. In contrast, the distributed game is a short-term process played once per each period of T TTIs. This implies that the duration of the period T cannot exceed a few hundreds frames, which corresponds to a few seconds during which traffic and channel conditions remain practically unchanged.

From a high level perspective, the algorithm works as follows. At the end of each period of T TTIs, the controller gathers from the base stations the performance resulting from the M_i values (and the corresponding ABSF pattern) used during the period. The metric chosen to represent the performance of a base station is given by the average user rate experienced by users of base station i in the period[3], i.e.:

$$\eta_i = \frac{1}{|\mathcal{U}_i|} \sum_{(u,t)\in\mathcal{U}_i\times\mathcal{T}} c_{u,t}. \tag{A.1}$$

The controller then uses the sum of the individual performance metrics, $\eta = \sum_{i\in N} \eta_i$, to keep track of the global system performance and drive M_i to the setting that maximises η. The algorithm to find such M_i setting follows an AIMD strategy: the M_i values are increased as long as performance is improved, and, when performance stops improving, then the M_i values are decreased. After each update of the M_i values, these are distributed to the base stations and used in the following period (i.e., the following iteration of game Γ).

The specific algorithm executed to calculate the new set of TTI bounds M_i is described in Algorithm 1. Each iteration of the algorithm is identified by an index k. At the initial step ($k = 0$), the controller initializes the system performance metrics η to 0 and assigns the initial TTI bounds $M_i^* = \lceil T/|\mathcal{N}| \rceil$ for every base station. This initial M_i^* setting has been chosen to allow base stations to schedule their users in disjoint portions of the period, which helps the convergence of the algorithm in case of very high mutual interference between all base stations. The M_i^* also provide a lower bound for M_i.

At each step, the controller collects the performance metrics η_i from base stations and checks whether the performance of this period, $\eta^{(k)}$, has improved with respect to the previous period, $\eta^{(k-1)}$ (line 3). If this is the case, this means that system performance is raising and the controller increases TTI bounds M_i as follows. The controller increases by 1 unit the M_i of the base station with the smallest η_i whose M_i is below T (lines 8-9). Once one M_i value is increased, step k of the algorithm terminates (line 10).

If no M_i can be increased, which means that all base stations are active in all TTIs, then no adjustment of the M_i values is made as long as the system performance does not degrade. In case performance degrades, i.e., $\eta^{(k)}$ decreases, (line 13), the controller drastically reduces the M_i. Specifically, the controller looks at the base station i with the largest η_i whose M_i is above M_i^*. It sets the new M_i value of this station equal to the minimum between the half of the current M_i value and the lower bound M_i^* (lines 17-18). If $M_i = M_i^*$ for all i, no change is carried out.

The rationale behind using AIMD to adjust the M_i values is that, similar to what happens with TCP, increasing the utilization of the system (i.e., increasing M_i values) may lead to congestion (in our case, this corresponds to excessive interference), which causes user rates to drop. In this case,

[3]Note that, since user allocation is carried out according to Problem `LOCAL`, the max-min objective tends to assign rates with limited variance; as a consequence, the average user rate and the rate of the worst-off user are likely to be similar.

Algorithm 1 Resource Sharing Algorithm: Adaptive algorithm to dynamically design M_i. Called at the end of $(k-1)^{th}$ ABSF pattern

Input: $\mathcal{N}, T, M_i^*, \eta^{(k-1)}$
Initialization: $\eta^{(k)} \leftarrow 0; M_i \leftarrow M_i^*, \forall i \in \mathcal{N}$
Procedure

$\mathcal{V} \leftarrow \{\eta_i, \forall i \in \mathcal{N}\}$
Order $\mathcal{V}$ non-increasing
$\eta^{(k)} = \sum_{i \in N} \eta_i$
if $\eta^{(k)} > \eta^{(k-1)}$ **then**
  while $\mathcal{V} \neq \emptyset$ **do**
    e = pop($\mathcal{V}$)
    Consider index i of element e
    if $M_i^{(k-1)} < T$ **then**
      $M_i^{(k)} = M_i^{(k-1)} + 1$
      break
    end if
  end while
else
  while $\mathcal{V} \neq \emptyset$ **do**
    e = pop($\mathcal{V}$)
    Consider index i of element e
    if $M_i^{(k-1)} > M_i^*$ **then**
      $M_i^{(k)} = \max\left\{M_i^*; \left\lceil M_i^{(k-1)}/2 \right\rceil\right\}$
      $\eta^{(k)} = 0$
      break
    end if
  end while
end if

a quick reaction is required by the controller to drive the system to a safe point of operation, by properly adjusting TTI bounds M_i. Also similar to TCP, the additive increase of TTI bounds M_i allows to gracefully approach the optimal utilization of the system. Furthermore, since the problem may admit more than one local maximum, using multiplicative decrease for the TTI bounds M_i helps our heuristic to escape from a local maximum where the optimisation function may be trapped in.

As a side comment, we point out that the proposed algorithm could accommodate different goals, such as, e.g., maximum throughput or proportional fairness, by simply replacing the function that gives the global system performance, η, by another function that reflects performance according to the objective pursued.

A.1.3 Control overhead for $H_2(IC)_2$

We conclude the analysis with the evaluation of the control overhead introduced by $H_2(IC)_2$. To this aim, we identify two different *interfaces*: one between central coordinator and base stations, namely I_C, and one between distinct base stations, namely I_B. They may be both implemented using, e.g., the LTE $X2$ interface.

In the centralized solution, the central coordinator requires message exchanges over I_C only. In

Table A.3: Overhead of centralized and $H_2(IC)_2$ semi-distributed approaches

Interface	centralized approach	$H_2(IC)_2$ approach
I_C	$64 \cdot \lvert U \rvert \cdot \lvert \mathcal{N} \rvert + T \cdot \lvert \mathcal{N} \rvert$	$64 \cdot \lvert \mathcal{N} \rvert$
I_B	0	$T \cdot k \cdot \lvert \mathcal{N} \rvert$

particular, per each pair (*user, base station*), it requires the transmission of an average channel quality indicator (e.g., the *RSRP* value in the LTE-Advanced networks [19]) which can be encoded in double precision floating point format, e.g., 64 bits. Then, the controller issues a scheduling pattern (a string of T bits) per each base station.

In the $H_2(IC)_2$ mechanism, the controller requires to receive the average user rate η_i per base station over I_C at the end of each game Γ, consisting in a binary string of fixed length (e.g., 64 bits for a double precision floating point number). Regarding the interface I_B between different base stations, $H_2(IC)_2$ needs a sequential exchange of ABSF scheduling patterns (strings of T bits) during the interference coordination game Γ, until the game reaches a convergence state or the convergence deadline expires.

We can therefore summarize the total load in terms of bits for each interface as reported in Table A.3. In the table, k is the number of rounds the interference coordination game plays before reaching the convergence, and $|\mathcal{U}| = \sum_i |\mathcal{U}_i|$ is the total number of users in the system. We can easily observe that the overhead of $H_2(IC)_2$ is lower than that of the centralized mechanisms when the following inequality holds:

$$|U| > 1 + \frac{T}{64}(k-1) \cong \frac{T|\mathcal{N}|^2}{64}, \tag{A.2}$$

where we have considered that the number of rounds k in the worst case is a function of $|\mathcal{N}|$ (i.e., at most $k = |\mathcal{N}|^2$ iterations are enough to converge, when convergence exists, as proven mathematically in [51] and empirically shown in Appendix A.1.4) and both T and $|\mathcal{N}|$ are (much) greater than 1. Therefore, our semi-distributed approach is convenient as soon as the number of users exceeds a threshold that depends on T and $|\mathcal{N}|$ (i.e., the threshold is $O\left(T|\mathcal{N}|^2\right)$). For example, in an (sub-)urban environment with $T = 70$ and $|\mathcal{N}| = 7$, as in our simulations described later, $H_2(IC)_2$ results convenient with as few as 54 users or more, while in a dense-urban environment with $|\mathcal{N}| = 30$, our approach exhibits a practical implementation starting with $\sim$1000 users in the entire network. Those values are pretty low, revealing how our semi-distributed approach drastically reduces the signaling overhead for existing cellular network size.

A.1.4 Performance Evaluation of $H_2(IC)_2$

In this section of the appendix, we use numerical simulations to show that our proposal performs near optimally and boosts achievable rates in the whole network, not just for topologically disadvantaged users. All simulations are carried out by means of MATLAB with all parameters summarized in Table 2.2. The average quality of the user channel is computed as function of the distance from the base station (according to the propagation model suggested by 3GPP specifications, Table A.2.1.1-3 of TR.25.814 v7.1.0), and Rayleigh fading is considered. Based on user channel qualities, each simulated base station solves the local optimisation problem by means of a remote call to a commercial solver, i.e., IBM CPLEX OPL. Additionally, we show that game Γ quickly approaches its Nash equilibrium, which enables $H_2(IC)_2$ to easily follow changes occurring dynamically in the network.

A.1.4.1 Benchmarks

We benchmark $H_2(IC)_2$ against the optimal solution, obtained by solving Problem `CENTRAL` by means of an ILP solver. Additionally, we compare $H_2(IC)_2$ to the case of uncontrolled base stations using the same frequencies (No ICIC) and to a traditional frequency reuse 3 scheme, in which the available band is split into three orthogonal sub-bands. For the sake of completeness, we also compare $H_2(IC)_2$ with two existing approaches fully based on a power control schemes, showing how $H_2(IC)_2$ can achieve high network performance at a bargain price of complexity. In the first scheme, namely Utility-Based Power Control (UBPC) [53], base stations are allocated in all available TTIs by tuning properly the transmitted power to reduce interference. The algorithm suggested in [53] maximises the user net utility by ensuring that the signal-to-noise-ratio of each transmission is greater than a minimum threshold γ_i (in our simulations we assume γ_i as the minimum MCS with nonzero rate). However, UBPC allows for multiple transmissions to different users in the same cell, which is not doable in schedule-based cellular networks. Therefore, to force the scheduling of a single user per cell on a per-TTI basis, we simply modify the original UBPC algorithm by setting the interference to infinite when two or more users from the same cell are scheduled. We refer the reader to [53] for more details on UBPC. While UBPC provides a rigorous centralized solution for the power allocation problem at the expense of a huge amount of information exchanged, a second power control scheme recently developed, namely REFerence based Interference Management (REFIM) [54], proposes a low-complex distributed scheme by exploiting the notion of reference user. The authors of [54] aim at simplifying the analysis of the impact of neighbouring cells by replacing all of them with a single virtual user, selected as the user with the worst channel condition belonging to the surrounding cells. This abstraction leads to a drastic reduction of the control signal overhead resulting in a practical implementation of the power control solution, which exhibits a conservative behaviour.

We have shown in Section 2.1.2 the dynamic behaviour of $H_2(IC)_2$. Figure 2.7 in that section shows that it takes a few seconds for $H_2(IC)_2$ to reach its stable operating point, after which it follows quite fast the evolution of channel and traffic conditions. In particular, at time $t = 16$ s, the number of users in the network doubles abruptly, but it takes only a fraction of a second for $H_2(IC)_2$ to adapt. In general, $H_2(IC)_2$ largely outperforms the No ICIC scheme and achieves significant gain over frequency reuse 3. Indeed, $H_2(IC)_2$ halves the distance between the optimal performance and the one of frequency reuse 3. Notably, after the initial adaptation period, the utility achieved by $H_2(IC)_2$ lies within 85% and 90% of the one achieved with the optimal solution for Problem `CENTRAL`. REFIM and UBPC results show the real potentials of power control schemes. UBPC can even go slightly beyond the performance of the optimal solution without power control, although it requires higher complexity in terms both of execution and device hardware. REFIM, notwithstanding a low-complexity scheme, shows lower performance with respect to frequency reuse 3 for a particular set of user populations due to the conservative assumption taken on interfering cells. Therefore, our approach $H_2(IC)_2$ perfectly lies in between an impermissible efficient power control scheme and a practical doable distributed power control solution.

Besides utility η, we want to evaluate the fairness achieved by the different schemes. To this aim, Figure A.1 presents the CDF of achieved user rates (averaged over the time horizon T). The figure clearly shows that the optimal solution, $H_2(IC)_2$ and UBPC behave similarly and exhibit two main advantages: (i) they achieve user rates in a compact interval of possible values (which is symptom of fairness according to Jain fairness definition), and (ii) with high probability, they guarantee a minimum rate which is several times higher than the one guaranteed by No ICIC or frequency reuse 3 (which is symptom of max-min fairness). In addition, REFIM shows a similar behaviour to the optimal solution in terms of fairness, even though its curve stays on the left side of the graph due to the critical user rates experienced by the users. For instance, with 95% probability, $H_2(IC)_2$ guarantees 3.7 Mbps per user, REFIM guarantees 2.4 Mbps per user, while No ICIC

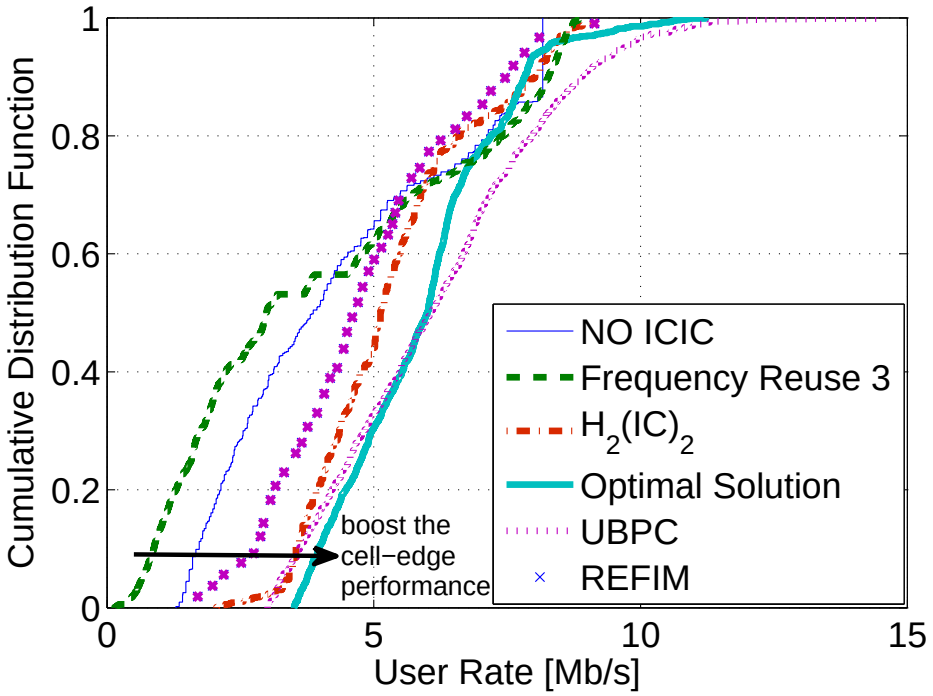

Figure A.1: CDF of average user rates with 7 base stations and 10 users per base station. The time horizon is set to $T = 70$ TTIs.

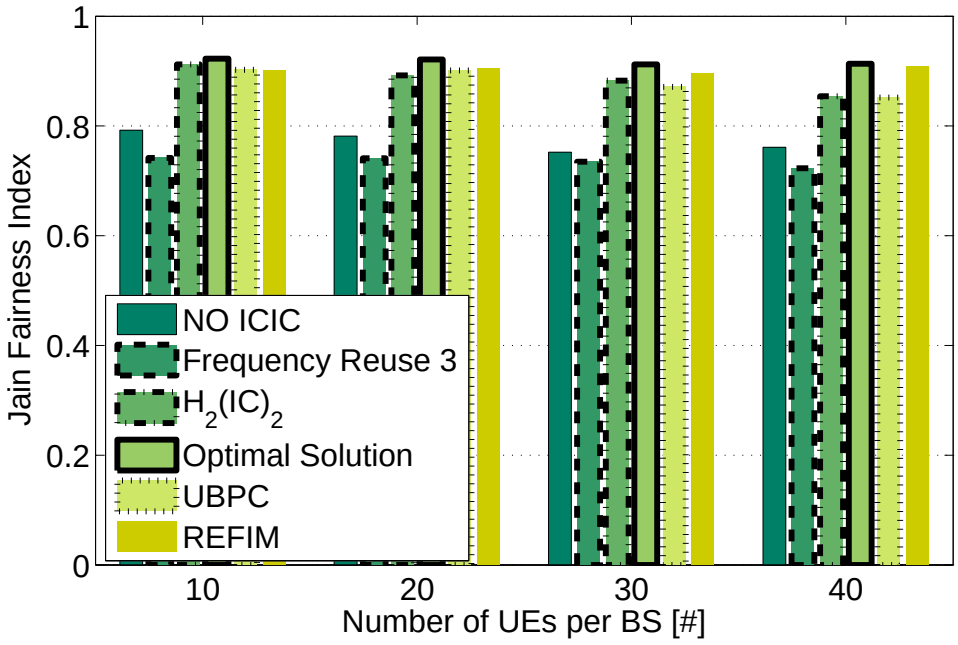

Figure A.2: Jain fairness indexes achieved with 7 base stations and a variable number of users per base station.

only guarantees 0.9 Mbps. Jain fairness achieved under the different schemes under evaluation is depicted in Figure A.2 as a function of the number of users per base station. As expected, REFIM presents a stable behaviour over different values of user population, due to its strong correlation with worst users. Also in this case, $H_2(IC)_2$ achieves near-optimal results, and it outperforms UBPC.

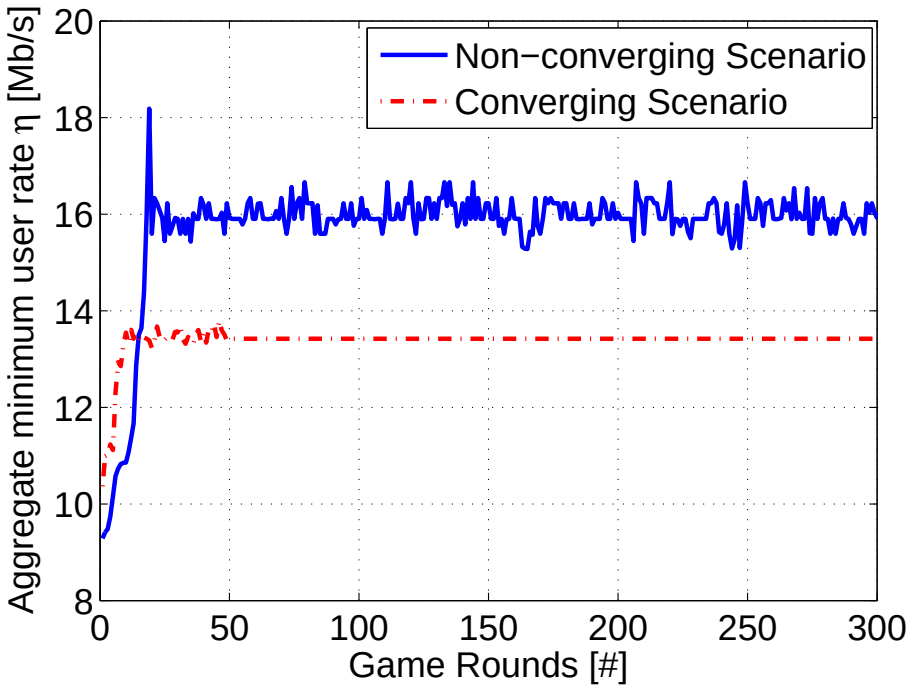

Figure A.3: Game convergence behavior considering two different cases with $|\mathcal{N}| = 7$ base stations.

A.1.4.2 Game convergence speed

A key feature of $H_2(IC)_2$ is its ability to adapt quickly to network changes. Such a feature relies on quick ABSF pattern computation, which follows the rule of the Interference Coordination Game Γ. The game evolves over time as illustrated in Figure A.3. In this example two different cases are considered: the dashed line represents a case of convergence, while the solid line is for a rare case in which the game does not converge to a Nash equilibrium point. In both cases, 7 base stations are considered, and the TTI bounds M_i are fixed.

In case of convergence, which occurs in about $|\mathcal{N}|^2$ rounds, it is clear that a few game rounds suffice to approximate the performance achieved at the Nash equilibrium with an error smaller than 3%. Notably, also in case the game fails to converge, after a few rounds the utility starts fluctuating around a stable value, with small oscillations (about $\pm 5\%$).

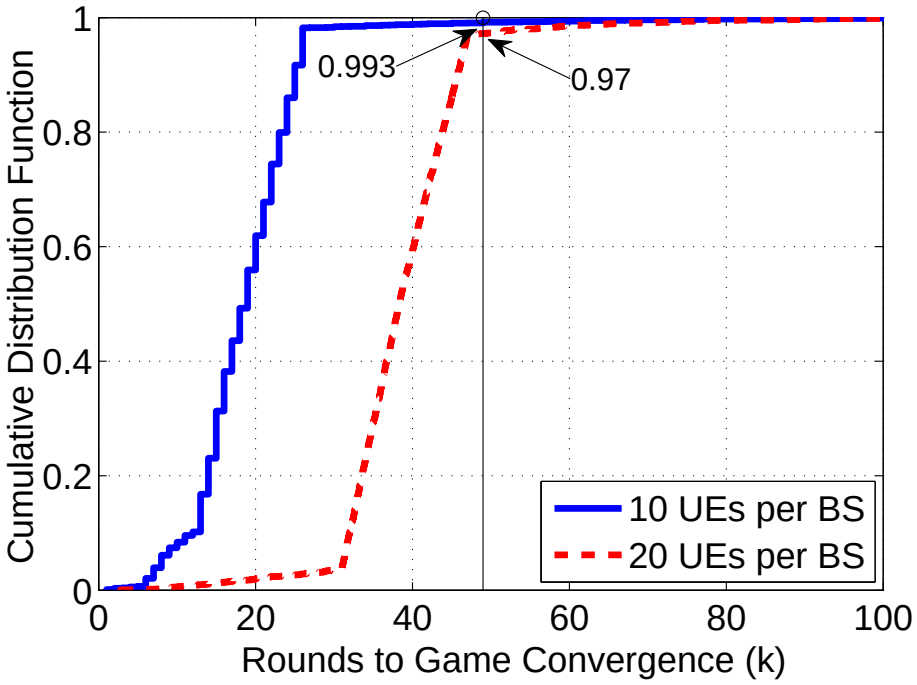

Figure A.4: CDF of number of rounds needed for game convergence with 7 base stations and different user populations.

Figure A.4 illustrates the CDF of the number of rounds needed to converge for a few different user populations (note that we use the value 1 to indicate that the game did not converge). The figure shows how the majority of the games converge much before $|\mathcal{N}|^2$ rounds (vertical line in the figure), and very few cases do not converge at all. We have observed very similar behavior for the majority of the cases analysed in our experiments, so we conclude that reasonably high utilities can be achieved by stopping the game after a number of rounds comprised between $|\mathcal{N}|$ and $|\mathcal{N}|^2$.

Overall, our results show that $H_2(IC)_2$ not only achieves near-optimal results according to the definition of utility given in the formulation of Problem `CENTRAL`, but also achieves high levels of max-min and Jain fairness, and significantly boosts average rates in the entire cellular network.

A.2 A CPS model for wireless networks

In this appendix we present the theory behind the CPS-based modelling we have used to characterise the stability region and the key factors used to evaluate the performance of wireless networks in Section 2.1.3 and in Section 2.2.3.

A.2.1 Model and Assumptions

A.2.1.1 System Model

We consider a system of N queues, where each queue receives traffic from one or more fresh sources (i.e., residing out of the system). We assume such queues are served by work conserving schedulers, and that the served traffic leaves the system. The system *state* at time t is the array $\bar{a}(t)$, whose i-th element is equal to 0 if the $i-$th queue is empty, and 1 otherwise. We assume the service rate $R_i(t)$ of queue $i \in \{1, ..., N\}$ at time t is determined only by the system state, i.e., $R_i(t) = R_i(\bar{a}(t))$. We call such a system a CPS.

We consider *monotonic decreasing* CPSs, i.e., if $\bar{a}_1$ and $\bar{a}_2$ are two different states such that $\bar{a}_1 \leq \bar{a}_2$, [4] then $\forall i \in \{1, ..., N\}$, $R_i(\bar{a}_1) \geq R_i(\bar{a}_2)$. This class of CPSs includes many problems of practical importance and it has been widely studied in the past, in the context of wireless networks and of bandwidth sharing in packet networks [25]. In [27] we also briefly discuss how the presented results can be extended to non-monotonic CPS. Without loss of generality, we assume arrivals to be *packetized*, with a finite number of packet sizes. Finally, we consider that no losses occur in the system at queues (i.e., queues have infinite buffer capacity).

A.2.1.2 Basic Concepts and Definitions

In this section we introduce some definitions and basic results useful for our CPS analysis.

Network Calculus (NC) is a min-plus system theory for deterministic performance analysis of a queuing system. It provides tools for the derivation of bounds to backlog and packet delay in a network. Traffic in NC is typically characterised by means of *arrival curves* [35]. Let $\mathcal{F}$ represent the set of nonnegative wide-sense increasing functions, and $\forall t \geq 0$ let $A(t)$ be the cumulative traffic arrival function for the time interval $[0, t]$. Then $\alpha \in \mathcal{F}$ is an *arrival curve* for the considered flow if for any $[t_1,\ t_2]$, $A(t_2) - A(t_1) \leq \alpha(t_2 - t_1)$. As such a traffic characterisation translates into quite loose assumptions on traffic statistics (mainly involving tail probabilities), it applies to a large spectrum of practical settings. One of the most common classes of arrival curves is the *leaky bucket* arrival curve, where $\alpha(t) = \rho t + \sigma$. The non-negative parameters ρ and σ are the *leaky bucket rate* and the *burstiness*, respectively.

Let us now introduce the following notion of stability [35]. Let $q_i(t)$ indicate the backlog of queue i at time t.

