317959

Mobile Opportunistic Traffic Offloading

D3.3.2 – Design and evaluation of enabling techniques for mobile data traffic offloading (release b, public)

Grant Agreement No.	317959
Project acronym	*MOTO*
Project title	Mobile Opportunistic Traffic Offloading
Advantage	
Deliverable number	D3.3.2
Deliverable name	Design and evaluation of enabling techniques for mobile data traffic offloading (release b)
Version	V 3.0
Work package	WP3 – Offloading foundations and enablers
Lead beneficiary	INTECS
Authors	Vania Conan (TCS), Filippo Rebecchi (TCS), Raffaele Bruno (CNR), Andrea Passarella (CNR), Elisabetta Biondi (CNR), Chiara Boldrini (CNR), Giovanni Mainetto (CNR), Marcelo Dias de Amorim (UPMC), Engin Zeydan (AVEA), Ahmet Serdar Tan (AVEA), Eva Pierattelli (INTECS), Daniele Azzarelli (INTECS).
Nature	R – Report
Dissemination level	PU – Public
Planned Delivery date	31/07/2015 (M33)
Actual Delivery date	31/07/2015 (M33)

Table of Contents

List of Figures

Executive summary

This document is the fourth deliverable of WP3, and updates the previously published deliverable D3.3.1 [3]. To make the document self-contained, it also includes the same material already presented in D3.3.1. The new results, with respect to D3.3.1, are presented in new sections of this document (**specifically, Sections 2.1.4, 2.2.2, 5.4, 6.4** and **7**). **Section 1** and **Section 8** have also been updated to reflect these additions (in particular **Section 1.2.3** summarises the key advancements with respect to Y2 results). A new paper has been included in the **Appendix**, which is a reprint of [16], and another one has been updated, to include the extensions of [56] contained in a journal paper submitted to an international journal. While activities of WP3 have been formally concluded at M29 (March 2015), as suggested by the reviewers during the second project review, we have decided to informally extend the WP activities up to M33 (July 2015). This document presents, therefore, the final results of WP3 activities.

As in the case of D3.3.1, this deliverable reports on the activities and results obtained in the Tasks 3.2 and 3.3. Activities on T3.1 have been presented in a separate document, i.e. D3.2 [2], which is the logical output of T3.1 (finished at M18). Activities in WP3 have progressed along the methodology discussed already in D3.1 [1]. As far as capacity assessment is concerned (T3.2) we have both analysed the performance of individual building blocks in isolation, and their performance when they are combined in a complete offload networking solutions. Results already presented in D3.3.1 related to these activities include: (i) analysing convergence issues in opportunistic networks; (ii) providing end-to-end delay guarantees in opportunistic networks with duty cycling; (iii) assessing the performance of LTE through modelling; and (iv) assessing the performance of complete offload networks (also using infrastructure WiFi components in addition to cellular and opportunistic) in presence of both synchronised and non-synchronised content requests. In this document, we extend this set of results by presenting (i) an analysis of necessary and sufficient conditions for convergence of opportunistic protocols, based on the features of users' mobility patterns; and (ii) an analysis of the effect of energy conservation techniques (duty cycling) on the detected inter-contact times between nodes, in case of general mobility patterns.

Moreover, the document also reports results from T3.3 about scheduling. We have analysed scheduling from multiple dimensions. We have analysed both intra-technology and inter-technology scheduling issues. From the first standpoint, we have considered joint scheduling of multicast and D2D transmissions to optimise offloading. As far as inter-technology scheduling is concerned we have developed a general optimisation framework based on TOPSIS, in order to optimise allocation of users to the various possible technologies based on different QoS performance indices and criteria. Last but not least, we have analysed how to schedule various architectural components of an LTE network (i.e., pico and macrocells) in order to reduce the LTE energy consumption without compromising the efficiency in terms of throughput perceived by the users. With respect to results already presented in D3.3.1, in this document we have advanced the activities, by (i) defining a learning framework to dynamically tune the share of traffic to be sent via LTE multicast and by D2D communications; (ii) including D2D communications (in addition to LTE and WiFi) in the set of choices available to the TOPSIS optimisation framework, and (iii) including D2D communications ad an additional technology to facilitate energy saving in the LTE access network.

In addition to presenting these results in detail, in Section 1 we remind the general strategy of activities in WP3 and how these results are aligned with it, and how they are synergic with the work undertaken in the rest of the project. At the end of the document, we draw the main conclusions of the WP activities.

1 Introduction

This deliverable reports on the main activities during the second and third years of the project in T3.2 and T3.3 of WP3. Task 3.1, focusing on the study of spatio-temporal contact patterns, was finished on M18, and the main outputs have been reported in D3.2 [2]. Activities described in this document, therefore, focused on two main broad topics. The first one is about assessing the capacity of offload networking solutions (with main focus on supporting terminal-to-terminal communication). The second one is about scheduling solutions, both inside single network technologies and across different technologies. Before summarising the key results achieved on these topics, let us recall the main objectives of the WP (related to T3.2 and T3.3), and how we addressed them through the results presented in this document. Formally, WP3 activities have concluded at M29. However, the project has decided to carry over some of the WP3 activities until M33 (July 2015), based on the explicit suggestion of reviewers during the second review meeting (see the review report). Specifically, reviewers encouraged the project to continue the work in WP3, as it is generating interesting theoretical and algorithmic results. We took this opportunity to better finalise some of the analyses that were scheduled for Y3, and the result is that these WP activities actually finished at M33 (instead of M29). This document reports the final results obtained in the WP also considering this additional timeframe.

1.1 Problem statement: Objectives of the WP and approach in addressing them

As described in the DoW, the objectives of WP3 related to T3.2 and T3.3 are as follows:

1. To quantify capacity improvements that can be achieved when offloading traffic across different wireless infrastructures and/or using terminal-based offloading, in both single- and multi-operator environments.
2. To characterize the impact on offloading efficiency of factors such as user mobility patterns, heterogeneity of network deployments, traffic loads, QoS application requirements, and variable terminal densities due to distributed duty cycling techniques.
3. To develop inter-technology scheduling algorithms allowing a more efficient synergy – in presence of offloading techniques – between multiple wireless infrastructures and opportunistic networks, which a special focus on high-load conditions.

The first two objectives are addressed by T3.2, while the third one is addressed by T3.3. As far as the first objective is concerned, the rationale of the work undertaken is as follows:

1.1. We have identified the main architectural blocks to be analysed as far as capacity improvements are concerned from the architecture definition in WP2. At the high level, these are (i) wireless broadband infrastructures and (ii) opportunistic networks.

1.2. We analyse capacity limits of wireless infrastructures, primarily focusing on LTE cellular networks. Although this is not the main focus of the WP, in some cases we also study modifications of LTE components that can overcome some of these limits.

1.3. We characterise the capacity that opportunistic networks can bring about. To this end, we consider a wide range of opportunistic networking protocols and users' mobility patterns, and assess capacity (in terms of throughput and/or end-to-end delay) as a function of these key elements.

1.4. We consider both architectural blocks together, i.e. we study the capacity of an integrated heterogeneous network composed of both a wireless broadband infrastructure and opportunistic network. We study the actual capacity gain that can be achieved when these networking environments are put together.

With respect to the second objective, the rationale of the work undertaken is as follows:

2.1. We have derived configurations for evaluation of the capacity of networks with offloading from the factors identified in the objective, primarily: (i) mobility patterns; (ii) heterogeneity of networks and of users' mobility; (iii) end-to-end delay requirements; (iv) contact patterns modifiers related to energy efficiency (duty cycling)

2.2. We analyse how these factors impact on the capacity limits of wireless infrastructures, and on the capacity improvements brought by opportunistic networks.

Finally, with respect to the third objective, the rationale of the work undertaken is as follows:

3.1. We have identified relevant scheduling problems for the different MOTO scenarios. As discussed in the first period review meeting, this has lead to re-focusing some of the activities, that now take into consideration also intra-technology scheduling. The resulting lines of activities are as follows: (i) intra-technology scheduling in LTE to jointly exploit multicast and D2D communications; (ii) intra-technology scheduling in LTE to improve energy efficiency in the LTE access network; (iii) inter-technology scheduling to optimize allocation of users to multiple wireless technologies available at the same time.

3.2. For each of the three lines of research, we identify specific research problems and address them. We propose algorithmic solutions to improve the efficiency of MOTO for each of them.

1.2 Enabling techniques for mobile data traffic offloading: A summary

Figure 1 provides a graphical representation of the main activities undertaken during Y2 and Y3 in WP3, and particularly in Tasks 3.2 and 3.3. Note that for each activity we also highlight the main methodological approach followed, consisting either of analytical modelling, definition of algorithms to improve capacity, or simulation-based analysis (or combinations thereof). Dashed boxes represent areas where new results are reported in this document, with respect to what already described in D3.3.1 [3]. The figure also shows how objectives presented in Section 1.1 are mapped to the activities presented in the rest of the document.

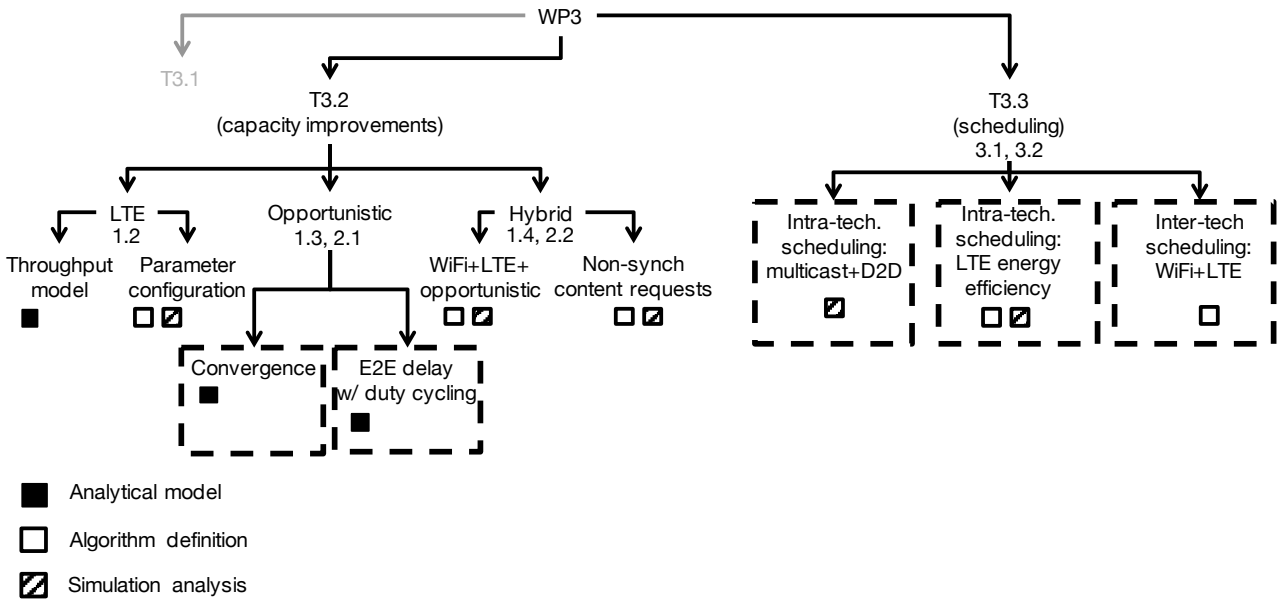

Figure 1. Schematic representation of WP3 activities.

The following subsections provide a summary of these activities, and highlight the key results achieved. Moreover we also present as a brief summary of the improvements for each activity with respect to the status at the end Y1 and Y2, respectively.

1.2.1 Task 3.2: Capacity limits and improvements in networks with offloading

Task 3.2 is devoted to better understand the capacity limits of LTE, and assess the capacity improvements that can be achieved through offloading. From a *methodological* standpoint, we take primarily an approach based on analytical modelling and simulation analysis. Simulation is used to explore the performance of specific systems or networking solutions. Analysis is also used for this purpose (in a pretty standard way with respect to the field of performance evaluation). In addition, it is also used to provide compact mathematical tools that can be used by network operators to plan how to dimension their network in presence of offloading. With respect to the specific *subjects of investigation*, we go step-by-step. At the high level, an offload network is made up of two main components, i.e. a wireless infrastructure part (primarily, LTE), and a mobile part (primarily, an opportunistic network). Therefore, several topics of research deal with assessing the performance and the capacity of these two building blocks in isolation. This is important, as there are still several open points in understanding the capacity limits of these types of network alone.

With respect to ***LTE networks***, while a number of studies have been carried out to characterise its capacity at the physical layer, little effort has been devoted in analysing the capacity perceived by the users, i.e. a number of network layers and functional blocks away from the basic physical layer. In T3.1 we contributed to fill this gap, and this document presents some results on this.

Specifically, in Section 3 we present the key results achieved with respect to this point, i.e.:

- ***we provide an initial model of the throughput experience by users of LTE networks, when we factor in not only parameters of the physical layer, but also the key mechanisms of retransmissions and data reliability***
- ***we exploit learning mechanisms to extend LTE algorithms for automatic configurations of transmission parameters, making them adaptive to changing network conditions.***

With respect to ***opportunistic networks***, research is still ongoing to model the end-to-end throughput (or, equivalently, the end-to-end delay, as explained below) of opportunistic networks. As summarised in the following of this section, this document presents results on convergence properties of opportunistic networks, and on modelling of end-to-end delay in presence of energy saving mechanisms such as duty cycling (we will come back on the energy saving dimension afterwards in this section). Specifically, in Section 2 we first focus on the issue of convergence of forwarding protocols in opportunistic networks, where a protocol is convergent if it yields expected finite end-to-end delay. While at a first sight this may seem a very theoretical problem, it has significant practical implications. Diverging, in practise, means loosing packets, and not be able to provide and end-to-end delay guarantee. Unfortunately, analysis of real mobility traces has shown that protocols may indeed diverge, depending on the contact patterns between nodes. Then, we study the end-to-end delay of forwarding protocols, in case of exponential contact patterns, when duty cycling is used at nodes to conserve energy. This piece of work directly exploits results previously presented in D3.1 (about how to model end-to-end delay in with exponential contact patterns) and D3.2 (about the effect of duty cycling on temporal contact patterns).

The key results achieved with respect to this point, are:

- ***we provide practical tools to select appropriate opportunistic networking protocols given a stochastic description of the contact patterns between nodes, such that convergence can be guaranteed. Moreover, we extend what presented in D3.1, by considering a much more vast family of routing protocols, and comparing them;***
- ***we provide sufficient and necessary conditions on the key features of mobility patterns such that convergence can be guaranteed by specific opportunistic networking protocols. Conditions are set on aggregate inter-contact time statistics, that are much more reasonable and easy to obtain with respect to pairwise inter-contact time statistics;***

- ***we provide probabilistic end-to-end delay guarantees when a certain duty cycling is used;***
- ***we provide practical tools to select the optimal trade-off between the delay that can be guaranteed and the maximum energy saving that can be achieved while still achieving that delay;***
- ***we analyse the effect of duty cycling on detected inter-contact times for general inter-contact time patterns (beside exponential mobility). This is instrumental for deriving general models of end-to-end delays in presence of duty cycling.***

Advancing what presented in D3.1, in this document we also provide some results on the ***capacity of offload networks, i.e. considering the two main building blocks working together***. Note that both level of analysis (considering individual blocks, and the two blocks together) are considered important outcomes of T3.2. Isolated analysis allows us to understand the performance of single technologies, and therefore provide tools to compare possible alternatives when configuring the two main building blocks. Joint analysis allows us to understand the interplay and the complementarity between the two blocks. We provide two main contributions along those lines, in Section 4. Remember that in D3.1 we have presented Push&Track (and its evolution, Droid), as one of the reference solutions in MOTO to implement offloading through opportunistic networks. In this document we extend Droid to use all the different networks that may be available in the MOTO context, i.e. cellular, WiFi and opportunistic. Then, in Section 4.2 we start investigating a relatively untapped problem in the offloading networking panorama, i.e. the offloading performance when the same content is requested in a non-synchronised way by users (note that results in D3.1 and in Section 4.1 focus on simultaneous requests, as the vast majority of the literature on offloading networks).

The key results achieved with respect to this point, are:

- ***through simulations based on real mobility trace and real WiFi deployments we compare the offloading performance when WiFi, cellular and opportunistic networks are used partly or all together at the same time (which clearly results in the best offloading performance).***
- ***in case of offloading with non-synchronised requests we find that even with unfavourable configurations of the opportunistic networking protocols (i.e., those where resources of mobile devices are used at the minimum possible level), offloading can be very effective, saving up to 90% of the total traffic from flowing on the cellular network.***

1.2.2 Task 3.3: Scheduling issues in networks with offloading

Task 3.3 focuses on scheduling policies at various levels. We have identified *three main threads of activities* within the scheduling topic. As discussed at the first review meeting, this was the outcome of a restructuring of the activities originally planned for T3.3, which should have focused exclusively on inter-technology scheduling. While this topic is still investigated in T3.3, we have started two more lines of research that were considered equally important.

First, we study how to **schedule multicast vs. terminal-to-terminal transmissions in an offload network** (Section 5). This is an example of solutions that can be seen both as intra-technology scheduling (e.g., if LTE-D2D is used) and as inter-technology scheduling (when opportunistic networking solutions are used for terminal-to-terminal communication). This piece of work is quite interesting, as it start clarifying one of the basic questions of the overall offloading approach, i.e. if multicast wouldn't be enough.

The key results achieved with respect to this point, are:

- ***our results show that, even in presence of synchronised content requests, multicasting alone is not as efficient as using multicast and D2D communication in an optimised way.***
- ***we have developed a learning framework to dynamically guide the scheduling of traffic between multicast and D2D, showing that it closely approximates an optimal solution, where the split of traffic is decided based on exhaustive search***

A second thread of activity is related to **scheduling picocells** in the emerging scenario of dense networks, **where LTE macro cells are complemented by picocells for increase capacity at the edge of the cell**. While this in principle an interesting and effective approach, it may significantly increase energy consumption in the core network, as multiple eNBs need to be powered. We therefore asked ourselves if (i) macrocells can be run in a more efficient way in presence of picocells, by reducing their transmit power, (ii) if picocells can be selectively switched off, concentrating traffic on fewer, better utilised, picocells, and (iii) what is the advantage of D2D communications in this case. Results presented in Section 6 show that the answer is positive in all cases.

The key results achieved with respect to this point, are:

- ***macrocells can be operated at reduced transmit power, without compromising the throughput perceived by the users thanks to the additional capacity provided by picocells (and the net power consumption is lower when the solution with picocells is adopted)***
- ***picocells can be switched off, without impacting on the throughput perceived by users, which remains at the same level***
- ***when D2D communications are used, additional picocells can be switched off, even for a higher throughput provided to end users; it is expected that even greater energy savings could be achieved for the same throughput level.***

Note that this activity is synergic with the work on end-to-end delay in presence of duty cycling, presented in Section 2.2. Specifically, the two works together provide initial results on how to operate offload networks in an energy efficient way, considering both the energy consumed in the access network, and the energy consumed on the users' mobile devices.

Finally, in Section 7 we present results on the third activity in Task 3.3, related to **inter-technology scheduling**. We used a standard optimisation framework (TOPSIS) to optimally allocate users to technologies, when multiple technologies are available at the same time. The framework is customised based on a number of QoS metrics, and modified to find the optimal allocation considering overall capacity performance goals (instead of maximising individual nodes benefit).

The key results achieved with respect to this point, are:

- ***our results show that by using the TOPSIS framework, together with appropriate global throughput maximisation functions, we can effectively schedule users across multiple wireless technologies (cellular and WiFi), avoiding cellular network overload;***
- ***we show how to extend the TOPSIS framework to also include D2D communications, and provide results about the further gain obtained when D2D is also available, in addition to cellular and WiFi technologies;***
- ***for a given population of users and traffic demands, we find the optimal operating points, where the LTE and WiFi networks are not saturated, and the users satisfaction is maximised.***

1.2.3 Progress with respect to Y1 and Y2 activities

To make it more clear the evolution of activities during the WP lifetime, we hereafter briefly recall the status of the work after Y1 (i.e., with respect to deliverable D3.1 [1]), and after Y2 (i.e., with respect to the first version of this document, D3.3.1 [3]).

After the end of the first year, the WP activities built upon the initial results presented in D3.1 [1]. This produced several new results during Y2, that have been presented in D3.3.1 [3]. Summarising:

- The overall methodology of the work in T3.2 was confirmed, and activities have extended what presented in D3.1 in the first place. This is the case of (i) results about convergence of opportunistic networks; (ii) results about the end-to-end delay in opportunistic networks; (iii) results about the

capacity of LTE networks in terms of throughput; (iv) results about the performance of Droid in joint cellular, WiFi and opportunistic networks

- In addition, we have analysed more comprehensively LTE networks by developing a user-oriented throughput model, and to analyse more in depth the capacity advantage of offloading in a complete offloading solution with non-synchronised content requests. Finally, we carried out and completed activities on scheduling producing several interesting results.

In the last phase of the WP, we completed the lines of research identified before, in order to achieve the objectives of the WP. With respect to the results already presented in D3.3.1 [3], this has resulted in several novel achievements both in terms of capacity assessment, and in terms of characterisation of scheduling solutions. In particular:

- We have derived sufficient and necessary conditions for convergence of opportunistic networking protocols for generic mobility patterns. These conditions are on aggregate inter-contact time statistics, as opposed to on pairwise inter-contact time statistics, already presented in D3.3.1. This is a major improvement from a practical standpoint. Checking conditions on pairwise statistics means characterising each and every contact process between pairs of nodes. On the other hand, in the case of aggregate statistics much fewer samples are required for each pair of nodes, and pairwise statistics cannot be – in general – traced back from aggregates. The resulting practical tools for assessing convergence become therefore more efficient (as they require less samples of contact events) and more privacy preserving (as only aggregate information is required). This is presented in Section 2.1.4 of this document.
- We have analysed the effect of duty cycling on detected inter-contact times for general mobility patterns, i.e. also for the case where resulting inter-contact times are not exponential. We have further characterised the detected contact process also when the duration of contact events is not neglected (as previously done in D3.1 [1]). This is also a significant step forward towards deriving complete end-to-end delay models in presence of duty cycling, which is the final goal of this part of WP3 activities (that will be achieved by M33). This is presented in Section 2.2.2 of this document.
- We have defined and evaluated a learning-based framework to guide the split of traffic between multicast and D2D, where these techniques are used at the same time for disseminating content within an LTE cell. We have shown that using both techniques significantly improves LTE resource efficiency, over using any single one. In addition we have also shown that the learning framework is able to dynamically learn the configuration of the split to use based on the population of nodes and the deadline for content delivery, well approximating an optimal (infeasible) solution whereby configuration is decided after an exhaustive search on the parameter space. This is presented in Section 5.4 of this document.
- We have refined the results about scheduling with the TOPSIS framework, by better analysing the resulting split of users across cellular and WiFi technologies. Furthermore, we have also extended the framework to take D2D communications into account, and we have derived results showing the beneficial effect on capacity gain. We have finally shown optimal operating conditions for scheduling traffic among the LTE, the WiFi and the D2D networks. This is presented in Section 7 of this document.

1.3 WP3 activities in the overall framework of the project

Due to the nature of how WP3 was planned, activities and results provide a library of solutions that can be composed together when needed, in order to assess the capacity of offload network either in operation or at the design stage. To give one concrete example of how such results can be combined and used, let us consider the case of an operator wishing to plan its network in case of opportunistic offloading with energy saving at the users' devices. Results presented in D3.2 tell how the process of inter-contact times between nodes is modified according to the use of duty cycling. Then, criteria presented in Section 2.1 of this

deliverable can be used to select which specific protocol should be used, to avoid divergence problems. Depending on the application to be supported, the operator can identify the requirements in terms of end-to-end delay that need to be guaranteed. Therefore, models in Section 2.2 can be used to tune energy saving to meet these constraints. Finally, the operator would use a Push&Track (or Droid) system to implement the offloading solution. Based on the guaranteed end-to-end delay, it will be clear which fraction of the traffic can be expected by the operator to be served through the opportunistic network, thus finally obtaining an estimate of the additional capacity that would be gained through the so-configured offloading process. The architectural framework and mechanisms to support these operations is presented in D2.2.2.

In addition, the WP3 work presented in this deliverable is well aligned with the overall flow of activities of the project, and specifically with WP2 (architecture), WP4 (protocols), WP5 (performance evaluation). Specifically, all capacity building blocks and proposed algorithms are totally inline with the MOTO reference architecture defined in D2.2.1 [4]. Specifically, in the architecture we have clearly separated blocks dealing with how to manage and use wireless infrastructure (LTE and WiFi), and opportunistic networks. The separation of the capacity work in T3.2 is aligned with this architectural separation, and results can thus be fed back to the modules corresponding to these architectural blocks. Moreover, the integrated studies of offloading network follow the general design features of Push&Track, which implements all the basic architectural building blocks identified in WP2. Finally, activities on intra-technology scheduling naturally fit into the corresponding architectural elements, while inter-technology scheduling solutions are amenable to be implemented in the control blocks of the architecture dealing with orchestration between multiple wireless technologies. With respect to WP4, WP3 provides identifies initial solutions and alternatives, to be then analysed more precisely in the framework of the specification of networking protocols in WP4. An example is the work undertaken in WP4 about resource limitations in opportunistic network. This is based on the basic algorithmic tools of Push&Track, whose performance are assessed in WP3, and on the selection of appropriate configurations for opportunistic protocols, available after the analysis in T3.2. Finally, with respect to WP5, in WP3 we evaluate either analytically or by simulation individual solutions, to better identify which technical solutions to integrate in the testbeds or in the integrated simulation platform. This is achieved often through simplified simulation models and scenarios, where we abstract (with respect to WP5) some characteristics, to be able to have quicker and “more agile” results about the performance of selected offloading building blocks.

In the rest of this document we expand on the concepts described above, and present the detailed results produced in WP3 during the second year and the first part of the third year of the project. Finally, in Section 8 we discuss the final goals of WP3, after the informal extension of the activities suggested by reviewers during the second review meeting.

As a note to the reader, each section starts with a summary of the content presented herein. Then, subsections present these activities in some more details. When appropriate, we omit technical details that will make the presentation too long. In these cases, Appendices are provided where all details are available. **Therefore, the document is structured so that it can be read at multiple levels of details**. The first parts of the sections are sufficient to understand the content in terms of the approach taken and the main results achieved. The rest of the sections go in more details presenting results and how they have been achieved. Appendices contain the rest of the details. **New results not already included in D3.3.1 are presented in Sections 2.1.4, 2.2.2, 5.4, 6.4 and 7, respectively.**

2 Capacity analysis: Assessing capacity of opportunistic networks

This section presents two main contributions, focused on the assessment of the capacity of opportunistic networks. Both contributions build upon, and significantly extend, results presented in previous documents, specifically in D3.1 (Initial results on offloading foundations and enablers) [1] and D3.2 (Spatiotemporal characterization of contact patterns in dynamic networks) [2].

In Section 2.1 we present results on convergence of opportunistic networking protocols. Remember that in our overall strategy, this is one of the necessary steps to characterise the capacity of the opportunistic network, because it allows us to identify configurations of the network (i.e., patterns of contacts between users) for which specific protocols do not converge. In these cases, the capacity gained by offloading would be zero, as protocols would yield infinite expected delay (more practically, messages will be lost and never delivered to the destination). Initial results on this topic were presented in Section 4.1 of [1]. In Section 2.1 we summarise the additional results we have obtained (complete details are available in Appendix A, which is a reprint of [12]). Specifically, we have fully characterised the convergence properties of both randomised (or social-oblivious) and social-aware routing protocols in case of Pareto distributed inter-contact times. Social-oblivious protocols do not use any contextual information about the behaviour (and thus resulting contact patterns) of users, while social-aware protocols are built to exploit such knowledge. The set of protocols considered cover the vast majority of forwarding algorithms defined in the opportunistic networking literature, while considering Pareto inter-contact times is inline with well-established results in the literature about the analysis of real mobility traces. Our results allow us to draw a pretty interesting and useful set of conclusions. Within each class of routing protocol (social-oblivious, social-aware) we are able to identify best solutions, i.e. those that guarantee convergence in most of the cases. Comparing best solution of each class, we find that there is not a unique winner, but the best overall choice actually depends on the pattern of contacts between nodes. It is particularly interesting to find that using social-oblivious protocols may yield convergence in some cases where social-aware would not.

In Section 2.2 we deal with end-to-end delay models in opportunistic networks in presence of duty cycling, for the case of exponential inter-contact time patterns (complete details are provided in Appendix A, which is a reprint of [9]). Specifically, we provide stochastic guarantees on the end-to-end delay as a function of the duty cycle period used by nodes. In other words, our results allow an operator to set the optimal duty cycling so that the end-to-end delay is below a given threshold with a given probability. This model provides an analytical tool to set the trade-off between the additional capacity that can be provided and the corresponding energy cost (in terms of battery of users' mobile devices). This result derives from two previous results achieved in the project. On the one hand, we exploit the results on how duty cycling modifies the patterns of *useful* contact between nodes in case of exponential contact patterns, where useful contacts are those that can be used to forward messages, i.e. those that occur when both nodes are active (not sleeping). These results have been presented in Section 4 of [2]. On the other hand, we exploit the models of end-to-end delay in case of exponential inter-contacts, presented in Section 4.2 of [1].

2.1 Convergence of opportunistic networking protocols

Modelling the performance of social-oblivious and social-aware forwarding protocols for opportunistic networks is still an open research issue. Knowing the distribution of intermeeting times and the rules applied by the forwarding algorithm used in the network, one could - in principle - model the distribution of the delay experienced by messages and compute its expectation. In practice, modeling analytically the delay of the various forwarding protocols for general distributions of inter meeting times is very hard, and models exist only for some specific cases, typically assuming exponential intermeeting times [65][66][41][67][54][31]. A related modelling challenge is to assess the convergence of routing protocols, i.e. whether a specific protocol yields finite or infinite expected delay. Assessing convergence allows us to understand whether a particular protocol can be safely used or not given a pattern of intermeeting times and how to configure it so that it converges, if possible. Although less informative than a complete delay

model, convergence models can be derived for a large class of routing protocols releasing the exponential intermeeting time assumption.

The convergence of the expected delay is not guaranteed in all cases in which the expectation of the inter-meeting times may diverge. In fact, being the delay the result of the composition of the time intervals between node encounters, depending on the convergence of inter- meeting times, the expectation of the delay itself might diverge. This can happen, for example, when intermeeting times feature a Pareto (also known as power law) distribution, as first highlighted in [22]. The problem with Pareto distributions is that their expectation is finite only for certain values of their exponent α. More specifically, the expectation is finite if $\alpha > 1$, while for $\alpha \leq 1$ it diverges to infinity. The first to postulate the existence of Pareto intermeeting times in real mobility scenarios (i.e., analyzing real traces of human mobility) were Chaintreau et al. in their seminal work in [22]. The relevance of Pareto intermeeting times in opportunistic networks is both theoretical and empirical. Cai and Eun [20] have mathematically derived that heavy-tailed intermeeting times can emerge depending on the relationship between the size of the boundary of the considered scenario and the relevant timescale of the network, showing that, at least in principle, Pareto intermeeting times are something that one may be faced with when studying opportunistic networks. Empirical evidence for the presence of Pareto intermeeting times was first suggested by [22], but it has been later criticised, arguing that the tail of the distribution is in fact exponential (e.g., [37]). Typically, these results are derived focusing on the aggregate inter-contact time distribution, while convergence depends on pairwise distributions. As proved in [51], the aggregate and pairwise distributions can be in general very different, and therefore analysis of pairwise inter-contact times are necessary, which are however mostly missing in the literature. To address this issue, we have performed a pairwise hypothesis testing on three popular publicly available contact datasets (Cambridge, Infocom'05, and RollerNet) and we have found that the Pareto hypothesis for intermeeting times cannot be rejected for 80%, 97%, and 85.5% of pairs, respectively. We believe that these results provide a strong case for Pareto intermeeting times in opportunistic networks and substantially motivate analyses like the one presented in this paper.