Definition A.2.1 (Stability)**.** *A system of N queues is deterministically stable if $\forall i \in [1, ..., N]$ it exists a $\Gamma_i < \infty$ such that*

$$sup_{t \geq 0}\ q_i(t) \leq \Gamma_i. \tag{A.3}$$

Practical sufficient conditions for stability typically imply some constraints on source traffic (e.g., on arrival statistics, or on their arrival curve parameters) and/or on the network (i.e., on some form of service guarantees at queues). For instance, for queues in isolation, a sufficient condition for deterministic stability is that at each queue the leaky bucket rate for the arrivals is less than the minimum service rate [55].

[4] Note that for binary vectors $\bar{a}_1$ and $\bar{a}_2$, the *component-wise* inequality $\bar{a}_1 \leq \bar{a}_2$ implies that the set of queues active in state $\bar{a}_1$ is a subset of the set of queues active in state $\bar{a}_2$.

Definition A.2.2 (Continuous Data Scaling Block)**.** *For any time interval* $[t_1, t_2]$*, with* $t_2 \geq t_1 \geq 0$*, let* $A(t_2 - t_1)$ *be the amount of bits arrived at a node in the time interval. The node is a continuous data scaling block, with scaling value* $S \in \mathbb{R}^+$*, if the amount of bits at its output during the same time interval is* $SA(t_2 - t_1)$*.*

Scaling blocks have been introduced in NC in order to model transformation processes which alter the total amount of traffic (lossy channels, data processing, encoding/decoding, discard of non-conformant traffic).

Definition A.2.3 (Policer)**.** *A policer with rate* Q *is a processing node that, for any arbitrary input traffic, forces* Q *as the maximum instantaneous departure rate at its output.*

In what follows we consider *unbuffered* policers, which discard non-conformant traffic.

Finally, we recall the concept of Generalized Processor System (GPS) node [56]. In general, a GPS node is composed by a server having a fixed service rate R, and M different queues. Each queue is characterised by a weight w_i. At any given time, R is split among the non-empty queues proportionally to their weights.

A.2.2 A new approach to CPS Analysis

One of the main reasons that accounts for the complexity in studying the performance of CPSs is the fact that correlations between service rates are mutual, which implies circular dependencies. To make the problem tractable, we propose a method to break circular dependencies in the analysis. Specifically, we suggest to order the CPS queues and to model the service rate impairment caused by each queue on the queues following in the ordered list, starting from the top of the list (position 1). With this, we model inter-queue dependencies in one direction (top-down in the list). To model the dependencies in the other direction (bottom-up in the list) without incurring in complex calculations, we consider that a queue in position j in the list is considered as always active by all the queues listed in position $k < j$. This is clearly a worst case approach, which will lead to the identification of performance bounds. Moreover, all possible orderings needs to be considered. In the following we describe the network abstraction that allows to study analytically the performance of a CPS according to the above described methodology (Section A.2.2.1) and derive the conditions under which such network abstraction yields performance bounds for a CPS (Section A.2.2.2).

A.2.2.1 Feed-forward networks

Our method is based on the derivation, from a given CPS, of a set of feed-forward networks, i.e., loop-free networks of queues. In such feed-forward networks, dependencies between queues, while still present as in the original CPS, are not mutual. Indeed, the feed-forward structure allows to analyse the network by stages: the activity of a given node, representing a CPS queue, affects only the service rates of those CPS queues which are mapped in following stages of the feed-forward network. Such *rate impairment* is modeled with an exchange of traffic from the affecting queue to the affected one(s). The effect on queues at preceding stages is modeled instead through a constant penalty on service rates at those stages. This structure allows to apply standard methods for performance analysis, such as classical queuing theory or basic network calculus results.

The structure of the feed-forward networks we propose is the following. Each of such networks has N stages and a two-queue work-conserving GPS node at each stage. GPS nodes are in a one-to one mapping relation with queues of the CPS. Specifically, let $j \in \{1, ..., N\}$ be the label of the j-th stage of a feed-forward network, as well as of the GPS node in it, and let us label each of the queues of the CPS from 1 to N. Let $\bar{n} = (n_1, ..., n_j, ..., n_N)$ indicate one of the $N!$ possible permutations of the labels of the CPS queues. To each permutation $\bar{n}$ it corresponds a specific mapping which

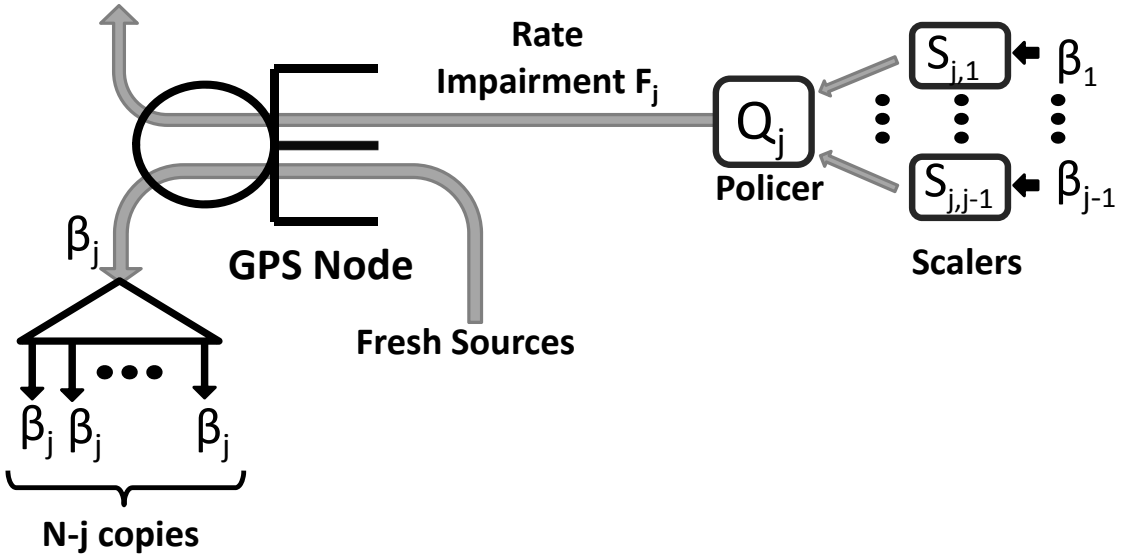

Figure A.5: Structure of the j-th stage of a feed forward network associated to a N-queue CPS.

associates the j-th GPS node to the n_j-th queue of the CPS.

Figure A.5 describes the general structure of the j-th stage of such feed-forward networks, whereas Figure A.6 shows an example of a three-queue CPS and one of the possible feed-forward networks associated to it. For stage j, each traffic flow β_k coming from one of the $j-1$ previous stages is fed to a dedicated scaling block, with scaling coefficient $S_{j,k}$. The aggregate output of all the $j-1$ scaling blocks is then fed to a policer with rate $Q_j = (R^{up}_{n_j} - R^{sat}_{n_j})$, where $R^{up}_{n_j}$ is the service rate of queue n_j of the CPS when the active queues at the CPS are $n_j, n_{j+1}, ..., n_N$, and $R^{sat}_{n_j}$ when all queues of the CPS are active.

The output of the policer is finally fed to a dedicated queue of the GPS. The other queue of the GPS is dedicated to traffic from fresh sources. At any time t, we assume arrivals from fresh sources at stage j are the same as at the corresponding queue n_j at the CPS. The total capacity of the GPS node is $R^{up}_{n_j}$, which means that the serving rate of the GPS in j-th stage of the feed-forward network is the minimum possible rate with respect to the state of queues $n_{j+1}, ..., n_N$ at the CPS. The GPS weights are $w = R^{sat}_{n_j}/R^{up}_{n_j}$ for fresh traffic, and $1-w$ for traffic from the policer. At the output of the GPS node, traffic coming from stages $1, ..., j-1$ exits the network. The remaining traffic is fed to a block which produces $N-j$ exact replicas of the same traffic, introducing no delay. Each replica is fed to one of the following stages.

Let us indicate with $F_j(t)$ the *rate impairment* at node j. It is the output of the policer at stage j, and models the effect of activity at nodes in stages $1, ..., j-1$ on the fresh traffic service rate at node j. Due to the network structure, we have:

$$F_j(t) = \min\left(\sum_{k=1}^{j-1} \beta_k(\bar{b}(t)) S_{j,k}, R^{up}_{n_j} - R^{sat}_{n_j}\right). \tag{A.4}$$

It is easy to prove that the rate impairment traffic depends only on the activity of the N queues dedicated to fresh traffic. $\bar{b}(t)$ is a binary vector which j-th element is zero if at time t the queue for the fresh traffic at the j-th GPS node of the feed-forward network is empty, one otherwise. The impairment traffic is therefore modeled at the fluid limit as $\beta_k(\bar{b}(t)) = R^{up}_{n_k} - F_k(\bar{b}(t))\ \forall k \in 1, ..., j-1$.

It is clear that the proposed feed-forward networks are also networks of coupled queues, where coupling translates into rate impairments.

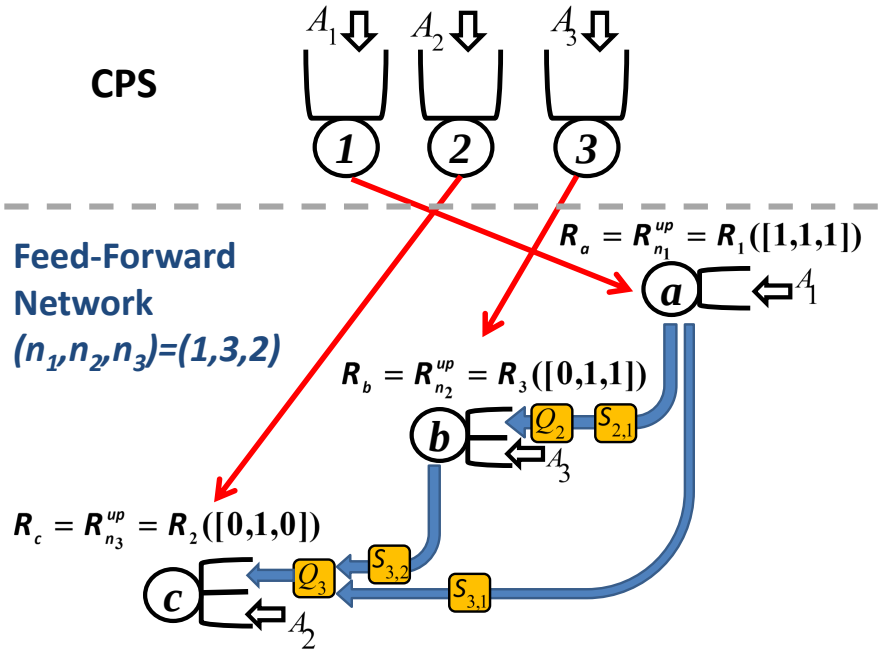

Figure A.6: An example of a three-queue CPS, and of a network associated to it, corresponding to the mapping (1,3,2).

A.2.2.2 Upper Bounding Conditions

In the networks we have described, the scaling coefficients $S_{j,k}$ are free parameters. Hence, if $\bar{S}$ is the matrix of the scaling coefficients of the network, then for a given CPS, each choice of the pair $(\bar{n},\ \bar{S})$ identifies a specific network associated to that CPS. In what follows we present a set of sufficient conditions on rate impairments (and therefore on scaling coefficients) that a feed-forward network must satisfy in order to allow deriving, from its stability, the stability of its associated CPS.

Theorem 3. *Given a CPS and a $(\bar{n}, \bar{S})$ feed-forward network associated to it, the stability of this network implies the stability of the CPS (and we say that the network upper bounds the CPS) if the following holds:*

$$F_j(t) \geq R^{up}_{n_j} - R'_{n_j}(t),\ \forall t \geq 0\ \forall j \in \{1, ..., N\}; \tag{A.5}$$

where $R'_{n_j}(t)$ is the service rate at queue n_j of the CPS when the queues of the CPS which are active are all those associated to corresponding active queues in the feed-forward network at t.

Proof. We begin with the following result, which defines a sufficient condition for a class of networks to upper bound a CPS:

Lemma A.2.1 (Upper Bounding Network)**.** *Consider a network with $N' \geq N$ queues, such that there is a one-to-one mapping between the queues of the CPS and a subset of N queues of the network. The mapping is such that each queue j in the subset has the same arrivals at any time t as its corresponding queue n_j in the CPS. Let $R_{n_j}(t)$ and $R_j(t)$ be the service rates at time t, respectively, at queue n_j and queue j. If at any time $t \geq 0$, for each queue n_j of the CPS, it holds $R_j(t) \leq R_{n_j}(t)$, then the network upper bounds the CPS.*

Let us indicate with $O_j(\bar{b}(t)) = R^{up}_{n_j} - F_j(\bar{b}(t))$ the instantaneous service rate for fresh traffic at the j-th GPS node. Then Equation (A.5) can be written as:

$$O_j(\bar{b}(t)) \leq R'_{n_j}(t) = R_{n_j}(\bar{b}(t)). \tag{A.6}$$

Here $R_{n_j}(\bar{b}(t))$ is the service rate of queue n_j of the CPS when the queues active at the CPS at time t are *only* all those which correspond to active queues for fresh arrivals at the network at time t. Equation (A.6) ensures that the service rate of each queue for fresh traffic of the network is always inferior to the one of the correspondent queues in the CPS, when all active queues for fresh arrivals in the network are associated to active queues at the CPS.

We now prove that such property is sufficient for the network to respect Lemma genTheorem. That is, we prove that if Equation (A.6) holds at any node and any time $t \geq 0$, then $O_j(\bar{b}(t)) \leq R_{n_j}(\bar{a}(t))$ for any node j and any time $t \geq 0$. In the following, indeed, we prove that $\bar{b}(t) \geq \bar{a}(t)$ at any time t. If so, from Equation (A.6), it follows that

$$O_j(\bar{b}(t)) \leq R_{n_j}(\bar{b}(t)) \leq R_{n_j}(\bar{a}(t)).$$

hence proving the theorem.

In the following, we prove that $\bar{b}(t) \geq \bar{a}(t)$ by contradiction. Considering that the arrivals are the same at the nodes of the network and at the corresponding queues of the CPS, we can assume that at $t = 0$, $\bar{b}(t) = \bar{a}(t)$.
Then, we assume by hypothesis that t^* is the first moment where a queue n_k that is not empty at the CPS turns empty in the corresponding node k of the network. In other words, t^* is the first moment where $\bar{b}(t) \geq \bar{a}(t)$ does not hold. In order to satisfy the hypothesis $\bar{b}_k(t) < \bar{a}_{n_k}(t)$, it exists at least a moment $t' \in [0, t^*)$ where $O_k(\bar{b}(t)) > R_{n_k}(\bar{a}(t))$. Since the CPS is monotonic, this is possible just if at t' one of the queues for fresh traffic is empty at the network, while the corresponding queue at the CPS is not. The existence of t' contradicts the definition of t^*, proving the theorem. □

The proof of this result is based on the fact that Equation (A.5) implies that at any moment, and at any node of the network, the instantaneous service rate for fresh traffic is always not larger than at the corresponding queue of the CPS. Hence, bounds on backlog and delay for fresh traffic queues in any feed-forward network hold also for the corresponding queue in the CPS.

For a given CPS and a given mapping $\bar{n}$, Theorem 3 identifies a set of feasible arrays $\bar{S}$, and hence of networks ($\bar{n}$,$\bar{S}$). From the structure of the feed-forward network, it can be easily verified that for every mapping $\bar{n}$ there are at least two choices of scaling coefficients which always satisfy Theorem 3.
The first one consists in setting all scaling values to a value larger than the largest possible service rate in the CPS. Note that such a choice is suboptimal, as it is equivalent of studying the system in saturation. A second choice corresponds to setting $\forall j \in \{1, ..., N\}$, $\forall k \in \{1, ..., j-1\}$,

$$S_{j,k} = \frac{R_{n_j}^{up} - R_{n_j}^{k-up}}{R_{n_k}^{up}},$$

where $R_{n_j}^{k-up}$ is the service rate at the CPS queue n_j when the CPS queues $n_k, ..., n_N$ are active. Such values of scaling coefficients make the contribution to interfering traffic from stage k equal to the one we have when stages $k, ..., j-1$ are active.

A.2.3 Sufficient Conditions for Stability

In the previous section we have seen how to derive a set of upper bounding networks for a given CPS. In what follows we present some sufficient conditions for stability of a CPS.

In what follows, we assume fresh traffic is constrained by arrival curves. More specifically, we consider leaky bucket arrival curves, despite similar results can be derived through our method for other types of arrival curves. The following theorem defines a set of sufficient conditions on leaky bucket rates for the deterministic stability of the CPS.

Theorem 4. *Given an N-queue CPS, where at each node $j = 1, ..., N$ fresh arrivals are constrained by leaky bucket arrival curves, with parameters (ρ_j, σ_j), the CPS is deterministically stable if there exists at least one associated network $(\bar{n}, \bar{S})$ satisfying Theorem 3, and such that at each stage $j = 1, ..., N$, ρ_{n_j} satisfies*

$$\rho_{n_j} \leq \max \left(R^{sat}_{n_j}, R^{up}_{n_j} - \sum_{k=1}^{j-1} S_{j,k} \rho_{n_k} \right). \tag{A.7}$$

Proof. Theorem 4 can be proved by induction on the index j of the stages of the feed forward network. Let us consider a network $(\bar{n}, \bar{S})$ associated to the given CPS, and satisfying Theorem 3. Let us consider stage $j > 1$. As stages 1 to j-1 are stable, $\beta_k(t)$, $k \in 1, ..., j-1$ is constrained by a leaky bucket arrival curve with rate ρ_{n_k}. Hence, from Equation (A.4), the rate impairment is constrained by a leaky bucket arrival curve, with rate equal to $\min(R^{up}_{n_j} - R^{sat}_{n_j}, \sum_{k=1}^{j-1} \rho_{n_k} S_{j,k})$. $(rho_c onstraint)$ derives from imposing that the sum of the leaky bucket rates of all arrivals at the GPS node should be less than its total service rate $R^{up}_{n_j}$. □

From the above, we can see that a trivial sufficient condition for stability is $\rho_j \leq R^{sat}_{n_j}$, $\forall j$, which corresponds to assuming the whole system is in saturation (all queues active). The derived expression for the achievable rates ρ in (A.7) takes into account the fact that queues are not always active, and they are function of the bounds to fresh traffic at those nodes which affect the service rate of the considered node.

A.3 CPS characterisation of LTE D2D inband and optimisation

In this appendix we provide more details on the CPS-based characterisation of LTE D2D inband mechanisms presented in Section 2.1.3. A basic introduction to CPS modelling has been presented in Appendix A.2.

A.3.1 System Model

We consider a base station (or eNB, in LTE) belonging to an LTE access network, serving U cellular users. Under the coverage area of the considered eNB, there are also D D2D transmitter-receiver pairs. We consider a static scenario, in which D2D pairs do not change over time and in which users do not move. We assume D2D transmissions happen on the uplink channel. The main reason for this choice is that in uplink the eNB is the receiver of all cellular transmissions and has complete knowledge of the sensed interference. We consider a log-distance path loss model. Specifically, path loss (in dB) between a transmitter and a receiver at a distance r is given by $L(r) = L_0 + 10\eta \log_{10}\left(\frac{r}{r_0}\right) + X_S$, where L_0 is the path loss at a reference distance r_0, η is the path loss exponent, and X_S is a Gaussian random variable with standard deviation σ_S, modelling the effects of shadowing [57]. As medium access technology, we assume that OFDMA is used.

Transmission time is split into slots of fixed duration, while multiple and independent sub-carriers are obtained over a wide channel bandwidth. A RB is the smallest resource that can be assigned to a particular user.[5] We model capacity through the Shannon formula. Therefore, if a generic transmitter i (either cellular or D2D) and its receiver j are at a distance $r_{i,j}$, the amount of bits $b(r_{i,j})$ transmitted per RB is:

$$b(r_{i,j}) = \min\left(\tau B \log_2\left(1 + \frac{P_T(i)10^{\frac{L(r_{i,j})}{10}}}{N_0 B + \sum_{k\neq i} P_T(k)10^{\frac{L(r_{k,j})}{10}} + I_C}\right), b_M\right), \tag{A.8}$$

where B is the bandwidth of the RB, τ is the duration of a time slot, $P_T(i)$ is the per-RB transmission power of i, N_0 is the noise spectral density, I_C is the inter-cell interference, while b_M is the maximum amount of bits that can be transmitted in a RB when the best modulation and coding scheme is used. The summation of the interference at the denominator goes over all the active transmitters in the RB, both cellular and D2D. We assume that the transmission power used by the devices is fixed over the RBs and over time.

A.3.2 A D2D in-band underlay scheme for LTE

In in-band underlay D2D transmission, a D2D pair can be scheduled by the eNB on the same RBs that are assigned to cellular transmitters, or to other D2D pairs. Figure A.7 represents an example of radio resource utilization of in-band underlay D2D communication.

The particular resource scheduling policy we consider is a variation over the one proposed in FlashLinQ [26]. FlashLinQ is a state-of-the-art PHY-MAC architecture for D2D that allows the scheduling of different transmitters (D2D or cellular) in the same time and frequency resource, through an OFDMA-like access selection mechanism. The scheduling of the transmitters is performed at RB level. Furthermore, FlashLinQ does not distinguish between cellular and D2D users. We believe indeed that the UEs have to be considered as primary users of the cellular access network and the D2D transmissions have to be scheduled without preventing satisfactory performance

[5] The analysis holds also for Single Carrier-FDMA (SC-FDMA), where a set of RBs (and not one) is the minimum quantum of resource.

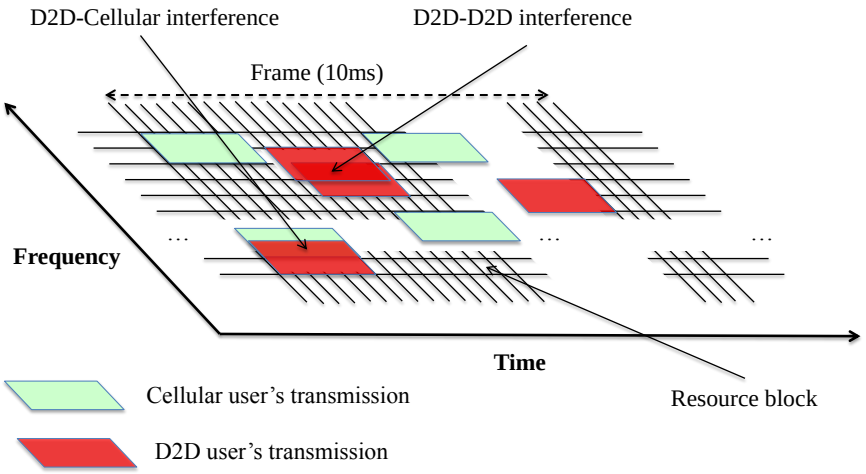

Figure A.7: Use of resources in D2D in-band underlay systems.

for cellular users. We introduce therefore a two-tier scheduling policy of the transmitters, where first cellular transmitters are scheduled, and where subsequently the FlashLinQ scheduling policy is applied to the D2D transmitters:

- **Cellular UEs transmitters**: The scheduling policy at the eNB is Equal Time [58], which is interference-unaware.[6] In order to have a full understanding of the effect of the D2D transmissions on the cellular ones, we also assume that cellular UEs are in saturation, i.e., they always have a packet ready to send. In practice, the eNB assigns RBs to each cellular user with the same probability.
- **D2D transmitters - FlashLinQ policy**: For each RB, all the D2D transmitters having at least one packet to send are considered. Those who are scheduled for transmission are all those (i) whose interference on the cellular transmitter scheduled in the same RB is below a given limit, and (ii) who are able to achieve a given minimum Signal to Interference Ratio (SIR), considering the cellular UE scheduled in the particular RB under analysis and all the D2D transmitters already scheduled. This ensures that the impact of D2D transmissions on cellular ones is kept below a given threshold and, furthermore, that D2D users achieve a target minimum throughput.

The order in which D2D transmissions are considered by the above scheduling procedure determines the set of D2D scheduled transmissions in a given RB. Therefore, in order to maximise fairness among D2D transmitters, as in FlashLinQ, every time a scheduling decision has to be made, the order of candidate D2D transmitters is picked at random, following a uniform distribution.

A.3.3 A CPS for D2D Schemes

In this section we show how to tackle the analysis of in-band underlay D2D schemes. First we introduce a model that is able to characterise the main features of such D2D schemes. Then, we present the main tools we use to analyse such a model.

In order to study the performance of the cellular and D2D systems, we use a CPS [25]. A CPS is a set of parallel queues, whose service rates at any time t is determined by the set of active queues at that time. More formally, if at time t the set of active queues is $I(t)$, the service rate of queue

[6] Any scheduling scheme, even interference aware, for which it is possible to determine the scheduling probability of a cellular user can also be used.

$i \in I(t)$ is $R_i(t) = R_i(I(t))$. Therefore, at each time t, the set of active queues $I(t)$ univocally determines the state of the system. Note that the set of system states $\mathcal{I}$ is finite, and represents all the possible subsets of active queues of the CPS.

The D2D system can be mapped into a CPS where D queues are present (the number of queues exactly corresponds to the number of D2D transmitters under analysis). Each of the "coupled" queues models one of the D2D transmission queues, where the coupling arises from mutual interference with the other D2D pairs scheduled in the same RB.

Considering that the number of bits transmitted in a RB by the D2D pair $d \in \mathcal{D} = \{1, ..., D\}$ depends on (i) the set of active D2D pairs, i.e., the system state $I \in \mathcal{I}$, (ii) the cellular user $u \in \mathcal{U} = \{1, ..., U\}$ scheduled in the same RB, and (iii) the particular order I_O in which D2D users are listed in the scheduling algorithm, the average service rate of d when I is the system state is expressed by:

$$R_d(I) = \frac{\phi}{U} \sum_u \frac{\sum_{I_O} b_d^u(I_O)}{|I|!}, \tag{A.9}$$

where $|I|$ is the number of active D2D pairs when I is the system state, $|I|!$ are the possible sorted lists of active D2D pairs, ϕ is the number of available RBs per second and $b_d^u(I_O)$, computed applying (A.8), is the amount of bits per RB transmitted by d when the cellular device u is active and I_O is the particular sorting used to schedule the active D2D pairs. Here we have used the fact that all $|I|!$ possible permutations of active D2D pairs are equally likely to be chosen, as well as the probability of having a cellular user u scheduled in a RB is $\frac{1}{U}$, due to the Equal Time scheduler adopted.

The approximation we make with this model consists in considering the instantaneous service rate of the CPS queues as the average rate of the corresponding D2D transmitters in a given system state. Furthermore, while the cellular transmissions are scheduled sequentially, on a per-RB basis, and the change of system state can happen only when a new RB is scheduled, the equivalent CPS model works at the *fluid* limit, i.e., each queue would serve his traffic as it were infinitely divisible, and it could change state at any time t. In Section A.3.6 we validate numerically that those approximations are acceptable, and that the CPS models accurately the dynamics of the original D2D system.

We summarized how to study a CPS queuing system in Appendix A.2.

A.3.4 Proportionally Fair Optimisation

By exploiting the knowledge of the conservative estimate of the stability region of the system, here we show how to achieve proportional fairness among D2D transmissions. Given a CPS with D queues, each representing one of the D D2D pairs in the scenario, we first formalize a proportional fairness throughput optimisation problem that exploits the feed-forward networks introduced in Appendix A.2, and then we present a heuristic which searches for the optimum by characterising just a suitably small subset of networks.

We assume to have D D2D transmitters, that their demand is given and that arrivals can be described through equivalent leaky bucket characterisations $\{\rho_i, \sigma_i\}$, $i \in \{1, ..., D\}$. The goal of the optimisation is to introduce leaky bucket shapers at the transmitters, having long term rates $\rho_i^* \leq \rho_i$, in order to (i) ensure stability at the transmission queues of the D2D users, (ii) maximise a weighted sum of the logarithms of the rates, thus achieving *proportional fairness* of user's throughputs. The solution of the optimisation problem and the distribution of the long term rates of the shapers to the D2D transmitters can be easily performed by the eNB where the D2D transmissions are taking

place. Formally, we express the optimisation problem as follows:

$$\begin{aligned}
&\underset{\rho^*}{\text{maximise}} \quad \max_{d_O} \sum_{i=1}^{D} w_{d_i} \log\left(\rho^*_{d_i}\right), \\
&\text{subject to:} \quad \forall i = 1, \ldots, D, \\
&\rho^*_{d_i} + \min\left(R^{up}_{d_i} - R^{min}_{d_i}, \sum_{k=1}^{i-1} \frac{R^{up}_{d_i} - R^{k,up}_{d_i}}{R^{up}_{d_k}} \rho^*_{d_k}\right) \leq R^{up}_{d_i}, \\
&\rho^*_{d_i} \leq \rho_{d_i},
\end{aligned} \tag{A.10}$$

where the maximisation is performed over any possible ordering d_O of $\mathcal{D}$ tuning the long term rates ρ^* of the shapers we want to introduce in the network and where w_{d_i} is the weight assigned to the transmitter d_i in order to achieve proportional fairness. The first constraint ensures that the long term rate $\rho^*_{d_i}$ we pick leads to a stable transmission queue for the generic transmitter D2D d_i (directly from Theorem 4), while the second ensures that the long term rate we pick for the shapers is at most the one of the demand of the generic transmitter d_i.

When dealing with a particular sorting d_O, the presence of the min function in the constraints leads to a piece-wise linear feasibility region. To deal with the problem, we use the Big-M transformation [59], so that the resulting optimisation problem has a convex objective function and belongs to the mixed-binary programming family. We achieve the global optimum of the optimisation through *branch-and-bound* ($B\&B$) method [60].

A.3.5 Derivation of a Heuristic

In order to reduce the computation required by the optimisation problem over the rates ρ_{d_i} when the sorting d_O is analysed, we stop the $B\&B$ evaluation when the intermediate solution is at most at ϵ from the optimum ($B\&B$ gives, at each step, a higher and a lower bound for the optimum). This parameter can then be tuned according to the desired trade-off between computational cost and performance of the heuristic.

The second and most important approximation we introduce is in the set of networks that we evaluate. Instead of choosing all the sortings d_O, we try to identify the subset of feed-forward networks that most probably contains the stable set of long term rates ρ^* that maximise our proportional fair problem. We limit to those networks the evaluation of (A.10).

To choose such subset we use the following reasoning. Due to the cascade structure of the feed-forward network, for each CPS queue, the earlier is the stage to which the corresponding GPS queue belongs, the larger the set of queues whose coupling with the considered one is modelled more accurately, i.e. through feed-forward traffic rather than via a penalty on the service rate which holds for any time t.

Obviously, the higher the weight of the D2D pair in the proposed optimisation, the higher is the contribution to the utility function of the system. Therefore, the objective of the heuristic we propose is to maximise the achieved throughput of the D2D pairs associated to the largest weights. Accordingly, the small subset of networks we decide to evaluate is initialized by probabilistically mapping D2D transmitters onto GPS nodes sequentially, starting from the earliest stages of the feed forward network, with probabilities proportional to the weights w_i of CPS transmitters in the objective function of the optimisation (A.10). We repeat the probabilistic mapping process to populate a small set of starting points for the heuristic. For each starting point network, we identify the set of $D-1$ *neighbouring* networks each obtained by swapping two adjacent GPS nodes. Within this set of neighbours, the network with the higher proportional fairness for the system is taken as the new reference network, and the evaluation of neighbour feed-forward networks is

Table A.4: Simulation Setup

LTE Carrier	2.45 GHz
Uplink Bandwidth	20 MHz
B	200 kHz
Subframe Duration τ	1 ms
b_M	468 $bits$/RB [7]
Size Scenario	300×300 m
Max Dist. D2D pair	40 m
r_0	3 m
η	2.5
σ_S	8 dB
$\mathcal{N}$	$3.98 * 10^{-18}$ W/Hz
P_{TX}^{MAX}	200 mW
packet length	12 Kb

repeated until a local maximum for the proportional fairness is found. Obviously, the trade-off among computation and accuracy of the solution is given by the number of starting points we choose.

A.3.6 Numerical Evaluation

In this section, we show the advantage achieved exploiting our analysis when a proportional fair assignment of the resources is the goal. For the whole set of proposed simulations, Table A.4 summarizes the values of the parameters we used. We also used power control mechanisms, aiming to achieve a per-RB SNR of 50.

Now, we show the results of the optimisation problem we presented in Section A.3.4. In the following, we set $\epsilon = 10\%$, a value that was showing a good trade-off among precision and complexity. First of all we present how close the heuristic performs if compared against a brute force approach, i.e., an approach that solves the optimisation problem over all the possible upper bounding networks w_O, with $\epsilon = 0\%$.

In order to show the scalability of the approach we propose, here we choose larger scenarios. In Table A.5 we present the results obtained when 20 cellular transmitters and 3, 4, 5 and 6 D2D transmitters are present. The choice of the number of transmitters is such that we can solve the optimisation problem with the brute force approach.

The position of the devices, the demand of each of the D2D transmitters and each weight in the utility computation is picked at random for any of the optimisation problems performed. Table A.5 shows that the presented heuristic performs as good as the brute force approach in most of the cases, although it does not need to explore the entire set of demands.

In order to evaluate the complexity of the heuristic proposed, Table A.6 shows the number of networks evaluated by the heuristic and the complexity of each of the optimisation problems solved (one for each network). Such complexity is evaluated in terms of different branches the $B\&B$ algorithm requires before reaching an intermediate solution that is at most $\epsilon = 10\%$ far from the optimum. In this case we also evaluate a set of larger optimisation problems, with 8 D2D transmitters. In particular, Table A.6 shows that the complexity of the solution proposed by the heuristic presented in Section A.6.2 remains computationally feasible, even though the complexity of the brute force approach grows fast.

[7] Given $64 - QAM$ as modulation and 0.93 as approximated coding rate.

Table A.5: Heuristic vs. Brute force: Utility

D2D transmitters	Heuristic		Brute force	
	Mean	95% Conf. Int.	Mean	95% Conf. Int.
3	2.5359	2.38-2.69	2.5360	2.38-2.69
4	2.5796	2.45-2.71	2.5843	2.45-2.72
5	2.3451	2.21-2.47	2.3503	2.22-2.47
6	2.2191	2.02-2.41	2.2369	2.04-2.43

Table A.6: Heuristic vs. Brute Force: Complexity

D2D transmitters	Mean # of Networks (Heuristic)	95% Confidence Interval	Networks Available
3	3.44	3.36-3.52	6
4	4.60	4.36-4.84	24
5	9.56	8.95-10.18	120
6	11.70	10.49-12.92	720
8	22.70	20.52-24.88	40320

D2D transmitters	Mean # of Branches (Heuristic)	95% Confidence Interval	Branches Available
3	2.00	2.00 - 2.00	2
4	1.99	1.90 - 2.07	4
5	2.74	2.56 - 2.93	8
6	2.85	2.55 - 3.15	16
8	4.82	4.10 - 5.53	64

Table A.7: Optimisation vs. Saturation: Utility

D2D transmitters	Mean Sat.	95% Conf. Interval	Mean Opt.	95% Conf. Interval
3	2.3289	2.11-2.54	2.5303	2.37-2.68
4	2.2870	2.06-2.50	2.5750	2.44-2.70
5	1.9881	1.80-2.17	2.3424	2.21-2.47
6	1.7460	1.39-2.09	2.2153	2.01-2.41
8	1.5185	1.25-1.77	2.0537	1.91-2.19

In order to show the gain that the knowledge of the conservative estimate of the stability region brings, we also simulate all the scenarios for which the optimisation was performed. We simulate, for each of the cases, the operation of the in-band underlay D2D system, considering or not the presence of shapers at the D2D transmitters set as the output of the optimisation problem we solved. In both cases we compute the throughput achieved from the D2D transmitters, and then we compare the corresponding log-utility (see (A.10)).

As it easy to see from Table A.7, even if we achieve just a conservative estimate of the whole stability region, the shapers improve sensibly the value of the utility achieved. In particular, when the scenarios get larger, the utility improves up to 35.2% on average and 131.53% in the best case. Please note that the simulations performed when the shapers are present reach almost the same utility values that the optimisation problems were giving in output.

A.4 Design and performance of ADOS

In this section we show how to derive the theoretically optimal operational values for access probabilities p_i and transmission rate thresholds $\bar{R}_i$ in a WLAN with distributed opportunistic scheduling enabled. We have used the results presented in this appendix in the design of the ADOS mechanism presented in Section 2.2.1. Here we also detail how ADOS works and report on its performance assessment.

A.4.1 Optimal p_i configuration

To compute the optimal p_i configuration, we start by expressing the throughput r_i as a function of $\mathbf{p}$. Let l_i be the average number of bits that station i transmits upon a successful contention and T_i be the average time it holds the channel. Then, the throughput of station i is

$$r_i = \frac{p_{s,i} l_i}{\sum_j p_{s,j} T_j + (1 - p_s)\tau}$$

where $p_{s,i}$ is the probability that a mini slot contains a successful contention of station i and p_s is the probability that it contains any successful contention,

$$p_{s,i} = p_i \prod_{j \neq i} (1 - p_j), \quad p_s = \sum_i p_{s,i}$$

Both l_i and T_i depend on $\bar{R}_i$. Upon a successful contention, a station holds the channel for a time $\mathcal{T} + \tau$ in case it transmits data and τ in case it gives up the transmission opportunity. In case the station uses the transmission opportunity, it transmits a number of bits given by $R_i(\theta)\mathcal{T}$. Thus, T_i and l_i can be computed as

$$T_i = Prob(R_i(\theta) < \bar{R}_i)\tau + Prob(R_i(\theta) \geq \bar{R}_i)(\mathcal{T} + \tau)$$

and

$$l_i = \int_{\bar{R}_i}^{\infty} r\mathcal{T} f_{R_i}(r) dr$$

where $f_{R_i}(r)$ is the pdf of $R_i(\theta)$.

Similarly as in [61], let us define w_i as

$$w_i = \frac{p_{s,i}}{p_{s,1}} \tag{A.11}$$

where we take station 1 as reference. From the above equation, we have that $p_{s,i} = w_i p_s / \sum_j w_j$; substituting this into (A.11) yields

$$r_i = \frac{w_i p_s l_i}{\sum_j w_j p_s T_j + \sum_j w_j (1 - p_s)\tau}$$

In a slotted wireless system such as the one of this paper, the optimal access probabilities satisfy $\sum_i p_i = 1$ (see [62]), which gives an optimal success probability p_s approximately equal to $1/e$,

$$p_s = \sum_i p_i \prod_{j \neq i} 1 - p_j \approx \sum_i p_i e^{-\sum_j p_j} = e^{-1} \tag{A.12}$$

With the above, the problem of finding the $\mathbf{p}$ configuration that maximises the proportionally fair rate allocation is thus equivalent to finding the w_i values that maximise $\sum_i \log(r_i)$, given that $p_s = 1/e$. To obtain these w_i values, we impose $\frac{\partial \sum_i \log(r_i)}{\partial w_i} = 0$ which yields

$$\frac{1}{w_i} - N \frac{p_s T_i + (1 - p_s)\tau}{\sum_i w_i p_s T_i + \sum_j w_j (1 - p_s)\tau} = 0$$

Combining this expression for w_i and w_j, we obtain

$$\frac{w_i}{w_j} = \frac{p_s T_j + (1-p_s)\tau}{p_s T_i + (1-p_s)\tau}$$

Given that $w_i/w_j \approx p_i/p_j$ and $p_s = 1/e$, the above can be rewritten as

$$\frac{p_i}{p_j} = \frac{T_j + (e-1)\tau}{T_i + (e-1)\tau} \tag{A.13}$$

Furthermore, the probability that a given mini slot is empty can be computed as follows,

$$p_e = \prod_i 1 - p_i \approx e^{-\sum_i p_i} = e^{-1} \tag{A.14}$$

We use a different approximation than [61]'s in order to remove any dependency with the number of stations, a result that we will exploit to design an algorithm that works well under non-saturation conditions too. Our simulation results show a very small performance impact for using this approximation instead, practically negligible for scenarios with $N > 4$ stations.

With the above, we compute the solution of the optimisation problem by finding the $\mathbf{p}$ values that solve the system of equations formed by (A.13) and (A.14).

Hereafter, we denote the unique solution to the system of equations by $\mathbf{p}^* = \{p_1^*, \ldots, p_N^*\}$. Note that determining $\mathbf{p}^*$ requires computing $T_i \ \forall i$, which depend on the optimal configuration of the thresholds $\bar{\mathbf{R}}$. In the following section we address the computation of the optimal $\bar{\mathbf{R}}$, which we denote by $\bar{\mathbf{R}}^* = \{\bar{R}_1^*, \ldots, \bar{R}_N^*\}$.

A.4.2 Optimal $\bar{R}_i$ configuration.

In order to obtain the optimal configuration of $\bar{\mathbf{R}}$, we need to find the transmission rate threshold of each station that, given the $\mathbf{p}^*$ computed above, optimises the overall performance in terms of proportional fairness. To this aim, we rely on Theorem 1 in [61] to find that the optimal configuration of the transmission rate thresholds is given by $\bar{R}_k^* = \bar{R}_k^1$, where $\bar{R}_k^1$ is the transmission rate threshold that optimises the throughput of station k when it is alone in the channel and contends with $p_k = 1/e$ (under the assumption that different channel observations are independent).

This is done in [28], which uses *optimal stopping theory* and finds that the optimal threshold can be obtained by solving the following fixed point equation:

$$E\left[R_i(\theta) - \bar{R}_i^*\right]^+ = \frac{\bar{R}_i^* \tau}{T/e} \tag{A.15}$$

Note that the above allows computing the threshold $\bar{R}_i^*$ of a station based on *local information* only, as (A.15) does not depend on the other stations in the network and their radio conditions. In particular, the optimal threshold configuration is *independent of the access probabilities* $\mathbf{p}$, which is crucial as it allows decoupling the algorithm that adjusts the configuration of $\bar{\mathbf{R}}$ from the one that adjusts $\mathbf{p}$. In the following, we present two independent adaptive algorithms to bring the system to the optimal point of operation: one that drives the access probabilities $\mathbf{p}$ to their optimal values $\mathbf{p}^*$, and another that drives the threshold of each station $\bar{R}_i$ to its optimal value $\bar{R}_i^*$.

A.4.3 The ADOS Mechanism

We next present our mechanism, named ADOS, to adapt the operation of the opportunistic scheduling to the conditions in the WLAN.

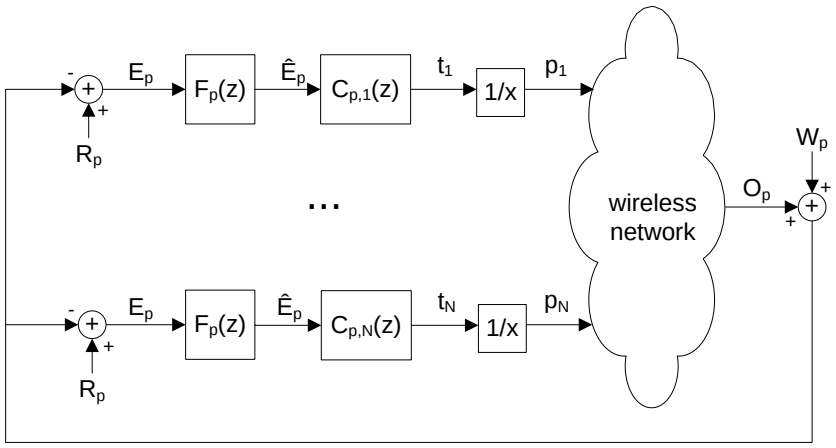

Figure A.8: Adaptive algorithm for p_i.

A.4.3.1 Adaptive algorithm for p_i

Following the first approximation above, with ADOS each station implements an adaptive algorithm to configure the access probability p_i, with the goal of driving the channel empty probability to $1/e$, as given by (A.12) in Section A.4.1.

Driving the channel empty probability toward a constant optimum value fits well with the framework of *classic control theory*. With these techniques, we measure the *output signal* of the system and, by judiciously adjusting the *control signal*, we aim at driving it to the *reference signal*. A key advantage of using such techniques is that they provide the means for achieving a good trade-off between the speed of reaction and stability while guaranteeing convergence, which is a major challenge when designing adaptive algorithms.

Figure A.8 depicts our algorithm to adjust $\mathbf{p}$, where each station computes the error signal E_p by subtracting the output signal O_p from the reference signal R_p (the functions in the figure are given in the z domain). The output signal O_p is combined with a noise component W_p of zero mean, modelling the randomness of the channel access algorithm. In order to eliminate this noise, we follow the design guidelines from [63] and introduce a low-pass filter $F_p(z)$. The filtered error signal $\hat{E}_p$ is then fed into the controller $C_{p,i}(z)$ of each station, which provides the control signal t_i, defined as the average time between two transmission of station i. Station i then computes its access probability as $p_i = 1/t_i$. With the p_i of each station, the wireless network provides the output signal O_p, which closes the loop.

In the above system, we need to design the reference and output signals R_p and O_p, as well as the transfer functions of the low-pass filter and the controller, $F_p(z)$ and $C_{p,i}(z)$. In the following we address the design of these components with the goal of ensuring that the empty probability p_e is driven to $1/e$.

In our system, time is divided into intervals such that the end of an interval corresponds to a transmission in the channel (either a success or a collision). Given that the target empty probability is equal to $1/e$, the target average number of empty mini slots between two transmissions (i.e., our reference signal) is equal to $R_p = 1/(e-1)$. In this way, after the n-th transmission, each station computes the output signal at interval n, denoted by $O_p(n)$, as the number of empty mini slots between the $(n-1)$-th and the n-th transmission. The error signal for the next interval is then computed as

$$E_p(n+1) = R_p - O_p(n). \tag{A.16}$$

With the above, if p_e is too large then $O_p(n)$ will be larger than R_p in average, yielding a negative error signal $E_p(n+1)$ that will decrease t_i for the next interval, which will increase the transmission

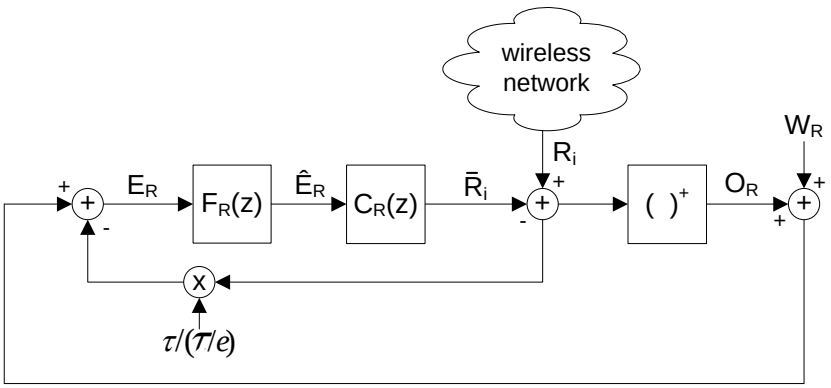

Figure A.9: Adaptive algorithm for $\bar{R}_i$.

probability p_i and therefore reduce p_e (and vice-versa). This ensures that p_e will be driven to the optimal value. For the low-pass filter $F_p(z)$, we use a simple exponential smoothing algorithm of parameter α_p [64], given by the following expression in the time domain,

$$\hat{E}_p(n) = \alpha_p E_p(n) + (1 - \alpha_p)\hat{E}_p(n-1),$$

which corresponds to the following transfer function in the z domain:

$$F_p(z) = \frac{\alpha_p}{1 - (1 - \alpha_p)z^{-1}}.$$

For the transfer function of the controllers $C_{p,i}(z)$, we use a very simple controller from classical control theory, namely the Proportional Controller [65], which has already been used in a number of networking problems (see e.g. [66]):

$$C_{p,i}(z) = K_{p,i},$$

where $K_{p,i}$ is a per-station constant. In addition to driving the empty probability to $1/e$, we also impose that the access probabilities satisfy the fairness constraints (A.13) reported in Section A.4.1. Since we feed the same error into the different stations, and the proportional controller simply multiplies this error by a constant to compute p_i, the following equation holds for all i, j:

$$\frac{p_i}{p_j} = \frac{K_{p,j}}{K_{p,i}}.$$

Therefore, by simply setting $K_{p,i}$ as

$$K_{p,i} = K_p \left(T_i + (e-1)\tau\right)$$

we ensure that (A.13) is satisfied.

A.4.3.2 Adaptive algorithm for $\bar{R}_i$

The goal is to adjust the threshold $\bar{R}_i$ to drive the threshold of all stations to the optimal value given by (A.15) in Section A.4.2. Note that (A.15) is equivalent to the following equation:

$$E\left[(R_i(\theta) - \bar{R}_i^*)^+ - \frac{\bar{R}_i^*\tau}{\mathcal{T}/e}\right] = 0. \tag{A.17}$$

In the following, we design an adaptive algorithm that drives $\bar{R}_i$ to the value given by the above equation. The algorithm is depicted in Figure A.9. Similarly to the adaptive algorithm for p_i, we

base the algorithm design on control theory. The key difference between the two algorithms is that, since the optimal value of threshold of a station depends on local information only and hence does not depend on the threshold value of the other stations, we can consider each station separately (in contrast to Figure A.8).

In order to ensure that the configuration of $\bar{R}_i$ satisfies (A.17), we design the output signal of the algorithm, O_R, equal to the term $(R_i - \bar{R}_i)^+$, and the reference signal, R_R, equal to the term $\bar{R}_i\tau/(\mathcal{T}/e)$. Thus, by driving the difference with these two terms (i.e., the error signal) to zero, we ensure that (A.17) is satisfied.

Following the above, upon its n^{th} successful contention, a station measures the channel transmission rate $R_i(n)$ and computes the output signal as

$$O_R(n) = \begin{cases} R_i(n) - \bar{R}_i(n), & \text{if } R_i(n) >= \bar{R}_i(n) \\ 0, & \text{otherwise} \end{cases}$$

From the above output signal, it then computes the error signal as

$$E_R(n+1) = O_R(n) - \frac{\bar{R}_i(n)\tau}{\mathcal{T}/e}$$

Due to the randomness of the radio signal, the output signal carries some noise W_R. In order to filter out this noise, we apply (like in the previous case) a low pass-filter $F_R(z)$ on the error signal, which yields

$$\hat{E}_R(n) = \alpha_R E(n) + (1 - \alpha_R)\hat{E}_R(n-1)$$

Also like in the previous case, the error signal is introduced into a proportional controller,

$$C_R(z) = K_R$$

where K_R is the constant of the controller.

The controller gives the threshold configuration $\bar{R}_i(n)$ as output. As mentioned above, by driving the error signal $\hat{E}_R(n)$ to 0, the controller ensures the threshold value satisfies (A.17) and thus achieves the objective of adjusting the threshold to the optimal value $\bar{R}_i^*$.

A.4.3.3 Control Theoretic Analysis

With the above, we have all the components of the ADOS mechanism fully designed. The remaining challenge is the setting of its parameters, namely the parameters of the adaptive algorithm for p_i (K_p and α_p) and the adaptive algorithm for $\bar{R}_i$ (K_R and α_R). For the derivation of these parameters, we refer the interested reader to our publication [7].

A.4.4 Performance assessment figures for ADOS

In Section 2.2.1, we have presented the basic performance figures for ADOS and assessed its behaviour with respect to other state of the art mechanisms. Here we add more performance figures achieved via simulation.

A.4.4.1 Stability of ADOS

The setting of the parameters $\{K_p, \alpha_p\}$ and $\{K_R, \alpha_R\}$ achieves a good tradeoff between stability and speed of reaction. This is verified by the results presented in this and the following two sections.

To verify stable behavior, we first analyse the evolution over time of the access probability p_i of a station for the proposed $\{K_p, \alpha_p\}$ setting and for a configuration of these parameters 10 times

larger, in a homogeneous scenario with $N = 5$ saturated stations and $\rho = 4$. Figure A.10 shows the evolution of p_i for both cases, sampled over $10^5\tau$ intervals. We observe from the figure that with the proposed setting (labeled "K_p, α_p"), p_i shows minor deviations around its average value, while for a larger setting (labeled "$K_p * 10, \alpha_p * 10$"), it shows unstable behavior with drastic oscillations.