Under the Pareto intermeeting times assumption, in this work we derive the stability region (i.e., the Pareto exponent values of pairwise intermeeting times for which finite expected delay is achieved) of a broad class of social-oblivious and social-aware forwarding protocols (single- and multi-copy, single- and multi-hop). The starting point of our paper is the work by Chaintreau et al. [22], where such conditions have been studied for the two-hop scheme under the assumption of homogeneous mobility (i.e., i.i.d. intermeeting times across all pairs). However, measurement studies [23] [22] have shown that real networks are intrinsically heterogeneous. Thus, in this work, we investigate whether heterogeneity in contact patterns helps the convergence of the expected delay of a general class of social-oblivious and social-aware forwarding protocols, and whether convergence conditions can be improved using multi-copy strategies and/or multi-hop paths.

Overall, the key findings we obtain from this analysis are as follows:

- For *social-oblivious strategies*, if convergence can be achieved, *two hops are enough* for achieving it.
- Using n *hops* can help *social-aware schemes*, and make them converge in some cases when all other social- aware or social-oblivious schemes diverge.
- In both the social-oblivious and the social-aware case, we find that *multi-copy strategies* can achieve a finite expected delay even when single-copy strategies cannot.
- Comparing *social-oblivious and social-aware multi-copy solutions*, we are able to prove mathematically that there is no clear winner between the two, since either one can achieve convergence when the other one fails, depending on the underlying mobility scenario.

A concise presentation of these findings is provided in the following subsections. All details are available in Appendix A.

2.1.1 Social-oblivious protocols

The analysis of convergence for social-oblivious protocols was presented already in [1], and is briefly summarized here for the reader's convenience.

To accurately represent the different variants in this class, we identify three main groups, differing in the number of hops allowed between source and destination, the number of copies generated, and whether the source and relay nodes keep track of the evolution of the forwarding process or not. First, forwarding strategies can be single-copy or multi-copy. In the former case, at any point in time there can be at most one copy of each message circulating in the network. In the latter, multiple copies can travel in parallel, thus in principle multiplying the opportunities to reach the destination (we assume that all copies are generated by the source node). Second, forwarding protocols can be classified based on the number of hops that they allow messages to traverse, or, in other words, based on a TTL computed on the number of hops. When the number of allowed hops is finite, the last relay can only deliver the message to the destination directly. Third, the amount of knowledge that each agent in the forwarding process can rely on (or is willing to collect and store) is an additional element for classifying forwarding strategies. Focusing on the source node, there can be social-oblivious strategies in which the source node does not keep track at all of how the forwarding process progresses. In this case, considering the configuration in which the source node can generate up to m copies of the message, the m copies might end up being all distributed to the exact same relay, thus eliminating the potential benefits of multi-copy forwarding. A memoryful source, instead, is able to guarantee to use distinct relays. A similar problem holds for intermediate relays. Memoryless relays can forward the message to the same next hop more than once, because they are not at all aware of what happened in the past. On the other hand, memoryful relays possess this knowledge, and are able to refuse the custody of messages that they have already relayed.

The following conditions are found for convergence of these protocols

	1 hop		2 hops		n-hop	
	1 copy	m copies	1 copy	m copies	1 copy	m copies
memoryless	$\alpha_{sd} > 2$	-	[C1,C2]	[C1,C2]	[C1,C2]	[C1,C2]
memoryful source	-	-	-	[C3,C4]	-	[C1,C2]
memoryful relays	-	-	-	-	[C1,C2]	[C1,C2]

Table 1. Convergence conditions.

Specifically, conditions C1 through C4 in the table are defined as follows

C1 $\sum_{j\in\mathcal{P}_s} \alpha_{sj} > 1 + |\mathcal{P}_s|$, *where* $\mathcal{P}_s$ *denotes the set of all nodes that can be encountered by node* s;

C2 $\alpha_{jd} > 2, \forall j \in \mathcal{P}_s - \{d\}$.

$C3 = m \leq m^*$

$C4 = \sum_{j=N-m}^{N-1} \alpha'_j > 1 + m$

where *s* and *d* denote the source and destination nodes, respectively, *m* denotes the number of copies generated by the source, *m** is defined as follows

$$m^* = \begin{cases} 0 & \text{if } \sum_{j\in\mathcal{P}_s} \alpha_{sj} \leq N \\ \arg\max_m \{m + \sum_{i=m}^{N-1} \alpha_i^* > 1 + N\} & o.w. \end{cases}$$

and α_i^* denotes the *i*-th largest α_{sj} with $j \in P_s$.

It is useful to focus on one specific case, in order to clarify how these conditions can be interpreted. Let us consider the 2-hop 1-copy memoryless scheme, which converges iff conditions C1 and C2 are met. The physical meaning of the conditions is quite intuitive. Recall that in the 2-hop 1-copy scheme the source hands over the only copy of the message to the first encountered node, which then has to relay it directly to the destination. Condition C1 guarantees that the first phase occurs with a finite expected time. Specifically, the source node encounters the first possible relay with a time that is distributed according to a Pareto law with shape $\sum_{j\in P_s} \alpha_{sj} - |P_s|$. Therefore, the first phase "converges" if the average value of this time

is finite, which leads to condition C1. Condition C2 guarantees that whatever relay is chosen by s, it encounter the destination within a finite expected time (note that the time for such relay to meet the destination is the residual of their intermeeting time, as the process of encounter between nodes is asynchronous, and therefore node s meets the relay at a random point in time with respect to the meetings between the relay and the destination).

Note that conditions C3 and C4 are needed only in case of multi-copy forwarding. The value m^* is a threshold on the number of copies, such that if the source generates up to m^* copies, all of them are handed over to m* distinct relays with finite expected delay, while if m exceeds m^* the additional copies cannot be handed over with finite expected delay. Condition C3 thus imposes that the source can actually relay m distinct copies of the message, while condition C4 guarantees that the destination meets at least one of the used relays with finite expected delay.

2.1.2 Social-aware protocols

2.1.2.1 Definition of social-aware forwarding protocols

Due to the variety of social-aware schemes available in the literature, here we only consider an abstract social-aware protocol that measures how good a relay is for a given destination in terms of its *fitness*. The fitness fit^d_i is assumed to be a function of how often node i meets the destination d, thus fit^d_i can be taken as proportional to the rate of encounter between node i and the destination. $1/E[M_{id}]$.

Under this abstract and general social-aware strategy, upon encounter, a node i can hand over the message to another node j only if its fitness is lower than the fitness of the peer, i.e., if $fit^d_j > fit^d_i$ holds (in the following we drop superscript d). The fitness function considered here uses only information on contacts between nodes, which have a direct dependence on the intermeeting time distribution. This lets us clearly show what is the impact of the contact dynamics on the performance of opportunistic forwarding protocols. How such simple fitness function can be extended to more complex forwarding strategies has been discussed in [12]. In the following, we denote with R_i the set of possible relays for node i, i.e., the set of nodes whose fitness is greater than that of node i. Therefore, with social-aware forwarding, nodes can hand over a message only to nodes with higher fitness. This means that the set of potential relays shrinks as the message is handed over from hop to hop towards the final destination. This is pictorially represented in Figure 2.

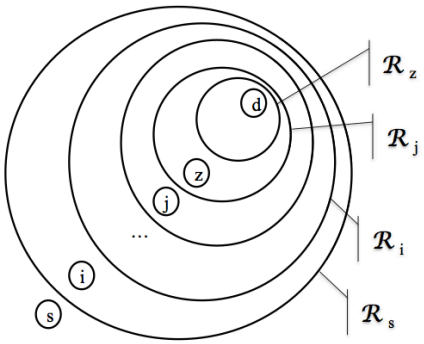

Figure 2. Schematic representation of social-aware forwarding.

2.1.2.2 Convergence conditions for social-aware protocols

Based on this definition of social-aware forwarding protocols, we have been able to analyse different families of protocols, again distinguishing between single- and multi-hop protocols, and between single- and multi-hop protocols. In all cases we consider memoryful protocols, as it does not make much sense to assume memoryless protocols when they have already to store information about contact patterns between nodes.

The corresponding convergence conditions are provided in Table 2 (note that the definition of conditions is slightly modified with respect to [12] for the sake of clarity, in order to match conditions for social-oblivious protocols in Table 1).

	1 hop		2 hops		n-hop	
	1 copy	m copies	1 copy	m copies	1 copy	m copies
social-oblivious	$\alpha_{sd} > 2$	-	[C1,C2]	[C3,C4]	[C1,C2]	[C1,C2]
social-aware	$\alpha_{sd} > 2$	-	[C5,C6]	[C9]	[C7,C8]	[C7,C8]

Table 2. Convergence conditions for both social-oblivious and social-aware schemes

We hereafter only sketch the intuitive meaning of one such condition, by focusing on the 1-copy 2-hop scheme (and refer the reader to Appendix A for the details). This protocol converges if and only if conditions [C5,C6] are met, which are defined as follows

C5 $\sum_{j \in \mathcal{R}_s} \alpha_{sj} > 1 + |\mathcal{R}_s|$

C6 $\alpha_{jd} > 2, \forall j \in \mathcal{R}_s - \{d\}.$

where Rs denotes the set of nodes encountered by the source with higher fitness (than itself) towards the destination. Conditions C5 and C6 maps exactly conditions C1 and C2 of the social-oblivious case. Specifically, they state that the expected delay in the 2-hop 1-copy case if finite iff the source encounters one of the possible relays (i.e., one node in Rs) in a finite amount of time (condition C5), and each of these nodes encounters the destination in a finite amount of time (condition C6). Conditions C7,C8 and C9 can be derived for the other protocols using a similar line of reasoning (see Appendix A).

As in the social- oblivious case, multi-hop schemes do not benefit from the use of multiple copies, and in fact the 1-copy n-hop scheme and the m-copy n-hop scheme share the same convergence conditions. Similarly, the difference between 2-hop schemes mirrors that between the corresponding social-oblivious versions. Thus, the 1-copy 2-hop scheme is effective when $\alpha_{jd} > 2$ for all $j \in$ Rs, since it allows us to save resources by sending a single copy. However, when conditions [C5,C6] do not hold, the only chance to achieve convergence is to exploit multiple copies.

If we focus on single-copy schemes, it is interesting to note that, differently from the social-oblivious case in which using additional hops did not provide any advantage, 1-copy social-aware schemes may benefit from multiple hops. In fact, for the 1-copy 2-hop scheme we need to impose that all intermediate relays j meet the destination with $\alpha_{jd} > 2$, (conditions C6) which is a quite strong condition. On the other hand, if we use multiple hops (1-copy n-hop case), conditions C7 and C8 are required, which are milder than C5. Their definition requires several steps, and therefore we don't report them here for the sake of simplicity (see Appendix A for the details). Basically, the only constraint for the 1-copy n-hop case is that there must be at least one node z (the one with the highest fitness) meeting the destination with $\alpha_{zd} > 2$.

Finally, we compare the m-copy 2-hop case with the 1-copy n-hop case (which is equivalent to the m-copy n-hop scheme). There is no clear winner here, as each scheme can provide convergence when the other one cannot. For example, consider the case in which the source node is not able to send more than one copy within a finite amount of time. In this case, the m-copy 2-hop scheme becomes effectively a 1-copy 2-hop scheme, which fails to achieve convergence if some intermediate hop j does not have exponent α_{jd} greater than 2 (condition C6). Instead, exploiting multiple hops pays off in this case, as it allows us to rely on more intermediate relays, which may not meet the destination within a finite expected time but can bring the message "closer" to nodes that do meet d with $\alpha_{jd} > 2$. Vice versa, when the source node can hand over multiple copies (m>1) within a finite delay, the cooperative delivery of the multiple copies can overcome the presence of intermediate relays for which conditions C8-C9 do not hold. For example, when there is not even one relay j with $\alpha_{jd} > 2$, then the m-copy 2-hop case is the only possible choice.

2.1.3 Comparing social-oblivious and social-aware schemes

In the following we take the champions of each class and we investigate whether there is a clear winner between social-oblivious and social-aware strategies when it comes to the convergence of their expected delay.

Let us first consider the case $\alpha_{sd} > 2$. With this configuration the Direct Transmission scheme is the best choice from the convergence standpoint. In fact, with social-oblivious schemes using more than one hop, "bad" relays can be selected even starting from a source that is already able to reach the destination with a finite expected residual intermeeting time. This does not happen with social-aware strategies. In fact, assume that the source is the only node with $\alpha_{sd} = 2 + \varepsilon$, while all other nodes meet the destination with $\alpha_{jd} = 1 + \varepsilon$, with ε being a very small quantity. In the social-aware case, Rs contains only the destination, as all other nodes are clearly worse than the source node as relay. This shows the adaptability of social-aware schemes: the additional knowledge that they exploit makes them able to resort to simpler approaches (in this case, Rs = {d} is equivalent to the Direct Transmission) when they realize that additional resources in terms of number of copies or number of hops would not help the forwarding process. This implies that one can safely use the m-copy 2-hop or the 1-copy n-hop social-aware protocols because in the worst case they will do no harm (they will downgrade to simpler strategies, without exploiting wrong paths), while in the best case they are able to improve the convergence of the forwarding process.

When $\alpha_{sd} \leq 2$ and $\alpha_{jd} > 2$ for all nodes j in the relay set (i.e., $j \in Rs - \{d\}$ for the social-aware case and $j \in Ps - \{d\}$ for the social-oblivious case), the strategy of choice is the 1-copy 2-hop for both the social-oblivious and social-aware category. However, the 1-copy 2-hop social-aware scheme is overall more advantageous than its social-oblivious counterpart. More specifically, when the source node is the worst relay for the destination (i.e., $\min_i\{\alpha_{id}\} = \alpha_{sd}$), the social-oblivious and the social-aware approaches are equivalent (given that Ps = Rs). In all other cases, instead, $Rs \subset Ps$, thus, for the set of nodes in $Ps - Rs$, social-aware forwarding does not impose any constraint, while social-oblivious forwarding needs to impose constraints, thus resulting in stricter conditions for convergence.

Let us now focus on the remaining cases, namely i) when $\alpha_{sd} \leq 2$ and not all intermediate relays have exponent greater than 2, and ii) when $\alpha_{jd} \leq 2$ for all nodes j. In the first case, the social-aware m-copy 2-hop, the social-aware 1-copy n-hop, and the social-oblivious m-copy 2-hop can achieve convergence. In the second case, the only options for convergence are the social- aware m-copy 2-hop and the social-oblivious m-copy 2- hop. We first highlight the differences between the n-hop approach and the 2-hop approach by discussing when the social-aware 1-copy n-hop outperforms the other two strategies in terms of convergence (which can only happen in case i), then we focus on the social-aware and social-oblivious m-copy 2-hop strategies, thus covering both case i and ii.

Assume that there exists at least one node z that meets the destination with $\alpha_{zd} > 2$. The m-copy 2-hop strategies send multiple copies to a set of relays, which in turn can only deliver the message to the destination directly. This implies that intermediate relays must have collectively the capability of reaching the destination, for all subsets with size m of possible relays. Here, only meetings with the destination are relevant, and if all relays but z have very low exponent for encounters with the destination, convergence may not be achieved. Differently from the 2-hop strategies, the social-aware n- hop scheme do not rely exclusively on the capabilities of meeting with d, but it is able to generate a *path* towards the destination in which intermediate nodes may not be good relays for d but good relays towards nodes with high fitness (in the extreme case, only $\alpha_{zd} > 2$ can hold). Thus, in the n-hop case, as long as the message can leave intermediate relays within a finite expected time, this could be enough for convergence. When all three strategies achieve convergence, the one to be preferred can be chosen based on resource consumption considerations. With the m-copy 2-hop strategies there can be up to 2m transmissions, while with the 1-copy n-hop scheme there are n. Hence, when $n < 2m$, the single-copy scheme should be preferred.

Let us finally compare the social-oblivious and the social-aware m-copy 2-hop schemes. Since they seem to cover similar mobility scenarios (as discussed in the previous section) and to be based on similar mechanisms (the mini and maxi quantities, whose relation with m determines the convergence), it may be

difficult to intuitively evaluate which one performs better in terms of convergence. As shown in Appendix A, it may happen that either the social-oblivious m-copy 2-hop scheme achieves convergence when the social-aware m- copy 2-hop scheme does not, or vice versa, depending on the underlying mobility process. An example of the first case is when there are a lot of nodes that meet the source with high α_{sj}; if those relays have very low α_{jd}, they will not be used by the social-aware scheme, and this may hinder convergence of the second hop. It is easy to construct a corresponding example for the other case.

2.1.4 Assessing convergence through aggregate statistics

As stated several times, a number of forwarding approaches for opportunistic networks are available in the literature. Two very large families are typically identified, distinguishing *randomised* protocols from *utility-based* protocols. In randomised protocols nodes do not exploit any contextual information about the encountered nodes, and forward according to simple, sometime probabilistic, rules. In utility-based protocols nodes exploit context information about the encountered peers to compute a utility score which tells how suitable the peers are to deliver the message to the destination.

It is clear that the performance of whatever type of forwarding protocol fundamentally depends on the properties of the contact patterns between nodes. A critical aspect is to assess whether a forwarding protocol converges or not, depending on the contact patterns between nodes. By *convergence* for a given source-destination pair, we mean that the forwarding protocol yields finite expected delay for that pair. By *global convergence* we mean that the protocol converges for any pair.

The overall goal of this part of WP3 activities is to characterise how the properties of inter-contact times in heterogeneous opportunistic networks determine the convergence of randomised forwarding protocols. Note that randomised forwarding protocols are not only relevant for homogeneous networks where contact processes between pairs of nodes are all identical, and therefore all nodes necessarily have the same utility towards all the other nodes. Also in heterogeneous networks, where contact processes are different across pairs, computing utility scores requires to collect statistics about the behaviour of nodes, such as their rate of encounter or their social network properties. Having reliable estimates for such parameters may require a number of contact events and thus significant time, during which using utility-based forwarding may be imprecise or not feasible at all. In addition, such statistics might not be available at all due to privacy and trust issues, or may be too unreliable if the system is highly dynamic and non stationary. Assessing convergence is important at least for two key reasons. On the one hand, assessing global convergence tells whether a given protocol can “safely” be used in a specific network. On the other hand, as we show in this section, highlighting the properties of the contact patterns that hinder global convergence indicates the nodes that should be avoided in the forwarding process when using a particular protocol.

Convergence of randomised protocols has been analytically characterised in the seminal work by Chaintreau et al. [22] for *homogeneous* networks where inter-contact times follow a Pareto distribution, and ICTs are independent and identically distributed (iid). In D3.3.1 [3] we have already extended this analysis to the case of heterogeneous networks where ICTs are Pareto with different parameters. However, to the best of our knowledge, a more general characterisation of this dependence for *heterogeneous* networks with *generic* ICT distributions is still missing. We think that filling this gap is relevant because analysis of ICT in human mobility traces has shown that ICT distributions vary across traces and, inside traces, across pairs of nodes. In addition, inside the same trace different pairs can meet with different distributions. Existing results about convergence, thus, do not completely address these scenarios.

More in detail, the contribution of this part of the document is threefold. The first contribution consists in extending the results in [22] and those already presented in D3.3.1, and provide sufficient and necessary conditions on the distributions of individual ICTs for convergence of randomised forwarding protocols in heterogeneous networks for a very large family of ICT distributions.

The second contribution is investigating whether it is possible to assess convergence by considering *aggregate* ICT statistics, instead of studying the ICT distributions of each and every individual pair. We

consider two aggregate statistics. The first one is the distribution of *aggregate ICTs*, i.e. the distribution obtained by considering altogether the ICT samples of all possible pairs of nodes. The second one is the distribution of the *contact rates*, where the contact rate of a pair is defined as the reciprocal of the average inter-contact time. Aggregate statistics are customarily used in the literature, although sometimes this is not fully justified (as we have shown in [51]). Being able to assess convergence by looking at aggregate statistics is highly desirable from a number of standpoints. Aggregate statistics disclose much less information about the behaviour of the users than the ICT distributions of individual pairs. It is thus reasonable to assume that, while individual ICT distributions might hardly be available in real scenarios due to privacy concerns, aggregate statistics could be distributed more easily. In addition, collecting enough ICT samples for each pair to reliably characterise individual ICT distributions might be prohibitive in terms of time, storage and network traffic overhead. This is the reason why, even in rather large datasets, it is difficult to identify with statistical confidence "the" distribution that fits ICTs of a given pair, and a certain level of uncertainty among a set of candidate distributions always remains [23]. The same number of samples can, instead, characterise more reliably aggregate statistics. In this part of the WP activities we have formally proved that convergence of randomised protocols can be assessed also in heterogeneous networks through aggregate statistics, specifically through the distribution of aggregate ICTs or contact rates. Beyond that, our results show that, in certain cases, aggregate statistics "contain" the same information of individual statistics. In particular, they can be used to entirely characterise the individual contact processes between nodes, and thus any property that depends on them (not only convergence).

Finally, the third contribution is defining distributed algorithms to (i) compute aggregate statistics at individual nodes without central control, and (ii) tune the behaviour of a range of randomised protocols to guarantee convergence.

Hereafter, we provide an intuitive sketch of each of these contributions. More details are available in [52].

2.1.4.1 *Convergence conditions on individual inter-contact times for general heterogeneous mobility patterns*

In this section we analyse the relationship between the convergence of randomised protocols and the properties of the ICT distributions. For all of them but one, the conditions we find on ICT distributions are sufficient and necessary for convergence. Specifically, we define this class as *naïve* forwarding protocols According to the analysis already presented at the beginning of the section, we consider randomised protocols where the source can generate m copies of a message, each copy can travel up to h hops before being delivered directly to the destination, and the source can (memoryless) or cannot (memory-full) re-use the same relay for different copies. Naïve forwarding protocols are defined as all such protocols, but the one that uses m > 1 copies, 2 hops, and is memory-full (m-copy 2-hop memory-full protocol).

The convergence conditions derived hereafter hold for ICT distributions whose density belongs to set Ψ, defined as the set of functions that, for large x, can be either upper bounded by $x^{-(3+\beta)}$ for some β > 0, or lower bounded by x^{-3}, i.e.

$$\Psi = \left\{ h(x) \; s.t. \; \exists \beta, x_0 > 0 \,|\, \forall x > x_0 \right. \\ \left. h(x) \leq x^{-(3+\beta)} \vee h(x) \geq x^{-3} \right\}$$

Figure 3 shows a graphical representation of functions belonging and not belonging to set Ψ (the axes are in log scale). Note that functions not belonging to Ψ must have a very particular behaviour. For any β > 0, they must lie in the band between $x^{-(3+\beta)}$ and x^{-3}, and they cannot "jump out" of this band for all values of x greater than some x_0>0 (or, in other words, if they "jump out" after x_0, they must "jump in" again after a *finite* interval). Note that, as β can be *any* small positive value, this band can be made arbitrarily narrow. More concretely, the densities of all the ICT distributions typically considered in the literature (e.g., exponential, Pareto, Pareto with cut-off) and of many more popular distributions (e.g., Gamma, Normal, Log-normal, all densities that can be approximated by a power law or that cannot be lower-bounded by a power law) all belong to set Ψ.

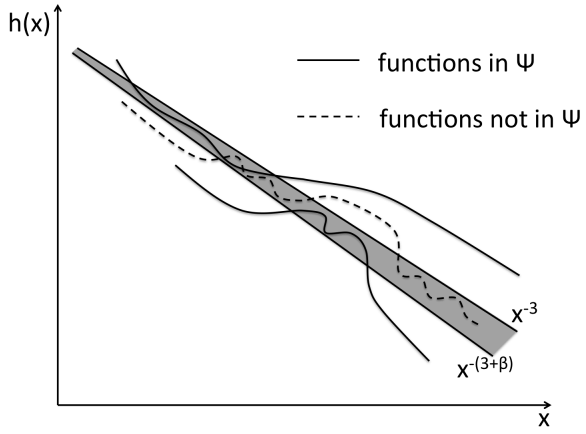

Figure 3. Functions belonging or not to set Ψ

The convergence conditions for naïve forwarding protocols we are able to obtain are described in Theorem 2 of [52], which is reported hereafter:

***Theorem 2**: If, for any pair of nodes, the density of ICTs $f_X(x, \lambda)$ belongs to Ψ, naïve forwarding protocols achieve global convergence iff the second moment of all the distributions is finite.*

Theorem 2 shows that, in order to assess global convergence, when the densities of ICTs belong to set Ψ, it is sufficient to consider the second moments of the distributions. From a complementary standpoint, it shows that when the second moment of even a single pair (n1,n2) diverges, naïve forwarding protocols yield infinite expected delay for a set of source-destination pairs (all pairs that can possibly communicate using a path that includes the (n1, n2) hop), and, therefore, using them without scrutiny, to support communication between *any* pair, would be risky.

2.1.4.2 Divergence of naïve forwarding protocols and aggregate inter-contact times

The importance of Theorem 2 is to derive conditions on generic distributions of ICTs, i.e. conditions that are not dependent on the specific pairwise contact distribution. However, as anticipated, we are able to extend the conditions of Theorem 2 to *aggregate* inter-contact time statistics. This is the key result of Theorem 3 in [52].

***Theorem 3**: If, for any pair of nodes, the density of ICTs fX(x, λ) belongs to Ψ, any naïve forwarding protocol achieves global convergence iff the second moment of the aggregate distribution is finite.*

The importance of Theorem 3 is to formally prove that, for a very broad class of contact patterns (whenever the densities of individual ICTs belong to set Ψ), studying the second moment of the aggregate distribution is sufficient to assess global convergence of naïve protocols. As discussed before, this result is very welcome, as characterising the aggregate distribution, instead of each and every individual pair distribution, is much easier and practical for a number of reasons. Moreover, Theorem 3 provides - for the first time - a theoretical foundation to the very common approach used in the literature, consisting in studying the aggregate distribution and not the distributions of individual pairs to assess convergence of forwarding protocols.

Another conceptually similar result can be obtained considering another aggregate statistic, i.e. the distribution of contact rates in the network. In order to derive these conditions, we exploit a well-known result about the approximation of generic random variables with appropriate Coxian random variables [34][35][59]. Specifically, we obtain Theorem 4.

***Theorem 4**: Irrespective of the distribution of inter-contact times, a density of contact rates allowing values arbitrarily close to 0 is sufficient condition for naïve forwarding protocols not achieving global convergence*

$$\lim_{\lambda \to 0^+} f_\Lambda(\lambda) > 0 \quad \Rightarrow \quad \text{naïve forwarding protocols do not achieve global convergence.}$$

Moreover, if individual ICT distributions are such that the distribution of the number of stages of their Coxian approximations always have finite coefficient of variation, then a density of contact rates allowing values arbitrarily close to 0 is sufficient and necessary condition for naïve forwarding protocols not achieving global convergence:

$$\forall \lambda \; C_K(\lambda) < \infty \;, \quad \lim_{\lambda \to 0^+} f_\Lambda(\lambda) > 0 \quad \Longleftrightarrow \quad \text{naïve forwarding protocols do not achieve global convergence.}$$

Based on Theorem 4, the divergence of global forwarding may even not depend on the specific distribution of ICTs, but only on the distribution of the contact rates. In other words, the properties of naïve forwarding protocols in terms of global convergence may be fully determined by the type of heterogeneity in the network, irrespective of the type of distributions of individual pair ICTs.

2.1.4.3 Exploiting aggregate statistics

The third contribution we highlight is how the above theoretical results can be exploited to tune the behavior of opportunistic forwarding protocols in case of risk of divergence. In [52] we first show how aggregate statistics can be computed in a distributed way. According to this algorithm, nodes exchange a number computed out of locally sampled inter-contact times (denoted as $C_j[0]$ for node j). If the second moment of ICT of a given pair (j, i) is infinite, both Cj [0] and Ci [0] will be infinite. For any other node in the network that receives Cj[0] and Ci[0], the fact they are infinite tells that j and i meet at least another (unknown) node with a contact pattern whose second moment is infinite. According to Theorem 2, this means that global convergence cannot be achieved. This property suggests a first way to extend the standard mechanisms of naïve protocols to achieve global convergence also under these conditions. This is described in Algorithm 3 of [52] ("Blacklisting nodes"), also reported below. Whenever node k receives an infinite value from node i for Ci[0], node i is blacklisted (line 4). When considering whether an encountered node is a suitable next hop for a message, node k checks whether this node is blacklisted (line 8), and considers it for forwarding only if it is not (line9). Thanks to this modification, naïve protocols only use the subset of nodes for which C[0] is finite. This guarantees that only pairs with finite individual second moment are used in the forwarding paths, and, based on Theorem 2, that the expected delay between any source and destination is finite.

Algorithm 3 Blacklisting nodes

```
1:  Initialise BL_k                          ▷ BL_k = nodes blacklisted by k
2:  procedure BLACKLIST(C_i[0])              ▷ Upon meeting node i
3:      if C_i[0] = ∞ then
4:          add i to BL_k
5:      end if
6:  end procedure
                                  ▷ Decides to use i or not as next hop
7:  procedure FORWARD(msg, i)
8:      if i ∉ BL_k then
9:          naïve_forward(msg, i)
10:     end if
11: end procedure
```

We have tested the effectiveness of using Algorithm 3 to avoid divergence in case where plain naïve protocols would diverge. Specifically, In Figure 4 we show the CCDF of the delay for various naïve protocols when the contact pattern configuration is as follows. We draw the parameters α_{ij} for the inter-contact times of each node pair i,j from a uniform distribution defined in [1, 2.5], except for one unique node u, for which we draw α_{uj} from a uniform distribution defined in [2.2,2.5]. Hence, for node u we guarantee that

$C_u[0]$ is finite, while for all the other nodes it is infinite. We can see in Figure 4 that, as expected, all the protocols considered do not converge (indicated by the non-null probability of infinite delay). Applying Algorithm 3 (Figure 5), we can instead achieve global convergence for each protocol but for Direct Transmission. In this configuration, plain naïve protocols diverge due to the fact that all nodes but one have at least a "divergent" link towards another node. Using Algorithm 3 all these nodes are filtered out, and only the node with "non-divergent" links is used. With Direct Transmission no intermediate relay is exploited, so filtering out nodes has no effect.

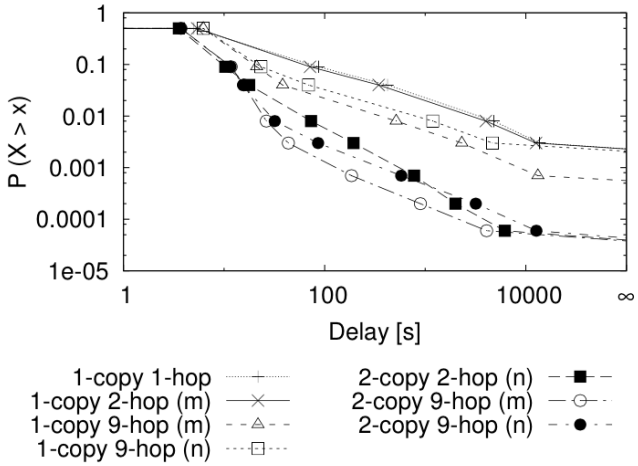

Figure 4. CCDF of delay for representative plain naïve randomised protocols.

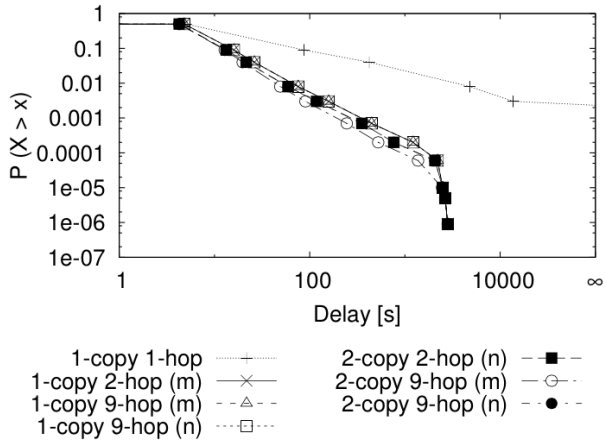

Figure 5 CCDF of delay with Algorithm 3.

Note that Algorithm 3 is just one example of possible extensions of naïve protocols based on Theorems 2, 3 and 4. [52] presents additional results, with Algorithm that drops nodes more selectively with respect to Algorithm 3.

2.2 End-to-end delay in opportunistic networks in presence of duty cycling

A possible roadblock in using opportunistic networking for offloading is the fact that direct communications may consume significant energy. To address this, nodes are typically operated in duty cycling mode, by letting their WiFi (or Bluetooth) interfaces ON only for a fraction of time. The joint effect of duty cycling and mobility is that, even if the network is dense, the resulting patterns in terms of communication opportunities is similar to that of conventional opportunistic networks, as devices are able to directly communicate with each other only when they come in one-hop radio range *and* both interfaces are ON.

The net effect of implementing a duty cycling scheme is thus the fact that some contacts between nodes are missed because the nodes are in power saving mode. Hence, *detected* intercontact times, defined as the time between two consecutive contact events during which a communication can take place for a pair of nodes, are longer than intercontact times determined only by mobility, when a duty-cycling policy is in place. This heavily affects the delay experienced by messages, since the main contribution to message delay is in fact due to the intercontact times. In [10] (presented in D3.2 [2]) we have focused on exponentially distributed intercontact times and we have studied how these are modified by duty cycling, obtaining that intercontact times remain exponentially distributed but their rate is scaled by the inverse of the duty cycle. Building upon this result, we have then investigated how the first moments of the end-to-end delay vary with the duty cycle for a number of opportunistic forwarding schemes. In addition, we have found that energy saving and end-to-end delay both scale linearly with the duty cycling. Therefore, for a single message delivery, the same energy saved through duty cycling is spent because the network must stay alive longer. Thus, the main advantage of duty cycling is enabling the network to carry more messages by being alive longer (rather than improving the energy spent for each single delivery).

Our work in [10] assumed that the value of the duty cycle was given and studied its effects on important performance metrics such as the delay, the network lifetime, and the number of messages successfully delivered to their destination. More in general, the duty cycling can be seen as a parameter that can be configured, typically, based on some target performance metrics. To this aim, in this part of the work we develop a mathematical model that allows us to tune the duty cycle in order to meet a given target performance, expressed as a probabilistic guarantee (denoted as p) on the delay experienced by messages. Considering probabilistic, instead of hard, guarantees, allows us to cover a very broad range of application scenarios also beyond best-effort cases – all but those requiring real-time streaming. Specifically, we study the case of exponential, hyper-exponential and hypo-exponential delays (please recall that any distribution falls into one of these three cases, at least approximately), deriving the optimal duty cycle for each of them. For the simple case of exponential delays we are able to provide an exact solution. For the other two cases, we derive an approximated solution and the conditions under which this approximation introduces a small fixed error ε (which is always below 0.14) on the target probability p. Specifically, in the worst case, the approximated duty cycle introduces an error on the target probability p of about 0.1 (hyper-exponential case) and 0.14 (hypo-exponential case), while in the other cases the error is well below these thresholds.

2.2.1 Optimal duty cycling settings

In this section we discuss how to derive the optimal duty cycle Δ_{opt} such that the delay of a tagged message remains, with a certain probability p, under a target fixed threshold z or, in mathematical notation, $\Delta_{opt} = \min\{\Delta : P\{D_\Delta < z\} \geq p\}$. Since the delay increases with Δ, the latter is equivalent to finding the solution to the following equation:

$$\Delta_{opt} = \{\Delta : P\{D_\Delta < z\} = p\} \quad (1)$$

In order to find the solution to this Equation, the distribution of the delay D_Δ should be known. To this end, we can exploit the results presented in [10] (see D3.2 [2]) as follows. From [10] we know that when inter-contact times are exponential the detected inter-contact times are also exponential. Therefore, we can use existing models to derive the moments of the end-to-end delay. For example, we can use the model presented in [12] and reported in D3.1 [1]. Based on the first and second moments, we can use well-known distribution approximation techniques for deriving D_Δ. Specifically, if the resulting coefficient of variation is 1, we can approximate D_Δ with an exponential distribution. When it is greater than 1 with a hyper-exponential distribution, while when it is lower than 1 with a hypo-exponential distribution. Based on these approximations, we can find approximate solutions for the optimal duty cycling by solving the Equation above.

The case when D_Δ can be approximated with an exponential distribution is very easy. In this case, D_Δ is distributed exponentially with a rate $\lambda\Delta$ where λ is the rate of the delay when no duty cycling is used. Therefore, the optimal duty cycling can be easily obtained as

$\Delta = -\log(1 - p) / \lambda z$

For the hyper- and hypo-exponential cases, the solution is not as straightforward, because by substituting the expression of $P\{D_\Delta < z\}$ in Equation 1, the resulting expression cannot be inverted to find Δ. However, it is possible to use approximate expressions, that prove to be within a very reasonable margin of error (below 0.14 in all cases, see Appendix A).

Based on these analytical results, we can therefore set the tradeoff between energy saving and throughput (i.e., end-to-end delay) as needed. Figure 6 illustrates this tradeoff when the end-to-end delay can be approximated with an exponential distribution. Similar results can be obtained in the other cases.

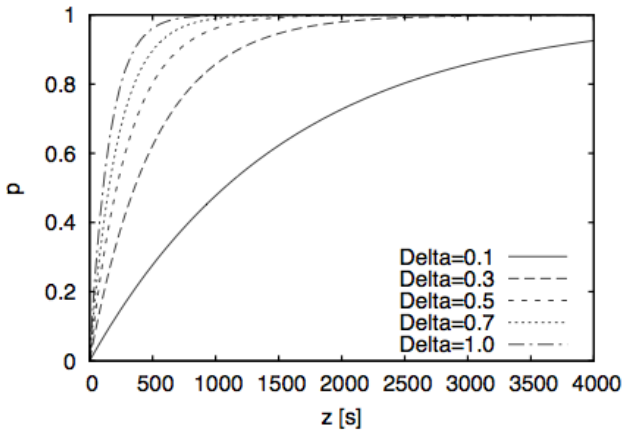

Figure 6. Example of delay/energy tradeoff.

Figure 6 shows, for various possible values of Δ, the delay that can be guaranteed (z) with a given probability (p). Clearly (i) for a given Δ, the higher the delay, the higher the probability with which it can be guaranteed and (ii) for a given probability p the delay that can be guaranteed with lower duty cycling is higher than the delay that can be guaranteed with lower duty cycling (i.e., curves move "from right to left when Δ increases). Even more importantly from a network configuration standpoint, if we want to guarantee a certain delay z with a certain probability p, we can identify the corresponding point (z,p) in the graph. Remember that Δ is the fraction of time during which nodes are active, and therefore we aim at the minimum possible value of Δ compatible with the point (z,p). Therefore, the optimal duty cycle is the one corresponding to curve with the minimum Δ that remains "on the left" of the point (z,p) in the graph.

Graphs like those in Figure 6 (that can be derived through our analytical results) thus provide simple tools for network operators to set the maximum energy saving that can be achieved, given a constraint in terms of extra capacity that they need to obtain from the opportunistic network, or, complementary, the maximum capacity that they can obtain if they have a constraint in terms of maximum energy consumption that users can tolerate.

2.2.2 Effect of duty cycling on inter-contact times in general mobility settings

Opportunistic networks are able to provide additional capacity for free to a cellular network exploiting localised communications between users. Unfortunately, these direct, opportunistic, communications consume significant energy, and this either discourages nodes to take part in the network or it rapidly depletes smartphones' batteries, leading, in both cases, to the failure of the opportunistic network. To address these problems, nodes are typically operated in duty cycling mode, by letting their WiFi (or Bluetooth) interfaces ON only for a fraction of time. With duty cycling, messages can be exchanged only when two nodes are in one-hop radio range *and* they're both in the active state of the duty cycle. So, power saving effectively reduces forwarding opportunities, because contacts are missed when at least one of the devices is in a low-energy state that does not allow it to detect the contact. Since some contacts are missed, the **measured intercontact times**, defined as the time interval between two consecutive *detected* encounters between the same pair of nodes, are, in general, larger and this clearly **affects the delay experienced by messages**.

In the following, we develop a mathematical model that enables us to derive, in **closed form**, the probability distribution of measured intercontact times under a generic distribution for the underlying intercontact times (i.e. intercontact times without duty cycling). In contrast, in our previous work reported in D3.2, we were able to derive a closed form only for exponential intercontact times, leaving out the very important case of Pareto intercontact times (which often emerge in traces of real human mobility). In addition, here we also consider the case in which contacts have a non-negligible duration and we discuss how this changes the effect of duty cycling on the measured contacts. *This work is instrumental for deriving closed form expressions for the end-to-end delay, and thus the additional capacity, that can be obtained by offloading part of the traffic in an opportunistic network. This is the final result of this activity, that we will present at the end of WP3.*

2.2.2.1 Preliminaries

We use duty cycling in a general sense here, meaning any power saving mechanism that hinders the possibility of a continuous scan of the devices in the neighbourhood. We assume that **nodes alternate between the ON and OFF states**. In the ON state, nodes are able to detect contacts with other devices. In the OFF state (which may correspond to a low-power state or simply to a state in which devices are switched off) contacts with other devices are missed. Using this generalisation, we are able to abstract from the specific wireless technology used for pairwise communications. The model we derive in the next sections assumes that the duration of ON and OFF states is fixed. However, it can be shown that the results hold, on average, also when their duration is stochastic. In the following, we assume that the duty cycle process and the contact process are independent and, considering a tagged node pair, we denote with τ the length of the time interval in which both nodes are ON, and with *T* the period of the duty cycle. Thus, *T*-τ corresponds to the duration of the OFF interval and $\Delta = \frac{\tau}{T}$ is the actual duty cycle (i.e., the percentage of time nodes are in the ON state).

Similarly to the related literature, we assume that, from the mobility standpoint, node pairs are independent. When two nodes meet, they remain in contact for a certain time, then they separate for a while, than they come into contact again. In many real scenarios, contact duration is orders of magnitude smaller than the time between contacts, thus it can often be neglected. For this reason, we first consider a case in which the contact duration of contacts is negligible with respect to the time interval between consecutive contacts. Next, we discuss the impact of contact duration and how it changes the modelling.

When contact duration is negligible, the **contact process** of each pair can be described **as a *renewal* process**. Focusing on a tagged node pair, we denote the time between the (*i-1)*-th and *i*-th contacts as S_i. By definition of renewal process, the intercontact times S_i for this pair of nodes are independent and identically distributed. When contact duration is not negligible, the **contact process** of each pair can be modelled **as an *alternating renewal* process**. In this case, the node pair alternates between the CONTACT state in which the two nodes are in radio range, and a state in which they are not. The time interval between the beginning and the end of the *i*-th contact is denoted as C_i. The time interval between the end of a contact and the beginning of the next one again corresponds to the intercontact time and it is denoted as S_i. Hence, the alternating renewal process corresponds to the independent sequence of independent random variables $\{C_i, S_i\}$, with $i \geq 1$.

2.2.2.2 Measured intercontact times when contact duration is negligible

We now discuss how the *measured* contact process depends on the contact process described above, starting from the negligible contact duration case. To this aim, we approximate **the measured contact process as a renewal process**. This approximation is very accurate when the PDF of the intercontact time *S* does not vary much inside an ON or OFF interval. Under this assumption, the measured intercontact time $\tilde{S}$ can be express as a random sum of the i.i.d. random variables *S*, according to the following definition.

Definition 1: *The measured intercontact time $\tilde{S}$ can be obtained as $\tilde{S} = \sum_{i=1}^{N} S_i$, where N is the random variable describing the number of contacts needed in order to detect the next one.*

Thus, in order to derive $\tilde{S}$, we need to characterize the distribution of *N*. In our previous work, we provided an exact model for *N*, which however could only be solved in closed form in the case of exponential intercontact times. Here we have derived two approximate models that yield the PMF of *N* in closed form. These models are denoted as Approx1 and Approx2 in the following. Model Approx2 is a refinement of model Approx1 in which further approximations are incorporated in order to derive a simpler closed form. Specifically, in model Approx2 *N* is approximated with a geometric distribution with parameter $\Delta = \frac{\tau}{T}$. The validation of the two models Approx1 and Approx2 against simulation results is shown in Figure 7 and Figure 8 for two representative cases of Pareto intercontact times (the Pareto exponent is hereafter denoted with α and the scale with *b*) and for two duty cycling process with parameters $\tau = 2, T = 10$ and $\tau = 8, T = 10$. It can be seen that Approx1 is always very accurate, while Approx2 suffers in some configurations in which $T \ll b$.

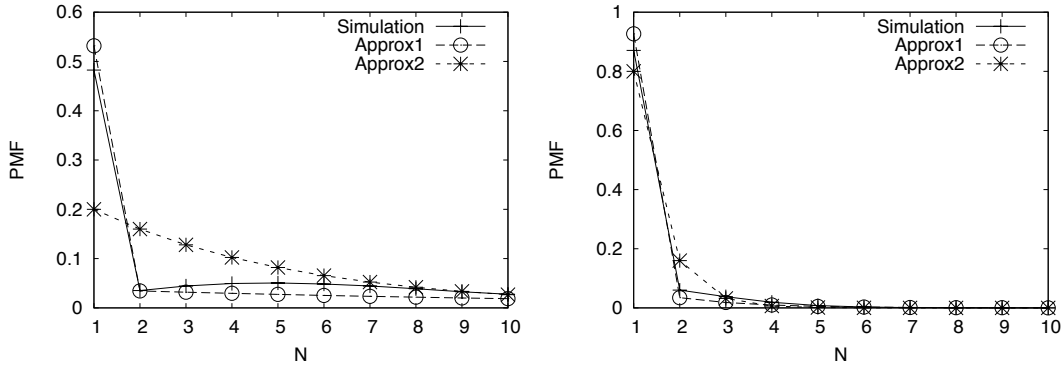

Figure 7: PMF of N, predictions of the approximate models VS empirical distribution when intercontact times are Pareto. On the left, $\alpha = 1.01, b = 1, \tau = 2, T = 10$. On the right, $\alpha = 1.01, b = 1, \tau = 8, T = 10$.

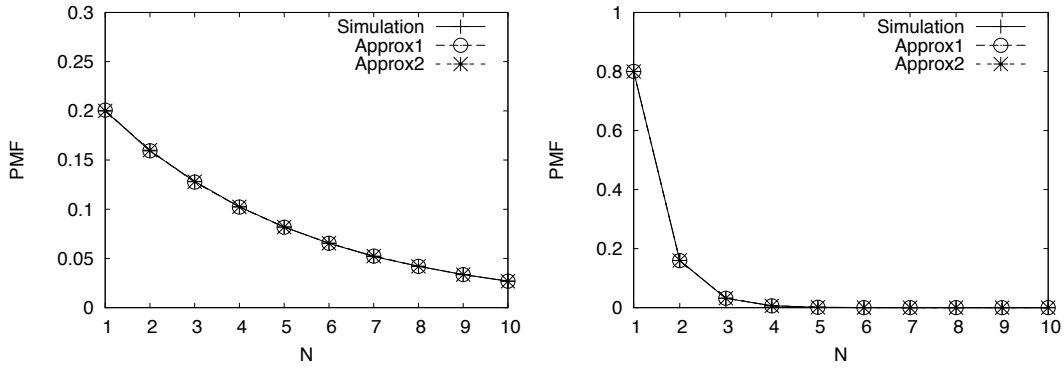

Figure 8: PMF of N, predictions of the approximate models VS empirical distribution when intercontact times are Pareto. On the left, $\alpha = 2.01, b = 101, \tau = 2, T = 10$. On the right, $\alpha = 2.01, b = 101, \tau = 8, T = 10$.

Using the above characterization of N, the moments and coefficient of variation of $\tilde{S}$ can now be derived, exploiting the well-known properties of random sums of random variables whereby the following holds:

$$\begin{aligned} E[\tilde{S}] &= E[N]E[S] \\ E[\tilde{S}^2] &= E[N^2]E[S]^2 + E[N]E[S^2] - E[N]E[S]^2 \\ cv_{\tilde{S}}^2 &= \frac{cv_S^2}{\mathbb{E}[N]} + cv_N^2 \end{aligned} \qquad (1)$$

The first two moments of $\tilde{S}$ can then be used to derive its approximate distribution using, e.g. the moment matching technique.

More generally, the above formulas can also be exploited to answer the following question: can duty cycling transform an intercontact time with low variability (e.g. hypo-exponential) into a high variability (e.g. Pareto) measured intercontact time? Is the opposite also possible? These questions are relevant to the delay experienced by messages, since measdeured intercontact times with high variability tend to also increase the variability of the delay (up to a point at which there is no convergent expectation anymore), and vice versa. Thus, defining the two classes of low variability (coefficient of variation *cv* smaller than 1) and high variability (coefficient of variation *cv* greater than 1) distributions, using the above formulas we have derived the conditions whereby the measured intercontact time does or does not inherit the class of its underlying intercontact time. Table 3 summarises these results.

Table 3: Summary of high/low variability regions depending on cv_S^2, g, p, where $\omega(g) = 3/2(g-1) + 1/2\sqrt{9-10g+g\hat{}2}$ and $\xi(g,p) = 1 - \frac{2g}{p} + \frac{2}{1-g+p}$

	$cv_{\tilde{S}}^2 \leq 1$	$cv_{\tilde{S}}^2 > 1$
$cv_S^2 > 3$	-	always
$1 < cv_S^2 <= 3$	$p > g \wedge cv_S^2 \geq \xi(g,p)$	$g > p$ $p > g \wedge cv_S^2 < \xi(g,p)$
$0 < cv_S^2 <= 1$	$p > g$ $g > p > \omega(g) \wedge cv_S^2 \leq \xi(g,p)$	$p < \omega(g)$ $g > p > \omega(g) \wedge cv_S^2 > \xi(g,p)$

Exploiting Defintion 1, we are also able to prove that when intercontact times are Pareto with exponent α, the CCDF of $\tilde{S}$ decays as a Pareto random variable with exponent α. Hence, if there are convergence problems with *S*, the same problems will show up also with the measured intercontact times, since they decays as a Pareto random variable with the same exponent.

2.2.2.3 Measured intercontact times when contact duration is not negligible

When contact duration is non negligible, the contact process can be modelled as an alternating renewal process. There are now two types of events: the beginning of a contact and the end of a contact. Since a contact is not anymore a punctual event but lasts for some time, contacts are detected more easily when their duration is not negligible. In fact, also contacts starting in an OFF interval can be detected, as long as they last until the next ON interval. Actually, a contact can cover also two or more ON intervals, and in this case pseudo-intercontact times are introduced. In fact, long contacts are split into smaller measured contacts separated by intervals of size $T - \tau$ in which no communication is possible (because the two nodes, despite theoretically in each other's radio range, are in the OFF interval of their duty cycle).

The contact process "filtered" by the duty cycle results in a **measured contact process characterised by a sequence of measured contacts $\tilde{C}_i$ that alternates with measured intercontact times $\tilde{S}_i$**. The dependence of $\tilde{C}_i$ and $\tilde{S}_i$ on *C* and *S* is more complicated than in the previous case. Specifically, we prove that the relationships below hold.

$$\tilde{C} = \begin{cases} C_{short} & p_1 \\ Z^{ON} & p_2 \\ \tau & 1 - p_1 - p_2 \end{cases} \qquad \tilde{S} = \begin{cases} R + \sum_{i=1}^{N} S_i + \sum_{i=2}^{N} C_i^{miss} + R & p_3 \\ T - \tau & 1 - p_3 \end{cases} \qquad (2)$$

Thus, $\tilde{C}_i$ and $\tilde{S}_i$ are mixture distributions. Probabilities p_1, p_2, and p_3 are a function of *T*, τ, and *C*. C^{short} represents the length of a contact when it is fully contained in an ON interval, while C^{miss} that of a contact fully contained in an OFF interval. Z^{ON} is a uniformly distributed random variable in [0, τ]. R takes values from a uniform distribution in [0, T-τ] with probability $1 - \Delta$, while it is equal to zero otherwise. The

distribution of *N* can be derived following the same line of reasoning as that used for the negligible contact duration case.

It is interesting to note that now $\tilde{S}_i$ is a bimodal distribution: there is an accumulation at $T - \tau$ due to the pseudo-intercontact times introduced by a long contact spanning more than one ON interval, while all other samples are due to the actual contribution of missed contacts. As far as measured contact duration is concerned, it can now last at most τ seconds, hence reducing greatly the amount of data that can be exchanged during a contact. Summarising our results for the non-negligible contact duration case, we observe that the **duty cycling** process **affects both the capacity** of the opportunistic network (because contacts are shorter, hence less data can be exchanged) **and the QoS** experienced by the users (because the time between consecutive contact opportunities is longer and the delay increases).

3 Capacity analysis: Assessing the capacity of LTE

To achieve high throughput performance, in addition to an advanced physical layer design LTE exploits a combination of sophisticated radio resource management functionalities, such as Channel Quality Indicator (CQI) reporting, link rate adaptation through Adaptive Modulation and Coding (AMC), and Hybrid Automatic Retransmission Request (HARQ) [1]. Specifically, a base station (eNB) can simultaneously serve multiple users on orthogonal subcarriers that are grouped into frequency resource blocks (RBs). Then, each user (UE) periodically measures channel state information that is fed back to the eNB in terms of CQI reports. Typically, only aggregate CQI values are reported to reduce channel feedback information. CQI measurements are used by eNBs for scheduling and link rate adaptation on the downlink [21]. For instance, the modulation and coding scheme (MCS) is typically selected in order to maximise the data rate to the scheduled UE subject to a constraint on the error probability. How CQI values should be computed by the UE using channel state information (e.g., SINR measurements) is implementation dependent. Unfortunately, past research has shown that it is difficult to derive accurate link performance predictors under realistic channel assumptions. Automatic retransmission protocols with channel coding (HARQ) are also exploited to mitigate errors at the physical layer. More precisely, HARQ procedures use the classical stop-and-wait algorithm, in which the eNB decides to perform a retransmission based on the exchange of ACK/NACK messages with the UE. Then, UEs try to decode the packet by combining the retransmitted copies.

Since user, cell and radio link throughputs are among the most important performance indicators that the operators adopt to asses the QoS in an LTE system [71], an extensive literature exists that investigates LTE throughput performance based on analytical models [25][71][42], simulation tools [21] or field tests [27][19]. However, it is evident that a complex interplay exists among the various mechanisms that operate at the MAC layer to improve communication reliability and to increase data rates. This makes accurate LTE throughput analysis notably difficult. Thus, most studies limit the analysis only to the radio link throughput or consider single MAC functions in isolation [33]. Furthermore, simplified error models are typically considered that only allow deriving upper bounds for the LTE throughput [71].

The contribution of this section is twofold. The first contribution is the development an initial model of the **user-perceived MAC-layer throughput on the downlink channel**. Our model is valid for homogeneous cells [62] and Rayleigh-distributed fading. Our model simultaneously caters for CQI feedback schemes that use spectral efficiency to generate CQI, as well as AMC and HARQ protocols. Furthermore, we include in the analysis an accurate link layer abstraction model that uses the Mean Mutual Information per coded Bit (MMIB) metric to derive the physical error probability [14]. The throughput estimates of our model are accurate, as validated using the ns-3 simulator extended with the LENA module for LTE [80]. The second contribution of this section is a proposal for increasing the LTE downlink capacity by enabling **automatic tuning of AMC parameters**. Specifically, we have developed a new AMC scheme that exploits a reinforcement learning (RL) algorithm to adjust at run-time the MCS selection rules based on the knowledge of the effect of previous AMC decisions. The salient features of our proposed solution are: i) the low-dimensional space that the learner has to explore, and ii) the use of direct link throughput measurements to guide the decision process. Simulation results obtained using ns3 demonstrate the robustness of our AMC scheme that is capable of discovering the best MCS even if the CQI feedback provides a poor prediction of the channel performance. Therefore, this result provides a possible capacity enhancing tool for the LTE part of an offloaded network, to be used in addition (and orthogonally to) offloading solutions.

3.1 LTE MAC-layer Throughput

3.1.1 LTE MAC Model

We now briefly describe relevant details of the LTE downlink, with special attention to frame structure, CQI feedback mechanisms and HARQ protocols. We also introduce the system model and notation, and we

discuss the main assumptions that underlay our analysis. Complete details are available in Appendix A, which include a reprint of [16].

In LTE, each DL frame is 10 ms long and it consists of ten transmission time intervals (TTIs). Furthermore, each TTI consists of two 0.5 ms slots. Each slot contains seven OFDM symbols. In the frequency domain, the system bandwidth, W, is divided into several orthogonal subcarriers. Each subcarrier has a bandwidth of 15 kHz. A set of twelve consecutive subcarriers over the duration of one slot is called a physical Resource Block (RB). Let q denote the total number of RBs available over the system bandwidth.

Since the RB bandwidth is only 180 kHz, it is reasonable to assume that the channel response is frequency-flat across all the twelve subcarriers of the RB[1]. Then, let us denote with $\gamma_{i,k}$ the SNR of the i^{th} RB of the k^{th} UE. Clearly, the statistics of the SNR depend on the channel model and the multi-antenna diversity mode of operation. As commonly adopted in other LTE models, e.g. [25], in this study we assume that the fading from the eNB to the UEs is *Rayleigh distributed*. This implies that *the SNR of each RB is an exponential* random variable (RV) [33]. Furthermore, we also assume an *homogeneous cell model* [62], i.e. the SNR is independent for different users and RBs. This also means that the SNRs of all RBs are *uncorrelated* in frequency and space, and $\gamma_{i,k}$ can be regarded as independent and identically distributed (i.i.d.) RV.

Popular methods (e.g., EESM and MIESM) that are typically used in LTE to compute CQI values rely on the concept of *effective* SNR. Basically, the UEs map the SNRs of multiple subcarriers/RBs into a single value by applying complex non-linear transformations. Then, the effective SINR is used to estimate the BLER experienced by a user and to determine the appropriate MCS, i.e. the MCS that allows the UE to decode the transport block with an error rate probability not exceeding 10%. However, the statistics of the effective SNR generated by EESM and MIESM techniques are not known in closed-form. Thus, they must be approximated or computed numerically, which makes performance analysis difficult [63][25]. An alternative approach proposed in [75] implement AMC capabilities is based on the *spectral efficiency*. Specifically, let us denote the with $\eta_{i,k}$ the spectral efficiency of the i^{th} RB of the k^{th} UE. Then, it holds that [60]

$$\eta_{i,k} = \log_2\left(1+\frac{\gamma_{i,k}}{\Gamma}\right) \tag{4.1}$$

where $\Gamma = -\ln(5/\beta)/1.5$ and β is BLER upper bound. Now a static mapping can be determined between the spectral efficiency and the CQI index, as well as between the CQI index and the MCS value [75]. More formally, let us denote with $C_{i,k}$ the CQI index for the i^{th} RB of the k^{th} UE. Typically the value of CQI can range between 1 and L. Then, $C_{i,k} = j$ ($j = 1,\ldots,L$) if $S_j \le \eta_{i,k} \le S_{j+1}$, with $S_0 = 0$ and $S_L = \infty$. In other words the CQI value is a quantised version of the spectral efficiency[2]. Furthermore, in the 3GPP-LTE standard the available MCS indexes are 32 but a 4-bit CQI allows selecting only 15 MCS. Thus, in practical LTE systems only a subset of available MCS is typically used.}. Closely related to the MCS selection is also the transport block (TB) size determination. More precisely, let n_k the number of RBs allocated to the k^{th} UE during a frame. Then, the number B of bits that can be delivered in those RBs, which is called transport block, is a function of the MCS index[3]. Furthermore, if $B > Z$ (with $Z = 6144$ bits in 3GPP-LTE) the transport block is *segmented* into a number C of *code blocks* (CBs) that are independently encoded. Note

[1] This assumption will not hold for highly dispersive channels with long delay spread.

[2] Note that in the 3GPP-LTE standard, L=15 and the S_j thresholds are specified in Table 7.2.3-1 of [75].

[3] See Table 7.1.7.2.1-1 of [75] for the static mapping between TB size, MCS and number of RBs allocated to the UE.

that the CB size highly impacts the actual BLER performance for a given MCS [14]. Figure 9 exemplifies the transport block segmentation.

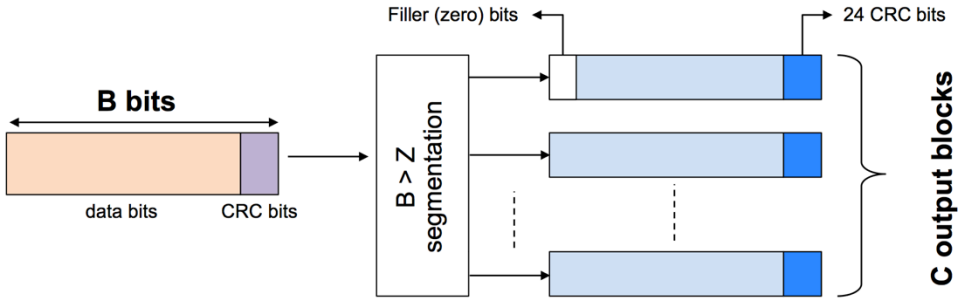

Figure 9: Transport block segmentation.

Regarding the HARQ protocol, LTE employs two types of HARQ schemes. In HARQ type-I, each encoded data frame is retransmitted until the frame passes the CRC test or the maximum number of retransmissions is reached. Erroneous frames are simply discarded. In contrast, in HARQ type-II, each transmission contains incremental redundancy (IR) about the data frame. Thus, consecutive transmissions can be combined at the receiver to improve error correction. Although our model is valid for all HARQ types, in the following we only consider HARQ type-II that is the most widely used in LTE. Note that in LTE systems *retransmissions typically use the same MCS index as the initial transmission*. It is also important to point out that the transmission of HARQ feedbacks (i.e. ACK/NACK messages) is not instantaneous but each received packet experiences a processing delay. According to the LTE standard, the processing delay at the receiver is about 3ms. Thus, assuming the same delay to process data transmissions and ACK/NACK messages, the HARQ round trip time, say τ_{ARQ}, is 7 TTIs, as shown in Figure 10. For this reason, an eNB must support up to 8 parallel HARQ processes for each UE to enable *uninterrupted* communications. In this way, an eNB can continue to transmit new TBs while the UEs are decoding already received TBs.

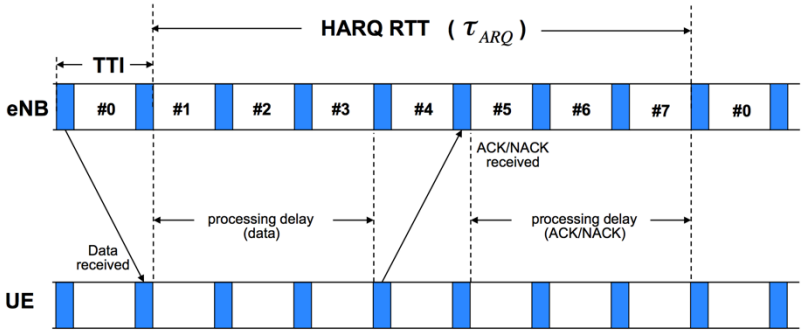

Figure 10: HARQ processes and timing in FDD-LTE DL.

3.1.2 Throughput Analysis

In this section we develop the mathematical model of the MAC-layer throughput for the LTE downlink. First of all, let us assume that n UEs are randomly distributed in the cell, and let d_k be the distance of the k^{th} UE from the eNB. As discussed in Section 3.1.1 we develop our analysis in the case $\gamma_{i,k} \sim Exp(\lambda_k)$, where the rate parameter λ_k of the exponential distribution depends on the UE position. Under this assumption the statistics of the spectral efficiency for each RB can be expressed in a closed-form as given by the following Theorem (unless otherwise stated, all proofs are reported in [18]).