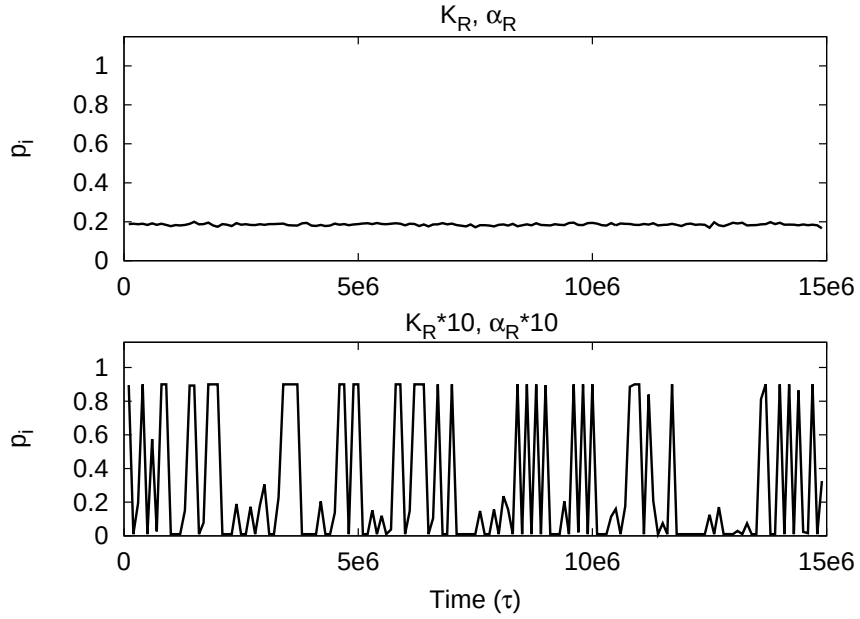

Figure A.10: Stability of p_i.

Similarly, we also analyse the evolution over time of the threshold $\bar{R}_i$ of a station for the proposed $\{K_R, \alpha_R\}$ setting and for a configuration of these parameters 10 times larger in the same scenario. The results, depicted in Figure A.11 confirm that the proposed setting for these parameters is stable while a larger setting is highly unstable.

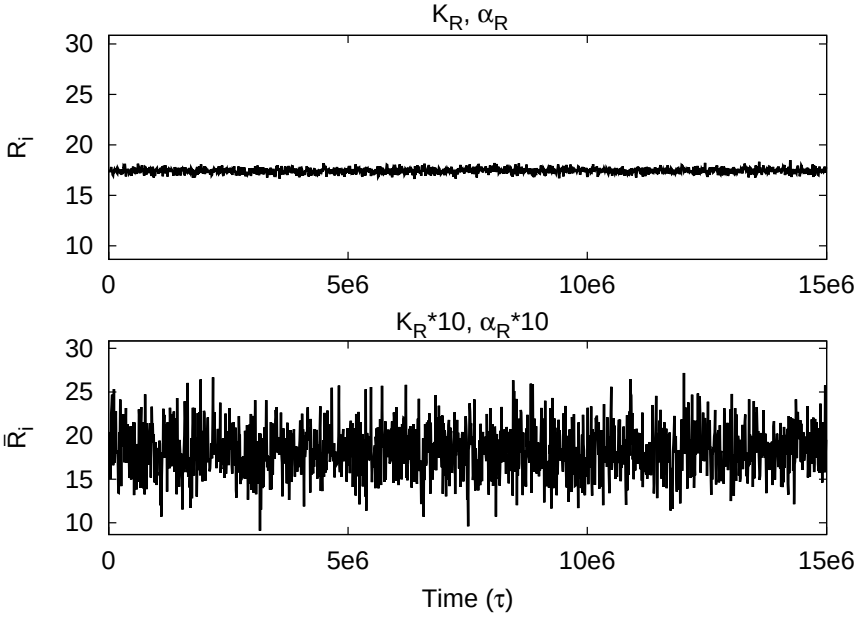

Figure A.11: Stability of R_i^*.

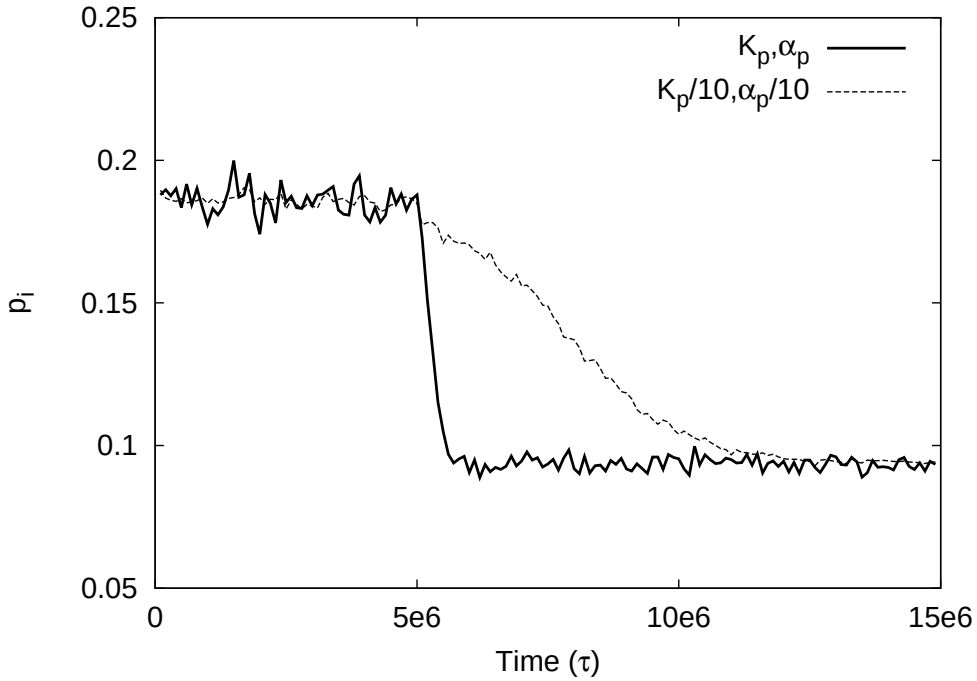

Figure A.12: Changing number of stations.

A.4.4.2 Changing number of stations

We next investigate the speed with which the system reacts to changes in the number of stations of the wireless network, which triggers the adjustment of the access probabilities p_i. To this aim, we consider a wireless network with initially 5 stations, where 5 additional stations join the network after a time $5 \cdot 10^6 \tau$. All stations have $\rho = 4$. Figure A.12 shows the evolution of the access probability of one of the initial stations sampled over $10^5 \tau$ intervals. We observe from the figure that with our setting (labeled "K_p, α_p"), the system quickly adapts the p_i of the station to the new value. In contrast, for a setting of these parameters 10 times smaller (labeled "$K_p/10, \alpha_p/10$"), the reaction is very slow and the system only converges after $5 \cdot 10^6 \tau$.

The results confirm that the proposed configuration for the parameters of the algorithm that adjusts p_i, $\{K_p, \alpha_p\}$, provides a good tradeoff between stability and speed of reaction, since with a larger setting of these parameters the system suffers from instability (as shown in the previous section), while with a smaller setting it reacts too slowly (as shown here).

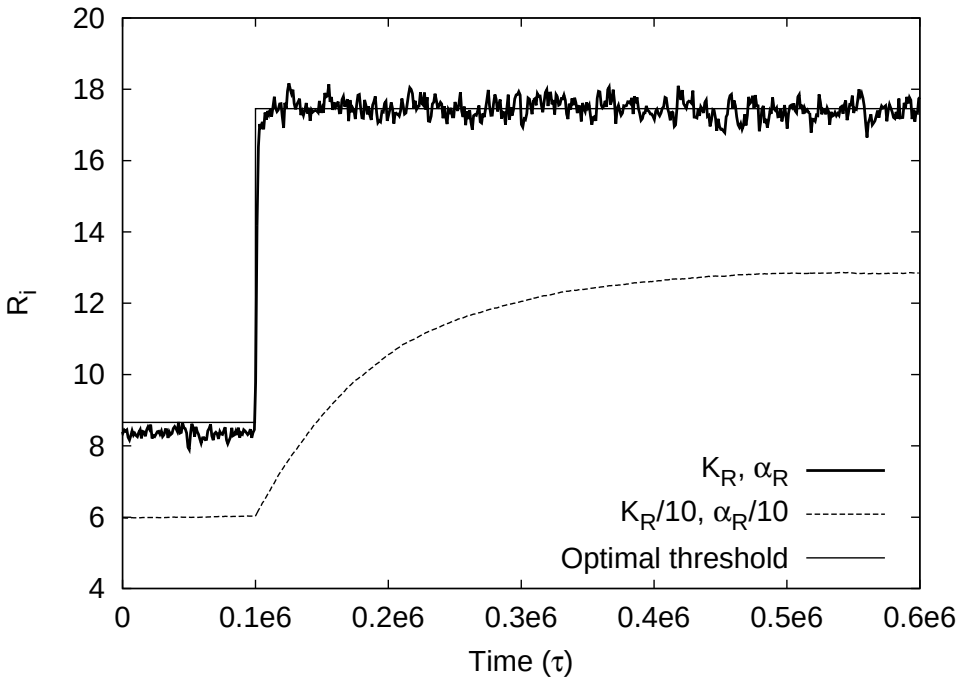

Figure A.13: Drastic change of ρ.

A.4.4.3 Changing radio conditions

To analyse the speed of reaction to changing radio conditions, we consider that in a wireless network with $N = 2$ stations, both of them with a normalized SNR $\rho = 1$, one of the stations changes its normalized SNR to $\rho = 4$ after a time $5 \cdot 10^5 \tau$. Figure A.13 shows the evolution over time of the $\bar{R}_i$ of the station whose normalized SNR has changed, for the proposed setting of the $\{K_R, \alpha_R\}$ parameters as well as for a setting of these parameters 10 smaller. As a benchmark, the figure also shows the optimal setting of the threshold as given by the analytical results. The results show that: (i) with our setting of the parameters, the system reacts quickly and closely follows the benchmark, while the reaction is much slower for a smaller setting, and (ii) the steady state error with our setting is negligible, whereas with a smaller setting of the parameters it is much larger. The latter effect is caused by the fact that the steady error with a proportional controller increases as its proportional gain (K_R) is reduced. Therefore, by choosing a too small value for K_R, we do not only worsen the speed of reaction of the system but also its steady error.

A.5 The SOLOR mechanism for opportunistic packet relay with WiFi Direct

In this appendix we detail how a SOLOR node operates. The SOLOR mechanism is based on the availability of WiFi Direct for D2D communications, and it has been presented in Section 2.2.2. In that section, we have introduced the basic properties and operation of SOLOR. This operation is summarised in Figure A.14. Next we provide a description of SOLOR as a MAC protocol.

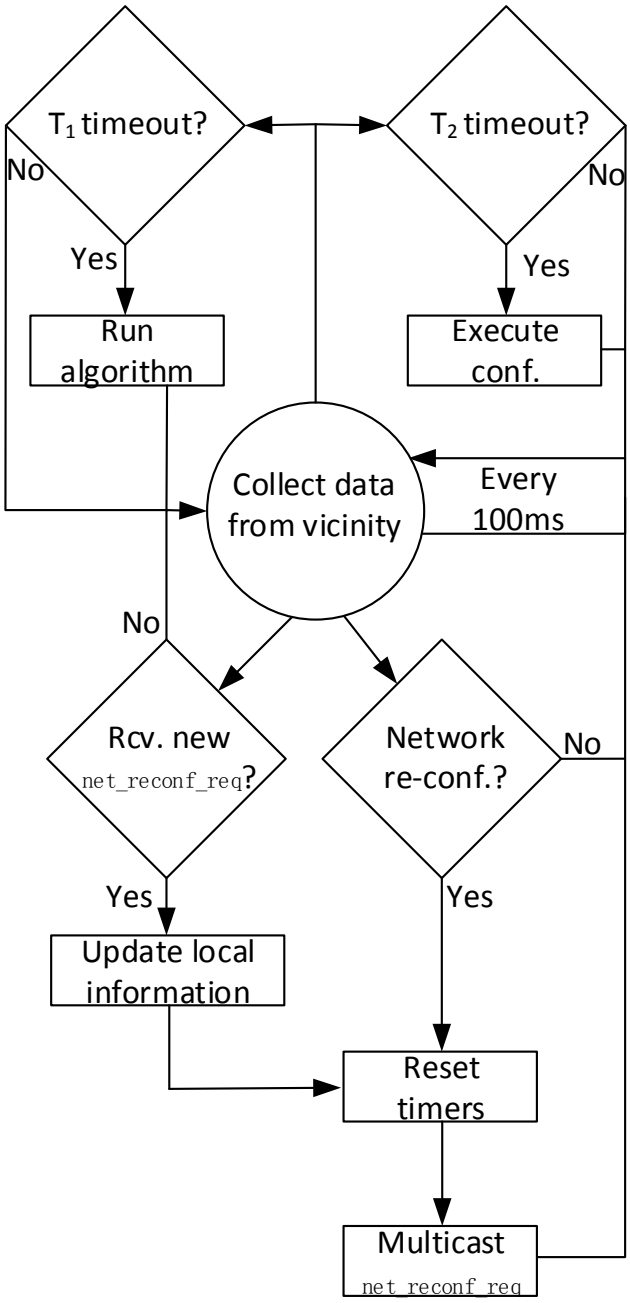

Figure A.14: Operation of a SOLOR node.

A.5.1 Protocol overview

When powered on, a SOLOR node multicasts its presence to the rest of SOLOR nodes (if any), following the communication scheme detailed below. Then, to estimate the topology of the network, it continuously snoops the transmissions from all nodes (legacy and SOLOR), and collects the MCS used by each node to transmit and the estimated SNR using an exponentially weighted moving average to filter out small fluctuations. Following [67], the SNR information serves to estimate the MCS a node will use when transmitting to the SOLOR node, which completes the estimation mechanism of the network conditions.

Based on the above mechanism, a SOLOR node compares the network conditions vs. the information utilised in the last re-configuration. In case conditions change (or when it is the first time the node is powered on), it multicasts a re-configuration message, which is extended by the other SOLOR nodes as they forward it with additional information (as described next). The re-configuration is triggered with a `net_reconf_req` message containing: (*i*) SOLOR ID, (*ii*) the estimated network conditions (i.e., estimated MCS between each pair of nodes known), (*iii*) the SOLOR operation parameters (α values), both from the node and clients, based on a default set of parameters or an estimation of the type of device (e.g., based on their MAC addresses, or the "Device Type" attribute of the Wi-Fi Protected Setup), (*iv*) its per-station power consumption figures ($\rho_{tx/rx/id/s}$) and those of the legacy clients it can hear (again, using a pre-defined set of parameters, or after an estimation), (*v*) the channel list where the relay can operate, and (*vi*) the timestamp when the re-configuration is issued.

A SOLOR node that receives a *new* `net_reconf_req`, updates its local database, updates the `net_reconf_req` by adding its local data, and multicasts this updated message with its own SOLOR ID. This simple controlled flooding protocol allows the SOLOR nodes to have a global view of the scenario, i.e., each relay knows the MCS for all potential links, and the individual preferences and per-state power consumption figures of the nodes (α_i, ρ^i_{tx}, ρ^i_{rx}, ρ^i_{id}, ρ^i_s), to run the algorithm with the same shared information. SOLOR relays record the timestamp of the initial re-configuration message, but do not immediately initiate the computation of the optimal configuration; instead, they wait T_1 seconds with no new messages to trigger the computation. This configuration is committed T_2 seconds after the timestamp, which guarantees synchronisation between SOLOR nodes. Note that T_2 has to be longer than the time it takes for the re-configuration message to reach all relays, plus the time to compute the optimal configuration.

A.5.2 Communication between SOLOR nodes

The operation of SOLOR relies on a mechanism to reliably deliver messages across all relays. To this aim, in our experiments we leverage the default multicast operation, as we found that it results extremely reliable due to the use of a robust MCS (i.e., 100% delivery rate). Still, for harsher network conditions, we could easily extend SOLOR with one of the mechanisms from the Group Addressed Transmission Service described in the recent 802.11aa standard, which specifies more reliable multicast services, as there is an implementation readily available [68].

The direct communication between SOLOR nodes, when one is acting as a parent for the other, results immediate, as they share the same schedules and therefore the transmitter knows when the intended destination can receive the data. However, when SOLOR nodes communicate through the (legacy) AP, they need to be associated with the AP long enough, so the multicast transmission is successfully forwarded from one SOLOR node to the other. To this aim, we fix a minimum amount of time that all clients have to be simultaneously connected to their parent, i.e., $F^{A_n}_{\mathcal{C}_{A_n}} = 10$ ms $\forall n \in \mathcal{S}$, and schedule multicast messages at the beginning of this time fraction.

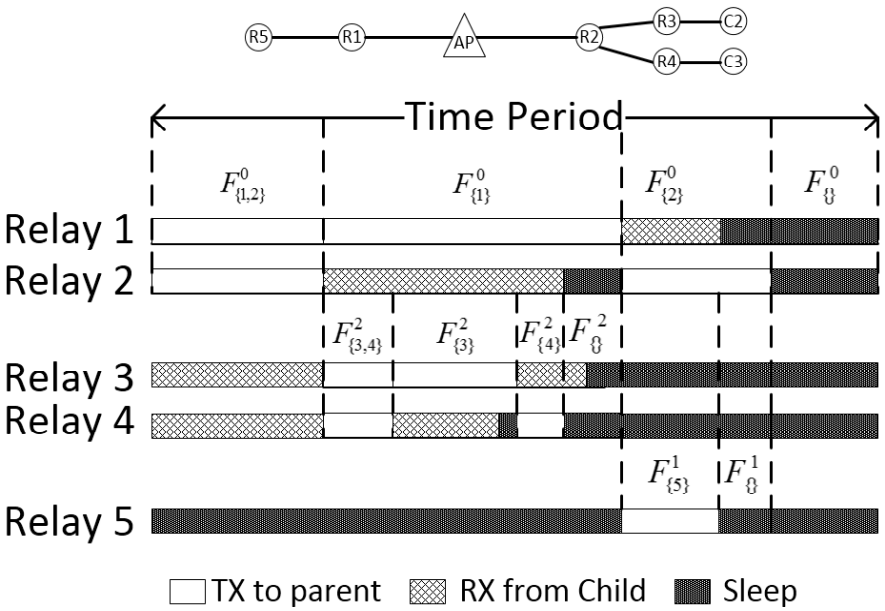

Figure A.15: Relay schedules computation for a 2-hops topology.

A.5.3 Computing a feasible schedule

To find a feasible schedule for the optimal configuration, we start with the relays one hop away from the AP, and then move one hop at a time (the schedule of the relays at the same number of hops from the AP can be computed in any order). For each relay $s \in \mathcal{S}$, we impose a deterministic ordering of the sets in $\mathbb{W}^s$, based on the size of the set (note that $\mathbb{W}^s$ does not include the empty set) and using the smallest node identifier (its MAC address) as a tiebreaker. We use this ordering of $\mathbb{W}^0$ to arrange the fractions $F^0_{\mathcal{V}}, \mathcal{V} \in \mathbb{W}^0$, which specifies the time periods when the children of the AP have to contend for access. Next, for each relay s one hop from the AP, we determine the rest of its schedule by splitting the time that s is not sending to the AP into the time fractions $F^s_{\mathcal{V}}$, ordered after the set $\mathbb{W}^s$ as well. The time left at the end of the schedule is the fraction of time the relay spends in sleep mode. Following this methodology, we find a feasible schedule that fulfils the requirements of the solution to the optimisation problem. Figure A.15 illustrates the above schedule computation for a scenario with five relays (R1–5) and two legacy clients (C1–2) with a 2-hops topology.

A.5.4 Applying the new configuration

Once the optimal configuration is found, the links between nodes must be configured. To force legacy nodes to disassociate from the AP and associate to the relay, we use a simple scheme based on the behaviour of most wireless network managers, which consists on the relay *forging* a disassociation message as if it were sent from the AP, thus forcing the legacy node to re-scan the network to look for the *best* AP announcing the same Service Set ID (SSID) to associate with. This AP should be the relay node, as it supports the use of better MCS and therefore has better link quality. For simplicity, in our experiments the client obtains a new Internet Protocol (IP) address after the re-association, but this could be prevented if the SOLOR relay sends a "gratuitous ARP" to the AP [69].

Finally, we need to ensure that legacy nodes go to sleep or, at least, do not transmit while the relay is not available (either sleeping or sending data to its parent). For simplicity, we use the *Notice of Absence (NoA)* protocol [70], specified for WiFi Direct and already present in many

current devices (e.g. Android phones), which allows the relay node to send a unicast packet to its attached clients with the relay's sleep schedule. We confirmed that other schemes also work, e.g., sending `null data` frames with the Network Allocation Vector set to the time the AP is not available, which enables the node to sleep for that period of time (we confirmed that old network interfaces overhearing all traffic do not go to sleep, but do not transmit neither).

A.6 CPS characterisation of WiFi D2D inband and optimisation

In this appendix we provide more details on the CPS-based characterisation of WiFi D2D inband mechanisms presented in Section 2.2.3. A basic introduction to CPS modelling has been presented in Appendix A.2.

A.6.1 Analysis of ad hoc wireless networks

Here we present an analytical method for the analysis of ad hoc networks under non saturation. First we introduce the *coupled processor* model, by which we capture the coupling in performance which characterises ad hoc networks. Then we present our main results, based on a worst case approach to coupled processors systems. Finally, we describe a practical method for deriving performance bounds, based on an approximation technique which adapts to the characteristics of the system under study and of the specific performance problem to solve.

A.6.1.1 The underlying CPS model for ad hoc wireless networks

A crucial aspect of performance analysis in an ad hoc network is to being able to capture the effect of traffic of a given user on the performance experimented by other users. In what follows, we address this issue by modelling these interaction by means of a *coupled processors* model [25]. A CPS is a set of parallel queues (i.e., queues which do not exchange traffic among them) served by work conserving schedulers, and whose service rates at any time t is completely determined by the set of active queues at that time. We define as the *state* of the system at a given time t the array $\mathbf{I}(t) = (I_1(t), I_2(t), ..., I_N(t))$ where for each node i, $I_i(t)$ is a binary variable which is equal to 0 if the queue at the $i-$th node is empty at time t, and 1 otherwise. Then at time t the service rate of the i-th queue in $\mathbf{I}(t)$ is $R_i(t) = R_i(\mathbf{I}(t))$, i.e., it is only function of the state of the system at time t.

In what follows, we model our N-node ad hoc network as a N-queue CPS, with one queue per transmitter. In such system, the coupling is in users transmission rates. Such coupling arises from sharing the same transmission medium, and is mediated by the Carrier Sense Multiple Access (CSMA)/Collision Avoidance (CA) algorithm. The state $\mathbf{I}(t)$ of such CPS is given by the set of active ad hoc transmitters at time t. In modelling our ad hoc network as a CPS, we assume that for each ad hoc transmitter i at time t, the service rate $R_i(t)$ is completely determined by the state of the system. For each state of the system, we characterise the underlying CPS through the saturation throughput of the subset of active nodes, derived in [30]. That is, $\forall t \geq 0$, if $\mathbf{I}(t)$ is the set of active transmitters at time t, the instantaneous service rate at the i-th active transmitter is given by:

$$R_i(t) = \frac{P_s P_{tr} E(P)}{(1 - P_{tr})\delta + P_{tr} P_s T_s + P_{tr}(1 - P_s) T_c}. \tag{A.18}$$

Here, T_s is the average time the channel is sensed busy because of a successful transmission. T_c the average time the channel is sensed busy by each station during a collision; P_s is the probability that a transmission occurring on the channel is successful, which is given by the probability that exactly one station transmits on the channel, conditioned on the fact that at least one station transmits. P_{tr} is the probability that there is at least one transmission in the considered slot time. $E(P)$ is the average packet payload size, and δ is the duration of an empty slot time. All these parameters can be computed directly from the parameters of the CSMA/CA protocol, and they refer to the particular set of active transmitters $\mathbf{I}(t)$. For the expressions of each parameter, please refer to [30].

Indeed, in an ad hoc network the instantaneous service rate is determined by the CSMA/CA algorithm, and for a same system state (set of active queues) it generally varies over time. Therefore, as it is common in the study of such systems [30], in adopting a CPS model for such ad hoc

network we are assuming that those system dynamics due to the CSMA/CA mechanisms take place on a smaller time scale than traffic dynamics (more specifically, the events of queues getting empty or full), so that they can be adequately modeled through their average effect on the system. Furthermore, while the ad hoc transmissions are scheduled sequentially, via contention, and the change of system state can happen only when a new host is scheduled, we assume the equivalent CPS model works at the "fluid" limit, i.e., each queue serves its traffic as it were infinitely divisible, and it can change state at any time t. In Section A.6.3 we assess the validity of such assumption, showing numerically that these approximations model accurately the performance of our ad hoc network. The analysis of the presented underlying CPS for ad hoc transmissions is presented in Appendix A.2.

A.6.2 Derivation of the optimal operating point

A.6.2.1 Problem Formulation

As already discussed, choosing the leaky bucket parameters of traffic sources in our networks allows tuning the operating point of the system, possibly in order to maximise some utility function. For instance, in order to guarantee that some form of fairness is maximised. In the present work, the utility function that we choose to optimise is a weighted fairness function, which is one possible way of balancing some notion of fairness among users with, for instance, different classes of service. Its expression is:

$$U = \sum_{i=1}^{N} w_i \log\left(\frac{\rho_i}{\rho_0}\right), \tag{A.19}$$

where ρ_0 is the minimum bit rate for an acceptable performance for the application. The feasible set of leaky bucket rates over which to optimise such utility is given by the set of inequalities in Theorem 4 in Appendix A.2, as they define the set of rates for which the ad hoc network is able to serve the traffic load with a finite maximum packet delay and backlog at each node.