Theorem 4.1: If $\gamma_{i,k} \sim Exp(\lambda_k)$ then the cumulative distribution function (CDF) of the spectral efficiency $\eta_{i,k}$ in equation~(4.1) is computed as:

$$F_\eta(x;i,k) = \begin{cases} 1 - e^{-\lambda_k \Gamma\left(2^x - 1\right)} & \text{if} \quad x \geq 0 \\ 0 & \text{if} \quad x < 0 \end{cases} \tag{4.2}$$

LTE specifies different types of CQI reporting: *wideband* and *subband*. Specifically, the wideband CQI represents the SNR observed by the UE over the whole channel bandwidth, while the subband CQI represents the SNR observed by the UE over a collection of adjacent RBs. Note that a vector of CQI values should be transmitted to the eNB when using the latter feedback scheme. Thus, the subband-level feedback scheme ensures a finer reporting granularity but it also generates a higher overhead. In this study, we focus on the wideband feedback scheme and we assume that the CQI reported by the k^{th} UE, say $\hat{C}_k$ is the *arithmetic mean* of the CQI values computed over all RBs[4]. Then, we use the spectral efficiency to generate the CQI values from the SNR measures of all RBs. The statistics of the wideband CQI are mathematically derived below.

Claim 4.1: The probability mass function (PMF) of the CQI value for the i^{th} RB assigned to the k^{th} UE is given by

$$g_{i,k}[j] = F_\eta(S_{j+1};i,k) - F_\eta(S_j;i,k) \tag{4.3}$$

Claim 4.2: The probability mass function (PMF) of the CQI value for the i^{th} RB assigned to the k^{th} UE is given by

$$g_k[j] = \sum_{l=qj}^{q(j+1)-1} g_{i,k}^{(q)}[j] \ , \tag{4.4}$$

where $g_{i,k}^{(q)}[j]$ is the q-fold convolution of $g_{i,k}[j]$.

As described in Section 3.1.1 a static mapping is typically established between the CQI value received at the eNB and the MCS for the downlink transmissions. For simplicity of notation we indicate with $m(j)$ the MCS that the eNB uses when the wideband CQI reported by a UE is equal to j.

Before proceeding with the throughput analysis, we need to introduce the physical layer error model. In this study we adopt the general approach initially proposed in [75] to accurately approximate the BLER curves of OFDMA-based wireless systems, and later specialised for the LTE case in [47]. Specifically, we assume that the *mutual information per coded bit* (MIB) of MCS m, as defined in [30], can be accurately approximated by a combination of Bessel functions of the SNR γ as follows

$$I_m(\gamma) \approx \sum_{h=1}^{H} \alpha_h J\left(\psi_h \sqrt{\gamma}\right) \ , \tag{4.5}$$

where H, α_h and ψ_h parameters are empirically calibrated as a function of the MCS index. Subsequently, the *mean* MIB (MMIB) value for each UE is computed by averaging the corresponding mutual information of all RBs allocated to that UE. Specifically, let $\Omega(k)$ be the set of RBs that are allocated to the k^{th} UE by the scheduler. Then, the MMIB value over the vector of SNR values for each RB assigned to the k^{th} UE when m is the adopted MCS is simply given by

[4] Note that an alternative solution would be to report the worst CQI value over all (or a subset of) RBs as in [71].

$$\hat{I}_{m,k} = \frac{1}{\omega(k)} \sum_{i \in \Omega(k)} I_m(\gamma_{i,k}) \ , \tag{4.6}$$

where $\omega(k)$ is the cardinality of the $\Omega(k)$. The non-linear nature of (4.5) makes an exact analysis difficult. Thus, previous studies limit the computational complexity of deriving MMIB values in multi-user scenarios by considering a quantised version of the $I_m(\gamma)$ function (4.5) in order to *discretise the MIB metric [47]*. More precisely, let us define a set $V_m = \{\mu_m[0], \mu_m[1], \ldots, \mu_m[v_m]\}$ for each MCS m such that

$$\mu_m[v] = I_m(Q_{m,v}) \ , \tag{4.7}$$

where $(Q_{m,v+1} - Q_{m,v}) = \delta\gamma$ is the quantisation step size, and $Q_{m,0}$ is the minimum usable SNR for MCS m. Now, let us denote with $H_{i,m,k}$ the discrete MIB value for the i^{th} RB scheduled to the k^{th} UE when \$m\$ is the adopted MCS. Similarly to the approach adopted for CQI mapping, we assume that $H_{i,m,k} = \mu_m[v]$ ($v = 0, \ldots, v_m$) if $Q_{m,v} \le \gamma_{i,k} \le Q_{m,v+1}$. In other words the discrete MIB value is associated to a *range* of SNRs. It is straightforward to derive the statistics of the discretised MIB metric as follows.

Claim 4.3: The probability mass function (PMF) of the CQI value for the i^{th} RB assigned to the k^{th} UE is given by

$$h_{i,m,k}[v] = \int_{Q_{m,v}}^{Q_{m,v+1}} f_\gamma(x; i, k) dx \ , \tag{4.8}$$

where $h_{i,m,k}[v] = \Pr\{H_{i,m,k} = \mu_m[v]\}$.

Similarly, we introduce a discrete MMIB metric, say $\hat{H}_{m,k}$, computed over the set of RBs allocated to the k^{th} UE when m is the adopted MCS. In particular, $\hat{H}_{m,k}$ can be obtained as the mean of the $H_{i,m,k}$ values over the set $\Omega(k)$. Thus, the statistics of the discretised MMIB value are derived using the same technique of Claim 4.2.

Claim 4.4: In an homogeneous cell the PMF of $\hat{H}_{m,k}$ is given by

$$h_{m,k}[v] \approx \sum_{l \in \Phi_v} h_{i,m,k}^{(\varpi(k))}[l] \ , \tag{4.9}$$

where $h_{i,m,k}^{(\varpi(k))}[l]$ is the $\omega(k)$-fold convolution of $h_{i,m,k}[l]$. The definition of the Φ_v set is quite involved and is given in [18].

Once the MMIB value is given, a direct MMIB to BLER mapping can be used to obtain the *code block error rate*, without necessarily defining an effective SINR. Following the approach proposed in [75], the empirical BLER curve for MCS m can be approximated with a Gaussian cumulative model as follows

$$CBLER_m(y, e) = \frac{1}{2}\left[1 - erf\left(\frac{y - b_e}{c_e}\right)\right] \ , \tag{4.10}$$

where y is the MMIB value, while b_e and c_e are parameters used to fit the Gaussian distribution to the empirical BLER curve[5]. These parameters depend on the Effective Code Rate (ECR), i.e. the ratio between

[5] Empirical BLER curves can be obtained through field measurements or detailed link-level simulations.

the number of downlink information bits (including CRC bits) and the number of coded bits. Intuitively, the ECR value is a result of the selected TB size, MCS, and $\Omega(k)$. Then, the overall error probability for a transport block transmitted as a combination of C code blocks, each one associated with a MMIB and ECR value, can be computed as

$$TBLER_m(y,e) = 1 - \frac{1}{2}\prod_{i=1}^{C}\left(1 - CBLER_m(y_i, e_i)\right) \quad . \tag{4.11}$$

However equation (4.11) does not take into account the impact of an IR-HARQ mechanism that combines retransmissions to improve error correction. To generalise equation (4.11) for a system with incremental redundancy we adopt the same approach as in [78]. In particular, we introduce an *equivalent* MMIB metric as the average of the mutual information values per HARQ block received on the total number of retransmissions. More precisely, let us assume that the original transport block has been retransmitted r times. Then, let $\left(\hat{I}_{m,k}^{(0)}, \hat{I}_{m,k}^{(1)}, \ldots, \hat{I}_{m,k}^{(r)}\right)$ be the vector of MMIB values for each of these transmissions. The equivalent MMIB for the r^{th} retransmission can be computed as follows

$$\hat{I}_{m,k,r} = \frac{1}{r+1}\sum_{i=0}^{r} I_{m,k}^{(i)} \quad . \tag{4.12}$$

Then, the PMF of the equivalent MMIB value for the r^{th} retransmission is $h_{m,k}^{(r)}[v] = \Pr\left\{\hat{I}_{m,k,r} = \mu_m[v]\right\}$. This PMF can be obtained using the same technique as in Claim 4.4 and it is not reported here for the sake of brevity. Similarly, we compute the effective ECR after r retransmissions, say $e^{(r)}$, by dividing the number of information bit of the original transmission with the sum of the number of coded bits of each retransmission. Finally, by applying the law of total probability the *average* TB error probability at the r^{th} retransmission for the k^{th} UE when m is the adopted MCS can be computed as

$$P_e(m,k,r) = \sum_{v=0}^{v_m} TBLER_m\left(\mu_m[v], e^{(r)}\right) \cdot h_{m,k}^{(r)}[v]dy \quad . \tag{4.13}$$

To conclude the MAC-layer throughput analysis we have to model the operations of the packet scheduler at the eNB, which is responsible for allocating RBs to UEs every TTI. In this study we consider the Round Robin (RR) scheduler that works by dividing the available resources among the UEs in a fair manner. In particular the scheduler allocates a set of consecutive resource blocks, called resource block groups (RBGs), whose size P depends on the system bandwidth [76]. Consequently the number of RB assigned to each UE is simply given by

$$n_k = \max\left\{P, \left\lfloor \frac{q}{n} \right\rfloor\right\} \quad . \tag{4.14}$$

The scheduler also interacts closely with the HARQ protocol. Typically, a non-adaptive HARQ mechanism is implemented in LTE systems, which implies that the scheduler should maintain the same RBG and MCS configuration of the original TB when scheduling the retransmissions.

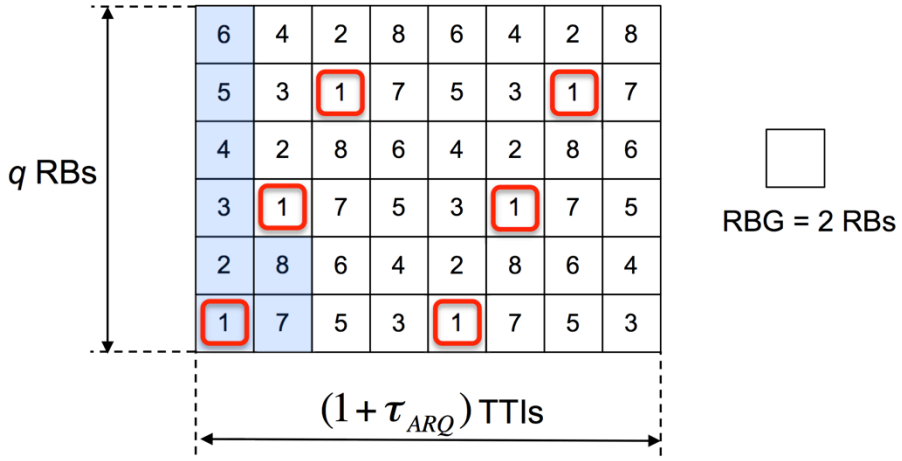

Figure 11: RR operations with q=12, P=2 and n=8.

As discussed in Section 3.1.1, the scheduler can control up to 8 HARQ processes for generating new packets and managing the retransmissions. However, the actual number of HARQ processes that are activated by the scheduler is bounded by the number of times the same UE is scheduled during the HARQ period τ_{ARQ}. In turn, this depends on the total number of available resources during a time window of duration τ_{ARQ}, the RGB size and the number of UEs. To illustrate this dependency in Figure 11 we exemplify the scheduling decisions that are cyclically performed by the RR scheduler during an HARQ period with q=12, P=2 and n=8. As shown in the figure, each UE is scheduled six times. In general, the average number of times each UE is scheduled in $\left(1+\tau_{ARQ}\right)$ TTIs is simply given by

$$n_{RR} = \frac{q\left(1+\tau_{ARQ}\right)}{nP} \quad . \tag{4.15}$$

However, not all the transmission opportunities allocated by the eNB to an UE result into a successful transmission. In particular let us denote with $P_s\left(m,k,r\right)$ the probability that the k^{th} UE correctly decodes a TB after r retransmissions when m is the adopted MCS. It holds that

$$P_s\left(m,k,r\right) = \left[\prod_{i=0}^{r-1} P_e\left(m,k,i\right)\right] \times \left[1 - P_e\left(m,k,r\right)\right] \quad . \tag{4.16}$$

The above equation is easily derived by observing that the r^{th} retransmission is a success only if the previous $(r-1)$ transmissions were TBs received erroneous and the r^{th} transmission is correctly decoded.

To perform the throughput analysis we observe the system behaviour only during the HARQ period because the HARQ processes that define the occupancy pattern of the channel (i.e., new transmissions and retransmissions) regenerates after each such period. Then, it follows that the MAC-layer throughput for the k^{th} UE is

$$\rho_k = \frac{n_k^{succ} \cdot E[TB]}{1+\tau_{ARQ}} \quad , \tag{4.17}$$

where n_k^{succ} is the average number of HARQ processes for the k^{th} UE which terminate with a success in the HARQ period, while $E[TB]$ is the average number of information bits that are delivered with a successful transmission. To compute n_k^{succ} we can note that in an HARQ period there are at most n_{RR} active HARQ

processes. Since RR equally distributes transmission opportunities to each UE, then all UEs have the same number of active HARQ processes. Note that only a fraction $P_s(k)$ of the n_{RR} HARQ process that are on average active in each HARQ period terminates with a successful transmissions. Hereafter we derive exact expressions for the unknowns in (4.17)

Theorem 4.2: By assuming an homogenous cell with Rayleigh-distributed fading, and a RR scheduling policy

$$E[TB] = \sum_{j=0}^{L} TBS\left(m(j), n_k\right) g_k[j] \quad , \tag{4.18a}$$

$$n_k^{succ} = n_{RR} \sum_{j=0}^{L} \left[\sum_{r=0}^{r_{\max}} \frac{\left[P_s(m,k,r)\right]^2}{\left[1 - P_e(m,k,r)\right]} \right] g_k[j] \quad . \tag{4.18b}$$

3.1.3 Validation

In this section we show a preliminary validation or our modelling approach using the ns-3 simulator with the LENA module for LTE. The main simulation parameters are reported in Table 4. Specifically, we consider an *Urban Macro* scenario, in which path loss and shadowing are modelled according to the COST231-Hata model [24], which is widely accepted in the 3GPP community. The fading is Rayleigh distributed. To limit the computation complexity of the simulator pre-calculated fading traces are included in the LTE model. Given the downlink system bandwidth (see Table 4) a RBG comprises two RBs [76], i.e., P=2. Regarding the network topology, we considered a single cell with a varying number of static UEs, chosen in the range [10,50]. One *tagged* UE is positioned at a fixed distance from the eNB, while the other UEs are randomly deployed within the cell. The cell radius is 2 Km. Note that, in our settings a maximum number of 96 (i.e., 8q/P) unique UEs can be scheduled within an HARQ period. Indeed, if n>96 the RR period is longer than the HARQ period. All results presented in the following graphs are averaged over multiple simulation runs with different fading traces and topology layouts. Confidence intervals are generally very tight and are not shown in the figures when below 2%. Each simulation run lasts 300 seconds.

Table 4: Simulation parameters

Parameter	Value
Carrier frequency (GHz)	2.14
DB bandwidth (MHz)	5
q	25
eNB TX Power (dBm)	43
CQI Processing time (TTI)	2
CQI transmission delai (TTI)	4
Antenna scheme	SISO
PDCCH & PCFICH (control ch.)	3 OFDM symbols
PDSCH (data ch.)	11 OFDM symbols
n	[10,50]

The accuracy of our modelling approach is validated considering the throughput of the tagged UEs. Specifically, Figure 12 shows the MAC throughput obtained by the tagged UE by varying its distance from the eNB and the number of competing UEs. The shown results have been obtained by setting the maximum

number of transmission equal to one. In other words, transport blocks that are not correctly received are discarded without being retransmitted. The plots clearly indicate that our analysis is very accurate in all the considered settings. At the time of the writing of this deliverable an extensive validation is in progress.

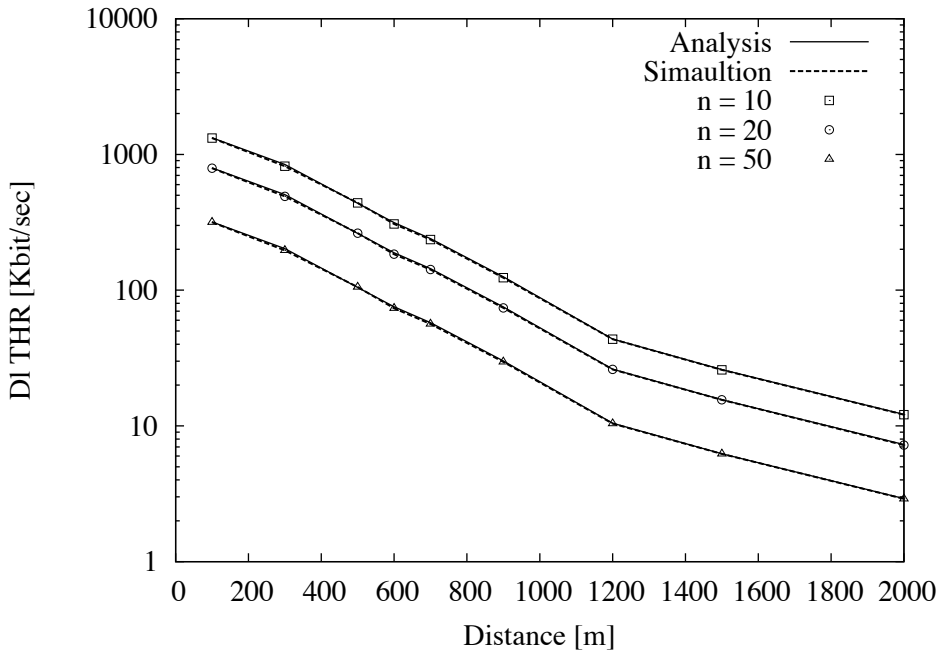

Figure 12: Comparison of simulation and analytical results versus the distance of the tagged UE.

3.2 Robust AMC in LTE using Reinforcement Learning

Adaptive Modulation and Coding (AMC) in LTE networks is commonly employed to improve system throughput by ensuring more reliable transmissions. Traditional AMC schemes rely on CQI feedbacks that are periodically reported by UEs to their eNBs. AMC schemes typically exploit static mappings between these link quality metrics and the BLER performance of each MCS to select the best MCS (in terms of link throughput). In other words, for each MCS a range of LQM values is associated via a look-up table, over which that MCS maximises link throughput. Either link-level simulations or mathematical models can be used to generate such static BLER curves under a specific channel model. Unfortunately, past research has shown that it is difficult to derive accurate link performance predictors under realistic channel assumptions [28][11][14][48]. Furthermore, a simulation- based approach to derive the mapping between LQM values and BLER performance is not scalable since it is not feasible to exhaustively analyse all possible channel types or several possible sets of parameters [40]. The second main problem with table-based AMC solutions is that a delay of several transmission time intervals (TTIs) may exist between the time when a CQI report is generated and the time when that CQI feedback is used for channel adaptation. This is due to process- ing times but also to the need of increasing reporting frequency to reduce signalling overheads. This mismatch between the current channel state and its CQI representation, known as CQI ageing, can negatively affect the efficiency of AMC decisions [39][6].

To deal with the above issue we have propose a new flexible AMC framework, called RL-AMC, that autonomously and at run-time decides upon the best MCS (in terms of maximum link-layer throughput) based on the knowledge of the outcomes of previous AMC decisions. To this end we exploit reinforcement learning techniques to allow each eNB to update its MCS selection rules taking into account past observations of achieved link-layer throughputs. In this section we outline the general design principles of the proposed solution and we show the main results of the performance evaluation performed in ns3 (complete details are available in Appendix A, which is a reprint of [15]). Overall, our RL-based scheme not only improve the LTE system throughput compared to other schemes that use static mappings between

SINR and MCS, but it is also capable of discovering the best MCS even if the CQI feedback provides a poor prediction of the channel performance.

3.2.1 Protocol design

In order to apply the Q-learning approach to the MCS selection problem it is necessary to define: i) the state space of the problem, ii) the feedbacks that the decision agent receives from the LTE network, and iii) the admissible actions for the agent with the action selection strategy. In our RL-based AMC framework, the problem state consists of CQI feedbacks and their evolution trends. The reward is the instantaneous link throughput obtained by a user after each transmission. Finally, an action is the selection of a correction factor to be applied to each CQI feedback to identify the best MCS under the current channel conditions.

Intuitively, a straightforward approach to define the state of the MCS selection problem would be to use the SINR values of received segments of data as state variables, as in [17]. However, the SINR is a continuous variable and it should be discretised to be compatible with a discrete MDP formulation. The main drawback is that a fine discretisation leads to a large-dimensional state space, which increases convergence and exploration times. To avoid this problem, we directly use CQI-based metrics for the state representation. Specifically, we adopt a two-dimensional space to characterise the LTE communication channel. The first state variable represents the CQI value (called CQI^m) that the UE should select using the internal look-up table that associates BLER and MCS and received SINR. The second state variable represents the ΔCQI^m value, which is defined as the difference between the last two consecutive CQI^m estimates. In other words, ΔCQI^m provides a rough indication of the trend in channel quality evolution. For instance, $\Delta CQI^m < 0$ implies that the channel quality is temporarily degrading.

Since the objective of the MCS selection procedure should be to maximise the link throughput it is a natural choice to define the reward function as the instantaneous link-layer throughput achieved when taking action a in a given state. Thus, a key aspect in the design of the Q-learning algorithm is represented by the set A of admissible actions. In our learning model we assume that an action consists of applying a correction factor to the CQI value that is initially estimated by means of the internal look-up table. As discussed above, the mapping relationship between SINR values and MCS may be inaccurate and the correction factor allows the agent to identify the best modulation and coding scheme (in the sense of maximising the link throughput) for the given channel conditions. For instance, it may happen that the SINR-to-MCS mapping is too conservative for the current channel conditions and an MCS with a higher data rate can be used without violating the target BLER requirement. In this case the correction factor should be positive. Furthermore, a correction factor is also needed to compensate eventual errors due to CQI feedback delay. More formally, we assume that an action taken by the AMC decision agent at time t is one possible choice of an integer number in the set $(-k,\ldots,-2,-1,0,1,2,\ldots,k)$, that we denote as a_t in the following. This index is added to the original CQI^m value to compute the CQI to be sent to the eNB. We argue that $\Delta CQI^m < 0$ we should prefer conservative MCS selections (and thus use values of at lower than 0) because the channel trend is negative, while if $\Delta CQI^m \geq 0$ we can try to use MCSs offering higher data rates (and thus positive values for at). Thus, the set of admissible actions is different whether the channel-quality trend is negative or non-negative. Before proceeding it is useful to point out that the choice of the k value determines how aggressively we want to explore the problem state space. In general, the selection of the k value could take into account the CQI difference statistics, i.e., to what extent a current CQI may be different from the reported CQI after a feedback delay [48]. Finally, a very important learning procedure is the action selection rule, i.e., the policy used to decide which specific action to select in the set of admissible actions. There is a clear trade-off between exploitation (i.e., to select the action with the highest Q-value for the current channel state) and exploration (i.e., to select an action randomly). In our solution we adopt a *softmax action-selection rule* [69] that assigns a probability to each action by applying a Boltzmann-like function to the current Q-value for that action (see Appendix A for more details).

3.2.2 Performance evaluation

The simulation setup is the same as the one already described in Section 3.1.3. Regarding the simulation scenarios, we consider two cases. In the first one, ten UEs are randomly deployed in the cell and they are static. Then an additional tagged user is moving with pedestrian speed from the centre of the cell to its boundaries. However, independently of the UE position the *CQI feedback is constant*. Then, Figure 13 shows a comparison of the throughput achieved by the tagged user with and without reinforcement learning. This is obviously a limiting case which is analysed to assess the robustness of our RL-AMC scheme even when CQI provides a very poor prediction of channel performance. As expected with fixed MCS the user throughput is constant when the MCS is over provisioned, while it rapidly goes to zero after a critical distance. On the contrary, our RL-AMC is able to discover the correction factor that should be applied to the initial CQI to force the selection of a more efficient MCS. In addition, the throughput performance of RL-AMC is almost independent of the initial CQI value. Note that in this case RL-AMC must explore the full range of CQI values and we set the k parameter for action selection equal to 15.

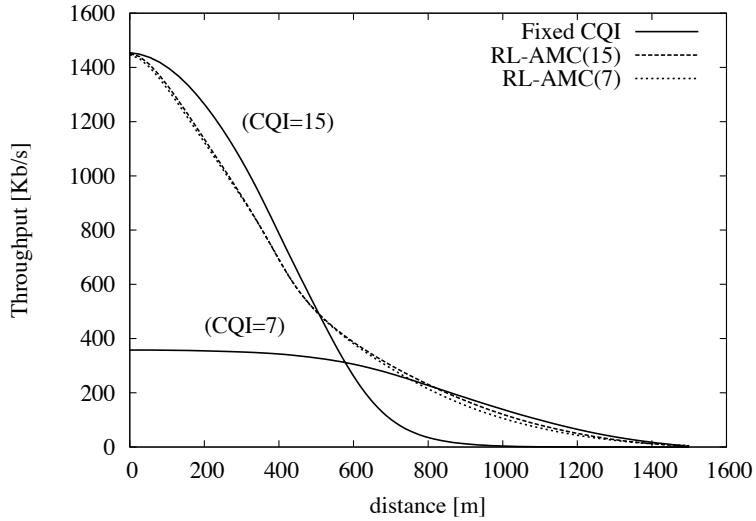

Figure 13: Average throughput as a function of the distance of the tagged user from the eNB in a pedestrian scenario.

In the second scenario, each UE implements the SINR to CQI mapping described in [76][75], called SE-AMC because it uses spectral efficiency to estimate the transport block error rates. Then, Figure 14 shows a comparison of the throughput achieved by the tagged user with both SE-AMC and RL-AMC schemes at different distances of the tagged UE from the eNB. We can observe that the MCS selection in SE-AMC is too conservative and this results in a throughput loss. On the contrary, RL-AMC method is able to discover the MCS configuration that can ensure a more efficient use of the available channel resources. This is more evident at intermediate distances from the eNB when short-term fading may lead to use more frequently low-rate MCSs. As shown in the figure, the throughput improvement varies between 20% and 55% in the range of distances between 200 meters and 800 meters.

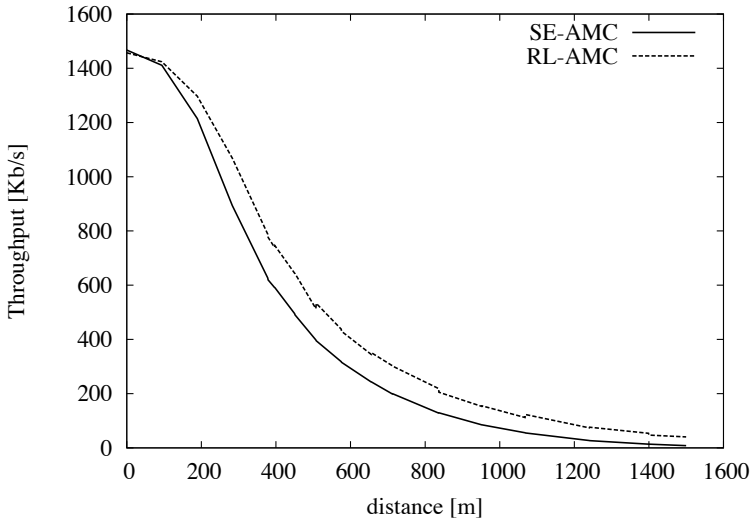

Figure 14: Average throughput as a function of the distance of the tagged user from the eNB in a pedestrian scenario

Finally, we also consider a more dynamic environment in which there is an increasing number of UEs in the cell, and all the UEs are moving according to a random waypoint mobility (RWM) model with speed 30 km/h and pause time equal to 5 seconds. Figure 15 shows a comparison of the aggregate cell throughput with both SE-AMC and RL-AMC schemes as a function of the network congestion (i.e., number of UEs). The results clearly indicate that the throughput improvement provided by RL- AMC is almost independent of the number of UEs and it is about 10%. We can also observe that the cell capacity initially increases when going from 10 to 20 UEs. This is due to two main reasons. First, RR is able to allocate RBs in a more efficient way when the number of UEs is higher. Second, the higher the number of UEs and the higher the probability that one of the UEs is close to the eNB and it can use high data-rate MCSs.

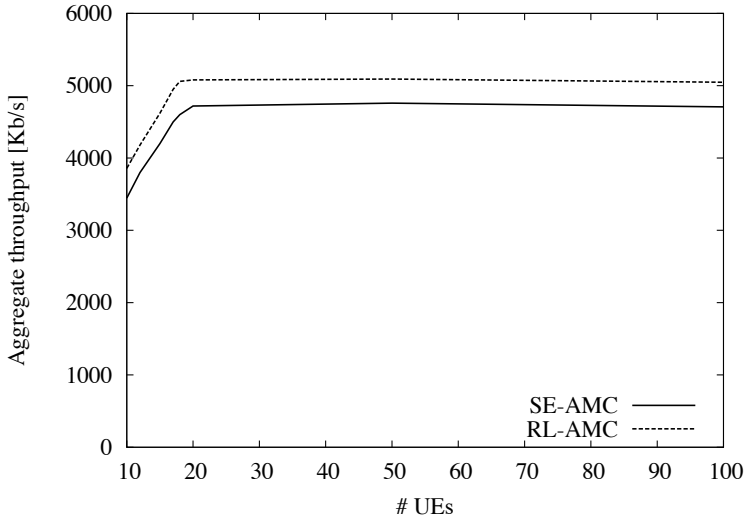

Figure 15: Average cell throughput as a function of the number of UEs in an urban vehicular scenario.

4 Capacity analysis: Assessing capacity of an integrated offloaded network

In this section we present the status of our work on the capacity gains that can be obtained in an integrated offloaded network. Also in this case, we build upon and extend the initial results presented in D3.1 [1], in particular those related to the Push&Track and Droid offloading systems, presented in Section 3 of D3.1.

As it was the case of solutions presented in D3.1, in WP3 we investigated in parallel several mechanisms and configurations for offloading. The rationale, as also described in the DoW of the project, is to design and quickly test such solutions in dedicated settings, in order to initially compare them, understand their feasibility and indicative performance. This part of the work in this WP is, therefore, preliminary to the more complex work of integration and comprehensive evaluation that is being carried out in WP5, and that will be fully developed during the third year of the project. Through these investigations we are able to already understand the most useful offloading mechanisms and solutions, that are then integrated in selected reference use cases, demonstrated and validated in WP5.

In Section 4.1 we start from the basic Droid system presented in D3.1 (and published in [55]), and extend its basic operations by considering the opportunity of using also a WiFi Access Point infrastructure in a city environment. This is one of the first studies – to the best of our knowledge – that takes into consideration all the three main content delivery enabling technologies that are nowadays considered for offloading, i.e. cellular networks, WiFi networks and opportunistic networks. We compare "vanilla" Droid (which uses only cellular+opportunistic) with content delivery solutions based only on WiFi APs, and on WiFi APs and opportunistic dissemination. Note that we consider a real WiFi AP development, i.e. the one currently available and open to the general users (managed by the municipality) in the city of Bologna. This analysis shows that (i) augmenting WiFi-based delivery with opportunistic networks is very important, as opportunistic delivery allows us to complement the limited pervasiveness of the WiFi AP deployment. However, comparing WiFi+opportunistic with Droid, we show that in case of stringent delivery deadlines, Droid possibility to exploit a more pervasive cellular network (as opposed to the WiFi network) to disseminate the content becomes a winning factor. Moreover, we also test the case of a fully combined solution based on Droid, and also exploiting WiFi APs, and therefore characterise the additional gain, in terms of offloading, of using it. Finally, we investigate the energy-capacity trade-off in this case. Complementary to the approach presented in Section 2.2, in this case we assume that the constraint in terms of maximum energy consumption is represented by a maximum number of copies of the message that each node can disseminate in the opportunistic network (i.e., a form of limited epidemic diffusion). Again, we show that, for a given such constraint, using an offloading solution integrating cellular and opportunistic network is winning, with respect to one that uses WiFi APs and opportunistic dissemination.