The feasible operating points which maximise the weighted fairness are therefore the solutions of the following optimisation problem, computed over the set of auxiliary networks and of fresh traffic leaky bucket rates:

$$\begin{aligned}
&\underset{\rho\geq\mathbf{0},\mathbf{n}\in\mathcal{N}}{\text{maximise}} \quad \sum_{i=1}^{N} w_i \log\left(\frac{\rho_{n_i}}{\rho_0}\right); \\
&\text{subject to:} \\
&\forall \mathbf{n},\ \forall j = 1,\ldots,N, \\
&\rho_{n_j} + \min\left(R^{up}_{n_j} - R^{sat}_{n_j}, \sum_{p=1}^{j-1} \frac{R^{up}_{n_j} - R^{p-up}_{n_j}}{R^{up}_{n_p}} \rho_p\right) \leq R^{up}_{n_j},
\end{aligned} \tag{A.20}$$

where the constraints derive from Theorem 4. $\mathcal{N}$ is the set of the $N!$ possible permutations of the labels of ad hoc users (i.e., CPS queues).
Solving this problem is challenging for two main reasons: (i) the presence of the min function in the constraints, which leads to a non-convex feasibility region; and (ii) the complexity of the problem, which scales factorially with N.

A.6.2.2 Heuristic Approach

In order to practically solve the above problem, we propose a heuristic approach that consists of two parts.

The first part aims at reducing the problem to a tractable problem that can be solved with standard tools. Specifically, since the presence of the min function in the constraints of the above problem leads to a non-convex feasibility region, we use the so called *big-M transformation* [59]. In such way, the two terms of the min in (A.20) are not *active* at the same time. Instead of that constraint, the method builds two constraints in which we add a binary variable multiplying a large constant value M. Whenever the binary variable is equal to one, the large constant makes the constraint useless because all the feasible sets of leaky bucket rates satisfy it, while when the binary variable is zero the constraint is active. By choosing a value for the binary variables we select a part of the feasibility region.

The problem obtained in this way belongs to the mixed-binary programming family. We solve it by means of the *branch-and-bound* method [60]. For each $\mathbf{n}$, we stop the branch-and-bound evaluation when the intermediate solution is at most at ϵ away from the optimum. Tuning ϵ allows achieving different tradeoffs between computational cost and optimality of the solution.

The second part of our heuristic aims at reducing the number of auxiliary networks over which to search for the optimum. It is based on running a set of greedy searches from a set of starting points, each of which has been derived as follows. To each node $j = 1, ..., N$ of the CPS, we associate the quantity $\frac{1}{w_j}$. Then, starting from the first stage of the auxiliary network, we assign a node of the CPS to each stage with a probability proportional to this quantity. The idea underlying such algorithm for the choice of the starting points is that nodes with higher weights w_i in the utility function need to be modeled more accurately than the others. This is achieved by assigning those nodes to the last stages of the auxiliary network. Indeed, due to the structure of the auxiliary network, the lower the stage a node belongs to, the larger the set of nodes whose coupling with the considered one is modeled through accurate rescaled traffic rather than via a conservative penalty on the service rate, holding for any time t, and therefore independent on traffic patterns at interfering nodes.

We describe now the elementary step of the search. From a network $\mathbf{n}$, we consider the set of $N - 1$ networks obtained by a swap of two contiguous nodes in $\mathbf{n}$. If the log utility value of $\mathbf{n}$ is lower than the max log utility among all these $N - 1$ networks, the network with the highest log-utility value among all the $N - 1$ is selected. Otherwise, the search stops. The largest of all local maxima computed from all starting points is the final output of our heuristic. By changing the number of starting points we can achieve different trade-offs between computational cost and optimality of the solution.

A.6.3 Numerical Evaluation

In this section we evaluate the fitting of the proposed CPS model for a WiFi network. In particular, we evaluate the impact of assuming the saturation throughputs in (A.18) as the service rates of the CPS equivalent model. The parameters used in the considered scenarios are presented in Table A.8. We have chosen the 802.11 b/g standard for WiFi communications. The values of the parameters are derived from [71].

A.6.3.1 CPS Model Validation

In this section, we evaluate the fitting of the CPS model for 802.11 communications. In particular, for different scenarios, we pick at random a stable set of arrival rates for the transmitters. Such decision ensures that the set of active transmitters present in the wireless scenario, i.e., the system state changes over time. Then, for each system state and for each transmitter, we evaluate the difference among the average throughput achieved during simulations and the service rate used in the CPS model, i.e., (A.18). Figure A.16 presents the probability density function of the differences

Table A.8: Setup of the Wireless Scenario

Available Channel Speeds (Mbit/s)	1;2;5.5;6;9;11;12 18;24;36;48;54
CWmin	16
Backoff stages	5
Preamble + PHY header(μs)	20
SIFS (μs)	16
ACK Time (μs)	24
DIFS (μs)	34
Slot time (μs)	9
MAC header(bits)	224
Chunk size (bits)	15000

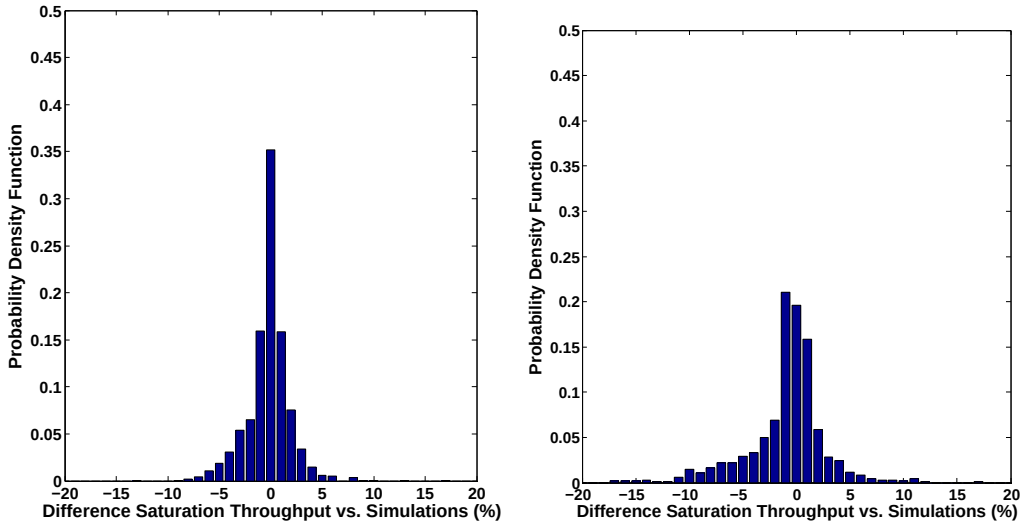

(a) 3 Nodes. (b) 7 Nodes.

Figure A.16: Probability density function, difference Simulations vs. CPS service rates.

for scenarios having 3 and 7 nodes. Similar results have been obtained also for different set-ups of the wireless scenario.

In both cases, the service rates used in the CPS modelling are close to the ones computed during simulations. In particular, in the 99.73% and the 97.26% of the cases, respectively, the absolute value of the difference among the service rates computed as in (A.18) and the corresponding ones achieved through simulations is less than the 10% of the ones achieved through simulations. Even if the system assumes a given state just for a short interval of time, i.e., even if the subset of active transmitters does not change just for a limited period, the above result shows that the saturation throughput (A.18) is reached fairly soon in almost all the cases. Therefore, we consider as negligible the impact of the approximations introduced during the modelling of the system as a CPS.

A.6.3.2 Proportional Fairness Optimisation Results

In this section we evaluate the results achieved through the optimisation problem (A.20) in a large set of scenarios. In each setting, we solve (A.20) exploiting our heuristic, computing for a given number of users the set of leaky bucket rates which maximise the log-utility function.
In the following, the weights w_i of the utility function, defined in Section A.6.2.1, are uniformly distributed in $[0, 1]$. Instead of considering a specific propagation model, and modelling its impact on the achievable channel speeds, for each user we assigned channel speed randomly, assuming speeds to be uniformly distributed among the set of available rates in Table A.8. We have set the parameter ϵ for the branch-and-bound algorithm to 5%. Furthermore, for a given number of users N, we considered a number of starting points for the heuristic, i.e., a number of upper bounding networks, which scales with N. Indeed, as N grows, it increases also the solution space, as well as it does the space of possible values of the weights of the utility function, and of the channel rates. Empirically, and for the number of users considered in our evaluations, we have found that scaling the number of starting points as $\lceil \frac{N}{2} \rceil$ brought acceptable results in terms of output of the optimisation and of computational complexity.

Overall, for each value of N we considered a total number of instances of our setting (i.e., a particular choice of starting points, set of weights and set of channel rates) sufficient to get a 95% confidence interval within the 15% of the value of the average utility U achieved. In any case, we never used less than 100 instances.

A first objective of our numerical evaluation has been to assess the performance of the proposed heuristic, which has been introduced as a computationally feasible approach to the problem of maximising the (weighted) proportional fairness in the allocation of leaky bucket rates among ad hoc users. More specifically, we have tried to give an idea of how far are, on average, the solutions of our heuristic from the optimal values, in order to evaluate the impact of the approximations on which the heuristic is based.

To this end, for scenarios with a small number of nodes (for which an exhaustive search still brings to an acceptable computational complexity), we have compared the average log-utility derived through our heuristic with the one derived through exhaustive search over all the possible upper bounding networks and choosing a $\epsilon = 0\%$. From Table A.9 we observe that, in the considered scenarios, the solutions from our heuristic bring a utility which is on average very close to the optimal values derived through exhaustive search. This suggests that the approximations on which our heuristic is based have an overall low impact on the optimality of the operating point derived. Log-utility and average throughput obtained through the presented simulations have been shown in Section 2.2.3.

In order to understand the feasibility of the proposed approach, we analysed the complexity of the heuristic we propose versus the exhaustive search of Table A.9.

Table A.9: Average Log-Utility: Heuristic vs. Exhaustive Search

Opt. technique	3 Nodes	4 Nodes	5 Nodes	6 Nodes
Ex. Search	3.4690	3.0688	2.8034	2.3933
Heuristic	3.4690	3.0590	2.7879	2.3630
Difference	0	-0.32%	-0.55%	-1.27%
Max Difference	0	-6.56%	-10.64%	-15.88%

Table A.10: Complexity of the Heuristic

	Average Number Network - Heuristic	Maximum Number Network ($N!$)
3 Nodes	6	6
4 Nodes	17.15	24
5 Nodes	41.04	120
6 Nodes	59.22	720
11 Nodes	409.05	$39.9\ 10^6$
	Av. Per-Network branches - Heu.	Av. Per-Network branches - Exhaustive Sear.
3 Nodes	4	4
4 Nodes	6.49	8.32
5 Nodes	9.47	12.87
6 Nodes	11.89	17.80
11 Nodes	29.11	-

In Table A.10 we present first the average number of networks the two methods analyse in order to get the final log-utility of the system. In case of the heuristic, we count the total number of upper bounding networks analysed, considering all the $\lceil \frac{N}{2} \rceil$ starting points. We also present the number of optimisations problems which the branch-and-bound method solves for each network. It can be easily proven that, in the worst case, the number of branches visited by the branch-and-bound for N nodes is $\sum_{i=1}^{N-1} 2^{N-i}$. Results are shown up to the point at which the comparison is computationally feasible (6 nodes), and for the largest scenario we analysed through our heuristic (11 nodes).

We can see how the heuristic requires a considerably inferior number of evaluations (both in terms of upper bounding networks, both in terms of branches), with a very limited impact on the optimality of the value of the log-utility derived (as seen in Table A.9). Please note that the different starting points of the heuristic are completely independent from each other. Therefore, the computational time can be reduced sensibly if the heuristic is evaluated in parallel.

Table A.11: Dependence of Results on ϵ and on Number of Starting Points (SP) heuristic

	4 Nodes	7 Nodes
U, $\epsilon = 5\%$, $SP = \lceil \frac{N}{4} \rceil$	-0.64%	-0.92%
U, $\epsilon = 0\%$, $SP = \lceil \frac{N}{2} \rceil$	$+0.02\%$	$+0.47\%$
U, $\epsilon = 10\%$, $SP = \lceil \frac{N}{2} \rceil$	-0.02%	0%
U, $\epsilon = 5\%$, $SP = N$	$+0.24\%$	0.74%
Weight. Av. Th., $\epsilon = 5\%$, SP $= \lceil \frac{N}{4} \rceil$	-1.99%	-5.54%
Weight. Av. Th., $\epsilon = 0\%$, SP $= \lceil \frac{N}{2} \rceil$	$+0.11\%$	$+0.44\%$
Weight. Av. Th., $\epsilon = 10\%$, SP $= \lceil \frac{N}{2} \rceil$	-0.08%	0%
Weight. Av. Th., $\epsilon = 5\%$, SP $= N$	$+1.10\%$	$+3.78\%$

Finally we evaluate how the results change varying the number of starting points and the value of ϵ. Also here, we evaluate all the scenarios shown antecedently when 4 and 7 nodes were present in the network. The results presented in Table A.11 use the setting having $\epsilon = 5\%$ and $\lceil \frac{N}{2} \rceil$ starting points as a benchmark.
Even though the utility U remains almost the same in every configuration of the heuristic, the weighted average throughput changes sensibly. The most affecting parameter of the heuristic is clearly the number of starting points used. Increasing the number of starting points indeed, U and the weighted throughput increases accordingly. Unfortunately, the number of upper bounding networks increases 67.24% and 106.10%, respectively, leading to a computationally expensive resolution of the proposed optimisation. On the other hand, considering as acceptable the complexity of the proposed heuristic, reducing the number of starting points results in an unnecessary lost in performance, that can become tricky when the number of nodes increases. The same reasoning applies when we evaluate the choice of ϵ: sometimes the difference of performance is too small to justify the decreasing of the value of ϵ. For instance, the reduction of branches analysed when $\epsilon = 10\%$ is of at most the 20.03%, while the accuracy of performance estimates remains similar. The presented heuristic represents, therefore, a good trade-off among complexity and performance.

A.7 Optimisation of LTE/LTE-A networks with macro and small cells

In this appendix we provide the reader with details on the analysis of the problem and the formal problem statement for the HetNet optimisation presented in Section 2.3.1.

A.7.1 Problem statement

We consider a HetNet which consists of macro and small cells. Denote M as the set of macro cells, S the set of small cells, and $K = M \cup S$ the set of all the cells in the HetNet. Let N_k be the set of eNBs forming the neighbouring an eNB $k \in K$. The CIO value from eNB i to eNB j is denoted by $CIO_{i,j}$. Each base station has T sub-frames in the time domain and F RBs in the frequency domain. The duration of time sub-frames is equal and the bandwidth of all RBs is also a constant. A physical resource block $PRB_{t,f}$ is denoted by its time sub-frame index t and RB index f, and we denote $B = T \times F$ as the total number of PRBs available at each eNB, where $t \in T$ and $f \in F$.

Consider that the sub-frames and RBs of base stations are synchronized. Let the binary variable $v_{k,t,f}$ specify the time and frequency pattern of each cell k. It indicates whether cell k is allowed to use $PRB_{t,f}$ or not. Each cell can allocate power in a plurality of discrete quanta $\Delta > 0$. We denote by $P_{k,t,f}$ the amount of power allocated by cell k for $PRB_{t,f}$. All cells operate under this constraint, $\forall k \in K$, $\forall t \in T$, we have

$$\sum_{f \in F} P_{k,t,f} \leq P_{k,\max}$$

where $P_{k,\max}$ is the maximum total transmitting power for cell k.

Denote U_k as the set of users who are associated with eNB k, where $U_k \subseteq U$. Let the binary variable $q_{u,k}$ indicate whether user u is in U_k. Each user can be served by only one cell, i.e.,

$$\sum_{k \in K} q_{u,k} = 1, \forall u \in U.$$

Each user has a throughput utility function. The problem to solve can be expressed as follows. Given a network state $\mathbf{s}$, which has three elements: the CIO, time and frequency pattern, and transmit power of each eNB, we look for the optimal values of those parameters so as to maximise a utility function $Ut(\mathbf{s})$:

$$Ut(\mathbf{s}) = \sum_{k \in K} Ut_k(\mathbf{s})$$

where Ut_k is the utility of the eNB k.

Using a Proportional Fairness (PF) approach, this utility can be expressed as follows:

$$Ut_k = \sum_{u \in U_k} \log(r_u)$$

where r_u is the throughput of the user u attached to eNodeB k.

As indicated by the Shannon-Hartley capacity, we choose the logarithmic function given by

$$r_u = C(SINR_u(\mathbf{s})) = c_0 \log_2(1 + SINR_u(\mathbf{s}))$$

where c_0 is the bandwidth of the channel in Hz. The SINR of user u when served by cell k on $PRB_{t,f}$, is equal to

$$SINR_{u,k,t,f} = \frac{v_{k,t,f} P_{k,t,f} G_{u,k,t,f}}{\eta + \sum_{l \neq k} v_{l,t,f} P_{l,t,f} G_{u,l,t,f}}$$

with η representing the thermal noise, $G_{u,k}$ the link gain from cell k to user u on PRB (t, f).

We now formulate the downlink resource allocation problem where the best way of allocating PRBs is to be determined subject to the different parameters of the eNBs: CIO value, time and frequency pattern, and transmit power level. Let the binary variable $x_{u,k,t,f}$ indicate whether $PRB_{t,f}$ is allocated to user u by its serving station or not, with the constraint

$$\sum_{u \in U_k} \sum_{f=1}^{f=F} x_{u,k,t,f} \leq F, \forall t.$$

The total throughput received by user u is given by:

$$\begin{aligned} r_u &= \sum_{k \in K} q_{u,k} \sum_{1 \leq t \leq T} \sum_{1 \leq f \leq F} x_{u,k,t,f} r_{u,k,t,f} \\ &= \sum_{k \in K} q_{u,k} \sum_{1 \leq t \leq T} \sum_{1 \leq f \leq F} x_{u,k,t,f} c_0 \log_2(1 + SINR_{u,k,t,f}). \end{aligned}$$

To summarize, our goal is to maximise the sum of logarithm of throughput utility of the users with the following decision variables in the scheduling and assignment problem:

- The assignment of users to cells is represented by the variables $q_{u,k}$;
- The time and frequency pattern $v_{k,t,f}$ to indicate whether the cell can transmit or not;
- The power control $P_{k,t,f}$ defines the power to be used by each cell on each PRB;
- The allocation of resource block to users within each cell is specified by the variable $x_{u,k,t,f}$.

We want to solve the following optimisation problem:

$$Maximise \quad \sum_{k \in K} \sum_{u \in U_k} \log(r_u)$$

subject to

$$\begin{aligned} & r_u = \sum_{k \in K} q_{u,k} \sum_{1 \leq t \leq T} \sum_{1 \leq f \leq F} x_{u,k,t,f} c_0 \log_2 \left(1 + \frac{v_{k,t,f} P_{k,t,f} G_{u,k,t,f}}{\eta + \sum_{l \neq k} v_{l,t,f} P_{l,t,f} G_{u,l,t,f}} \right), \\ & \sum_{k \in K} q_{u,k} = 1, \forall u \in U_k, \\ & \sum_{f \in F} P_{k,t,f} \leq P_{k,\max}, \\ & \sum_{u \in U_k} \sum_{f=1}^{f=F} x_{u,k,t,f} \leq F, \forall t. \end{aligned}$$

A.7.2 Proposed generic framework

An approach based on game theory has been proposed to determine the optimal configuration in iterative approach and enabling distributed scheme for both 3G [36] and 4G [42]. Each UE reports to its serving eNB long-term statistics (such as wideband CQI, Peak Signal to Noise Ratio (PSNR)) relative to the neighbour's eNBs. These statistics are then explicitly shared between the serving eNB and its neighbours so that each eNB can make, iteratively, decision that is not purely selfish but also considers the impact to its neighbours. The iterations consist in the successive selection of various eNBs. Once an eNB has been selected, a sampling (testing various options of CIO and ABSF) is

performed and one sample is selected (the best is selected either deterministically or probabilistically for example with respect to Gibbs distribution). After a number of iterations, utility function converges to a value. Depending on the selection method, this value can be guaranteed to be a local or global optimum. The framework proposed by the game theory is robust enough for the exploration needs. Recent paper shows that many enhancements are still possible for further advancement [39].

In game theory, the best response is the strategy which produces the most favorable outcome for a player, taking other players' strategies as given [72]. The concept of a best response is central to John Nash's well-known result, the Nash equilibrium, the point at which each player in a game has selected the best response to the other players' strategies [73]. To our optimisation problem, the idea is as follows:

- Starting with any arbitrary initialization;
- At each internal algorithm iteration, joint parameters are chosen that maximise utility function;
- Afterward, the system converges to a Nash equilibrium that could be a local optimal solution of the network utility.

While several algorithms are efficient, some algorithms can drive the system to a state of strict optimal solution by proper probabilistic relaxation [74]. Note that this approach can be implemented in a distributed fashion, however in practice it may be useful to support also more centralized architecture: distributed implementations require modification of the eNodeB code and also some exchange between the eNBs to converge to global optimum. To adapt to these practical constraints, we propose a more pragmatic framework that is a tradeoff between complexity and performance and also for today's C-RAN/V-RAN: a centralized coordinator using the best response.

In order to adapt this generic theoretical approach described in [39] to the particular case of eICIC deployment context, we use a specific framework that fits with the limitations of existing LTE cellular networks:

- Limited CPU capacity on eNB;
- Limited exchange capacity between the eNBs;
- Precise state knowledge known locally by each eNB;
- Introduction of a central entity in charge of performing the optimisation and called the "coordinator";
- Support of re-centralization of the optimisation computation.

In this model, the computation of the optimal CIO and ABSF values can be left to a central entity called the coordinator. In this case, the coordinator will perform the following Step 1 to 3:

- Step 1: Collect statistics from all eNBs (macro and metros), those statistics being collected from all the UE reports and processed by the eNB; for instance these statistics can be RSRP measurements from a group of UEs to a list of neighbouring (interfering) sectors. The key point is that the UEs can be grouped in groups (pools) having similar radio propagation, hence enabling lower data exchange between the local eNB and the central coordinator. We call this step the "UE grouping" and various examples of groupings are presented in the following Figure A.17. This approach leverages the eNB "local" knowledge since the eNB is the best point to qualify the users. At the central coordinator, this collection constitutes a database that reflects the state of the network. This Step 1 is therefore called the "collection

step". This 2-tiered approach is to enable a good balance between the local knowledge of the eNB (that is precise and real time but limited in scope to its attached UEs) and the global knowledge of the coordinator (that has access only to long-term statistics which are averaged in both time dimension and spatial dimension by the notion of UE grouping but in a system-wide scope).

- Step 2: Work on the database localized in the coordinator, the latter derives the best parameter values for CIOs and ABSF patterns using iterative algorithms; this step is therefore called the "optimisation step".
- Step 3: After the optimisation is done, the coordinator will send back to every eNB the optimised values. Each eNB will "execute" the selected parameters. This step is called the "execution step".

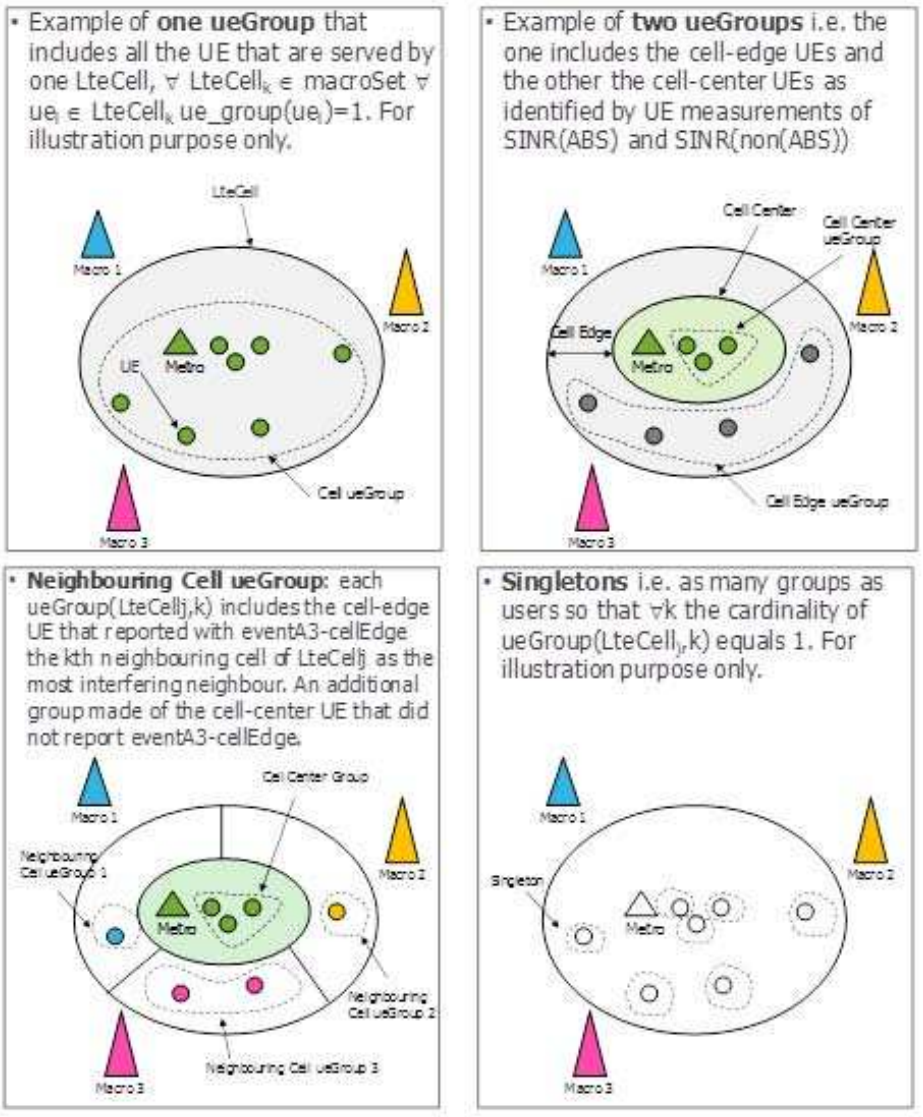

Figure A.17: Examples of various UE grouping policies

In the general case, the sampling may be partially distributed or centralized. In this case, a computing element function C is defined such that $C(c)$ gives the address of the element in charge

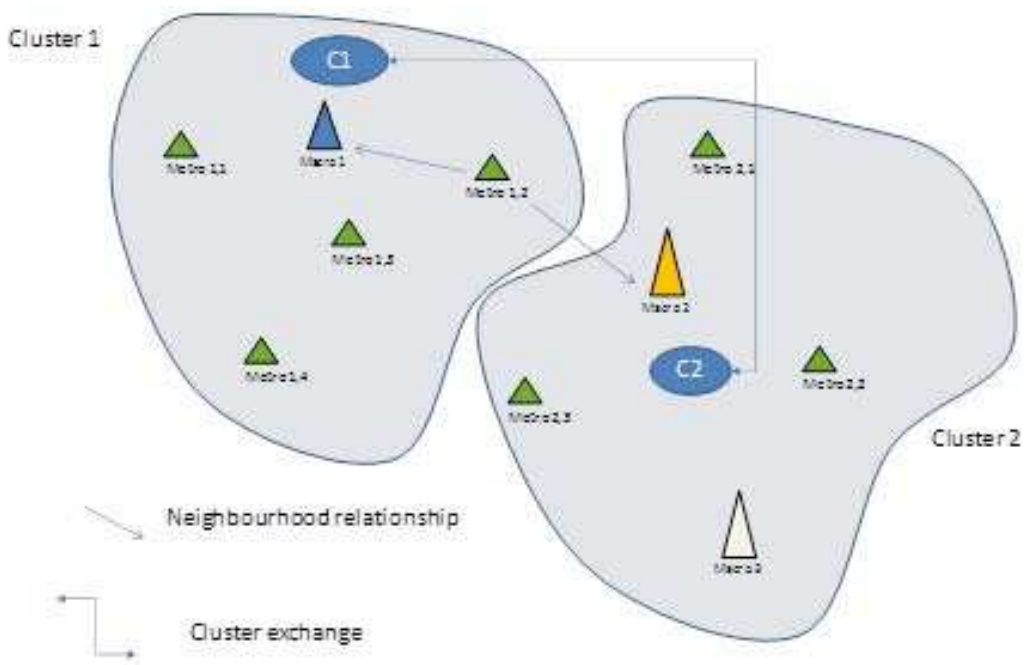

Figure A.18: Interaction between coordinator clusterization and neighbourhood patterns

of performing the sampling of cell c. Let us revisit the possible options:

- In the centralized case, $C(c)$ = coordinator whatever c is.
- In a partially centralized case, $C(m) = M(m)$ for all metro m belonging to its macro $M(m)$.
- In the fully distributed case, $C(c) = c, \forall c$.

In the case of distributed operation, whenever a cell $c1$ is selected, one has to check if there exists in the neighbourhood of cell $c1$ another cell $c2$ such that $C(c2) \neq C(c1)$. If this is the case, $C(c1)$ must give to $C(c2)$ the new parameters selected at the end of the iteration. Otherwise, there is no exchange needed. Therefore, a trade-off of the cluster size between the computing load on the cluster computing element and the level of message exchange between two iterations is required, as depicted in Figure A.18. Specifically, in the example shown in Figure A.18 we can see that whenever the Metro 1.2 gets updated during an iteration, its coordinator C1 must send its new parameters to the coordinator C2. Similarly, whenever the Macro 2 gets updated during an iteration, the coordinator C2 must send its new parameters to the coordinator C1. But no exchange has to occur for all other cells. The algorithm in the simpler case of the fully centralised and best response approach is represented in Figure A.19, which is organised as follows:

- At each iteration we pick up randomly a sector;
- Depending on the sector type (macro/metro), the parameters that can be sampled are: for macro sector CIOs and ABSF pattern can be modified whereas for metro sectors only the CIO can be modified;
- A sampling is performed on the admissible range of the variables; note that some sampling combinations may be discarded for various reasons (e.g. race condition detection, power limitations) therefore limiting the “sample space”;

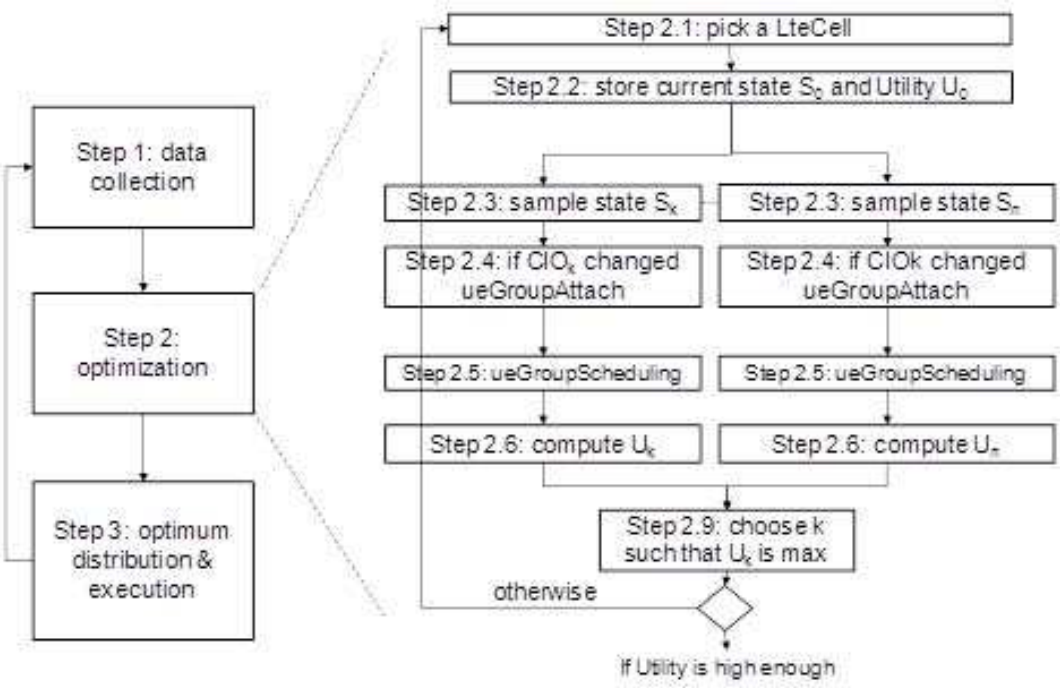

Figure A.19: Algorithm overview

- For each sampled case where the CIO has changed, we perform a "virtual handover" by calling the ueGroupAttach function. This functions tests if the user group has to support a handover to a neighbouring sector due to the change in the CIO;
- If some handovers have occurred and/or some ABSF ratios have changed, we perform the "virtual scheduling" via the ueGroupScheduling function call. This function renders the scheduling performed by each eNB (the one being sampled plus the neighbouring ones) in order to get an idea of the resulting bit rate. Several optimisation criteria options are available such as proportional fairness, absolute fairness, sum rate maximisation, etc. In order to keep low complexity, we need to schedule only a limited number of eNBs (the group constituted of the sampled and its neighbours). Defining the cluster size of the virtually scheduled eNBs is matter of research.

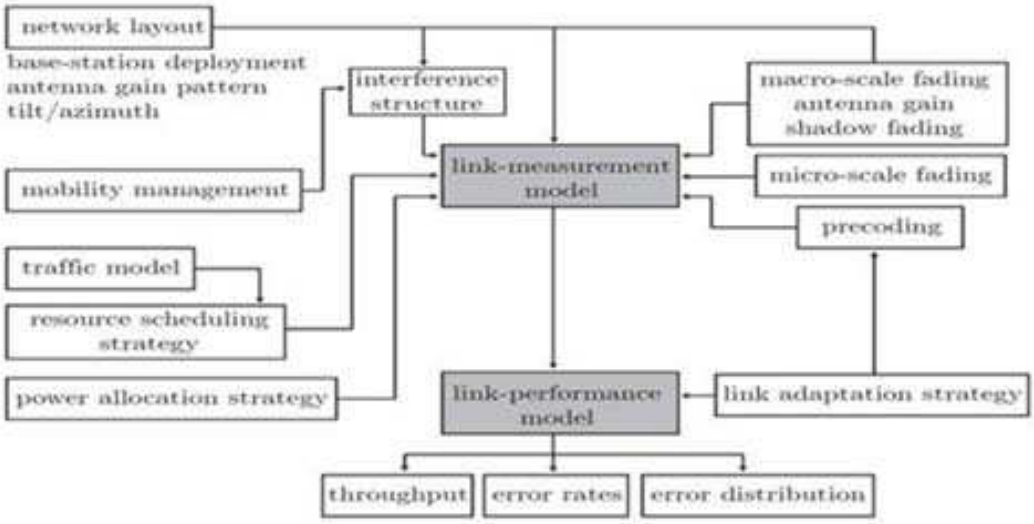

Figure A.20: LTE system level simulator

A.7.2.1 Simulation

For assessing the performance of the proposed optimisation scheme for HetNets, we have used a simulation-based approach. In particular, we have used a MATLAB-based LTE-compliant simulator developped by the TU Wien's Institute of Telecommunications [43]. Figure A.20 presents a scheme of the simulator, whereas Table A.12 presents the main simulation parameters used for producing the results previously shown in Section 2.3.1.

Table A.12: Simulation environment for the assessment of optimisation techniques for HetNets

Simulation parameters	
Waveform	10MHz LTE rel8 SISO
Propagation	TS 36.942 recommended pathloss models + shadow fading
Macro cell placement	Hexagonal 1 ring, ISD 500m
Macro cell antenna	Kathrein
Small cell placement	At fixed location(s) 0.5 ISD 30 SC
Small cell antennas	omnidirectional
Cell load	25 UEs per macro sector, 525 UEs in total
UE speed	fixed
Measurement error modeling	No error
Traffic	Full buffer model
Optimization sampling	$CIO \in \{0; 10\}$; $ABSratio \in \{0, 20, 30, 40\}$ Only 1 iteration per macro sector

A.8 D2D mode selection

In this appendix we provide the reader with details on the design of the D2D mode selection optimisation strategies that we have discussed in Section 2.3.2.

A.8.1 D2D modes

There have been extensive research efforts in both academia and industry to explore D2D techniques [21]. D2D communications have been considered for a large variety of use-cases such as multicasting [45], cellular offloading [75], mobile relaying [76], load balancing [77], and video streaming [78]. These studies indicate the potential outstanding gain of D2D communications in cellular networks. Indeed, the high performance gain motivated leading telecommunication companies such as Qualcomm to perform experimental studies on this paradigm using early stage prototypes [26]. Standardization bodies such as 3GPP have also joined this front by considering D2D communications as a public safety feature in the next release of LTE-A [48]. These efforts from academia, industry, and standardization bodies confirm that the society regards D2D communications as a crucial feature for next generation networks. Nevertheless, there is still no concrete agreement on D2D operational details such as which medium access control to adopt, or which spectrum allocation schemes, connection setup, and resource management protocols are to be implemented. Initial proposals for D2D communications aimed at re-using the same resources that are used for conventional cellular communications (i.e., *inband underlay* D2D mode) [47]. The significance of the D2D gain had led to proposals in which a part of the cellular resources is dedicated only to D2D communications (i.e., *inband overlay* D2D mode). Finally, the scarcity and the high price of cellular spectrum motivated some researchers to explore D2D communications over the unlicensed band (i.e., *outband* D2D mode). These D2D modes are schematically illustrated in Figure A.21.

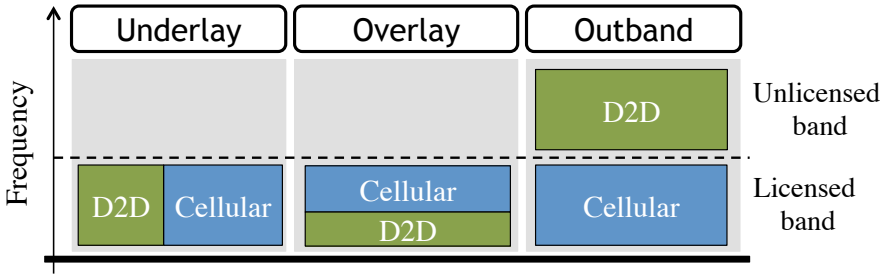

Figure A.21: Schematic representation of overlay inband, underlay inband, and outband D2D for cellular scenarios.

The majority of the existing studies on D2D communications select one of the aforementioned modes, then propose a method for resource allocation/interference management in order to handle the resulting complications, and finally illustrate the achievable performance improvement [79, 76, 47, 45, 75]. However, single mode D2D significantly limits the system performance to the interference profile of the network. Existing multi-mode D2D systems [80, 81] only focus on inband D2D modes, i.e., fully dependent on cellular spectrum. Other proposals focus on joint scheduling and mode selection [80, 81], although they are extremely complex (more complex than scheduling, which is already proven to be NP-hard for cellular systems such as LTE [49]) and introduce unnecessarily frequent mode selection decisions.

Interestingly, while some researchers limit D2D communications to cellular spectrum, the standards have a more liberal view of D2D. In fact, 3GPP defines D2D as "the communication between two users in proximity using a direct link between the devices without traversing the eNB(s) or the core network" [48]. We also remark that network-assisted outband D2D is accounted for in

3GPP Proximity-based (Prose) services [82]. Although both inband and outband D2D are considered valid options for Prose services, there is no indication on how to select between the two. Hence, given the fast-track emergence of D2D communications in cellular networks, the need for an adaptive D2D mode selection scheme is beyond question.

Here, we propose a flexible framework to adaptively select D2D mode or operating band and technology, which we name Floating Band D2D. In particular, we first discuss the practical implications of each D2D mode based on the latest standard releases of LTE-A and WiFi-Direct. This discussion clarifies that *there is no superior D2D mode* and the potential of each mode is highly scenario/use-case dependent. After discussing practical implementation issues of D2D-enabled networks, we provide analytical insights into the mode selection problem in an innovative multi-mode multi-band setup, which accounts for both achieved *throughput* and *energy costs*. We call such a novel approach *Floating Band D2D*, because D2D transmissions can occur on either inband or outband modes. The problem is formulated as a non-linear integer programming problem. Given the NP-hardness of the problem and time-stringent requirements of future cellular networks, e.g., 5G networks, we propose three practical heuristics with near-optimal performance and low complexity. Finally, we evaluate the performance of the proposed heuristics in a multi-cell scenario using a realistic setup designed based on the ITU-R guidelines for evaluating IMT-Advanced networks [83]. Our results confirm that the coexistence of D2D modes immensely ameliorates the performance of the system in terms of the key performance factors such as throughput and utility (up to one order of magnitude), and near complete fairness.

A.8.2 Practical limitations

In a D2D-enabled network, the eNB requires CSI between each pair of users (i.e., user-to-user CSI) in addition to user-to-eNB CSI in order to perform MCS assignment and scheduling. However, the existing cellular technologies do not have the means to obtain user-to-user CSI. Hence, we need a mechanism to obtain and send this information to the eNB efficiently because the addition of user-to-user CSI imposes high signaling overhead to the system.

CSI measurement. In LTE , the eNB-to-user CSI is estimated by active measurements from the received signal strength. However, there is no signaling message exchange between the users. Therefore, some researchers propose probing techniques to perform CSI estimation among users [80]. This approach imposes even higher signaling overhead to the system. In contrast, we propose an adaptive passive CSI estimation between users as explained in what follows. In LTE, each user has a unique ID (i.e., C-RNTI [84]) and this ID is included in the frame header. Thus, the users can detect the ID of the source of interference at each frame. Alternatively, the user can read C-RNTIs from the broadcasted scheduling map to identify the interfering user's ID. The latter does not require users to sniff and decode other users' frame headers. The CSI is then reported to the eNB. The eNB builds an interference table, whose elements $I_{n,m} \geq 0$ represent the interference caused by user n to user m ($\forall n, m \in \mathcal{N} \cup \{N+1\}$). In case two users do not detect each other for physical/timing reasons, the failure only causes an interruption on a millisecond scale. Once an interruption occurs, the user will report it to the eNB which will update the interference matrix. As for outband D2D, each user reports the last achieved rate over WiFi. In case of an inaccurate report due to long inactivity period, the users can send an updated report before the next mode interval.

Signaling overhead. The maximum number of CSI reports in LTE-A (with wideband CSI reporting [84]) is equal to N. This number increases to $N + 2|\mathcal{N}_d||\mathcal{N}_c| + |\mathcal{N}_d|\,(|\mathcal{N}_d| - 1)$ in a D2D-enabled network, where $\mathcal{N}_c$ and $\mathcal{N}_d$ are the sets of cellular and D2D users, respectively. For instance, a D2D-enabled network with 4 cellular and 6 D2D users may require up to 88 CSI reports, which is almost 9 times higher than its equivalent in a legacy system. Fortunately, the CSI feedbacks can be considerably reduced using the state-of-the-art feedback reduction techniques [85]. Moreover, we

will see in Section A.8.5 that the D2D signaling overhead is negligible as compared to the resulting gain. This overhead is further reduced by our proposal because we decouple scheduling from mode selection, hence the D2D related CSIs are obtained less frequently.

A.8.3 Floating Band D2D Framework

In this section we describe our proposed Floating Band D2D framework and formulate the problem of mode selection at the beginning of each mode interval j, i.e., each T seconds. The utility function in our problem formulation depends on throughput and energy costs, for which we provide a general model in which specific schedulers can be plugged in. Note that, although the general model can be used with specific schedulers to evaluate the performance of various strategies, the formulation of the problem does not depend on the scheduler actually implemented, and is not affected by resource allocation strategies for either cellular or D2D connections.

Throughput and energy costs. The transmitted data $\theta^i_{n,m}(j)$ for a connection (n,m) in mode $i \in \{0,1,2\}$ during mode interval j is formulated as follows:

$$\theta^i_{n,m}(j) = B^i_{n,m}(j) R^{i,\mathrm{CSI}}_{n,m}(j), \tag{A.21}$$

where $B^i_{n,m}(j)$ is the number of Resource Blocks (RBs) allocated to connection (n,m) in mode interval j. $R^{i,\mathrm{CSI}}_{n,m}(j)$ is the number of transmitted bits per RB of connection (n,m) in mode i during mode interval j, computed based on the channel gain between users n and m, and the interference matrix $\mathbf{I}$.

The energy consumption of a cellular user $E^0_{n,m}(j)$ and the energy consumption of a D2D pair $E^i_{n,m}(j)$ in inband mode $i \in \{1,2\}$ are given by:

$$E^0_{n,m}(j) = \beta_{lte} + p^{0,\mathrm{TX}}_n \cdot t_{B^0_{n,m}}(j) \qquad m = N+1, \tag{A.22}$$

$$E^i_{n,m}(j) = 2\left(\beta_{lte} + \beta^{\texttt{WiFi}}_{idle}\right) + \left(p^{i,\mathrm{TX}}_n + p^{i,\mathrm{RX}}_m\right) t_{B^i_{n,m}}(j), \tag{A.23}$$

where β_{lte} and $\beta^{\texttt{WiFi}}_{idle}$ are the baseline energy consumed in a mode interval by an active cellular interface and an idle WiFi interface, respectively. The WiFi interface is kept idle in inband modes to speed up WiFi connection setup. Here, $p^{i,\mathrm{TX}}_n$ and $p^{i,\mathrm{RX}}_m$ are the energy consumed for transmission and reception in one subframe, respectively. $t_{B^i_{n,m}}(j)$ is the duration of $B^i_{n,m}(j)$. Here, we do not calculate the energy per RB, because it is shown that the transmission/reception power mainly depends on time rather than bandwidth [86].

The expression of transmitted data $\theta^3_{n,m}(j)$ and the energy consumption $E^3_{n,m}(j)$ for connection (n,m) under outband mode (i.e., mode 3) in mode interval j is as follows:

$$\theta^3_{n,m}(j) = T \cdot R^{3,\mathrm{CSI}}_{n,m}(j), \tag{A.24}$$

$$E^3_{n,m}(j) = 2(\beta_{lte} + \beta^{\texttt{WiFi}}_{active}) + \left(p^{3,\mathrm{TX}}_n + p^{3,\mathrm{RX}}_m\right) \theta^3_{n,m}(j), \tag{A.25}$$

where $R^{3,\mathrm{CSI}}_{n,m}$ is the WiFi rate and $\beta^{\texttt{WiFi}}_{active}$ is the baseline WiFi energy consumed by a user in a mode interval. $p^{3,\mathrm{TX}}_n$ and $p^{3,\mathrm{RX}}_m$ are the energy consumed by user m per transmitted/received bit. Note that the energy consumption as defined here can incorporate both the consumption due to transmission/reception and packet processing (see [76]). The β_{lte} is due to the dependence of outband users to the eNB signaling.

We define a utility function for connection (n,m) under mode i in mode interval j as follows:

$$U^i_{n,m}(j) = \theta^i_{n,m}(j) - \alpha E^i_{n,m}(j), \tag{A.26}$$

where α is the relative cost of energy. The utility accounts for both throughput and energy consumption. The value of α determines whether the system is biased towards higher throughput or

lower energy consumption. In our model, the impact of schedulers is summarized in $B^i_{n,m}$ and $t_{B^i_{n,m}}$. Those parameters have to be computed in each mode interval and for each possible mode selection decision.

Problem formulation. Let $\mathcal{L}(j)$ be the set of all existing connections during mode interval, j, $\{Y^i_{n,m}(j)\}$ be the set of binary decision variables, and γ_n be the tolerable interference threshold that allows for a non-zero reception rate by user n. We formulate the problem of mode selection for mode interval j as a binary programming problem (the dependency on j is omitted for readability):

$$\begin{cases} \max \quad U_{sum} := \sum_{i=0}^{3} \sum_{(n,m)\in\mathcal{L}} U^i_{n,m} Y^i_{n,m} \ \forall (n,m) \in \mathcal{L} \\ \text{s.t. } \sum_{i=0}^{3} \sum_{n|(n,m)\in\mathcal{L}} Y^i_{n,m} \leq 1 \quad \forall m \in \mathcal{N} \\ \quad \sum_{i=0}^{3} \sum_{m|(n,m)\in\mathcal{L}} Y^i_{n,m} \leq 1 \quad \forall n \in \mathcal{N} \\ \quad \sum_{(n,m)\in\mathcal{L}} Y^1_{n,m} I_{n,x} \leq \gamma_x \quad \forall x \in \mathcal{N}_c \cup \{N+1\} \\ \quad \sum_{i\in\{0,1\}} \sum_{(x,y)\in\mathcal{L}\setminus\{(n,m)\}} Y^i_{x,y} Y^1_{n,m} I_{x,m} \leq \gamma_m \\ \quad \sum_{(x,y)\in\mathcal{L}\setminus\{(n,m)\}} Y^2_{x,y} Y^2_{n,m} I_{x,m} \leq \gamma_m \ \forall (n,m) \in \mathcal{L} \end{cases} \tag{A.27}$$

Problem (A.27) maximises the sum of utilities U_{sum} over all possible combinations of users and modes. Our assumption on single instantaneous connectivity is enforced with the first and second constraints (the eNB, which is labeled as $N+1$, is an exception). The third constraint ensures that the co-channel interference from underlay users to cellular users and to the eNB is kept below the threshold. The fourth constraint limits the interference from cellular and inband underlay users to other inband underlay users. The interference of overlay transmissions is limited by the fifth constraint.

Complexity. Problem (A.27) is NP-hard and non-linear since it can be reduced to the *longest path problem* (e.g., for a weighted directed and possibly disconnected graph), which is NP-hard [87]. This reduction is obtained when we consider Problem (A.27) for a single mode $i = 3$ (outband), in which the objective is to activate D2D pairs so as to achieve the maximum utility possible with the two restrictions on at most one incoming and at most one outgoing transmission for every user. Problem (A.27) requires the computation of $\{U^i_{n,m}\}$, which is based on Signal to Noise and Interference Ratios (SINRs) and its optimal solution can be achieved by brute force: exploring the consequences of assigning modes 1, 2, or 3 to any of the $\frac{|\mathcal{N}_d|}{2}$ D2D pairs. Hence, the resulting complexity is $O(N \cdot 3^{\frac{|\mathcal{N}_d|}{2}})$, which grows exponentially with the number of D2D pairs. The optimal solution to the above maximisation problem is computationally expensive and practically unfeasible in dense networks. However, we linearized the non-linear constraints so that the problem can be solved relatively efficiently by standard approaches, such as Branch & Bound [88]. The linearized problem is omitted for brevity, but can be found in [89]. Nevertheless, we deem such an approach impractical as the system requires a solution in milliseconds. Thereby, we propose efficient heuristics in what follows.

A.8.4 Heuristics

The exact solution to Problem (A.27) is computationally expensive and does not allow for a fast and scalable mode selection. Given the similarity of the problem to the longest path problem and the knapsack problem, we propose three practical heuristics. These heuristics explore the achievable utilities of the users in an iterative manner. Note that these utilities are computed assuming that the system is fully utilized (i.e., users' queues are fully backlogged) so that they do not require the knowledge of the actual user's offered load.

A.8.4.1 Heuristic 1. `Social`

The eNB iterates over the set of D2D transmitters $\mathcal{N}_{d,TX}$, and it selects the mode that maximises the aggregate utility (lines 7-13 in Algorithm 2). Note that the mode for user i is selected based on the modes selected for the precedent users. Initially, all D2D pairs are assigned to mode 3 (outband), to minimise the impact on cellular users. For better fairness [26], the order of users in $\mathcal{N}_{d,TX}$ is randomized at any mode interval. The mode selection repeats until the algorithm converges to a decision. We name this heuristic as `Social` because it decides based on social welfare. Since the utility of `Social` cannot decrease with mode selection decisions, the heuristic always converges. Algorithm 2 illustrates the pseudocode of this heuristic.