In Section 4.1 we focus on a typical publish/subscribe scenario, i.e. one where content becomes available at some point in time, and is automatically (and implicitly) requested by all interested nodes at that time. In Section 4.2, instead, we complement these results by considering the case where content is not requested by all mobile nodes in a synchronised way. Specifically, we again consider a vehicular scenario, and assume that content of interest is geo-localised, and becomes interested for users once they enter a specific geographical area. Therefore, requests for content are generated asynchronously from each other, and schemes such as Droid should be modified. We therefore define simple algorithms to support this case, compatible with the overall MOTO architecture described in D2.2.1 [4]. Then, we assess the feasibility of offloading, by showing how much capacity operators can gain if opportunistic dissemination is used also in this case. We show that, clearly, this depends on the features of the mobility patterns, on the amount of time during which users keep the content locally (after having received it), and by how many users are interested in it. Overall, the additional capacity that can be gained ranges between 20% and 90%. These results are presented in more detail in [17], also included as Appendix A.

4.1 Offloading with both opportunistic and Wi-Fi networks

In this section, we evaluate through simulation the performance of several offloading strategies. We compare opportunistic and AP-based offloading strategies under tight delays. We consider a location-based service in a vehicular context. Content with traffic information or some infotainment announcement must be distributed to a multitude of users within a given maximum reception delay (in order to guarantee a minimal QoS on a per-content basis). We assume that nodes are equipped with several wireless interfaces, so that they are able to communicate through multiple interfaces simultaneously. Possible combinations involve 3G and 4G to communicate with the cellular infrastructure, Bluetooth or Wi-Fi ad hoc to communicate with neighboring devices, and Wi-Fi in infrastructure mode to communicate with fixed APs.

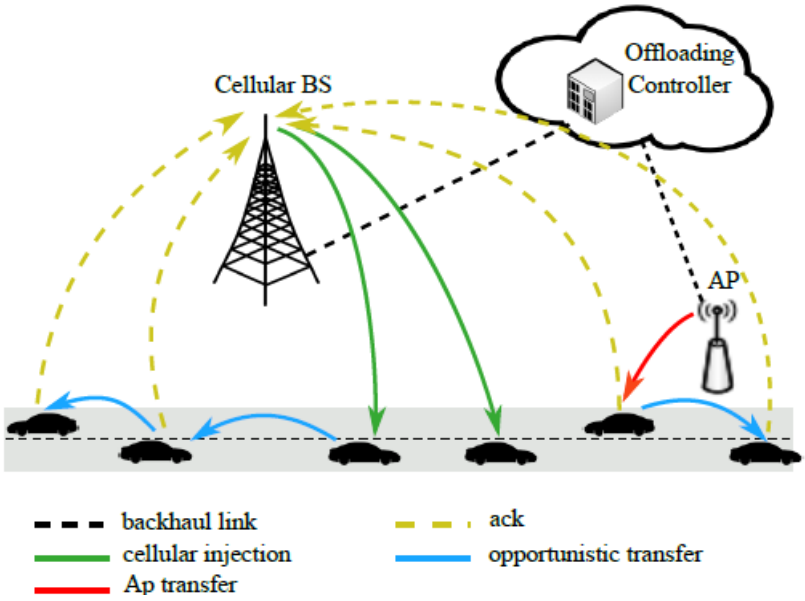

Figure 16: Offloading model: The dissemination process is kick started through cellular and/or AP transfers. Content is diffused among vehicles through subsequent opportunistic contacts. Upon reception, users acknowledge the offloading agent using the feedback cellular channel. The coordinator may decide at any time to re-inject copies through the cellular channel to boost the propagation. 100% delivery ratio is reached through fallback re-injections.

4.1.1 Mobility trace and AP position

We employ a large-scale vehicular mobility trace representing the city of Bologna (Italy). The Bologna dataset consists of 10,333 nodes, covering a total of 20.6 km^2 and 191 km of roads. The simulated traffic in the dataset mimics the everyday road activity in the two metropolitan areas. From the mobility trace, we derive a contact trace that features contacts between nodes when the distance between them is below a given threshold (we consider a range of 100 meters, in line with IEEE 802.11p specifications). The resulting trace has a duration of about one hour; on average, 3500 nodes are present at the same time (because some nodes leave while others join during the observation period). The advantage of using this large-scale trace is that, differently from other available datasets, we have a clear high turnover rate, due to vehicles entering and exiting the interest area, and no apparent social links between nodes. The distribution of contact durations is exponential. Most contacts are very short, confirming the highly dynamic nature of the trace. In addition, only few contacts last for more than a few minutes.

We extracted the location of an existing Wi-Fi Hotspot public deployment from http://www.comune.bologna.it/wireless. Figure 17 shows the position of 93 APs along with a map of the town center. We merged the location of APs with the vehicular mobility trace to extract a completely new dataset that includes vehicles mobility and AP positions. From this trace, we derived the connectivity traces between vehicles and fixed APs. The outcome is a completely new time-variant graph with unidirectional links connecting vehicles with the APs and the other vehicles.

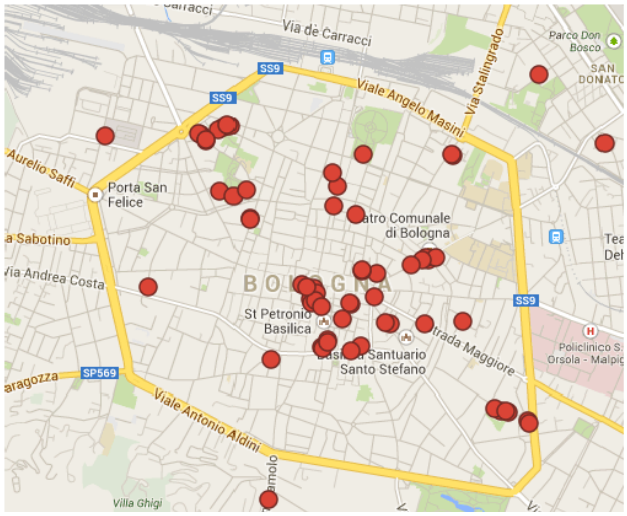

Figure 17: Bologna map with fixed AP positions.

To characterize this new dataset, we study the pairwise interactions among mobile nodes (vehicles) and fixed APs only. Figure 18 presents the distributions of contact and inter-contact times between vehicles and APs. Contact and inter-contact times follow a lognormal distribution.

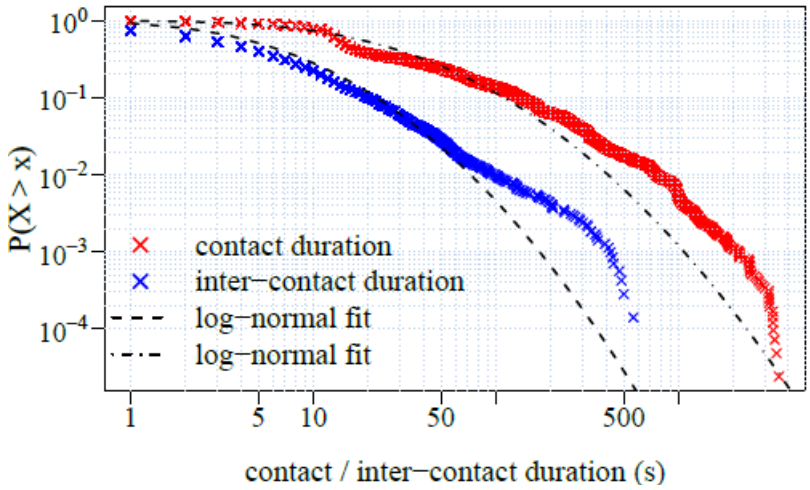

Figure 18: CCDF of contact and inter-contact times with fixed Aps.

Two anomalies make the study of the dataset interesting. First, a relevant number (around 20%) of inter-contact times are zero s, meaning that when a vehicle exits from the coverage area of an AP, it is already under the range of another. We may explain this aspect by noting that in the city center several APs are very close together. Nevertheless, the sole contact distribution does not tell us the whole story, as very few APs are located in southern and western town districts, and many vehicles passing there enter and exit the system without falling into the coverage zone of any APs.

Moreover, we note that many AP meetings occur in bursts. We infer a strong correlation between the geographic position and the expected duration of contact and inter-contact times with APs. In addition, we remark that around 80% of the contacts with APs last for more than 10 seconds. While this may be an acceptable duration for data transfers, short-lived contacts lasting less than that value could suffer from the duration of authentication and address granting procedures with an AP.

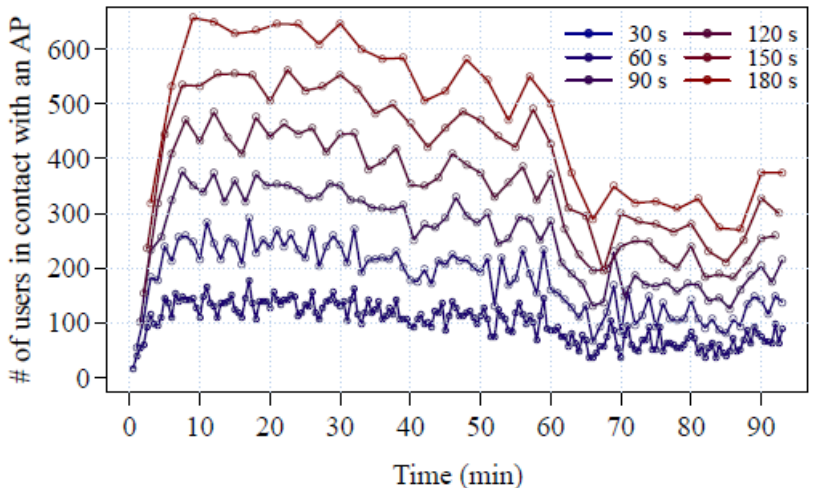

Figure 19: Time evolution of the total number of contact between vehicles and APs for different message lifetime. Note that, in average, 3500 users are active at the same time.

Finally, Figure 19 depicts the evolution of the number of vehicles in contact with at least an AP during a given time window (equivalent to the delay-tolerance of content reception in this case). The figure indicates the amount of vehicles that enter in the transmission range of least one AP during the considered distribution period. Augmenting the delay tolerance, the chances that a vehicle enters in the range of at least one AP increase. Still, the number of vehicles benefiting from this transfer opportunity is limited, if measured against their total number in the system, lying always between 5 and 25% of present users.

4.1.2 Simulation setup and scenario

In our implementation, we consider a simple contact-based ad hoc MAC model, where a node may transmit only to a single neighbor at a time. Transmission times are deterministic since we do not take into account complex phenomena that occur in the wireless channel such as fading. Communications consist of two different classes of messages (content and control). All transfers, including *ack* messages, may fail due to nodes moving out of each other's transmission range or exiting the simulation area. In addition, it is possible for the same message to be concurrently received through the two interfaces. In that case, we consider the one processed first. The ad hoc routing protocol employed by nodes to disseminate the content is the epidemic forwarding.

Parameters in simulation are set to mimic the functioning of communication technologies currently available to consumers. In each simulation run, the downlink bit-rate for the infrastructure network is set to 100 KB/s, while uplink is fixed at 10 KB/s. These values are in line with the average bit-rate experienced by users of a typical 3.5G network. The bit-rate for the ad hoc link is set to 1 MB/s, also in line with the advertised bit-rate of the IEEE 802.11p standard. The size of each content update is set at 100 KB. The size of the acknowledgement messages is 256 bytes, as it carries very little information (content and node identifiers).

4.1.3 Opportunistic or AP-based offloading?

How opportunistic offloading behaves when compared to most traditional AP-based offloading strategies?

To answer this question, we run network level simulations to benchmark the opportunistic offloading algorithm DROiD [55], which is at the base of MOTO, against other more conventional strategies based on direct offload from Wi-Fi hot spots. AP-based offloading takes advantage of the presence of fixed infrastructure that can serve to offload the cellular network. Nevertheless relying only on fixed deployment typically lacks of the flexibility of the pervasive cellular networks, since transmission range and spatial density are limited for physical reasons.

We exploit the Bologna dataset, considering also the AP deployment, with simulation parameters previously described in Section 4.1.2.

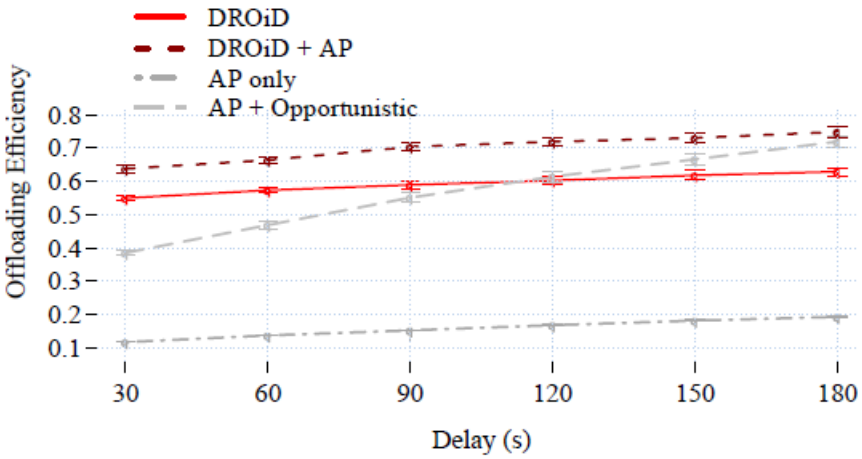

Figure 20: Bologna trace - offloading efficiency comparison between *DROiD*, *DROiD + AP*, *AP-based,* and *AP + opportunistic* distribution strategies. 95% confidence interval plotted.

In Figure 20 we may appreciate the efficiency of three alternative approaches to data offloading making use of fixed APs, simulated on the Bologna dataset and compared with the DROiD strategy as benchmark. We note that the ***AP only*** strategy gives extremely poor results. Even with larger tolerance to delays, this strategy is never capable of saving more than *20%* of traffic. Thus, this strategy turns out to be unable to relieve substantially a large fraction of the cellular load. As already hinted during the dataset analysis phase, the main problem in this case is that APs are not ubiquitously available in each town district. Vehicles traveling in areas without Wi-Fi coverage cannot download data from nearby APs, and will likely reach the panic zone without the content. An increase in delay-tolerance of the content only partially mitigates the issue. The analysis, carried out employing the real-world deployment in the city of Bologna, suggests that in order to offload a substantial part of traffic through fixed hot spots their deployment should be carefully planned, without black holes in spatial coverage.

As a second option, we consider that our nodes be able to communicate with both APs and other vehicles through direct ad hoc links, without considering any cellular re-injection (the ***AP + Opportunistic*** strategy). This proves to be very beneficial for the overall offloading performance, increasing efficiency up to 50% with respect to AP *only*. The possibility to exchange data directly among users, together with mobility, allows spreading the infection in many areas not covered by fixed APs through store-carry-forward routing. Gains, as expected, rapidly improve as the reception delay increases. Fixed APs, in this case, act as fixed (and free) infection source. Still, the benefits of re-injections through the cellular link emerge for shorter reception delays (up to 120s). For shorter deadlines, DROiD results always preferable than AP-based strategies. Once again, re-injections prove essential in the case of lagging infection evolution, and this is particularly true when opportunistic contacts among vehicles are scarce. For short distribution intervals, infected nodes hardly can carry the content far from AP location before its expiration. For longer delay tolerance instead, the continuous use of the APs to infect neighbor nodes gives the *AP + opportunistic* strategy an edge.

In order to take advantage of the fixed hot-spot infrastructure along with the feedback-based re-injection through derivative strategy, we evaluate ***DROiD + APs***. In this case, the APs pre-fetch the content at t_0. Cellular re-injections intervene only when the diffusion lags, so to overcome the difficulties encountered by users located away from APs range. This strategy emerges to be always the best, guaranteeing more than 65% of offloaded content. APs guarantee a steady infection rate to vehicles passing in their transmission range, letting the cellular infrastructure to target users in more isolated areas.

Further analyses of the traffic flowing on each interface reveal interesting and unexpected information. It turns out that in joint opportunistic/AP-based offloading strategies (***AP + opportunistic*** and ***DROiD + APs***), the fixed APs tend to kick-start the dissemination, which is then carried over with subsequent direct communications between mobile nodes. The aggregate amount of data flowing through APs in these cases

is roughly 10 times less than when hot spots are the only offloading options. Remarkably, this small fraction of data transferred through APs results very important to bootstrap the dissemination, offering an advantage compared to the non-AP based solutions.

4.1.4 Energy savings and fairness

A critical challenge to make mobile data offloading potentially attractive to end-users is to attenuate the impact of opportunistic communications on the battery of devices, concurring thus to increase their lifetime. For this reason, we analyze the impact that simple energy-saving methods have on offloading performance.

In our analysis, we compare ***DROiD*** [55] with the ***AP + opportunistic*** strategy, fixing the maximum number of possible opportunistic transmission that a node can do for each message. To put energy saving strategies in practice, we offer to users only a fixed amount of *tokens* for each content to be distributed. The local token count is decreased each time the content is forwarded by a node. When the token count is equal to zero, the node stops forwarding, and waits for the next content to appear.

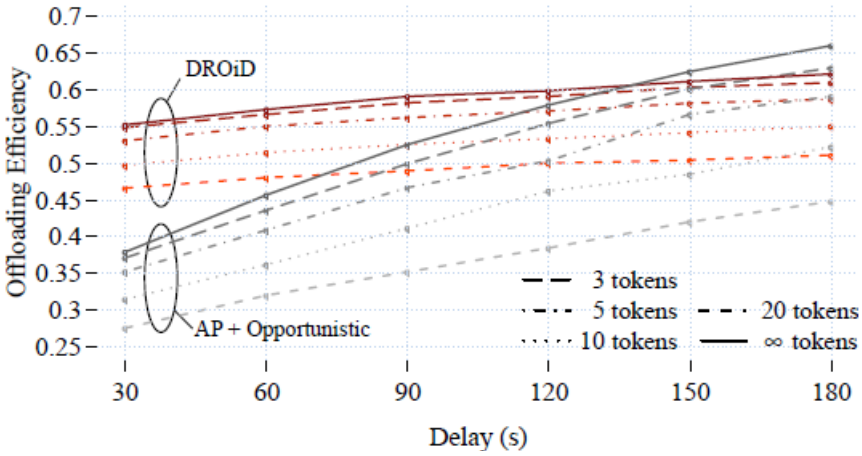

Figure 21: Offloading efficiency as a function of the number of transmission tokens for DROiD and AP + Opportunistic. Confidence intervals omitted for clarity.

Simulation results, presented in Figure 21, show that a limited number of tokens affects offloading performance, wasting possible contacts and lowering the aggregate capacity.

Energy saving strategies have a more pronounced effect on the AP *+ opportunistic* schema, which sees its performance highly lowered. The performance gap stretches as the delay-tolerance increases because nodes are more likely to run out of tokens. From the figure, we may appreciate that restricting the number of tokens to 20 does not bring a substantial performance hit for *DROiD*, while its influence is more pronounced in *AP + opportunistic.*

The energy saving scheme should trade off offloading efficiency for battery life, while ensuring at the same time to split the overall energy cost equally between all the nodes involved in the dissemination. However, in opportunistic networks, contacts are typically imbalanced between nodes, and it is common to find nodes that during the message lifetime sustain an important number of data forwarding.

	3	5	10	20	∞
30s	0.22	0.21	0.19	0.15	0.15
60s	0.27	0.25	0.23	0.2	0.17
90s	0.29	0.27	0.27	0.23	0.18
120s	0.3	0.33	0.3	0.25	0.21
150s	0.35	0.33	0.32	0.27	0.2
180s	0.39	0.37	0.34	0.29	0.21

(a) AP + Opportunistic

	3	5	10	20	∞
30s	0.43	0.4	0.34	0.29	0.28
60s	0.44	0.41	0.35	0.3	0.28
90s	0.46	0.42	0.36	0.31	0.29
120s	0.46	0.42	0.36	0.31	0.28
150s	0.47	0.43	0.37	0.31	0.28
180s	0.47	0.43	0.38	0.31	0.27

(b) DROiD

Figure 22: Jain's fairness index for different combinations of number of tokens and delay tolerance.

To evaluate **fairness**, we use the Jain's index to compare the fairness in the number of opportunistic transmission for the two schemes under evaluation, with different token values and content reception delay. Figure 22 shows that DROiD always presents better fairness indexes than AP-based strategy. The analysis of nodes forwarding reveals that in the *AP + opportunistic* strategy the number of users participating in content forwarding is sensibly lower than in *DROiD*. This depends on a mix of two factors: the efficiency is lower in general, so a greater number of nodes do not physically store the content to forward. The effect relies also on the fact that AP-based strategies tend to concentrate data forwarders among those nodes that receive the content first. This is the reason why in the 180 s scenario with infinite tokens, DROiD shows better fairness but lower efficiency.

The use of fixed Wi-Fi APs as the only data offloading strategy proves ineffective in the case of medium density AP deployment. Coupling it with opportunistic distribution could improves instead data offloading, namely to kick start the distribution process and to deliver free copies of the content to users located inside their coverage range. Nevertheless, if we consider tight delivery times, the use of the pervasive cellular infrastructure is still required to target isolated node. Offloading strategies relying on random re-injections result intrinsically more fair. This can improve the battery duration of mobile users with respect to AP-based strategies, which always target nodes in their spatial proximity.

4.2 Offloading with non-synchronised content requests

Most of the literature on opportunistic-based offloading investigates the scenario where a specific piece of content is generated, and the set of users to whom it has to be delivered is known already at that time and does not change subsequently, i.e. requests are *synchronised*. While significant, this scenario only partially captures relevant use cases. In particular, it does not cover cases where content demand is *dynamic*, i.e. users' requests for the same piece of content can arrive at different time instants. In the latter scenario offloading can still be applied: upon a request, content can reach the requesting user either through the opportunistic network, exploiting an ongoing dissemination process, or through the cellular network, in case the opportunistic dissemination does not reach the user in time. Offloading may even be more needed in case of dynamic requests, as synchronised requests could in principle be served also through multicast transmissions (although [56] – also presented in Section 5 of this document – shows that offloading is beneficial also when multicast is applied).

We have started investigating dynamic content requests, with a particular focus on vehicular scenarios. We deliberately use a very simple offloading scheme, whereby resources provided by mobile nodes are minimally used. Nodes interested in a content store it for a limited amount of time after receiving it. New requests from other users are satisfied either when the requesting user encounters another user storing a copy of the content, or through the cellular network upon expiration of the delivery deadline.

As opposed to most of the literature looking at offloading through opportunistic networks, in our scheme we do not use any epidemic dissemination mechanism. On the one hand, this allows us to test a minimally invasive offloading scheme from the mobile users' perspective. As additional resources spent by mobile devices are sometimes considered a possible roadblock for offloading, our results show the offloading

efficiency when this additional burden is extremely low. On the other hand, this simple scheme allows us to stress the efficiency of offloading in a particularly unfavorable configuration, thus providing a worst-case analysis, all other conditions being equal.

We focus on two complementary scenarios. In the first one, users move in a given physical area, and *all* request a piece of content, though at different points in time. This scenario is representative of users moving inside a limited area, and accessing very popular content, though not particularly time critical (i.e., content that does not generate a surge of requests immediately when it is generated). In the second scenario, users enter and exit (after a short amount of time) a given geographical area, and request content after a random amount of time after they entered the area. This complementary scenario is thus representative of users traversing a geographical area, as opposed to roaming there. Finally, in this scenario we also consider the case where content is requested only with a certain probability, i.e., when content has different levels of popularity.

We analyse the offloading efficiency in these scenarios, defined as the fraction of nodes receiving content through the opportunistic network. We characterise efficiency as a function of key parameters such as the number of users, the deadline of content requests, the time after which users drop the content after having received it, the popularity of the content. Even with an unfavourable opportunistic dissemination scheme, we find that offloading can be very efficient, as it is possible to offload up to more than 90% of the traffic. In other configurations, we find that the considered offloading scheme is less efficient, resulting in an offloading of only about 20%. In such cases, however, there is ample room for improvement, by further leveraging opportunistic networking resources, e.g., through more aggressive content replication schemes.

4.2.1 Offloading algorithms for non-synchronised requests

The algorithms we have defined are compliant with the general MOTO architecture presented in D2.2.1 [4]. We assume the existence of a Central Dissemination Manager (CDM), that can communicate with all nodes through the cellular network and keeps track of the dissemination process. The offloading mechanism is defined by the actions taken by requesting nodes and by the CDM, as described by Algorithms 1 and 2, respectively, shown in Figure 23.

Algorithm 1 Actions taken by requesting nodes

```
                                          ▷ Run by a tagged node k
 1: Upon request for content C
 2: content_received = false
 3: Send content_request to CDM
 4: if C not received immediately from CDM then
                                          ▷ try with opportunistic contacts
 5:     while content_timeout is not over do
 6:         request C to encountered nodes
 7:         if content received then
 8:             content_received = true
 9:             Send ACK to CDM
10:             break
11:         end if
12:     end while
13:     if content_received == false then
14:         Receive C from CDM
15:         content_received = true
16:     end if
17: end if
18: while sharing_timeout is not over do
                                          ▷ available for opportunistic sharing
19:     Send C to encountered nodes upon request
20: end while
21: Cancel content C
```

Algorithm 2 Actions taken by CDM

```
                                          ▷ Run by the CDM for content C
    Init #nodes_with_C = 0
 1: Upon request from node k
 2: k_served = false
 3: if #nodes_with_C == 0 then
 4:     Send C to k
 5:     #nodes_with_C++
 6:     Set sharing_timeout for node k
 7: else
 8:     while content_timeout is not over do
 9:         if ACK received by k then
10:             #nodes_with_C++
11:             k_served = true
12:             Set sharing_timeout for node k
13:             break
14:         end if
15:     end while
16:     if k_served = false then
17:         Send C to k
18:         #nodes_with_C++
19:         Set sharing_timeout for node k
20:     end if
21: end if

22: Upon sharing_timeout for node k over
23: #nodes_with_C = #nodes_with_C-1
```

Figure 23. Algorithms for offloading in non-synchronised requests.

Let us focus first on the actions taken by requesting nodes (Algorithm 1). When a request is generated at a node, the node sends it to the CDM via the cellular network (line 3). The node is guaranteed to receive the content within a given *content timeout*. During the timeout, the node tries to get the content from encountered nodes (lines 5-12). If the timeout expires, it receives it directly from the CDM (lines 13-16). Upon receiving the content, the node sends an ACK to the CDM (line 9 and, implicitly, line 14). In addition, it keeps the content for a *sharing timeout*, during which it can share the content with other encountered nodes (lines 18-20). After the expiration of the *sharing timeout* the content is deleted from the local cache. Note that requests and ACKs are supposed to be much shorter than the content size, and thus do not significantly load the cellular network.

Let us now focus on the actions taken by the CDM (Algorithm 2). Thanks to requests and ACKs, the CDM is always aware of the status of content availability in the network. Upon receiving a request, it checks whether some other node is already storing a copy of the content or not. In the latter case (lines 4-6) there is no chance that the user can get the content opportunistically through another node, and the CDM sends the content directly through the cellular network. In the former case (lines 7-21), it waits to receive an ACK during the *content timeout* (lines 8-15), indicating that the node has received the content. If this does not happen, it sends the content directly to the node (lines 16-20). Finally, upon expiration of the *sharing timeout* for a given node the CDM updates the view on the number of nodes with the content (lines 22-23).

4.2.2 Evaluation when all users request the content

In the first scenario we have considered, all users entering the physical area covered by the cell request the content, though they clearly do it at different points in time.

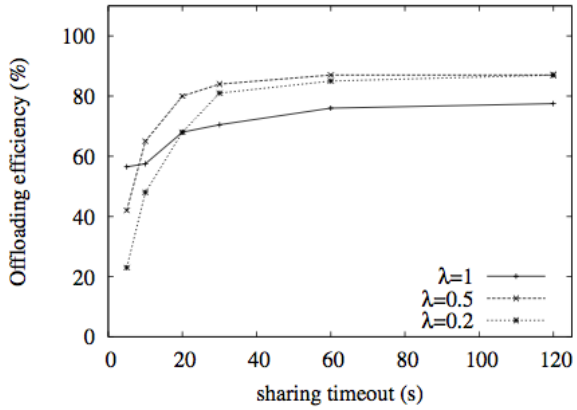

(a) $N = 20$, *content timeout* $= 60$s,

Figure 24. Example of results in the first scenario.

Figure 24 shows an example of the results we have obtained in this scenario. In the figure, λ represents the rate at which nodes generate the request for the content (i.e., two requests from two different users are spaced by an exponential interval with average 1/ λ). We observe two regimes. When the *sharing timeout* is low, higher request rates result in higher offloading. This is intuitive, because higher request rates results in requests being more concentrated in time. When nodes share the content only for very short amounts of time (see for example the case of 5s), concentrating the requests in time increases the probability of encountering other nodes sharing the content. Less intuitive is the behaviour for large sharing timeouts, where higher request rates results in *lower* offloading efficiency. Intuitively, when requests are more concentrated in time, *content timeouts* for nodes that do not get the content via the opportunistic network are also more concentrated. When a timeout expires and content is delivered via the cellular network, this kicks off a fast increase in the dissemination of content via the opportunistic network in the region of the node whose *content timeout* has expired. When expirations are less concentrated in time (i.e., when request rates are lower), the opportunistic diffusion process has more time to spread content, and therefore the offloading efficiency increases. Additional aspects are analysed in [17] (see Appendix A).

4.2.3 Evaluation when all users request the content

Figure 25 shows the offloading efficiency in the second scenario we have considered. In this case, when entering the area covered by the dissemination system, vehicles become interested in the content with a given probability p. If they are interested, they generate a request after a time interval uniformly distributed between the time when they enter and the time when they reach the centre of the area.

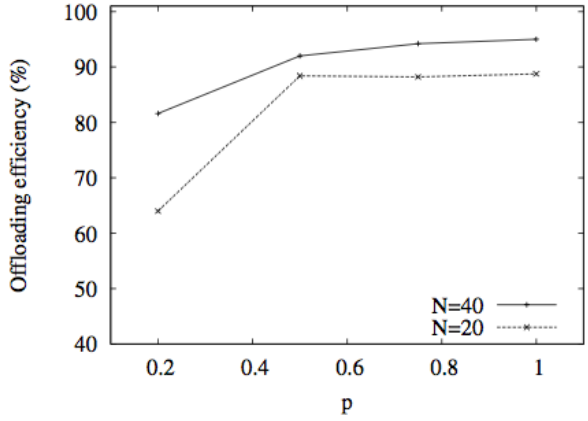

Figure 25. Example of results in the second scenario.

Figure 25 shows the offloading efficiency for two considered densities of nodes and the different content popularities (p). Results basically confirm previous observations. This is nevertheless important, as this scenario is more representative of a “steady state” behaviour of the offloading system, as nodes constantly enter and exit the area at a given rate, and continuously generate requests (with a given probability). Denser networks (N = 40) achieve higher offloading efficiency. The effect of the popularity parameter is similar to that of the request rate in the first: the higher the popularity, the higher the number of nodes sharing content, the higher the offloading efficiency. It is interesting to note, however, that, due to the mobility of the nodes, they stay within the area only for about 30s in total, and, on average, stay in the area for about 22s after having generated a request. This is the “useful time window” during which they can receive content via opportunistic dissemination. Even though this time window is rather short, offloading is very efficient, even at quite low popularities (p = 0.2). Additional aspects are analysed in [17] (see Appendix A).