Algorithm 2 `Social`

Input:
1: $\mathcal{N}_{d,\mathrm{TX}}$: set of D2D transmitters (randomized order).
2: $I_{n,x}$: interference between each pair of users.
Output: $Y^{i}_{(n,m)}, \forall n \in \mathcal{N}_{d,\mathrm{TX}}$
3: initialize: $\mathbf{Y}=\mathbf{Y}_{old}=\emptyset$; $Y^{0}_{(c,N+1)}=1, \forall c \in \mathcal{N}_c$; $Y^{3}_{(n,m)}=1, \forall n \in \mathcal{N}_{d,\mathrm{TX}}$; max $= U_{sum}$
4: **while** $\mathbf{Y} \neq \mathbf{Y}_{old}$ **do**
5: $\quad \mathbf{Y}_{old} = \mathbf{Y}$
6: $\quad$ **for** $n \in \mathcal{N}_{d,\mathrm{TX}}$ **do**
7: $\quad\quad$ **for** $j \in \{1,2,3\}$ **do**
8: $\quad\quad\quad$ Calculate: $U_{sum}|n$ is in mode j
9: $\quad\quad\quad$ **if** $U_{sum} >$ max **then**
10: $\quad\quad\quad\quad$ max $= U_{sum}$
11: $\quad\quad\quad\quad$ $Y^{j}_{(n,m)} = 1$; $Y^{k}_{(n,m)} = 0, k \in \{1,2,3\} \setminus \{j\}$
12: $\quad\quad\quad$ **end if**
13: $\quad\quad$ **end for**
14: $\quad$ **end for**
15: **end while**

A.8.4.2 Heuristic 2. `Greedy`

The `Greedy` heuristic is similar to `Social`. Unlike `Social`, `Greedy` selects the mode which maximises the user's individual utility (line 10 in Algorithm 3). The drawback of `Greedy` is that it might not converge. However, we can index each decision since the algorithm is running in the eNB. Once a duplicate index (stored in $\mathcal{D}$) is found, the algorithm stops the iteration. Algorithm 3 illustrates the pseudocode of the heuristic.

Algorithm 3 `Greedy`

Input:
1: $\mathcal{N}_{d,\mathrm{TX}}$: set of D2D transmitters (randomized order).
2: $I_{n,x}$: interference between each pair of users.
Output: $Y^i_{(n,m)}, \forall n \in \mathcal{N}_{d,\mathrm{TX}}$
3: initialize: $\mathbf{Y} = \emptyset$; $Y^0_{(c,N+1)} = 1, \forall c \in \mathcal{N}_c$; $Y^3_{(n,m)} = 1, \forall n \in \mathcal{N}_{d,\mathrm{TX}}$; $\max_i = U_{(i,m)}$; $exit = False, \mathcal{D} = \emptyset$
4: **while** $exit = False$ **do**
5: **for** $i \in \mathcal{N}_{d,\mathrm{TX}}$ **do**
6: **for** $j \in \{1,2,3\}$ **do**
7: Calculate: $U^j_{(i,m)} | i$ is in mode j
8: **if** $U^j_{(i,m)} > \max_i$ **then**
9: $\max_i = U^j_{(i,m)}$
10: $Y^j_{(i,m)} = 1$; $Y^k_{(i,m)} = 0, k \in \{1,2,3\} \setminus \{j\}$
11: **end if**
12: **end for**
13: **end for**
14: $dec =$ Index of current $\mathbf{Y}$
15: **if** $dec \in \mathcal{D}$ **then**
16: exit = True
17: **end if**
18: Add dec to $\mathcal{D}$
19: **end while**

A.8.4.3 Heuristic 3. `Ranked`

Both `Social` and `Greedy` operate on a list of D2D transmitters with a randomized order. In contrast, `Ranked` heuristic sorts this list based on the achievable utility of each user without considering the impact of other users (PHASE 1). In PHASE 2, the pre-ordered list $\mathcal{N}_{d,\mathrm{TX}}^{(\mathrm{ranked})}$ is evaluated using `Greedy`, which makes the heuristic *greedier* than `Greedy`. This helps to evaluate the ability of our approach to withstand unfair conditions. Algorithm 4 illustrates the pseudocode of the heuristic.

Algorithm 4 `Ranked`

Input:

1: $\mathcal{N}_{d,\mathrm{TX}}$: set of D2D transmitters (randomized order).
2: $I_{n,x}$: interference between each pair of users.

Output: $Y_{(n,m)}^{i}$

3: initialize: $\mathbf{Y} = \emptyset$; $Y_{(c,N+1)}^{0} = 1, \forall c \in \mathcal{N}_c$; $Y_{(n,m)}^{3} = 1, \forall n \in \mathcal{N}_{d,\mathrm{TX}}$

PHASE 1: Sorting D2D pairs based on their utility

4: **for** $i \in \mathcal{N}_{d,\mathrm{TX}}$ **do**
5: **for** $j \in \{1, 2, 3\}$ **do**
6: Calculate $U_{(i,m)}^{j}$
7: **end for**
8: $\mathrm{mode}_i = \arg\max\{U_{(i,m)}^{j}; j \in \{1, 2, 3\}\}$
9: **end for**
10: sort the $\mathcal{N}_{d,\mathrm{TX}}$ based on utilities $U_{(i,m)}^{\mathrm{mode}_i}$ & store in $\mathcal{N}_{d,\mathrm{TX}}^{(\mathrm{ranked})}$.

PHASE 2: Executing `Greedy` heuristic

11: Do `Greedy` with $\mathcal{N}_{d,\mathrm{TX}} = \mathcal{N}_{d,\mathrm{TX}}^{(\mathrm{ranked})}$.

A.8.4.4 Complexity Analysis

Our proposed heuristics compute $N - |\mathcal{N}_{d,\mathrm{TX}}|$ utilities $\{U_{n,m}^{i}\}$ for each mode and for every D2D transmitter in a sequential manner, i.e., $3\left(N|\mathcal{N}_{d,\mathrm{TX}}| - |\mathcal{N}_{d,\mathrm{TX}}|^2\right)$ utilities per round of evaluation. In each mode interval, the evaluation cycle is repeated r_i times, $r_i \geq 1$, until the algorithm converges to a decision. Therefore, the complexity of `Social` and `Greedy` is $O\left(3r_i N|\mathcal{N}_{d,\mathrm{TX}}|\right)$, $i \in \{1, 2\}$. `Ranked` has an additional sorting procedure before the mode selection in which the utility of each D2D pair is computed in isolation. Thus, the algorithm only needs to compute $3|\mathcal{N}_{d,\mathrm{TX}}|$ utilities in PHASE 1, which can be neglected with respect to the number of utilities to be computed in PHASE 2. Hence, the complexity of `Ranked` is $O\left(3r_3 N|\mathcal{N}_{d,\mathrm{TX}}|\right)$. Therefore, the three proposed heuristics have the same complexity, except for a constant factor r_i that we will quantify experimentally later.

A.8.5 Evaluation

In this section, we use numerical simulations to evaluate the performance of our proposed heuristics. The evaluation scenario consists of a hexagonal multi-cell network with a reference cell in the middle and its first-tier neighbours (see Figure 2.19). The results reported here pertain to the reference cell, and the neighbouring cells model the impact of inter-cell interference. Error bars in the results are the 95% confidence intervals. Although our approach can be tested with any scheduler, here we refer to the Proportional Fair (PF) scheme for scheduling cellular users, since it represents the state of the art for schedulers used in real implementations [49, 50]. In addition to our heuristics, we evaluate three benchmark schemes, namely, `Forced-LTE`, `Forced-WiFi`, and

`Optimal`. In `Forced-LTE`, D2D users are forced to use legacy cellular communications (i.e., mode 0). In `Forced-WiFi`, D2D users are forced to communicate over WiFi (i.e., mode 3). `Optimal` results are based on the exact solution to Problem (A.27). The benchmarks allow to compare our proposals with the legacy cellular system, to measure the gain due to extra WiFi bandwidth, and to see how far the heuristics are from the optimum.

A.8.5.1 Simulation setup

User placement follows the uniform distribution. The number of D2D users is on average 30% of the cell population. The simulation parameters are chosen according to the evaluation guidelines of ITU-R [83] which are reported in Table A.13. In the simulation, we show both the *packet simulation results* (i.e., performance under finite offered load and in the presence of probabilistic arrival processes) and the *achievable performance* (i.e., performance at capacity-level utilization, under infinite offered load conditions). Unless otherwise specified, the default values for α and overlay resource portion are those reported in Table A.13, with an aggregate D2D and cellular load of 30 Mbps and 90 Mbps, respectively. Since the D2D capacity is higher than the cellular one, due to proximity of D2D users and availability of outband resources, we deemed fair to impose higher load to D2D users. Note that the default value of α is selected based on a rough estimate of the current relative price of bit per Joule (b/J) in the market.

Besides the values of Table A.13, we investigate the impact of user density N, overlay resource portion, relative cost of energy α, and D2D load on the system performance. Moreover, we shed light on the convergence time of our heuristics and their flexibility in different environments.

Table A.13: The parameters used in the evaluation

Parameter	**Value**
Cellular	
Cellular uplink bandwidth	20 MHz
Cell radius	250 m
eNB, cellular user TX power	44 dBm, 24 dBm
Thermal noise power	-174 dBm/Hz
Mode interval length T	2 s
Fading, shadowing, pathloss	Reyleigh, 6 dB, UMa [83]
Buffer size	500 packets [90]
β_{lte}	1288.04 mW
WiFi	
WiFi bandwidth	22 MHz
WiFi effective range	150 m
WiFi TX power	20 dBm
$\beta^{\texttt{WiFi}}_{active}$, $\beta^{\texttt{WiFi}}_{idle}$	132.86 mW, 77.2 mW
D2D	
Underlay max bandwidth	20 MHz
Overlay resource portion	30%
D2D maximum distance	20 m
D2D inband TX power	10 dBm
Relative cost of energy α	1 bit/Joule

A.8.5.2 Simulation results

Wa have already discussed the analytical results on the impact of the number of users in Section 2.3.2.3, specifically in Figures 2.20 to 2.22. In what follows we discuss anlysis-based and simulation-based results that show the impact of more system parameters on the performance of D2D mode selection algorithms.

Impact of the number of users N. In Figures A.22 to A.24, we can observe the accuracy of our mode selection and its performance using packet simulation. Figure A.22 shows that cellular users have comparable throughput performance under all schemes due to PF scheduling. If the data rate of a cellular user degrades due to co-channel interference, the PF compensates for it by allocating more resources to that user. In Figure A.23, it is observed that the fairness among D2D users drops under `Forced-LTE` and `Forced-WiFi`. Under `Forced-LTE`, D2D users are scheduled as cellular users, hence they achieve similar fairness performance as cellular users (but not equal because their fairness is computed over a different set and their load is different). The fairness reduction under `Forced-WiFi` is due to topologically uneven distribution of contending outband users. In Figure A.24, we can observe that utilities of all D2D-enabled schemes grow until N reaches 50. The reason for this behavior is that the network operates under saturation up to this point. In fact, one can observe in Figure 2.20 that the achievable throughput with 50 users or less is below 120 Mbps which is equal to the total offered load (i.e., $30+90$) in the scenario of Figure A.24. For $N>50$, the utility in the packet simulation is limited by the adopted load.

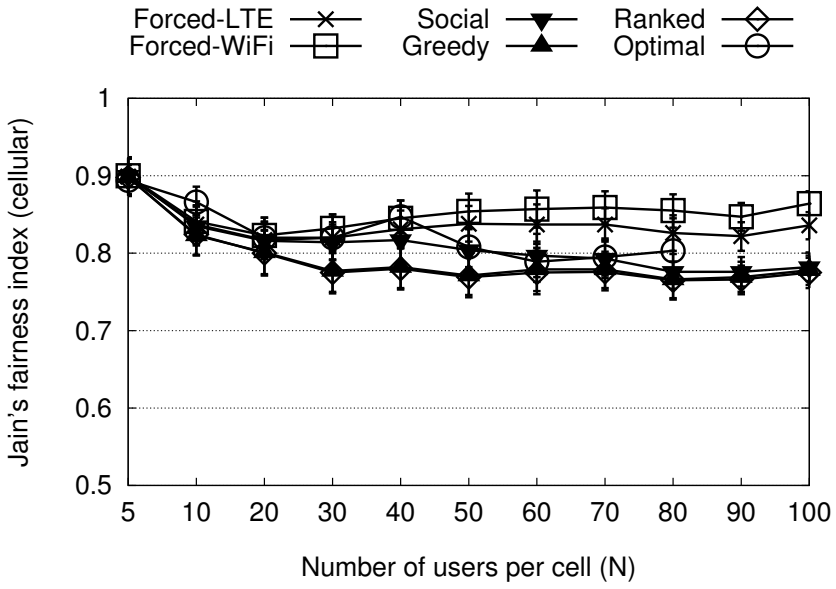

Figure A.22: Impact of user population on the system performance evaluated through packet simulation. Fairness among cellular users.

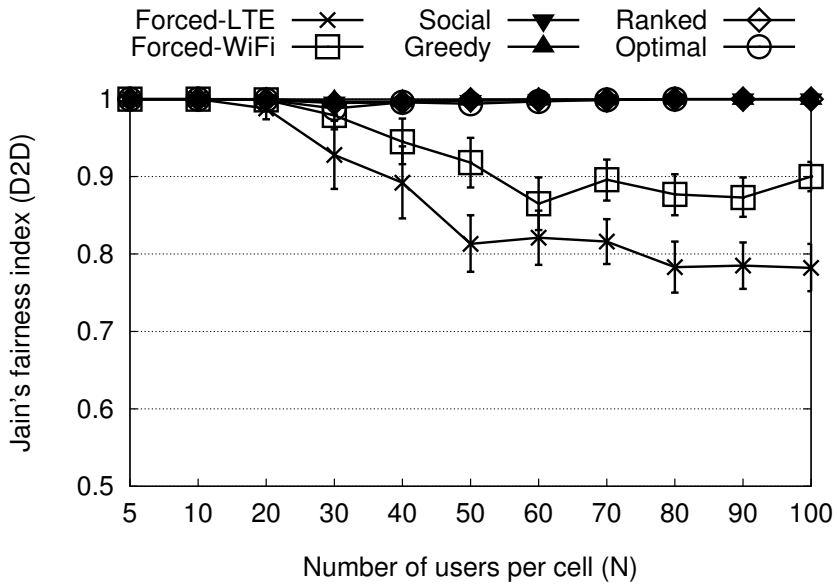

Figure A.23: Impact of user population on the system performance evaluated through packet simulation. Fairness among D2D users.

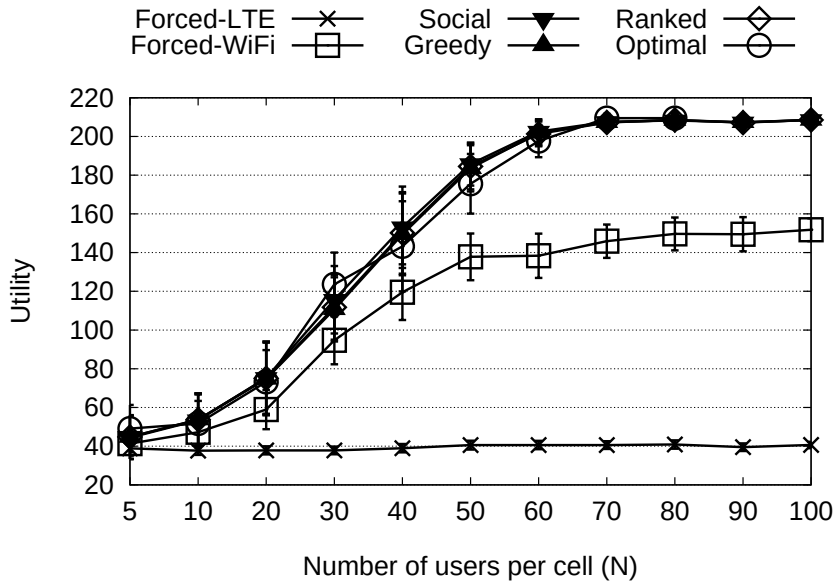

Figure A.24: Impact of user population on the system performance evaluated through packet simulation. Aggregate cell Utility.

Impact of overlay resource portion. Here, the number of users per cell is fixed to 50. Figure A.25 shows that the utilities of multi-band schemes increase with the overlay bandwidth. This increment is due to throughput improvement under mode 2. This implies that D2D users tend to receive more interference from cellular users than from other D2D users, hence, the spectral efficiency is higher in overlay than underlay. As mentioned, the overlay portion is given to modes 0 and 1 if there are no overlay users. As a result, the utilities of `Forced-LTE` and `Forced-WiFi` remain unchanged here.

Although the aggregate utilities are improved, we should also investigate the impact of overlay bandwidth on cellular users. Figure A.26 illustrates that the delivery ratio of cellular users degrades as the overlay bandwidth grows because there is less bandwidth at their disposal. Figure A.26 also sheds light on the differences among multi-band schemes. Cellular users experience higher packet delivery ratio with `Social`. Indeed, `Social` is the only scheme that aims to maximise the aggregate utility, which includes the utility of cellular users. Finally, Figure A.27 shows how the delivery ratio of D2D users approaches 1 with higher overlay bandwidths, as expected.

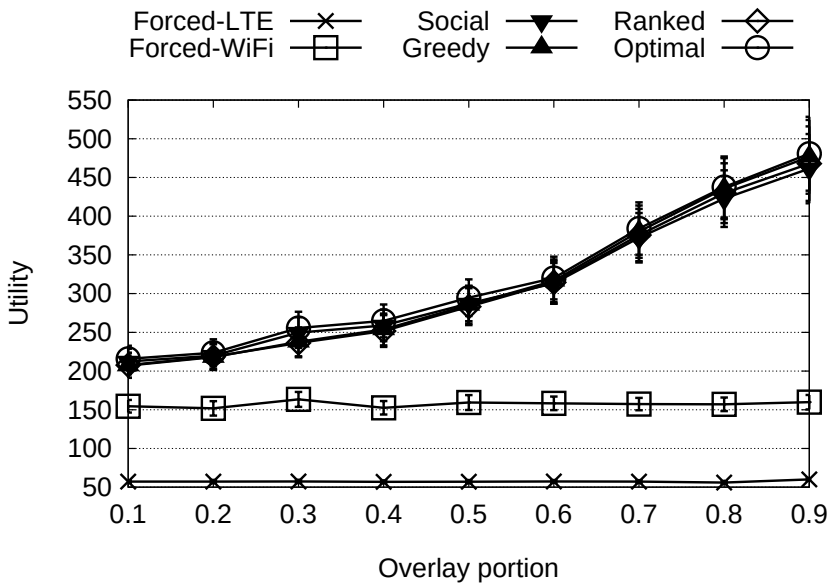

Figure A.25: Impact of overlay portion on system performance ($N = 50$). Aggregate achievable utility.

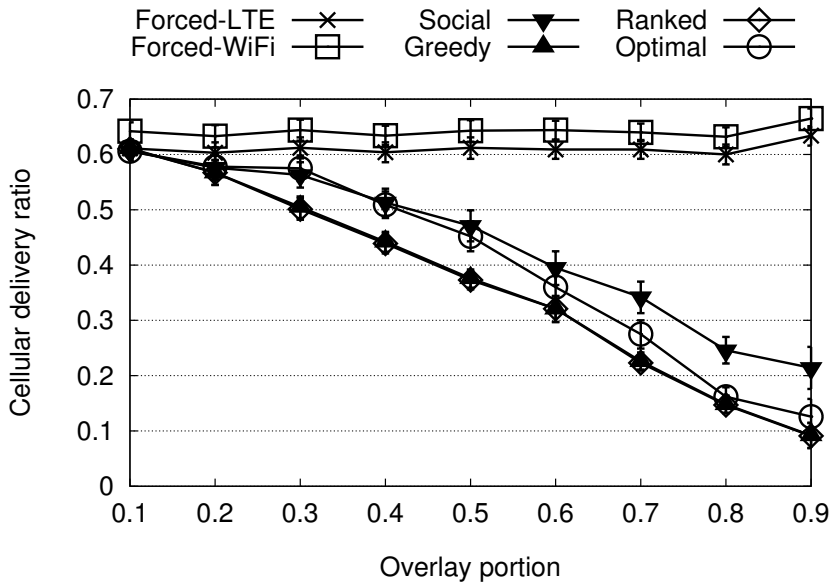

Figure A.26: Impact of overlay portion on system performance ($N = 50$). Average delivery ratio of cellular users.

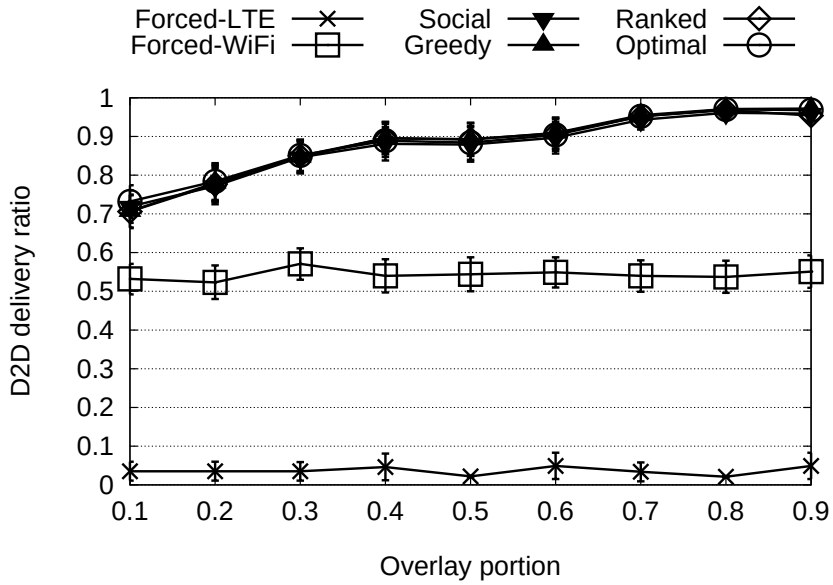

Figure A.27: Impact of overlay portion on system performance ($N = 50$). Average delivery ratio of D2D users.

Impact of energy cost α. Recall that with the current relative energy cost α our system is biased towards throughput. Hence, we investigate the impact of α in Figures A.28 and A.29, for $N = 50$. We start with Figure A.28, in which a 20% throughput reduction is observed at $\alpha = 10^6$ b/J. This shows the system's bias shifts towards energy minimisation as α increases. In Figure A.29, the utility reduces as α grows, although the behavior of the curves is not linear at all. In particular, for very large values of α, our system prefers `Forced-LTE` (i.e., mode 0) because it only powers one interface. Since we disallow multi-band schemes to assign mode 0 to D2D users, `Forced-LTE` might achieve utilities higher than that of `Optimal` when α is very large (e.g., with $\alpha = 10^6$ b/J, which is too unrealistic as of today and for the near future due to the high cost of electricity).

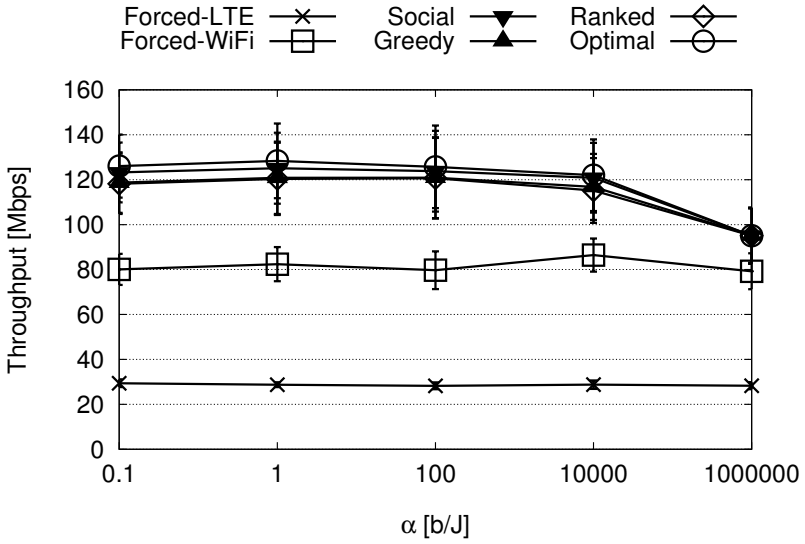

Figure A.28: Impact of α on system throughput and utility ($N = 50$). Aggregate throughput.

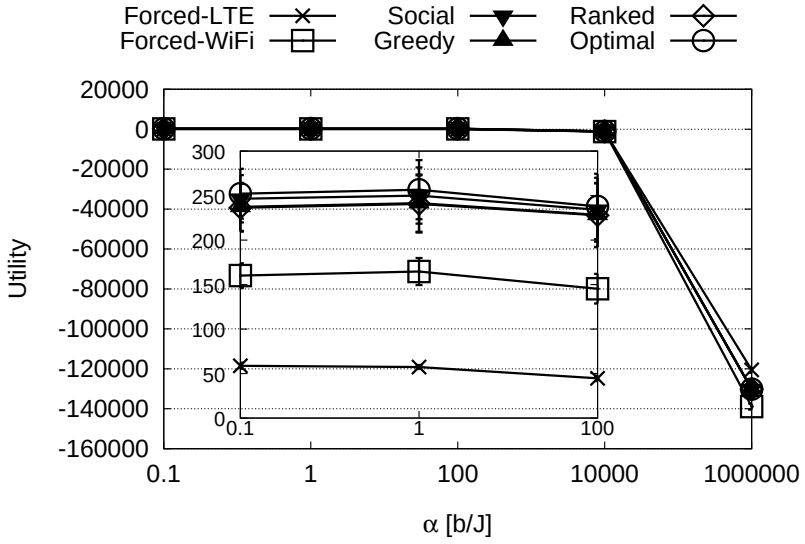

Figure A.29: Impact of α on system throughput and utility ($N = 50$). Aggregate utility.