5 Intra-technology scheduling: Joint use of multicast and D2D in cellular networks

This is the first section where we present results on scheduling. As discussed in Section 1, we explored multiple directions from this standpoint. In this section, we start addressing one of the fundamental questions for the practical applicability of offloading, i.e. whether using cellular multicast would not be used in most of the cases. This is particularly relevant for multimedia (non real-time) popular content, that may be required by multiple users simultaneously in the same physical area (i.e., covered by the same cell). In fact, among general multimedia services, some involve delivering the same piece of data to a community of interested users. Examples that fit this use case are software updates, on-demand videos, and road traffic information. When a multitude of co-located users is interested in the same content, two possible approaches could help operators to relieve their cellular infrastructures: **mobile data offloading** and **multicast**.

Multicast makes use of a single unidirectional link, shared among several users inside the radio cell, allowing, in principle, a more efficient use of network resources with respect to the case where each user is reached through dedicated bearers.[6] To ensure the coexistence between multicast and unicast services, operators must reserve a fixed amount of resources for multicast transmissions. Despite its attractive features, multicast presents intrinsic and still unresolved issues that limit its exploitation due to the difficult adaptation to radio channel conditions.

Mobile data offloading is an alternative low cost solution to reduce the burden on the infrastructure network. Direct device-to-device (D2D) communications help lowering the load on the infrastructure. The increase in the density of mobile users gives rise to an abundance of contact opportunities and represents a strong argument to support opportunistic offloading strategies. In order to encourage subscribers to offer their battery and storage resources to this end, mobile providers may offer monetary incentives and pricing discounts. As a counterpart, users should accept a delayed content reception.

In the present section*, we explore the combination of opportunistic traffic offloading with multicasting*. As we will see later, this allows significant reduction in the load on the access part of the cellular network. Multicast is not intended for retransmissions, and performance suffers and resources are wasted in the case of a single bad channel user inside the cell, due to trade-offs in coverage and efficiency. By including D2D communications into the picture, we obtain additional performance gains in terms of radio resources. Well-positioned users participate in mitigating the inefficiencies of multicast, by sharing their short-range resources to hand over content to users in bad cellular channel conditions. Depending on the number of participants requesting data, we find a break-even point that achieves a good trade-off in terms of covered users and reception delay.

To assess the performance of this joint multicast/D2D approach it is necessary to evaluate the amount of radio resources consumed at the base station. This leads us to introduce a finer model of radio resource consumption than previous works in the literature. Existing proposals do not consider heterogeneous channel conditions and assume that delivering a given amount of data to different users has always the same cost. Such an assumption does not hold in reality, as radio resources vary according to the channel condition experienced by each user. In other words, transmitting the same piece of content to users with different channel conditions do lead to uneven costs at the base station. To the best of our knowledge, we are the first to evaluate this aspect in the context of data offloading. Note that our results are one of the basis for designing cellular scheduling policies mixing together multicast and D2D transmissions. D2D should be intended in a broad sense, and includes both the standardized LTE-D2D technique, as well as solutions exploiting other technologies for D2D communications, such as WiFi and Bluetooth. Nevertheless,

[6] Note that a more precise terminology would be “multicast/broadcast”, because only a subset of nodes is concerned by the content (multicast), and the shared nature of the wireless medium (broadcast) is exploited to transmit data. For the sake of readability, in the following we will only employ the term “multicast”.

we consider this part of the work more related to intra- than inter-technology scheduling, as we see a more direct applicability to scheduling policies implemented by a cellular operator inside its network.

In the following of the section we provide an extended summary of the work undertaken on this topic. Additional details are available in [56], also included in Appendix A.

Moreover, we have extended these results by defining a Reinforcement Learning algorithm that automatically sets the optimal fraction of users that should be served via multicast, over the total number of users requesting a given content. We also provide a summary of this extension in Section 5.4, while details are available in [57].

5.1 Multicast in 4G networks

LTE proposes an optimized broadcast/multicast service through **eMBMS** (*enhanced Multimedia Broadcast Multimedia Service*), a point-to-multipoint specification to transmit control/data information from the cellular base station (eNB) to a group of user entities (UEs).

Cellular UEs can use different modulation and coding schemes (**MCS**) to deal with variable channel characteristics. Each UE experiences different radio conditions, depending on path loss, interference from other cells, and wireless fading. UEs that are closer to the base station are able to decode data at a higher rate, while others located near the edge of the cell have to reduce their data rate and use a degraded MCS. This heterogeneity (time-varying and user-dependent) reduces the effectiveness of multicast because the eNB uses a single MCS to multicast downlink data. The selected MCS for multicast should be robust enough to ensure the successful reception and decoding of the data-frame for each recipient inside the cell. Thus, the worst channel among all the receivers dictates performance. An increase in the number of UEs boosts the probability that at least one UE experiences bad channel conditions, degrading the overall throughput.

To quantify this effect, we simulate a 500 x 500 square meters single LTE cell with an increasing number of randomly located receivers using the ns-3 simulator. Figure 26 presents the average minimum channel quality, in terms of CQI (Channel Quality Indicator), reported at the eNB by UEs (static). The reported CQI is a number between *zero* (worst) and *15* (best). The CQI indicates the most efficient MCS giving a Block Error Rate (BLER) of *10%* or less.

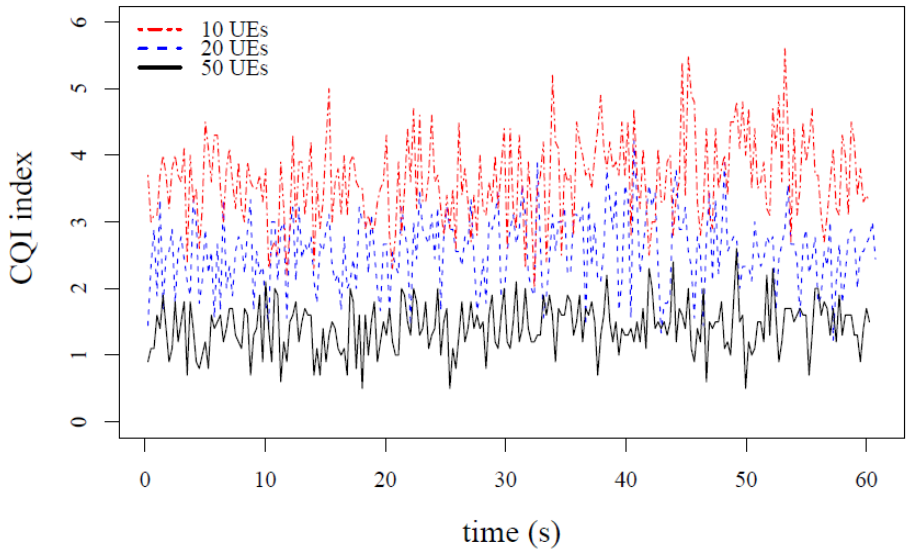

Figure 26: Minimum CQI for different multicast group sizes. 100 runs, confidence intervals are tight (not shown).

The average minimum CQI value decreases as the number of users in the multicast group increases. The result is that augmenting the number of multicast receivers clearly affects the attainable cell throughput. This greatly motivates us to investigate methods to cope with the inefficiencies of multicast.

5.2 Joint D2D / multicast offloading

We address the distribution of popular content to a set of *N* mobile UEs inside a single LTE cell. We want to transmit data to each UE with a guaranteed maximum deadline D at the minimum cost for the cellular infrastructure. We exploit D2D connectivity and store-and-carry forwarding at UEs.

A UE with good channel quality can obtain higher bit-rates with the same amount of resource blocks (RBs), while bad channel users consume more RBs in order to transmit the same amount of data. We capture the allocation expenditure, dynamically ranking the cost of transmitting the content to UEs according to their instantaneous CQI values.

The principles behind our approach are:

(1) at initial time, the eNB sends data to the I_0 UEs with the best radio conditions through a single multicast emission;

(2) the UEs that have received the data I_0 start disseminating it in a D2D (epidemic) fashion;

(3) Before the deadline, we define a time interval, a panic zone where all the nodes that have not yet retrieved the content receive it through unicast LTE emissions.

The proposed scheme allows all UEs to receive data by the deadline (as long as the panic zone is sufficiently large). It adapts to different deadlines -- the larger ones allowing for more D2D dissemination. Its performance relies essentially on one key parameter I_0 that characterizes the number of UEs reached by the initial multicast transmission. This immediately improves the usage of resources at the eNB, because it excludes the $N - I_0$ worst-channel UEs.

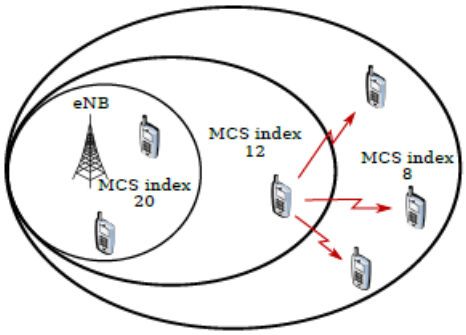

Figure 27: UEs can decode data with a maximum modulation schema depending on their position in the cell. The eNB may decide to multicast at higher rate (E.g., MCS index 12). UEs unable to decode data are reached through out-of-band D2D links.

Figure 27 offers a representative example of the proposed strategy with 6 UEs in the cell. In the D2D dissemination phase, *outaged* UEs benefit from nearby nodes, fetching data directly from them through D2D transmissions. This cooperative strategy is by far more efficient in terms of cellular resource consumption than multicast alone, given that the transmission rate increases and the D2D links typically exploit a much larger bandwidth than cellular communications.

5.3 Performance Evaluation

We compare the performance of the proposed joint distribution system with the one achieved by the classic cellular multicast alone. All the results presented in this section are averages over 25 independent simulation runs.

We consider a static number of UEs within the cell for each simulation run, to prove the validity of the concept. Future work will tackle the case where UEs can enter and exit the distribution area. Node mobility

is implemented according to the random waypoint model with speed fixed at 27 m/s and pause-time set at 0.5 s. We simulate UDP constant bit-rate downlink flows, each one with packet size s_k = 2048 bytes and a total load of 8 Mb.

We implemented our joint D2D/multicast strategy employing the MOTO simulator [5]. Since the MOTO simulator does not natively support cellular multicast, we implemented an additional module that interacts with the packet scheduler emulating single-cell multicast. The multicast module receives the CQI reports of UEs and decides the transmission rate following the steps explained in Section 5.2. Further parameters for the simulations can be found in Appendix A.

Reference Strategies: No D2D is the basic strategy, where UEs have no direct connectivity options, and multicasting through the cellular infrastructure is the only means of distributing content. We compare this base case to our joint D2D/multicast strategy. We assess the performance for three different values of ***N*** -- the number of users inside the cell -- respectively ***10***, ***25***, and ***50***, so to evaluate performance under different loads. We also consider various values for the parameter I_0 -- the number of direct multicast recipients. In order to be consistent with the notation, we evaluate this value as a percentage of ***N***.

Reception Methods: Figure 28 provides the fraction of packets partitioned by their reception method. For now, we focus only on their relative weight. As expected, the fraction of packets delivered through multicast follows I_0. The fraction of panic and D2D messages strongly depends on the parameters *D* and *N*. Tight service delays leave less time to opportunistic distribution to reach *outaged* UEs, resulting in a more intense use of panic retransmissions.

We can find a small amount of packet retransmitted during panic zone even in the ***No D2D*** strategy. These are packets incorrectly decoded by UEs during the initial multicast emission. In the other strategies, D2D allows not to make use of retransmissions where possible, because UEs can retrieve missing packets from other UEs. For instance, the strategies *No D2D* and *100%* have the same fraction of multicast reception, but differ on the amount of panic and D2D messages. We note also that for sufficiently long deadlines, panic zone is never triggered, and D2D transmissions meet the goal of guaranteeing total data diffusion.

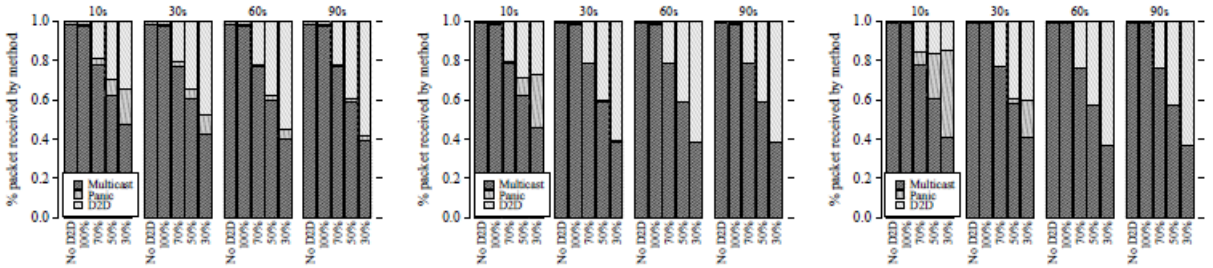

(a) 10 UEs (b) 25 UEs (c) 50 UEs

Figure 28: Data packet ranked by reception method. Multicast and Panic flows through the cellular infrastructure, D2D is on the Wi-Fi channel.

Cellular Resources: Mobile operators are primarily concerned about radio resource usage. This gives hints on the actual amount of RBs devoted to distribute data in the considered scenarios. Unlike the previous case, Figure 29 focuses on the amount of consumed radio resources at the eNB.

We note that the parameter *N* strongly affects the number of employed resources. This is even more evident if we consider very short deadlines. While the amount of resources devoted to multicast only slightly increases with the number of UEs, the impact of unicast re-injections heavily depends on the number of UEs in the cell. In some cases, a small fraction of unicast transmissions could translate into great resource usage. For large *N*, the choice of good values of I_0 becomes fundamental in order to avoid congesting the cell with too many panic retransmissions.

The interesting result is that for any possible value of N and D we may always find a joint strategy that offers better efficiency than No D2D.

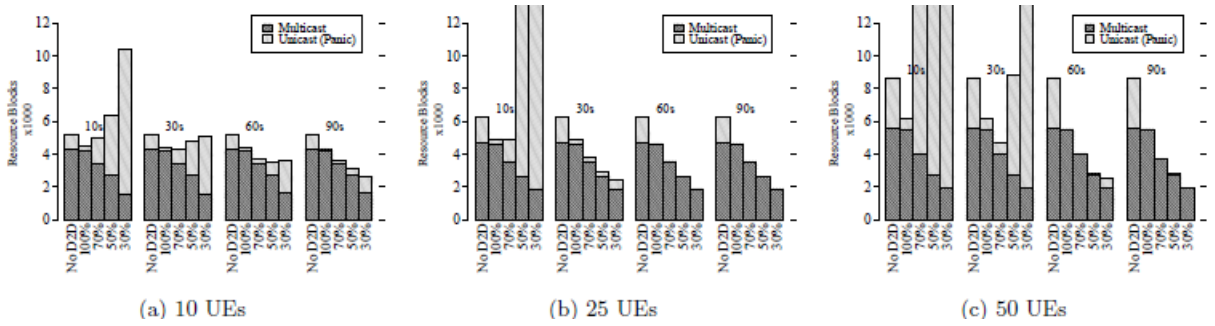

Figure 29: Average resource blocks employed at eNB to reach 100% dissemination. Note that even few panic zone retransmissions (in unicast) result very costly in resources.

In this section, we have presented a hybrid distribution system for popular content with guaranteed delays. Multicast is a valuable option to distribute popular data into a cellular network. However, performance is limited by the channel quality of the worst UE in the cell. In this context, we propose a framework that exploits D2D capabilities at UEs to counter the inefficiencies of cellular multicast.

The performance of a joint D2D/multicast strategy is evaluated by varying the number of UEs in the cell and the maximum reception deadline. Simulation results prove that the use of D2D communications allows increasing the multicast transmission rate, saving resources and improving the overall cell throughput.

5.4 A reinforcement learning strategy for joint multicast and D2D scheduling for data dissemination

In this section, we develop further one of the fundamental results on integrated scheduling initially evaluated in Section 5 of deliverable D3.1.1, where we explored the combination of opportunistic D2D offloading with cellular multicasting. A typical scenario in which a considerable amount of cellular resources can be saved is represented by many multimedia services requiring distributing the same piece of data to a community of interested users grouped in a limited geographical area (i.e., covered by the same cell). This is the case, for example, of software updates, on-demand videos, and road traffic information. The typical Zipf distribution of content interest makes this scenario highly relevant.

Despite the benefits of a hybrid distribution strategy were already evident from the analysis performed in the previous version of this document, we faced several challenges in designing a proper distribution strategy:

- Performance of opportunistic delivery is bounded to the mobility pattern of users. In addition, opportunistic networks can only guarantee a probabilistic assurance of data reception.
- Understanding the best fraction of users to reach through multicast and opportunistic transmissions is vital to guarantee users with a minimal QoS.

Since a truly optimal solution is not conceivable without precise knowledge of future contact patterns and channel qualities of UEs, we attack the problem from a more practical point of view. We apply a Reinforcement Learning (RL) approach that adapts the delivery strategy to the ongoing scenario. In our problem, a coordinator (the cellular base station) interacts with an unknown environment (the opportunistic network). In particular, for each packet to be delivered, the coordinator should decide the modulation and coding scheme (MCS) for the cellular multicast emission, which in turn affects the set of seed users capable of directly decoding data. This process, along with the opportunistic dissemination, entails a reward, i.e., a feedback to evaluate the action taken. The goal is to learn the best allocation of resources at the base station. Due to the many similarities in the formulation, we adopt the well-known multi-armed bandit RL technique to solve our problem.

5.4.1 Reinforcement learning strategy

We address the distribution of popular content to a set of N mobile UEs inside a single LTE cell. Each UE is a multi-homed device that embeds both a LTE interface and a short-range technology that allows D2D communications. In simulation, we consider IEEE 802.11g, however, the future integration of D2D capabilities within the LTE standard could be employed as well.

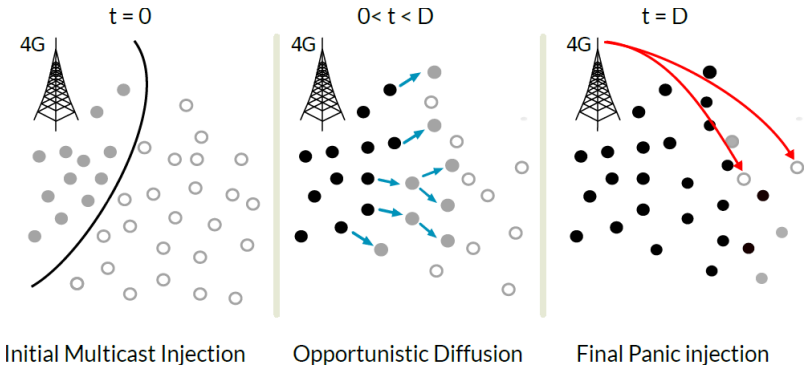

Figure 30. UEs can decode data with a maximum modulation schema depending on their channel quality. The eNB may decide to multicast at higher rate. UEs unable to decode data are reached through out-of-band D2D links.

We want to transmit data to each UE with a guaranteed maximum service delay D and at the smallest cost for the cellular infrastructure. In order to increase the efficiency of multicast, we exploit the possibilities offered by D2D connectivity and store-and-carry forwarding. The challenging issue is that opportunistic dissemination is, by definition, unreliable, as it depends on many factors that are outside of the control of the cellular infrastructure (e.g., movement pattern of nodes, variable density of opportunistic neighbors, or interference on the D2D channel). To achieve guaranteed delivery, we consider an acknowledgment mechanism, and panic zone retransmissions that are at the base of MOTO (D3.2). When the service delay reaches its maximum value D, the eNB pushes all the missing data to uninfected nodes using unicast transmissions. Of course, panic retransmissions are rather costly in terms of radio resources when compared to multicast. Different UEs could miss different packets leading to an inefficient use of resources. On the other hand, they represent the last opportunity to assure data reception.

Figure 30 offers a representative example of the proposed dissemination strategy. To avoid the penalty due to the presence of UEs experiencing severe channel conditions, the eNB emits at higher modulation, leaving them in outage. This is equivalent to restrict access to the multicast group only to the UEs in good channel conditions. In the opportunistic dissemination phase, outaged UEs benefit from nearby nodes, fetching data through out-of-band D2D transmissions. This cooperative strategy is by far more efficient in terms of cellular resource consumption than multicast alone, given that the cellular rate increases and the D2D links typically exploit a much larger bandwidth than cellular communications. Finally, panic injections assure data reception to all users. To make this strategy work, however, multicast and opportunistic dissemination must cover the majority of outaged UEs, leaving only a few of them uninfected at D. Problem definition

The problem we address is the following: *how to select the initial set of seed users to be injected using multicast transmissions with the objective of minimizing the total number of physical resource blocks (RBs) needed for content dissemination.*

We design our strategy as a network-based solution that can be deployed at eNB level, interacting with the packet scheduler. Let Γ be the set of seed belonging to the multicast group we build for each packet transmission. The users in the set are selected as follows. First, we rank the UEs according to their reported CQI. Then, we include in Γ at least I_0 UEs selecting first those with the best CQI. In this way, only the UEs with a reported CQI value greater or equal than the I_0th UE belong to the multicast group Γ. Note that by adjusting the I_0 value we implicitly control not only the number of initial seeds, but also the RBs amount associated with multicast emission. However, this is not the only cost we have to take into account. In fact, another conflicting cost is represented by the number of RBs allocated in the panic zone for unicast

transmissions. Specifically, if we consider a large I_0 threshold, we consume more RBs for the initial multicast emission because of the less efficient MCS employed. However, larger I_0 thresholds allow, in principle, to cover larger multicast groups, which potentially leads to a greater number of UEs infected at D. This means that fewer unicast transmissions are needed in panic zone. Conversely, smaller values of I_0 allow more contained costs for multicast transmission but dependently on the D2D dissemination effectiveness, could lead to more costly unicast transmission. While the number of RBs for the initial multicast emission is a deterministic and known value, the potential gain on the resource usage for the final unicast transmissions is a stochastic variable that depends on the effectiveness of the opportunistic dissemination, i.e. the more packets the D2D dissemination can deliver the less unicast transmission are needed.

Therefore, the problem we address is to try to dynamically estimate the function that relates the multicast cost to the unicast cost, in order to learn the best I_0 that permits to minimize the cost of the entire packet diffusion. We model this problem as a multi-armed bandit problem and we solve it through a Reinforcement Learning approach.

5.4.2 Background on multi-armed bandit algorithm

First, let us introduce the general formulation of a multi-armed bandit problem. In the simplest case, this kind of problem consists of a set of K probability distributions FD1 , ..., FDk with associated expected values μ1 , ..., μk and variances σ1 , ..., σk . For the sake of illustration, let us assume the distributions FDi describe the distribution of the outcomes of the i-th arm on a slot machine; the player is viewed as a gambler whose goal is to collect as much money as possible by pulling these arms over many turns. Initially, the distributions FDi are completely unknown to the player. At each turn, t = 1, 2,...n, the player selects an arm, with index j(t), and obtains a reward $r(t) \sim Dj(t)$. Since the player does not know in advance the distribution FDi, it has to test explicitly the i-th action with a trial-and-error search. Therefore, the player has two conflicting objectives: on the one hand, finding out which distribution has the highest expected value (or explore the distribution space); on the other hand, gaining as much rewards as possible while playing (or exploit its knowledge). Reinforcement Learning algorithms specify a probabilistic strategy by which the player should choose an arm j(t) at each turn. Clearly, the effectiveness of the bandit algorithm depends on how the gambler handles the exploration/exploitation dilemma when testing the different arms iteratively. Exploitation maximizes its reward at present time; at the same time, exploration may lead to a greater total reward in the future.

5.4.3 Learning Algorithm

The general multi-armed bandit formulation can be specialized as follows. First, in our problem each arm of the bandit corresponds to a different I_0 threshold. We consider I_0 as a percentage of requesting UEs to be covered by the multicast emission. Thus, K is the number of different thresholds chosen for multicast emission. It follows that FDi corresponds to the distribution of the amount of RBs that are used during the entire dissemination process when the i-th ranked UEs is used as I_0 threshold. More precisely, $Di = mi + Xi$, where mi is the fixed and known number of RBs that are used for a multicast transmission at the MCS of the i-th ranked UEs, and Xi is the random variable that models the total number of RBs used for the unicast transmissions during the panic zone. Note that Xi depends on many factors, including the set of seeds that are activated, the network topology and node mobility, as well as the dissemination strategy. In our case, each turn corresponds to the dissemination of a content (i.e., a packet). Once the nth packet with $I_0 = i$ is complete, the obtained reward is computed as:

$$\mu i\,(n) = 1/(mi + xi\,(n))$$

Where $xi(n)$ is the number of RBs that are used for the unicast transmissions in the nth panic zone. Note that the higher the number of used RBs and the lower the reward. To estimate dynamically the average reward $\mu i\,(n)$ for each value of I_0 we use a classical exponential moving average with rate α:

$$\mu i\,(n) = \alpha \mu i\,(n-1) + (1-\alpha)\mu i\,(n).$$

Now, we must define the policy to select at time n+1 the next I_0 value given the knowledge of the average rewards estimated at time n.

Different learning methods have been proposed in the literature for the armed bandit problems. The simplest one is the ε -greedy algorithm that selects with probability *(1 − ε)* the I_0 value with the maximum accumulated reward (greedy action), while it selects with probability one of the remaining I_0 values at random (with uniform probability) independently of the reward estimates (exploration action). More formally, let *πi(n)* be the probability to set $I_0 = i$ for the transmission of the nth packet, and *i(n) = argmax μi (n − 1)*. Then, in the -greedy algorithm it holds that *πi(n) = 1 − ε.*

Another class of learning algorithms is known as pursuit methods, in which the π probabilities are selected to strengthen the last greedy selection. Specifically, let *i(n)* be the greedy I_0 defined above. Then, just prior to selecting the CQI for the transmission of the nth packet, the greedy probability is reinforced as follows

$\pi i(n)\,(n) = \pi i(n)\,(n-1) + \beta[\pi_{MAX} - \pi i*(n)\,(n-1)].$

While all the non-greedy probabilities are updated as follows:

$\pi i(n)\,(n) = \pi i(n)\,(n-1) + \beta[\pi_{MIN} - \pi i(n)\,(n-1)].$

Here π_{MAX} , π_{MIN} are respectively the upper and the lower bound that the probability πi(n) (n) can take ∀i, n. In this way, the pursuit method is able to cope with the non-stationarity of the problem we considered, i.e. the distribution of rewards can change over time due to the underlying mobility, thus by restricting the range of the probability values from [0, 1] to [π_{MIN} , π_{MAX}] the exploration phase has always chance to be performed.

The principles behind the joint multicast/D2D approach are:

- At initial time, the eNB sends data to the best I_0 CQI-ranked UEs through a single multicast emission;
- The UEs that have received the data through the multicast emission start disseminating it in a D2D (epidemic) fashion;
- Before the maximum service delay D, we define a time interval, a panic zone where all the nodes that have not yet retrieved the content (either with the initial multicast emission or in D2D fashion) receive it through unicast cellular retransmissions.

The proposed scheme allows all UEs to receive data by the deadline (as long as the panic zone is sufficiently large). It adapts to different service delays -- the larger ones allowing for more D2D dissemination. Its performance relies on the multi-armed bandit algorithm that permits the cellular base-station to learn by experience the best transmission rate for each multicast emission.

Here resides the novelty of our approach: the eNB trades off the set of recipients that minimizes the multicast cost on the cellular network, while guaranteeing full coverage through D2D communications and panic re-injections when needed}. Next, we will determine the performance of the multi-armed bandit algorithm with the aid of simulations for different scenarios of utilization.

5.4.4 Performance Evaluation

5.4.4.1 Methodology

Synthetic mobility of users is implemented according to a Random-Waypoint model on a 200 × 200 sq.m. area. Nodes move in this space with speeds falling between 1 and 2.5 m/s (pedestrian speed). The synthetic mobility trace is the input of a packet level simulator. Indeed, we implemented the multi-armed bandit algorithm in the MOTO simulator (D5.2), which emulates the full LTE and Wi-Fi stack, allowing very realistic simulations. Since ns-3 does not natively support cellular multicast, we implemented an additional module that interacts with the packet scheduler to emulate single-cell multicast. The multicast module is installed at the eNB, and receives the CQI reports of UEs, deciding the transmission rate following the RL algorithm presented in Section 3.4. The network is composed of an eNB placed in the center of the interest area, a

remote server that provides the content, and multiple mobile devices. All mobile devices connect to the same eNB during the experiments.

We fix the bandwidth allocated for the multicast service at 5 MHz. LTE standard recommends not to reserve more than 60% of RBs to multicast, so the 5 MHz value could represent respectively the 50% or the 25% of RBs in a typical 10 or 20 MHz deployment. We simulate UDP constant bit-rate downlink flows, with packet size sk = 2048 bytes and a total load of 8 Mb. Additionally, we implemented DTN store-carry-forward routing mechanism at UEs to allow data forwarding on the Wi-Fi interface. Regardless of its reception method, an unexpired packet can be forwarded on the Wi-Fi interface upon meeting with neighbors. Neighbor discovery is implemented through a beaconing protocol triggered each 250 ms. UEs periodically broadcast beacon messages containing their identifier and the list of buffered packets. Upon beacon reception, UEs update their vicinity information and can transmit packets opportunistically.

All the results are averages over 10 independent simulation runs. Standard multicast implementation transmits data to all the UEs inside the cell using the MCS allowed by the lowest reported QI value. Even in that case, UEs have no assurance of reception. The radio channel could suddenly degrade during data reception (e.g., due to fast fading or mobility), preventing certain users to correctly decode data. For this reason, we consider an additional resilience layer in the form of panic zone retransmissions, which guarantee full dissemination at the cost of much higher resource consumption. The analysis will focus on aggregate metrics, as well as instantaneous values.

Implementation assumption: In simulation, we make the following simplification:

- HARQ-level retransmissions and RLC-level feedback are disabled in multicast. This is a reasonable assumption: otherwise, the eNB should merge the ack/nack messages received from all the UEs, and decide the best retransmission strategy. We guarantee data reception with panic zone retransmissions.
- The PUCCH channel is employed to acknowledge data reception towards the eNB. Panic zone retransmissions are then triggered looking at the list of received acknowledgments.
- The RL algorithm acts as a packet scheduler. The proposed approach employs a cross-layer design at the eNB. Exploiting signaling from physical layer (the amount of RBs utilized), it decides the modulation and coding of each multicast transmission. Moreover, CQI feedbacks are employed to dynamically rank of UEs.

We are aware that our simulation-based evaluation has some limits. First, we consider a simplified version of the eMBMS standard. The proposed approach requires deeper integration with the eNB scheduler. Moreover, we leave out the discussion on incentives that are vital to convince users to agree to spend their battery and storage resources to relay data to someone else.

5.4.4.2 Reference Strategies

Multicast-only is the basic strategy, where UEs have no direct connectivity options, and multicasting through the cellular infrastructure is the only means of distributing content. We compare this base case to our RL-based strategy. We assess the performance for three different values of N – the number of users inside the cell – respectively 10, 25, and 50, so to evaluate performance under different loads.

Fixed-best minimizes the number of RBs employed maintaining a fixed allocation of multicast users following the schema. Since the optimal fraction of multicast allocation is unknown, we ran several simulations to find experimentally this value.

ε-greedy method estimates the reward using the exponential moving average in Eq. 2. This simple algorithm selects the greedy value of I0 with maximum probability. In our implementation, we select = 0.05 and α = 0.5. We motivate this choice as a trade-off between different requirements. We need to maintain the exploration phase active in order to counter the non-stationarity of the underlying process. However, transmitting with a wrong CQI can lead to significant efficiency loss.

Pursuit method selects the I0 transmission probability following Eq. 3 and 4. In this case, the transmission probability pursues the greedy action. In the implementation, we choose β = 0.3, πMIN = 0.01, πMAX = 0.95.

5.4.4.3 Evaluation

UEs	D	Pursuit	Epsilon	Fixed–Best
10	30s	73	63	63
	60s	73	70	54
	90s	75	73	73
25	30s	69	12	70
	60s	80	80	71
	90s	83	84	83
50	30s	54	–28	55
	60s	55	–31	72
	90s	88	87	89

Figure 31. Steady-state levels of cellular offloading for the considered scenarios. Savings are referred to the multicast-only scenario (%).