Impact of D2D load. Finally, the impact of D2D load is shown in Figures A.30 to A.32 for $N = 50$. The packet delivery ratio for D2D users drops as the load increases, as shown in Figure A.30. This is the expected behavior of systems in saturation. However, as we see in Figure A.31, our schemes are designed in such a way that saturation of D2D users does not impact the cellular users. This shows that our proposal can be a candidate for distributed D2D mode selection implementations in which *cellular users are protected from mode selection decisions of D2D users*. It is observed in Figure A.32 that system utility approaches its achievable limit (220) when the D2D load is almost 250 Mbps (see Figure 2.22, $N = 50$). In Figure 2.20, we observe that the achievable capacity for $N = 50$ is almost 120 Mbps. Indeed, by multiplying the packet delivery ratios (see Figures A.30 and A.31) with the aggregate network load (250 Mbps for D2D users and 30 Mbps for cellular users), we observe that the achieved throughput is almost 120 Mbps (i.e., $250 \cdot 0.4 + 30 \cdot 0.6 = 118$). With similar calculations, one finds that `Forced-WiFi` saturates almost at 100 Mbps.

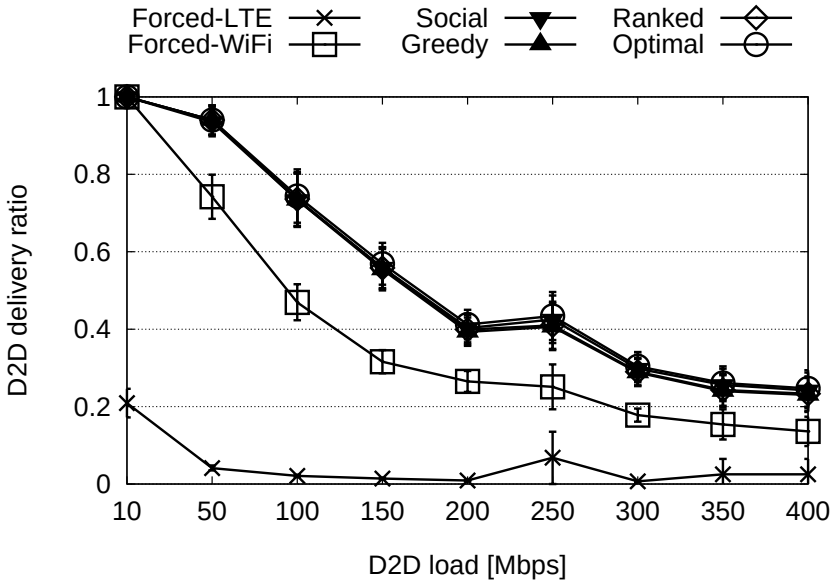

Figure A.30: Impact of D2D load on system performance evaluated through packet simulation ($N = 50$). D2D delivery ratio.

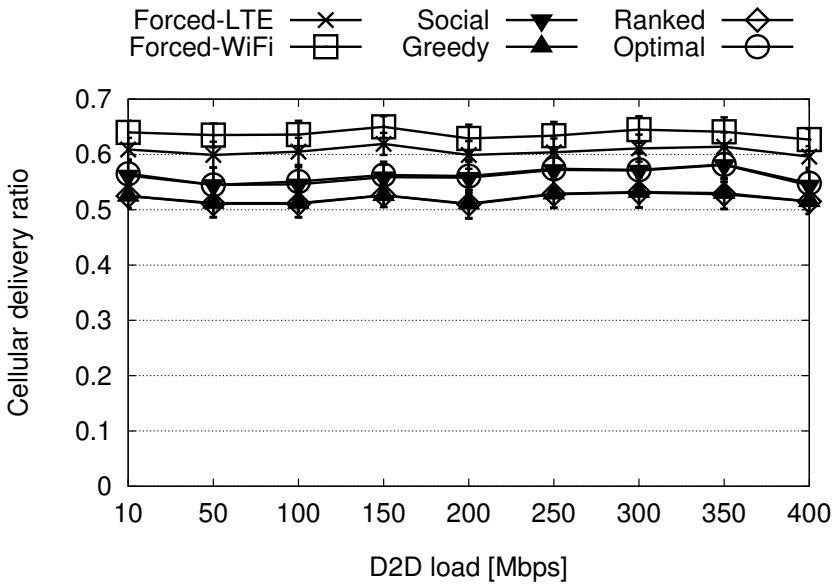

Figure A.31: Impact of D2D load on system performance evaluated through packet simulation ($N = 50$). Cellular delivery ratio.

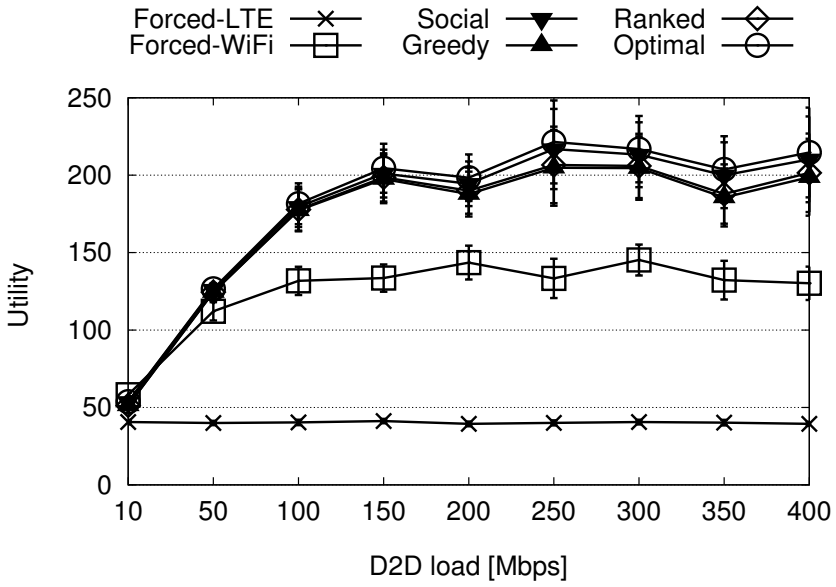

Figure A.32: Impact of D2D load on system performance evaluated through packet simulation ($N = 50$). Aggregate utility.

Convergence. In Table A.14, we report the convergence of our proposed heuristics in terms of time and the number of iterations. The heuristics are tested on Mathematica™ on a machine with a 3.6 GHz processor and 8 GB memory. `Greedy` and `Social` have a very similar convergence time. `Greedy` is slightly slower than `Social` due to decision indexing. Interestingly, notwithstanding its ranking operations, `Ranked` has a better performance. This happens because `Ranked` converges to a decision with less iterations (which is what we have indicated as factor r_i in Section A.8.4.4).

Table A.14: Convergence of the heuristics ($N = 100$)

	`Social`	`Greedy`	`Ranked`
Average convergence time [s]	1.61	1.62	1.43
Average number of iterations r_i	2.69	2.80	1.46

Flexibility. As mentioned, Floating Band D2D is key to flexible D2D architectures. We emphasize this fact by evaluating our proposal in various cellular environments, according to ITU-R guidelines [83]. Table A.15 shows that, moving from micro-cell to rural macro-cell, the system relies more on the cellular spectrum as density reduces. As a consequence, the number of underlay connections increases. For denser environments, we observe that a significant part of connections is served using outband D2D.

Table A.15: Percentage of each mode in different environments ($N = 100$)

	Urban micro-cell	Urban macro-cell	Suburban macro-cell	Rural macro-cell
Inband underlay	4%	8%	29%	31%
Inband overlay	63%	66%	66%	67%
Outband	33%	26%	5%	2 %

Bibliography

[1] Vincenzo Mancuso, Claudio Cicconetti, Vincenzo Sciancalepore, Arianna Morelli, Pablo Serrano, Sebastien Auroux, Erick Bizouarn, Antonio de la Oliva, Rohit Goupta, Pierre Seite, Maria Isabel Sanchez, and Martin Draexler. CROWD Deliverable D1.1: "Preliminary Architecture Design". 2013.

[2] Arianna Morelli, Claudio Bottai, Fabio Toninelli, Rohit Gupta, Vincenzo Mancuso, Arash Asadi, Vincenzo Sciancalepore, Christian Vitale, Sebastien Auroux, Martin Drexler, Antonio de la Oliva, and Pablo Serrano. CROWD Deliverable D1.3: "Final architecture design". 2014.

[3] Vincenzo Mancuso, Claudio Cicconetti, Vincenzo Sciancalepore, Arianna Morelli, Pablo Serrano, and Laurent Roullet. CROWD Deliverable D2.1: "Initial specification of wireless enhancements for very dense deployments". 2013.

[4] Simone Mattiacci, Claudio Bottai, Vincenzo Sciancalepore Vincenzo Mancuso, Arash Asadi, Christian Vitale, Pablo Serrano, Laurent Roullet, Engin Zeydan, and Maria Isabel Sanchez. CROWD Deliverable D2.2: "Final specification for the wireless enhancements functions and interfaces". 2014.

[5] Pablo Serrano, Pablo Salvador, Vincenzo Mancuso, and Yan Grunenberger. Experimenting with Commodity 802.11 Hardware: Overview and Future Directions. *IEEE Communications Surveys and Tutorials*, 2015.

[6] Arash Asadi, Peter Jacko, and Vincenzo Mancuso. Modeling D2D Communications with LTE and WiFi. *ACM SIGMETRICS Performance Evaluation Review*, 2014.

[7] A. Garcia-Saavedra, A. Banchs, P. Serrano, and J. Widmer. Adaptive mechanism for distributed opportunistic scheduling. *Wireless Communications, IEEE/ACM Transactions on*, 2015.

[8] Vincenzo Sciancalepore, Ilario Filippini, Vincenzo Mancuso, Antono Capone, and Albert Banchs. A semi-distributed mechanism for inter-cell interference coordination exploiting the absf paradigm. In *IEEE SECON*, June 2015.

[9] Christian Vitale, Vincenzo Mancuso, and Gianluca Rizzo. Modelling d2d communications in cellular access networks via coupled processors. In *IEEE COMSNETS*, January 2015.

[10] Christian Vitale, Gianluca Rizzo, and Vincenzo Mancuso. A coupled processors model for 802.11 ad hoc networks under non saturation. In *IEEE ICC*, June 2015.

[11] Arash Asadi, Vincenzo Mancuso, and Peter Jacko. Floating band d2d: Exploring and exploiting the potentials of adaptive d2d-enabled networks. In *IEEE WoWMoM*, June 2015.

[12] H. S. Dhillon, R. K. Ganti, F. Baccelli, and J. G. Andrews. Modeling and analysis of K-tier downlink heterogeneous cellular networks. *IEEE Journal on Selected Areas in Communications*, 30(3):550–560, 2012.

[13] Makoto Matsumoto and Takuji Nishimura. Mersenne twister: A 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Trans. Model. Comput. Simul.*, 8(1):3–30, January 1998.

[14] Q. Ye, B. Rong, Y. Chen, M. Al-Shalash, C. Caramanis, and J. G. Andrews. User association for load balacing in heterogeneous cellular network. *IEEE Communications Magazine*, 12(6):2706–2716, 2012.

[15] V. Sciancalepore, V. Mancuso, and A. Banchs. BASICS: Scheduling base stations to mitigate interferences in cellular networks. In *Proceedings of IEEE WoWMoM*, 2013.

[16] Imran Ashraf, Lester T. W. Ho, and Holger Claussen. Improving energy efficiency of femtocell base stations via user activity detection. In *WCNC*, pages 1–5, 2010.

[17] C. Desset, B. Debaillie, V. Giannini, A. Fehske, G. Auer, H. Holtkamp, W. Wajda, D. Sabella, F. Richter, M. J. Gonzalez, H. Klessing, I. Godor, M. Olsson, M. A. Imran, A. Ambrosy, O. Blume. Flexible power modeling of LTE base stations. *IEEE*, pages 1–5, 2012.

[18] V. Sciancalepore, V. Mancuso, and A. Banchs. BASICS: Scheduling base stations to mitigate interferences in cellular networks. In *IEEE WoWMoM 2013*, June 2013.

[19] Third Generation Partnership Project (3GPP). Physical Layer Measurements. 3GPP TS 36.214 v 11.1.0, December 2012.

[20] Olga Goussevskaia, Roger Wattenhofer, Magns M. Halldrsson, and Emo Welzl. Capacity of arbitrary wireless networks. In *INFOCOM'09*, 2009.

[21] A. Asadi, Q. Wang, and V. Mancuso. A survey on device-to-device communication in cellular networks. *IEEE Comm. Surveys &Tutorials*, 2014.

[22] Shaoyi Xu, Haiming Wang, Tao Chen, Qing Huang, and Tao Peng. Effective interference cancellation scheme for device-to-device communication underlaying cellular networks. In *Proceedings of IEEE VTC-Fall*, pages 1–5, 2010.

[23] Klaus Doppler, Mika Rinne, Carl Wijting, Cassio B Ribeiro, and Klaus Hugl. Device-to-device communication as an underlay to lte-advanced networks. *Communications Magazine, IEEE*, 47(12):42–49, 2009.

[24] Lei Lei, Ye Han, Zhangdui Zhong, and Chuang Lin. Performance analysis of device-to-device communications with frequency reuse using stochastic petri nets. In *Communications (ICC), 2013 IEEE International Conference on*, pages 6354–6359. IEEE, 2013.

[25] Sem C. Borst, Matthieu Jonckheere, and Lasse Leskelä. Stability of parallel queueing systems with coupled service rates. *Discrete Event Dynamic Systems*, 18(4):447–472, 2008.

[26] Xinzhou Wu, Saurabha Tavildar, Sanjay Shakkottai, Tom Richardson, Junyi Li, Rajiv Laroia, and Aleksandar Jovicic. FlashLinQ: a synchronous distributed scheduler for peer-to-peer ad hoc networks. *IEEE/ACM Trans. on Networking*, 21(4):1215–1228, 2013.

[27] Christian Vitale, Gianluca Rizzo, and Balaji Rengarajan. *Performance Bounds in Coupled Processor Systems*, 2013. Available at *http://publications.hevs.ch/index.php/attachments/single/665*.

[28] D. Zheng, W. Ge, and J. Zhang. Distributed opportunistic scheduling for ad hoc networks with random access: an optimal stopping approach. *IEEE Transactions on Information Theory*, 55(1), January 2009.

[29] Andres Garcia-Saavedra, Balaji Rengarajan, Pablo Serrano, Daniel Camps-Mur, and Xavier Costa-Perez. SOLOR: Self-Optimizing WLANs with Legacy-Compatible Opportunistic Relays. *Networking, IEEE/ACM Transactions on*, 2015.

[30] Giuseppe Bianchi. Performance analysis of the ieee 802.11 distributed coordination function. *IEEE JSAC*, 18(3):535–547, 2000.

[31] A. De la Oliva, A. Banchs, P. Serrano, and F. A. Zdarsky. Providing throughput guarantees in heterogeneous wireless mesh networks. *Wirel. Commun. Mob. Comput.*, 2012.

[32] D. Malone, K. Duffy, and D. Leith. Modeling the 802.11 distributed coordination function in nonsaturated heterogeneous conditions. *Networking, IEEE/ACM Transactions on*, 15(1):159–172, Feb 2007.

[33] Vern Paxson and Sally Floyd. Wide-area traffic: The failure of poisson modeling. *IEEE/ACM Transactions on Networking*, pages 226–244, 1995.

[34] Matthew Andrews. Instability of FIFO in the permanent sessions model at arbitrarily small network loads. *ACM Transactions on Algorithms*, 5(3), 2009.

[35] Jean-Yves Le Boudec and Patrick Thiran. *Network Calculus: A Theory of Deterministic Queuing Systems for the Internet*, volume 2050 of *LNCS*. Springer, 2001.

[36] Chung Shue Chen, Francois Baccelli, and Laurent Roullet. Joint optimization of radio resources in small and macro cell networks. In *Proceedings of IEEE VTC*, 2011.

[37] 3GPP TR 36.300. Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN), Release 10. v10.11.0, 2010.

[38] I-Hong Hou and Chung Shue Chen. An energy-aware protocol for self-organizing heterogeneous lte systems. *IEEE JSAC*, 2013.

[39] Chandramani Singh and Chung Shue Chen. Distributed downlink resource allocation in cellular networks through spatial adaptive play. In *Proceedings of 25th International Teletraffic Congress*, 2013.

[40] Nessrine Trabelsi, Laurent Roullet, and Afef Feki. A generic framework for dynamic eICIC optimization in LTE heterogeneous networks. In *Proceedings of IEEE VTC*, 2014.

[41] Ye Liu, Chung Shue Chen, and Chi Wan Sung. Joint optimization on inter-cell interference management and user attachment in LTE-A HetNets. In *Proceedings of WiOpt*, 2015.

[42] Virgile Garcia, Chung Shue Chen, YiQing Zhou, and JingLin Shi. Gibbs sampling based distributed OFDMA resource allocation. *Science China (Information Sciences)*, 2014.

[43] Mehlfuhrer et al. The vienna lte simulators - enabling reproducibility in wireless communications research. *EURASIP Journal on Advances in Signal Processing*, 2011.

[44] Arash Asadi and Vincenzo Mancuso. DRONEE: dual-radio opportunistic networking for energy efficiency. *Elsevier Computer Communications*, 2014.

[45] Bin Zhou, Honglin Hu, Sheng-Qiang Huang, and Hsiao-Hwa Chen. Intracluster device-to-device relay algorithm with optimal resource utilization. *IEEE Trans. on Vehicular Technology*, 2013.

[46] Xuan Bao, Uichin Lee, Ivica Rimac, and Romit R. Choudhury. DataSpotting: offloading cellular traffic via managed device-to-device data transfer at data spots. *ACM SIGMOBILE'10*.

[47] Klaus Doppler, Mika Rinne, Carl Wijting, Cassio B Ribeiro, and Klaus Hugl. Device-to-device communication as an underlay to LTE-Advanced networks. *IEEE Communications Magazine*, 2009.

[48] X. Lin, J.G. Andrews, A. Ghosh, and R. Ratasuk. An overview of 3GPP device-to-device proximity services. *IEEE Comm. Magazine*, 2014.

[49] Suk-Bok Lee, Ioannis Pefkianakis, Adam Meyerson, Shugong Xu, and Songwu Lu. Proportional fair frequency-domain packet scheduling for 3GPP LTE uplink. In *Proceedings of IEEE INFOCOM'09*, pages 2611–2615.

[50] Robert Margolies, Ashwin Sridharan, Vaneet Aggarwal, Rittwik Jana, N.K. Shankaranarayanan, Vinay A Vaishampayan, and Gil Zussman. Exploiting mobility in proportional fair cellular scheduling: Measurements and algorithms. In *Proceedings of IEEE INFOCOM'14*, pages 1339–1347.

[51] Tobias Harks, Martin Hoefer, Max Klimm, and Alexander Skopalik. Computing pure nash and strong equilibria in bottleneck congestion games. *Springer Mathematical Programming*, 141(1-2):193–215, 2013.

[52] Chunmei Liu and Eytan Modiano. On the performance of additive increase multiplicative decrease (AIMD) protocols in hybrid space-terrestrial networks. *Comput. Netw.*, 47(5):661–678, April 2005.

[53] Mingbo Xiao, N.B. Shroff, and E. K P Chong. A utility-based power-control scheme in wireless cellular systems. *Networking, IEEE/ACM Transactions on*, 11(2):210–221, Apr 2003.

[54] Kyuho Son, Soohwan Lee, Yung Yi, and Song Chong. REFIM: A practical interference management in heterogeneous wireless access networks. *Journal on Selected Areas in Communications*, 29(6):1260–1272, 2011.

[55] Cheng-Shang Chang. *Performance Guarantees in Communication Networks*. Springer-Verlag, London, UK, UK, 2000.

[56] Abhay K. Parekh and Robert G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: the single-node case. *IEEE/ACM Trans. Netw.*, 1(3):344–357, June 1993.

[57] John S Seybold. *Introduction to RF propagation*. J. Wiley & Sons, 2005.

[58] Oana Iosif and Ion Banica. Lte uplink analysis using two packet scheduling models. In *Telecommunications Forum (TELFOR), 2011 19th*, pages 394–397. IEEE, 2011.

[59] Aldo Vecchietti, Sangbum Lee, and Ignacio E. Grossmann. Modeling of discrete/continuous optimization problems: characterization and formulation of disjunctions and their relaxations. *Computers and Chemical Engineering*, 27(3):433–448, 2003.

[60] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, USA, 2004.

[61] A. Banchs, A. Garcia-Saavedra, P. Serrano, and J. Widmer. A game-theoretic approach to distributed opportunistic scheduling. *Networking, IEEE/ACM Transactions on*, 21(5):1553–1566, Oct 2013.

[62] P. Gupta, Y. Sankarasubramaniam, and A. Stolyar. Random-access scheduling with service differentiation in wireless networks. In *Proceedings of IEEE INFOCOM*, Miami, FL, March 2005.

[63] B. Kristiansson and B. Lennartson. Robust Tuning of PI and PID Controllers. *IEEE Control Systems Magazine*, 26(1):55–69, February 2006.

[64] A. K. Palit and D. Popovic. *Computational Intelligence in Time Series Forecasting: Theory and Engineering Applications*. Springer-Verlag New York, Inc., 2005.

[65] G. F. Franklin, J. D. Powell, and M. L. Workman. *Digital Control of Dynamic Systems*. Addison-Wesley, 2nd edition, 1990.

[66] A. Banchs, P. Serrano, and L. Vollero. Providing Service Guarantees in 802.11e EDCA WLANs with Legacy Stations. *IEEE Transactions on Mobile Computing*, 9(8):1057–1071, August 2010.

[67] J. Kim, S. Kim, S. Choi, and D. Qiao. CARA: Collision-aware rate adaptation for IEEE 802.11 WLANs. In *INFOCOM 2006. 25th IEEE International Conference on Computer Communications. Proceedings*, pages 1–11, 2006.

[68] Pablo Salvador, Luca Cominardi, Francesco Gringoli, and Pablo Serrano. A first implementation and evaluation of the ieee 802.11aa group addressed transmission service. *SIGCOMM Comput. Commun. Rev.*, 44(1):35–41, December 2013.

[69] P. Bahl, R. Chandra, P.P.C. Lee, V. Misra, J. Padhye, D. Rubenstein, and Yan Yu. Opportunistic use of client repeaters to improve performance of wlans. *Networking, IEEE/ACM Transactions on*, 17(4):1160 –1171, aug. 2009.

[70] D. Camps-Mur, A. Garcia-Saavedra, and P. Serrano. Device-to-device communications with wi-fi direct: overview and experimentation. *Wireless Communications, IEEE*, 20(3):96–104, June 2013.

[71] IEEE LAN/MAN Standards Committee et al. Ieee 802.11-wireless lan medium access control (mac) and physical layer (phy) specifications. In *IEEE, June*, 2007.

[72] Drew Fudenberg and Jean Tirole. *Game Theory*. Cambridge, MA: MIT Press, 1991.

[73] Drew Fudenberg and David K. Levine. *The Theory of Learning in Games*. Cambridge, MA: MIT Press, 1998.

[74] Chung Shue Chen and Francois Baccelli. Gibbsian method for the self-optimization of cellular networks. *EURASIP Journal on Wireless Communications and Networking*, 2012.

[75] Y. Zhang, E. Pan, L. Song, W. Saad, Z. Dawy, and Z. Han. Social network aware device-to-device communication in wireless networks. *IEEE Trans. on Wireless Communications*, 2014.

[76] Arash Asadi and Vincenzo Mancuso. On the compound impact of opportunistic scheduling and D2D communications in cellular networks. In *Proceedings of ACM MSWIM'13*, pages 279–288.

[77] J. Liu, Y. Kawamoto, H. Nishiyama, N. Kato, and N. Kadowaki. Device-to-device communications achieve efficient load balancing in LTE-Advanced networks. *IEEE Wireless Communications*, 21(2):57–65, 2014.

[78] Joongheon Kim, Feiyu Meng, Peiyao Chen, Hilmi E. Egilmez, Dilip Bethanabhotla, Andreas F. Molisch, Michael J. Neely, Giuseppe Caire, and Antonio Ortega. Adaptive video streaming for device-to-device mobile platforms. In *Proceedings of ACM MobiCom'13*, pages 127–130.

[79] Mohammad Zulhasnine, Changcheng Huang, and Anand Srinivasan. Efficient resource allocation for device-to-device communication underlaying LTE network. In *IEEE WiMob'10*, pages 368–375.

[80] K. Doppler, Chia-Hao Yu, C.B. Ribeiro, and P. Janis. Mode selection for device-to-device communication underlaying an LTE-Advanced network. In *Proceedings of IEEE WCNC'10*, pages 1–6.

[81] P. Phunchongharn, E. Hossain, and D.I Kim. Resource allocation for device-to-device communications underlaying LTE-Advanced networks. *IEEE Wireless Communications*, 20:91–100, 2013.

[82] 3GPP TR 23.303. Proximity-based services (ProSe);Stage 2 (Release 12). v. 12.2.0, September, 2014.

[83] ITU. Guidelines for evaluation of radio interface technologies for IMT-Advanced. 2009.

[84] C. Johnson. *LTE in Bullets*. Createspace Independent Pub, 2012.

[85] Jin-Hao Li and Hsuan-Jung Su. Opportunistic feedback reduction for multiuser MIMO broadcast channel with orthogonal beamforming. *IEEE Trans. on Wireless Communications*, 13(3):1321–1333, 2014.

[86] Bjoern Dusza, Christoph Ide, Liang Cheng, and Christian Wietfeld. An accurate measurement-based power consumption model for LTE uplink transmissions. *Proceedings of IEEE INFOCOM'13*.

[87] D. Karger, R. Motwani, and G.D.S. Ramkumar. On approximating the longest path in a graph. *Algorithmica*, 18(1):82–98, 1997.

[88] Eugene L Lawler and David E Wood. Branch-and-bound methods: A survey. *Operations research*, 14(4):699–719, 1966.

[89] Floating Band D2D: Exploring and Exploiting the Potentials of Adaptive D2D-enabled Networks (Extended version with proofs).

[90] Haiqing Jiang, Yaogong Wang, Kyunghan Lee, and Injong Rhee. Tackling bufferbloat in 3G/4G networks. In *Proceedings of ACM IMC'12*, pages 329–342.