Figure 31 provides a summary of the resource savings (aggregate over 1 hour, and at steady state) for the two considered algorithms, related to the basic Multicast-only approach. The RL solution to the joint multicast-D2D problem is an effective method to save resources at the eNB, approaching and even surpassing Fixed-best in more than one occasion.
Our system allows saving up to 88% of RBs for the 90 s scenario if compared to Multicast-only. This result confirms that the right synergy in the utilization of multicast and D2D resources allows for significant resource savings. Even with shorter deadlines, the pursuit method performs very well, saving at least 54 % of RBs. The main difference between RL-based methods and the benchmark represented by Fixed-best is that this latter is pre-computed with a fixed I_O selection made in advance. Its performance is stable over all the dissemination period, but this optimal value is the outcome of an extensive trial and error simulation phase. Conversely, a learning strategy is able, most of the times, to predict the most profitable values for I_O on its own, without relying on any prior knowledge and adapting to the network state.

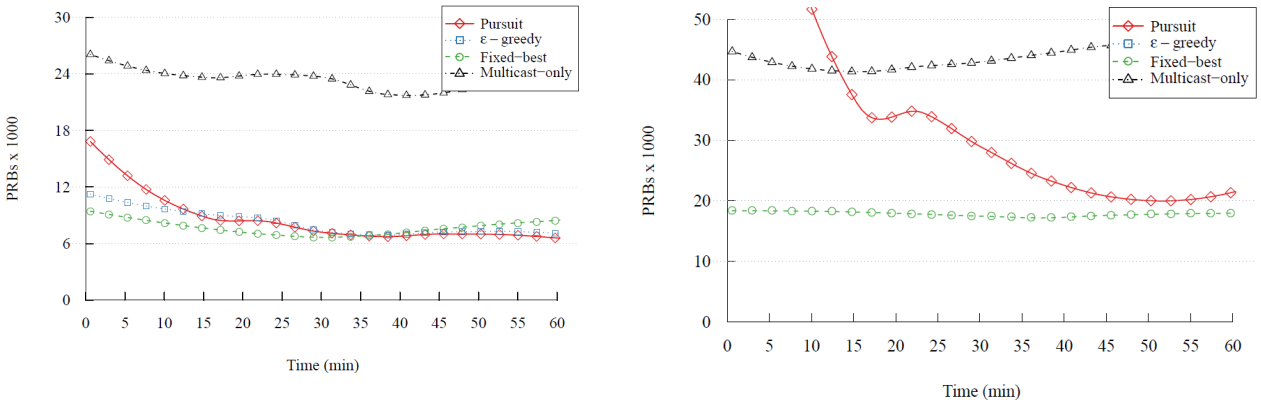

Figure 32. RBs usage for Multicast-only (black), ε -greedy (blue), Fixed-best (green), and pursuit method (red). 30 s. Left 10 UEs, right 50 UEs. Content is divided into 4000 packets of 2048 bytes. Plots are averaged over 10 runs, 95 % confidence intervals are not plotted but are knit.

We can observe this effect in Figure 32 for the tightest deadline considered (30 s). The behavior of RL methods (pursuit and ε-greedy) and Fixed-best are significantly different. Pursuit and ε-greedy are based on a learning algorithm; they therefore need time to learn the most appropriate value for I_0. Once trained, their performance is often on par or even better than the experimental optimal fixed-value strategy represented by Fixed-best. Another advantage is that even when those strategies are trained, they

continue to explore the solution space, being able to cope with the non-stationarity of the contact process that rules the opportunistic diffusion. Conversely, fixed-best is locked to a static value of the parameter I_0 and insensitive to variations in the mobility of UEs. One of the key advantages of the RL strategies are that they can autonomously find the optimal/suboptimal trade-off between multicast and D2D users in a reasonable time - without extensively search all the possibility for each time.
Given the heavy request for mobile data today, operators are mainly concerned about radio resource usage. To examine the impact on resource consumption, we fix the deadline at 30 s and vary the number of multicast UEs in the cell from 10 to 50. Intuitively, more UEs asking for the same content should require more infrastructure resources. On the other hand, the number of contact opportunities increases, offering more possibilities to offload the network.
Figure 32 gives hints on the actual amount of RBs devoted to distribute data in the considered scenarios. Unlike many other works in the literature, the use of the ns-3 simulator allows us to evaluate precisely the amount of radio resources consumed at the eNB.
In general, we note that when it converges, the ε-greedy method is faster than the pursuit method. In many cases, when ε-greedy fails to converge, pursuit reaches a nearly optimal value. This depends on the fact that ε-greedy always selects the value of I_O that maximizes the expected rewards. Instead, pursuit has an indirect selection method that better adapts to the temporal evolution of the system. The added complexity is however beneficial in most cases, as it results into an improved performance. The reinforcement allows smoothing out the inherent variations in epidemic diffusion that prevent the proper prediction in the ε-greedy method. The effect appears when the number of targeted UEs increases while keeping a tight deadline (i.e., 30 s). In those scenarios, the variability in performance of the opportunistic diffusion prevents the ε-greedy method to learn properly the best value for I_0.
We draw the lesson that the ε-greedy method owing to its simplicity does not fit well scenarios with a lot of variability. For those cases, the pursuit method is a better match. On the other hand, in scenarios where the variability of the opportunistic process is low -- i.e., when the deadline is large -- the ε-greedy approach allows for a quicker convergence time.

We plot in Figure 33 the fraction of packets partitioned by their reception method. Considering the same deadline and increasing the number of UEs has the effect of reducing the impact of D2D transmissions. While a larger number of UEs should multiply the contact opportunities, many of them are not adequately exploited because UEs can transmit only to one neighbor at a time. The result is that fixing the deadline, the contribution of the opportunistic diffusion is upper bounded. This is the inner process that the RL method should learn.

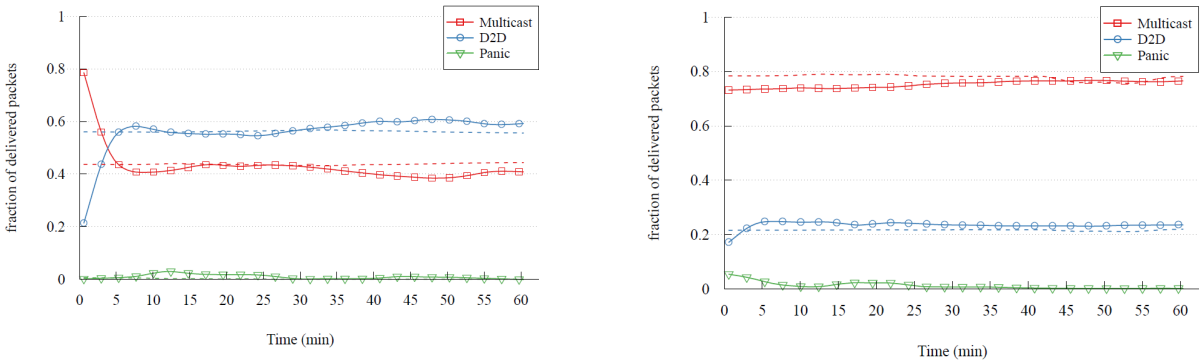

Figure 33. Pursuit method, reception method. 30 s, 10 UEs (left) and 50 UEs (right). Dashed lines are the objective ratio for Fixed-best. Content is divided into 4000 packets of 2048 bytes. Plots are averaged over 10 runs, 95 % confidence intervals are not plotted but are knit

We also take note of a peculiarity. Looking at the reception methods in Figure 33, the convergence time looks like always less than 10 minutes. Comparing this value to Figure 32 however, we realize that the actual convergence (in terms of RBs employed at the eNB) happens much later in time (around 40 min). This anomaly is motivated by the fact that even a small amount of unicast retransmissions in the panic zone

consumes many more resources than the multicast emission. The fine-tuning required to reach an optimal RBs usage level is thus the responsible for this longer convergence time. When the number of UEs is large, the choice of the appropriate values for I0 becomes fundamental in order to avoid congesting the cell with too many panic retransmissions.

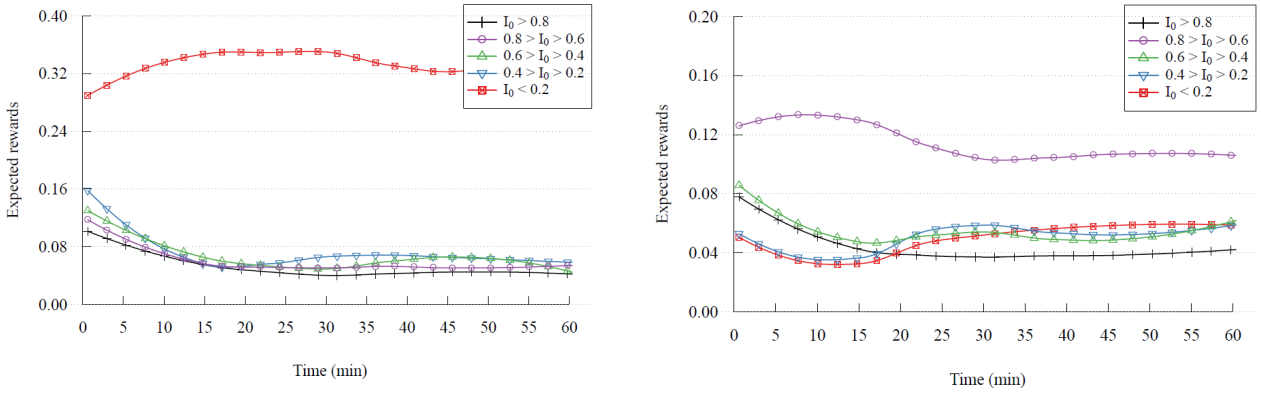

Figure 34. Pursuit method, average reward values for I_0. 30 s, 10 UEs (left) and 50 UEs (right). Content is divided into 4000 packets of 2048 bytes. Plots are averaged over 10 runs, 95 % confidence intervals are not plotted but are knit.

Considering the inner working of the pursuit method Figure 34 compares the rewards and emission probabilities respectively. In two cases out of three (namely for 10 and 50 UEs), there is a set of values for I_0 that performs clearly better than the others do. In the figures, we gathered the value of I0 to form five levels to improve the understanding of the curves. In the 25 UEs scenario instead, the best value for I0 fluctuates and no clear winner emerges. The best I_0 value is the one that is not affected too much by the loss in spectral efficiency due to the reduced multicast rate, but at the same time can guarantee a low penalty due to unicast panic re-injections. In the 10 UEs scenario, emitting with a multicast rate that targets one or two users is sufficient to achieve high efficiency. On the other hand, we note that increasing the multicast group size, the best value for I_0 increase. Intuitively, the penalty due to panic re-injections risk being extremely severe in these scenarios and the pursuit algorithm tends to allocate more seeds in the opportunistic domain.

The good functioning of the learning algorithm depends partially on the tuning of its parameters. This is the principal reason why in simulation we considered two different strategies, a simple ε-greedy and the more complex pursuit method. While pursuit highlights better overall performance, it needs a careful tuning of the parameters of learning. Beside α, which is the exponential moving average parameter used also in ε-greedy, pursuit needs to define also π_{MAX}, π_{MIN}, and β. In our application, we found that the value of π_{MIN} is particularly critical, as it affects the likelihood at which the algorithm explores the non-greedy strategies. Fixing too high values of π_{MIN} is most of the times unfavorable, as the learning algorithm emits at a wrong rate too many times. In particular, this results catastrophic when the optimal value of I0 is high, because of the too many panic retransmissions.

We tested different learning strategies, trying to learn various parameters beside the value of I_0 - the ratio of UEs to insert in the multicast group. We performed several tests trying to learn the best emission CQI value instead of the percentage I_0. After several failed attempts, it was realized that trying to learn the optimal CQI value is very complex, as its optimum value fluctuates following the instantaneous position of users and the experienced fading. Since the underlying opportunistic diffusion process depends strictly on the amount of seed UEs (tunable following I_0) and content deadline (fixed a priori), it is much more efficient to learn directly the percentage of seed UEs I_0.

5.4.4.4 Conclusion

In this work, we have presented a hybrid distribution strategy for popular content with guaranteed delays. Multicast is an advantageous option to distribute popular data into a cellular network. However, performance is determined by the UE with the worst channel quality inside the multicast group. We proposed a framework that exploits D2D capabilities at UEs to counter the inefficiencies of cellular multicast.

The proper balance of multicast and D2D transmissions is achieved using a multi-armed bandit learning strategy. We proposed and evaluated two different algorithms under variable multicast group size and reception deadline. Simulation results prove that D2D communications allow increasing the multicast transmission rate, saving resources and improving the overall cell throughput. At the same time, the analysis demonstrates that both algorithms have a reasonable convergence time.

6 Intra-technology scheduling: Towards energy efficiency in the LTE network

In the past decades, the design of cellular networks was primarily aimed at delivering the maximum performance in terms of coverage and/or capacity performance neglecting the energy efficiency of the solutions adopted in the core or the access parts of the network. Energy efficiency in cellular networks has become a central point in recent research efforts as it became evident that the operation of this sector is responsible of a non-negligible part of the energy bill of an operator. Furthermore, the access part of the cellular networks has drawn significant attention being the biggest contributor to this energy footprint. Consequently various approaches have been employed to improving its energy efficiency, including putting to sleep the electronic components in the eNB (or the entire eNB) whenever possible combined with dynamic coverage adjustment, data offloading through D2D communications and so on. However, the problem of improving the energy efficiency of eNBs is complicated by the recent adoption in LTE networks of a multi-tier architecture, in which a mix of macrocells and smaller cells (namely microcells, picocells and femtocells) coexist and cooperate. Indeed, the deployment of heterogeneous LTE networks is seen as a cost-effective solution to ever-growing traffic demands, because small cells have lower implementation costs, use less expensive equipment and consume less energy than traditional macrocells. On the other hand, the introduction of small cells may potentially lead to increased ICI (Inter Cell Interference) due to intense frequency reuse in neighbouring or overlapping cells. For this reasons, LTE standards envisage the use of various techniques, such as ICIC (Inter-Cell Interference Coordination) to mitigate the negative effect of the interference between overlaying macro- and microcells.

In this section, we use the ns3 simulator extended with LENA module to explore the trade-off between capacity increase and energy efficiency in a heterogeneous LTE cell in which a macrocell coexists with multiple outdoor picocells. Moreover, we have also considering the possibility of using D2D communications to further offload traffic from the cellular network. We assume that in each cell there is an *energy scheduler* that decides which eNBs to activate and when. Then, we investigate different strategies for switching off eNBs based on network status without degrading the overall cell capacity. Note that in the considered network scenario we also assume that the picocells employ Cell Range Extension (CRE) techniques to offload data from nearby macrocells. Such technique consists in adding a positive range expansion bias to the pilot downlink signal strength received from picocell so that more users connect to them. A novelty of our study is that we consider different requirements for different data types. In particular we assume the existence of **delay tolerant traffic**, which is amenable to be offloaded using terminal-to-terminal communications. Moreover, we evaluate the joint effect of an energy scheduler and D2D operations to save energy in the LTE core network.

In initial study we consider energy saving at three levels. On the one hand, we compare configurations with a single macrocell with others where the macrocell uses a lower transmit power and exploits picocells to serve additional UEs. On the other hand, we consider the effect of policies that switch off some picocells (those that are serving fewer users). In the latter case, we also consider the impact of using also D2D communications to disseminate traffic, to further reduce the the traffic carried by the pico-cells, and thus be able to switch them off more aggressively. Our results show that, in general, reducing the transmit power of the macro cell and complementing this with picocells is viable, as the overall throughput provided to the users is not negatively affected. The additional gain of switching off picocells is still to be better investigated. Interestingly, our results show that when D2D is used, the total throughput increases quite significantly, but this does not necessarily translates in a linear increase of the number of switched off pico-cells. Therefore, additional gains in energy savings are expected by dropping the throughput increase provided by D2D in favour of the possibility to switch-off more pico cells.

With respect to what already presented in D3.3.1 [3] and in the intermediate version of D3.3.2, we have significantly extended the preliminary simulation environment used for D3.3.1, and we have completed the simulation studies using D2D.

6.1 Power consumption model

We use the power consumption model of a BS with the e3F model introduced by Auer et al. [36]. This model provides a relation between P_{out} (the output power radiated by antenna) and P_{in} (the total power needed by the eNB to operate) for different types of eNBs in an LTE system with 10 MHz of bandwidth and 2x2 MIMO antenna configuration. The energy model is well approximated by this linear model:

$$P_{in} = \begin{cases} N_{TRX}(P_0 + \Delta_P P_{out}) & 0 < P_{out} < P_{max} \\ N_{TRX} P_{sleep} & P_{out} = 0 \end{cases}$$

where N_{TRX} represents the number of transceiver chains, P_0 represents the power consumption at the minimum non-zero output power (empty eNB), Δ_P is the slope of the load dependent power consumption, P_{max} represent the maximum transmission power achievable by the BS and P_{sleep} represents the power consumption of the eNB in sleep mode. Table 5 reports the typical values for the parameters of the power consumption model for different types of eNBs.

Table 5: Power model for different eNB types

eNB type	N_{TRX}	P_{max}[W]	P_0[W]	Δ_p	P_{sleep}[W]
Macro	6	20	130.0	4.7	75.0
RRH	6	20	84.0	2.8	56.0
Micro	2	6.3	56.0	2.6	39.0
Pico	2	0.13	6.8	4.0	4.3
Femto	2	0.05	4.8	8.0	2.9

6.2 Energy scheduler simulation set-up

The simulations were performed using ns3 simulator extended with the LENA module for LTE. The simulation parameters are reported in Table 6, while Figure 35 illustrates the network scenario that we have used in the simulations. Specifically, we consider a heterogeneous network in which there is one macrocell and a varying number of picocells are deployed within the macrocell. Furthermore, we consider two different types of deployment for the picocells. The first one is a *total random deployment* in which no constraint is set on the location of the picocell within the macrocell. The second one is a *planned deployment*. In particular, to avoid excessive interference from the macrocell to the picocels we avoid deploying picocells at a distance shorter than 200 meters from the centre of the macrocell. Furthermore, we also avoid deploying two picocells at a distance shorter than 150 meters. Note that in the simulations we have used the *Hybrid Building Propagation Loss* model that is used in ns3 to model propagation losses due to the presence of different types of buildings (i.e., residential, office and commercial). More details on the implementation details of this propagation loss model are provided in Section 6.2.2.

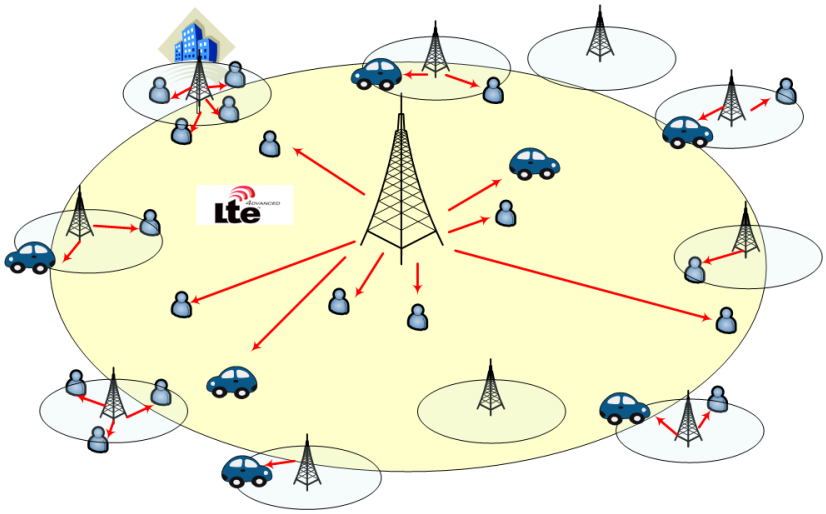

Figure 35: Network scenario

To evaluate the benefits of introducing an energy scheduler within a heterogeneous network we have chosen the following metrics.

- *Overall network throughput:* the total throughput measured over the entire network including macro- and small-cells
- *Network energy consumption:* the total energy consumed by the eNBs.

We remind that our ultimate goal is to design a scheduling policy that achieves a good comprise between power consumption reduction and network capacity degradation. Thus, before presenting the simulations results we describe in greater details: *a)* how UEs are connected to macrocells and picocells, and *b)* the algorithm that we have used to decide when a picocell should be switched-off and to which picocell to move its associated UE.

Table 6: Simulation parameters for energy efficiency in LTE core network

eNodeB type	Macro	Pico
Bandwidth	5Mhz	5Mhz
Tx Power	46, 43, 36 dbm	28, 24 dbm
Antenna	parabolic	omnidirectional
$bias_{CRE}$ value	0	10 db
Number of antenna	3	1
Cell layout	Hexagonal, ISD =500m	0,6,8,10 per macrocell sector
Min dist to picocells	200 m	150
Ue max speed	54 km/h	
Mobility Model	Random Waypoint	
Pathloss	Urban/Buildings	
BS Distruibution		Random
Data Rate	8.2Mb/sec	
UE Distribution	Random (density per m^2 0.0002, 0.00015, 0.0001, 0.00005)	

Packets dimension	1024 KBytes
Simualtion time	20 s

6.2.1 Switch-off procedures

At the beginning of each experiment each UE is connected to the closest eNB. Then, standard handover procedures implemented in the LENA module are activated and UEs automatically hand over towards the eNBs to which they receive the highest RSSI. After the initial configuration phase the switch-off algorithm decides which are the picocells to temporarily deactivate by looking at the network load distribution. Specifically, picocells with less than 4 users are switched off and their users are handed over to the closest picocell that has less than 25 users already connected. Note that after this tentative association of UEs to picocells based on network load, the automatic handover procedures implemented in ns3 are used to trigger additional handovers towards picocells with better SINR. For the sake of clarity, the various phases of the switch-off procedures are also illustrated in Figure 36. Note that in this scenario we have introduced a bias CRE value equal to 10 dBm, to enable picocells to capture more traffic and to favor macrocell offload.

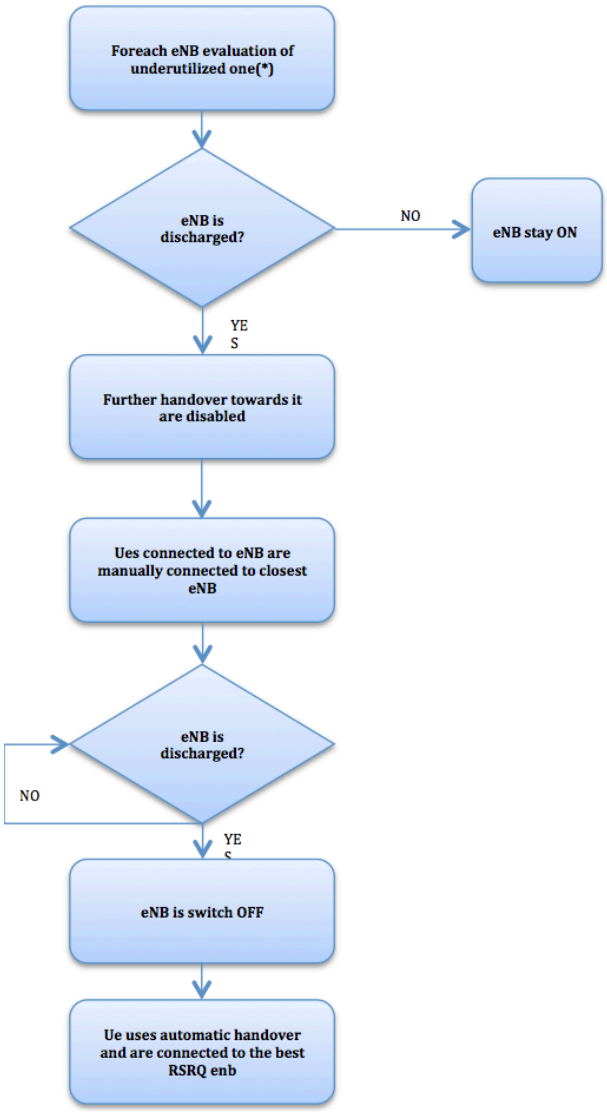

Figure 36: Switch-off procedures.

Next, we provide an outline of the implementation of the automatic handover procedures in the LENA module. Specifically, the RRC (Radio Resource Control) model implemented in the simulator provides handover functionality through the shared X2 interface. There are two ways to initiate a handover procedure:

- The handover could be triggered explicitly by the simulation program by scheduling an execution of the method *LteEnbRrc::SendHandoverRequest()*
- The handover could be triggered automatically by the eNB RRC entity. The eNB executes the algorithm Figure 37 for a UE providing measurements in its serving cell and the neighbour cells the UE measures:

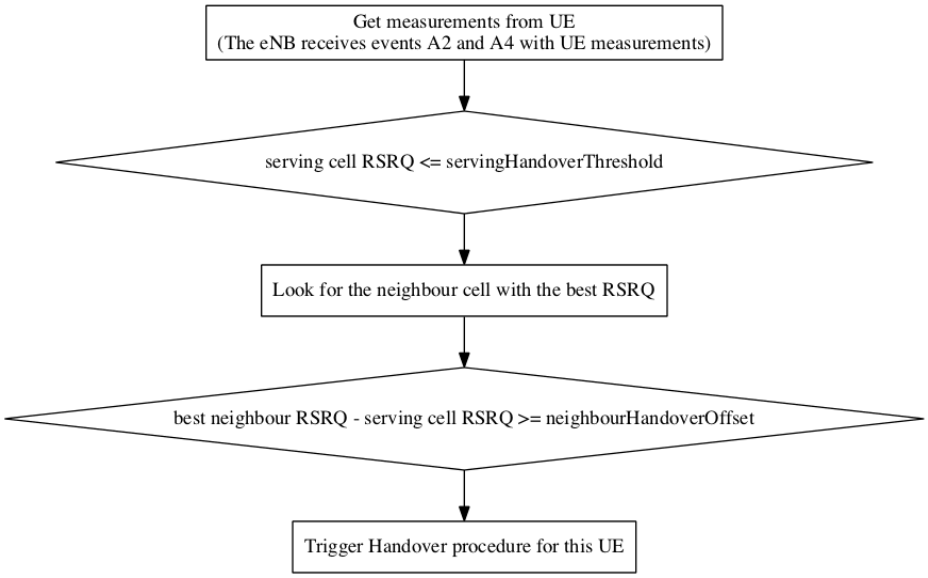

Figure 37:Algorithm to automatically trigger the Handover procedure

Furthermore, the following parameters can be adjusted to control the handover decision process:

- *servingHandoverThreshold*: if the RSRQ (Reference Signal Received Quality) value measured by the UE in its serving cell is less or equal to the *servingHandoverThreshold* parameter (i.e. the conditions of the UE in the serving cell are getting bad or not good enough), then the eNB considers this UE to hand it over to a new neighbor eNB. The handover will eventually be triggered depending on the measurements of the neighbor cells.
- *neighbourHandoverOffset*: if the UE is considered for handover, and the difference between the best neighbor RSRQ and the RSRQ difference between the neighbor and the serving cell is greater or equal to the *neighbourHandoverOffset* parameter, then the handover procedure is triggered for this UE.

The X2 interface interconnects two eNBs though a point-to-point link between the two eNB. The X2 interface implemented in the LENA module provides detailed implementation of the following elementary procedures of the Mobility Management functionality:

- Handover Request procedure
- Handover Request Acknowledgement procedure
- SN Status Transfer procedure
- UE Context Release procedure

The above procedures are involved in the X2-based handover. We note that the simulator model currently supports only the *seamless handover* while *lossless handover* is not supported. Generally speaking, seamless handover is used for channels transporting traffic that is very sensitive to delay and jitter and would rather accept retransmissions than delay, like VoIP, while the lossless handover is used for channels that transport traffic that does not care too much about delay but is sensitive to retransmissions, like FTP and HTTP. Figure 38 shows the interaction of the entities of the X2 model in the simulator.

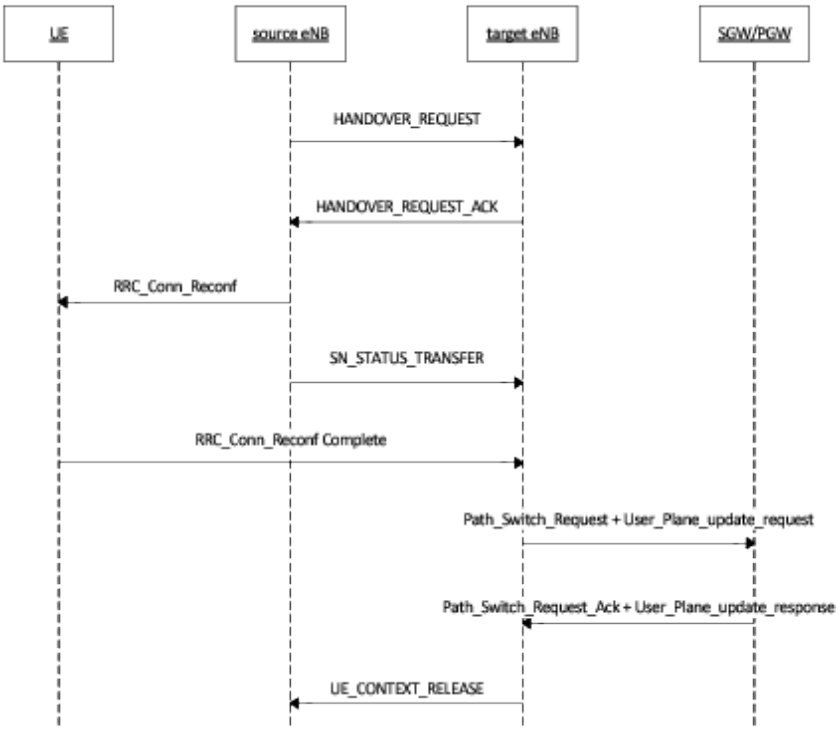

Figure 38: Sequence diagram of the X2-based handover

6.2.2 Buildings Module

The Buildings module used in our scenario and implemented in ns3, provides:

1. a new class (*Building*) that models the presence of a building in a simulation scenario;
2. a new class(*MobilityBulidingInfo*) that allows to specify the location, size and characteristics of buildings present in the simulated area, and allows the placement of nodes inside those buildings;
3. a container class with the definition of the most useful pathloss models and the correspondent variables called *BuildingPropoagationLossModel*.
4. a new propagation model (*HybridBuildingsPropagationLossModel*) working with the mobility model just introduced, that allows to model the phenomenon of indoor/outdoor propagation in the presence of buildings.
5. a simplified model working only with the Okumura Hata propagation model (*OhBuildingsPropagationLossModel*), which consider the phenomenon of indoor/outdoor propagation in the presence of buildings.

The models have been designed with LTE in mind, though their implementation is in fact independent from any LTE-specific code, and can be used with other ns-3 wireless technologies as well (e.g., wifi, wimax).

*HybridBuildingsPropagationLossMode*l pathloss model included is obtained through a combination of several well-known pathloss models in order to mimic different environmental scenarios such as urban, suburban and open areas. Moreover, the model considers both outdoor and indoor indoor and outdoor communication has to be included since HeNB might be installed either within building and either outside.

In case of indoor communication, the model has to consider also the type of building in outdoor <-> indoor communication according to some general criteria such as the wall penetration losses of the common materials; moreover it includes some general configuration for the internal walls in indoor communications.

OhBuildingsPropagationLossModel pathloss model has been created for simplifying the previous one removing the thresholds for switching from one model to other. For doing this it has been used only one propagation model from the one available (i.e., the Okumura Hata). The presence of building is still considered in the model; therefore all the considerations of above regarding the building type are still valid. The same consideration can be done for what concern the environmental scenario and frequency since both of them are parameters of the model considered.

6.3 Simulation Results

In the following figures we report a subset of the most interesting simulation results. Figure 39 shows the total throughput with 8 picocells in the following configurations:

1. 1 macrocell transmitting with a power equal to 43 dBm (red line)
2. 1 macrocell transmitting with a power equal to 36 dBm and 8 picocells with a *random* distribution and transmitting with a power equal to 24 dBm (green line)
3. 1 macrocell transmitting with a power equal to 36 dBm and 8 picocells with a *planned* distribution and transmitting with a power equal to 24 dBm (blue line).

In case 2 and 3 the throughput is lower than the one obtained in case 1 if the number of UEs is small. Indeed, in this case the bandwidth of the macrocell is sufficient to serve all the users and adding the smallcells negatively affects the level of interference in the cell. Furthermore, if the number of picocells is small the probability that a UE has a better connection with a picocell than with the macrocell is typically low. An important finding of Figure 39 is also that a planned deployment of picocell inside the macrocell is beneficial to reduce the ICI, and intuitively this effect is more evident for high number of UEs. For these reasons, in the following we only show results for planned picocell deployments and for larger number of picocells, which provide the most interesting and significant results.

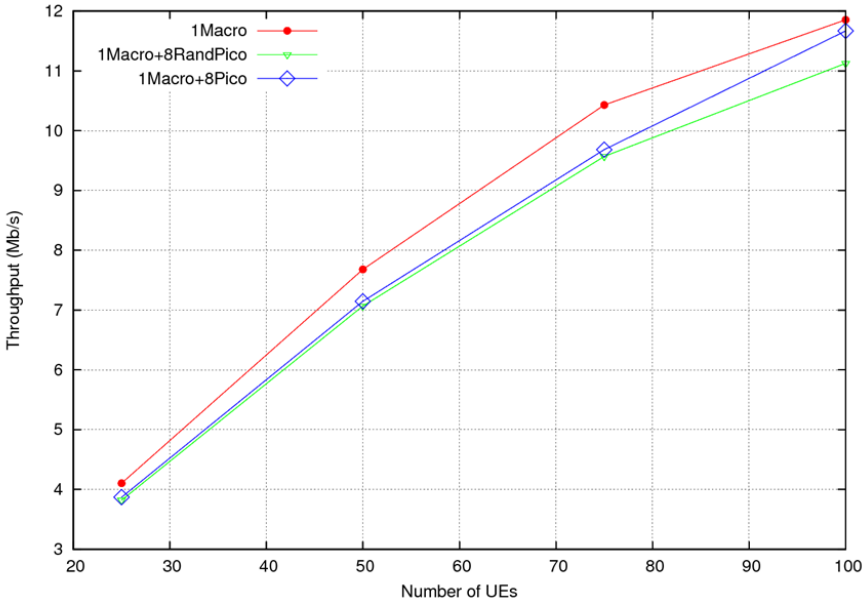

Figure 39: Throughput measured with 8 picocells.

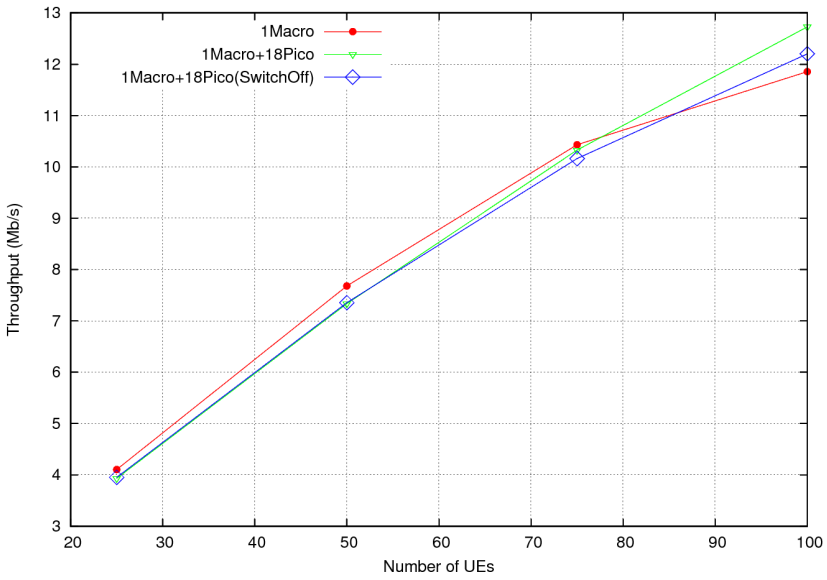

Figure 40: Throughput measured with 18 picocells with and without switch-off of lightly loaded picocells.

In Figure 40 we show the total cell throughput in a network scenario with 18 picocells with and without switch-off of lightly loaded picocells. We can observe that even if the transmission power of the macrocell has been reduced from 43dBm to 36 dBm the overall throughput provided to the users increases due to the additional capacity provided by the picocells. Moreover, in the considered scenario, switching off lightly loaded picocells (i.e., picocells with less than 4 UEs) has a negligible impact on the throughput. Finally, Figure 41 shows the energy consumption in the same network scenarios of Figure 40. Two important conclusions can be derived from the plots. The first one is that in a heterogeneous network with a macrocell and multiple picocells the energy consumption can be significantly less than in a network without picocells because in the former case the macrocell can use a significantly lower transmission power than in the latter case. The second observation is that the additional energy gain due to switching off picocells is way lower than that due to reducing the transmission power of the macrocell. This is due to the fact that the picocells are low power nodes and their energy consumption can be an order of magnitude less than the one of a macrocell. It is also important to point out that a more aggressive scheduling strategy might entail to also switch off the macrocell. This could provide a much higher energy reduction with respect to the case of switching off only the picocells, but at the cost of reduced coverage. On the other hand, when the macro cell cannot be switched off, and its transmission power cannot be further reduced, switching off picocells can still provide a noticeable benefit in terms of lower energy consumption.

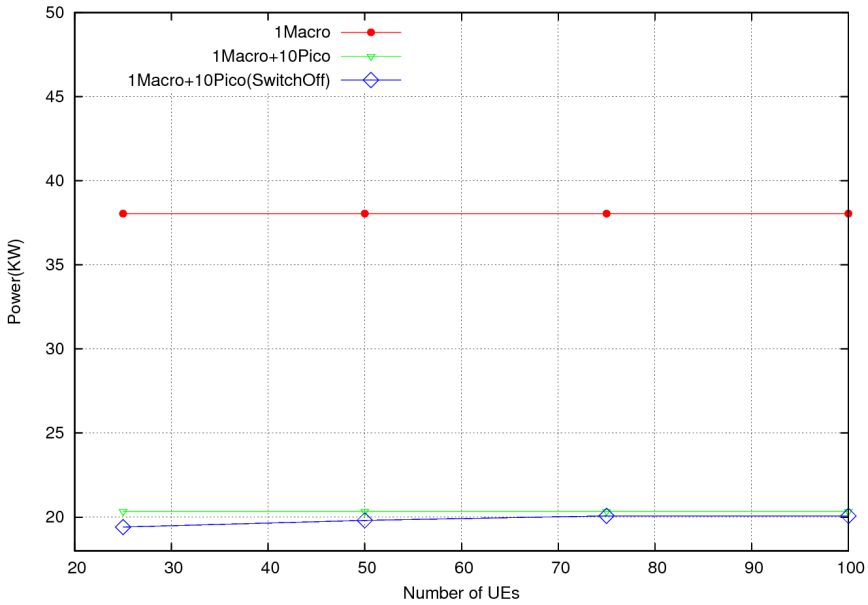

Figure 41: Power consumption with 18 picocells when switching off lightly loaded picocells

In conclusion, this part of the study has shown that it is possible to reduce the transmit power of macrocell without negatively affecting the throughput performance experienced by users, as the effect of a reduction in the macrocell transmit power is compensated by additional picocells. On the other hand, to switch off lightly loaded picocell also does no negatively affects the overall capacity of the LTE network. Furthermore, we have also shown that the throughput obtained by deploying picocells is heavily dependent on the deployment strategy of picocells (random vs. planned deployments) and the density of UEs. Clearly, the efficiency of the switch-off strategy depends on several parameters (e.g, CRE bias, macrocell transmission powers, thresholds to decide to switch off a picocell) and it would be possible to identify an optimal configuration of such parameters.

6.4 D2D Simulation scenario

In this set of simulations we compared the performance, in terms of energy gain and network efficiency, of the combined network with and without D2D offloading and with the energy scheduler.

In a real environment we couldn't think to have only MOTO users or DTN contents to diffuse using Wifi ad-hoc connections, therefore we configure our scenario, simulating the shopping mall use case, using two different traffic sources together:

- *Source A*: UDP constant bit-rate, only in downlink, where users use LTE to communicate with a remote host, with packet size of 1024 bytes and source rate 10240 B/s per user. The total amount of UDP traffic for a 20 seconds simulation is 10240*8*nUes*20 seconds = 1.64Mb*nUes, where nUes is the number of users receiving traffic from this source.
- *Source B*: The MOTO CDM data source generates a DTN flow to be spread using D2D connections between terminals, using Push&Track. Therefore, epidemic routing is used in the D2D network, while LTE is used for initial and periodic seeding, and to serve content in the panic zone. The traffic is made up of 8 packets of 25KB each, which have to be delivered to all the users. The total amount of traffic is again 1.6Mb*nUes.

For the simulations we consider a single LTE hexagonal cell with inter-site distance of 500meters, a variable number of UEs (respectively 25, 50, 75, 100) randomly distributed within the cell. The different number of UEs allows us to evaluate the performance of the algorithm under different loads. Each cell is served by a single central macrocell and 18 picocells randomly distributed but with distance constraints: the minimum distance between picocells is 90 meters and the minimum distance between macrocell and picocell is 200

meters). Users move according to the random waypoint mobility model, with variable speed from 0 to 54km/h and pause-time set at 2s. All the results presented in this section are averages over 15 independent simulation runs.

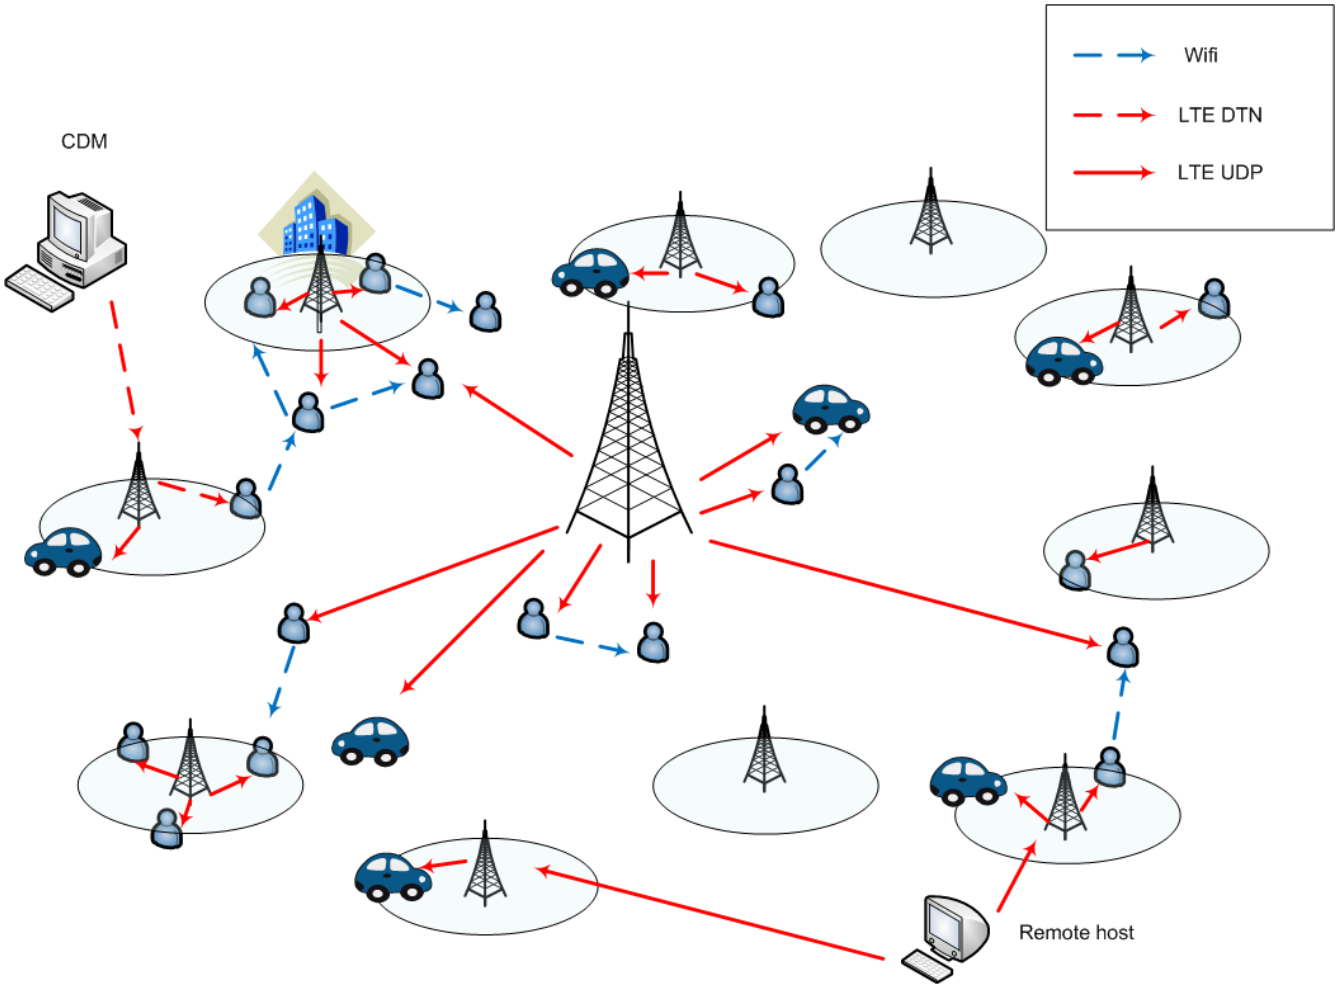

Figure 42: Simulation Scenario (D2D & energy scheduler)

We tested this scenario in four configurations, as follows:

- *First configuration*: NO energy scheduler - NO D2D. With this configuration we evaluate network efficiency as a function of the network load, without any of the mechanisms considered in MOTO. This is thus the benchmark to see the advantage of the proposed additional mechanisms (energy scheduler and D2D offloading).
- *Second configuration*: ENABLE energy scheduler - NO D2D. With this configuration we evaluate the throughput as a function of the network load, and the number of picocell that could be switch-off without influencing the perceived throughput. This basically corresponds to the mechanism described in Section 6.2.1.
- *Third configuration*: NO energy scheduler – ENBLE D2D with this configuration we evaluate the decrease of network load using the offloading algorithm.
- *Fourth configuration*: ENABLE energy scheduler - ENBLE D2D with this configuration we evaluate the decrease of network load, and the number of picocells that could be switched-off without influencing perceived throughput. In this last scenario users receive traffic from Source B through

D2D communications (assisted by LTE ad in the Push&Track algorithm), and the traffic from Source A through the LTE channel. In this way we want to test to what extent using D2D offloading allows us to switch off a greater number of picocells without influencing perceived throughput.

Note that in this set of simualtions picocells are switched off if they carry a traffic below a certain threshold. This threshold is inversely proportional to the total number of picocells.

6.4.1 Simulation results

Figure 43 shows the number of pico-cells that can be switched off when D2D is or is not used, respectively, as a function of the number of UEs. In Figure 43 we can see that the number of picocells that can be switched-off grows up in function of the number of Ues and this number always increases using the D2D ad-hoc communication with respect to the case with only LTE communication.

As we can see we have the best result with 25 users connected to the network using epidemic diffusion for Source B data in terms of energy efficiency, with 14 switched-off picocell. The fact that the number of pico-cells increases when the number of UEs is low is expected, as in this case there is a higher probability that the total traffic carried by the picocells is below the threshold.

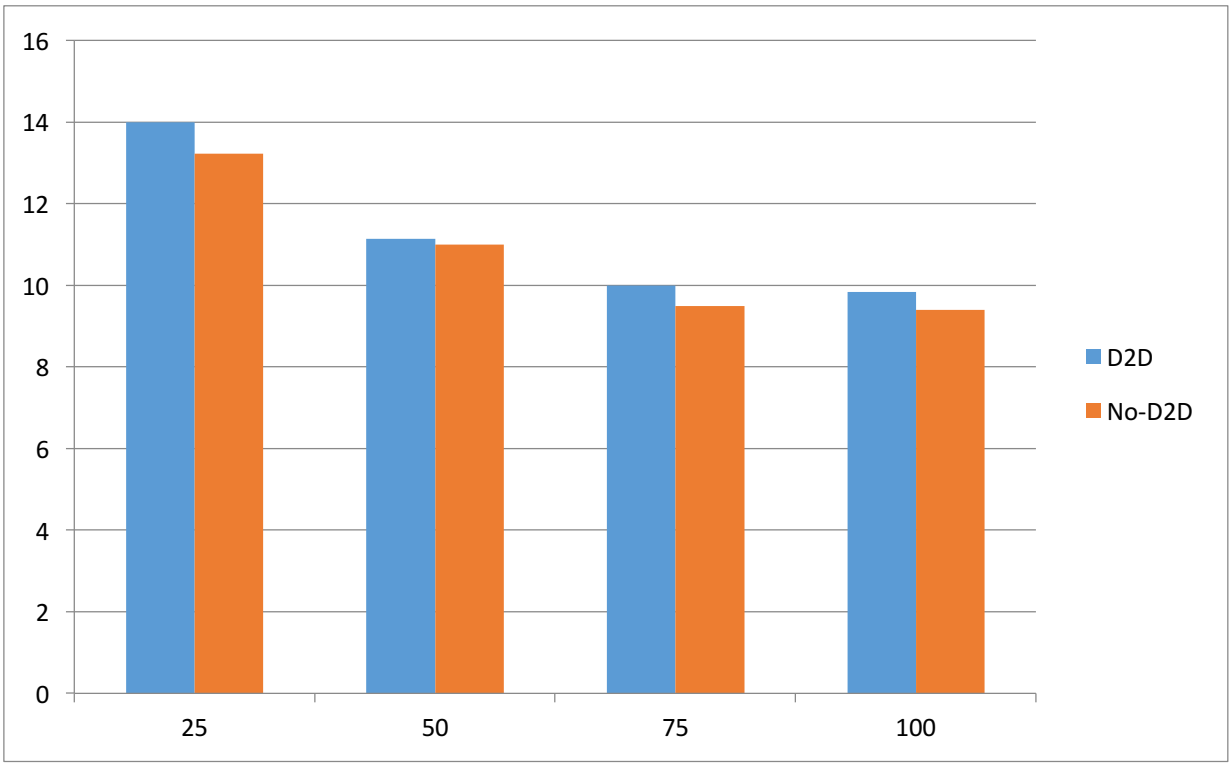

Figure 43: Number of switch-off picocells in function of the number of UEs

It might seem surprising that the additional number of picocells that can be switched off when D2D is used is relatively low. However, this is a side effect of the additional throughput received by users. Specifically, in Figure 44 we show the total network throughput considering the messages delivered using LTE network and D2D communications. Clearly, when D2D is used, the throughput increases quite significantly. This has an small effect also on the LTE traffic, because the additional traffic delivered through D2D also generates some control traffic over the LTE network (remember the Push&Track mechanisms). This results in an increase of the number of switched-off picocells lower than one could expect. However, we expect than by controlling the throughput received by users such that it is the same in both cases (with and without D2D), the traffic carried by the LTE network will reduce, and therefore the number of switched-off picocells could increase.

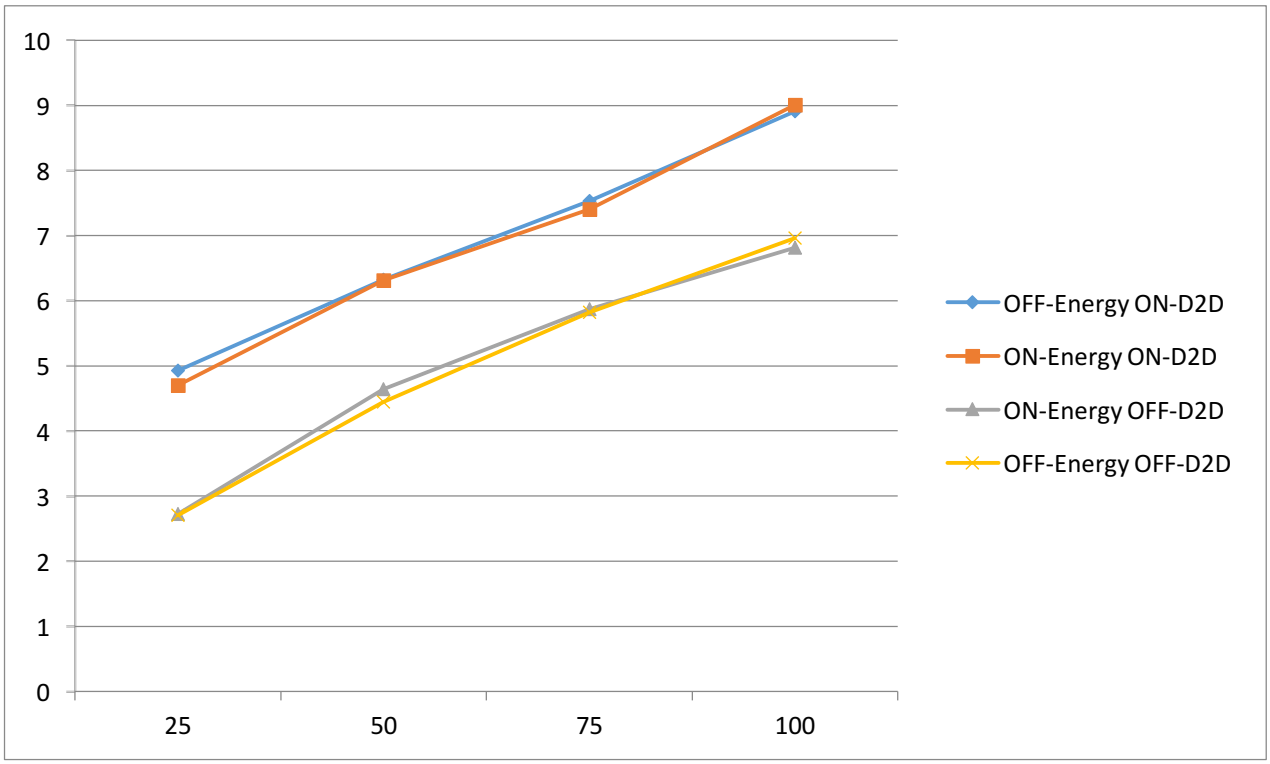

Figure 44: Network throughput (LTE+D2D) in function of the number of UEs (Mb/s)

Finally, in Figure 45 we show the results in terms of energy consumption (considering only network infrastructure consumption.

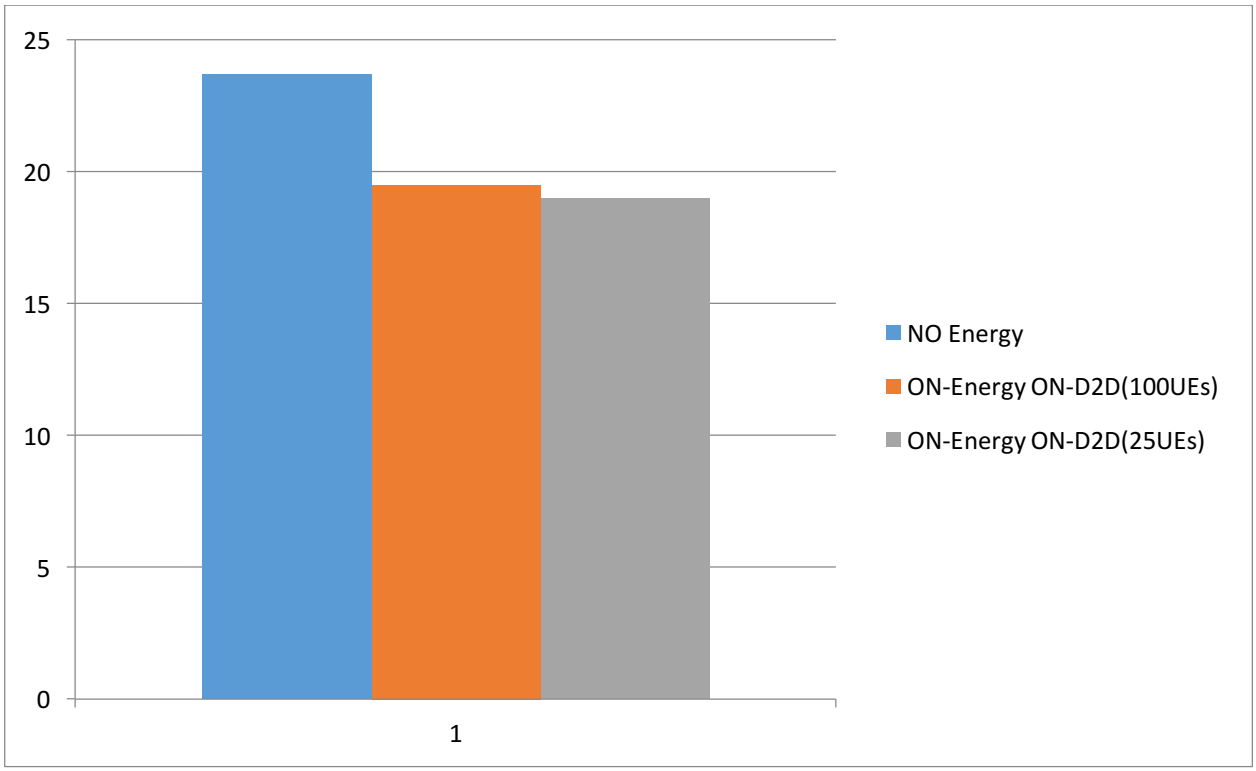

Figure 45: Power Consumption (KW) in 20 seconds simulation

Looking together at Figure 44 and Figure 45, we see that using D2D epidemic diffusion for an half of the traffic, the network throughput is increased of up to 70% (almost twice than the one without D2D, both with and without Energy saving algorithm). At the same time the "Energy Saving" algorithm doesn't affect the perceived throughput, but even manages to slightly reduce the energy consumption in the LTE core network. Also, a small saving for 20 seconds simulation (5KW) is nevertheless a good result, considering the length of the simulation. In fact, we can expect that the energy scheduler mechanism could result in a greater saving for traffic profiles spanning longer time intervals.

7 Inter-technology scheduling: Multi-user offloading in heterogeneous wireless network infrastructures

In this section we present the results of the work on intra-technology scheduling. This activity allowed us to better understand the option of using multiple wireless infrastructures to offload traffic from possibly congested cellular networks.

Specifically, we consider the case where users may either be served by an operator through a cellular or a WiFi infrastructure, or device-to-device communications and investigate how to best schedule users across these technologies. A sample generic scenario is shown in where two access technologies are depicted. In this basic scenario all users are assumed to possess D2D communication capability as well. The objective of this work is to investigate handover decision making algorithms targeting the users under multiple access technology coverage and point out the metrics and factors influencing data offloading in the presence of D2D communications.

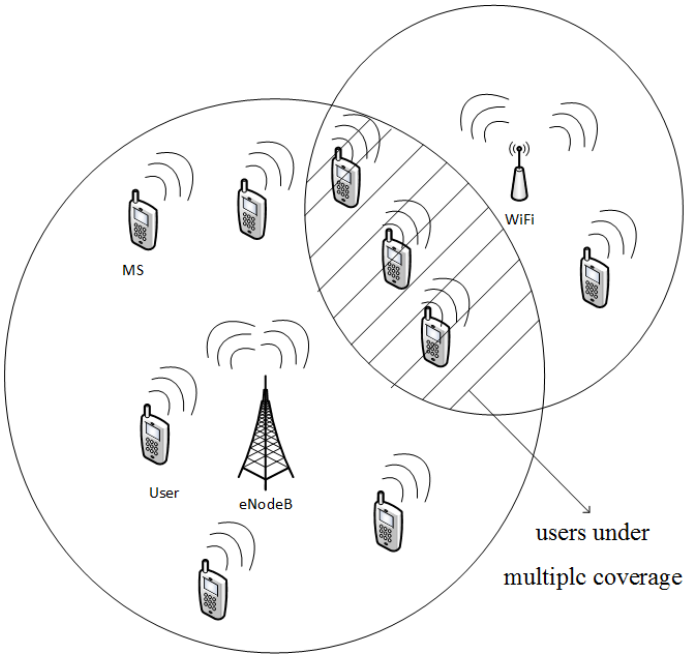

Figure 46. A sample user distribution map under multiple wireless technology coverage

To this extent, two multi-user multiple attribute decision making (MADM) algorithms have been developed. These algorithms are built on the TOPSIS framework [46] adapting it to the case of offloading across wireless infrastructure. This is basically an optimization framework where we can express decision criteria based on QoS metrics. We thus identify appropriate QoS metrics for our case, and show how to customize the optimization framework for our purposes, with specific reference to real-time traffic. Specifically, in the first algorithm we consider a straightforward application of TOPSIS, where we optimize individual users' performance. On the other hand, in the second algorithm we take into consideration network-wide performance by integrating the available capacity in WLAN and cellular access points (APs) into the TOPSIS framework and assume sequential handling of the users. In this context, we named the former algorithm as Standard TOPSIS (ST), and the latter as Capacity Aware TOPSIS (CAT).

The performance of two algorithms is compared in simulation environment under stadium scenario where the network is congested, user density is high and user mobility is low. The main purpose of the simulations is not to provide an assessment of the offloading algorithms in the presence of D2D communications in terms of user satisfaction metrics. The simulation results reveal the performance gains of the CAT algorithm compared to ST owing to the centralized approach considering the available network capacity.

Concurrently, the results also reveal how the user satisfaction fluctuates in the presence of D2D communications in dense and congested heterogeneous networks.

7.1 TOPSIS-based solutions in the context of ongoing research

Resource management in heterogeneous Wireless Networks or WiFi integrated cellular networks could be coordinated through user-centric models, network-centric models or collaborative schemes. User centric models offer ease if implementation and scalability, as opposed to the other two approaches, at the expense of reduced –system wide- efficiency. The network-centric models, on the other hand, provide more efficient solutions that improve "social welfare" (addressing both 3GPP and non-3GPP subsytems), at the cost of increased control overhead and risk of "single point of failure". Collaborative solutions, on the other hand, introduce a little bit complexity; however, in return, offer drastic performance difference with respect to network-centric solutions in terms of QoE. Basically, in a collaborative solution, UE data such as RSS or CQI along with access network metrics obtained from an operator server are combined in decision making phase and based on the implementation choice either on network-side or user-side the decision is executed.

The handoff decision algorithm aims at selecting a network for a particular service that can satisfy objectives based on some criteria (such as low cost, good Received Signal Strength (RSS), optimum bandwidth, low network latency, high reliability and long life battery) and taking into account the preferred access network of user. Some techniques used for network-centric solutions such as stochastic programming, game theory and utility function could be performed in this respect [53].

In network-centric approaches, the goal is often to acquire maximum total allocation in 3GPP and non-3GPP networks while minimizing cost of underutilization and demand rejection. Stochastic linear programming obtains maximum allocation in each network by using probabilities related to allocation, underutilization, and rejection in Heterogeneous Wireless Networks (HWNs). Game theoretic approaches take advantage of the bankruptcy game, and efficient bandwidth allocation and admission control algorithms are developed by utilizing available bandwidth in each network. In utility function, operator prioritizes users and classifies services to allocate bandwidth for the users [49][70][73].

In user-centric solutions, the users themselves (or their agents) make the decisions, often prioritizing the needs and objectives of the individual users. Analytical hierarchy processes help ranking the networks based on induced QoS indicators, by checking user's requirements and network conditions. Proposed approaches make use of the consumer surplus model and similar economic theory based techniques. Users are often modelled to have profit functions amounting to the difference between bandwidth gain and handoff cost for each network is computed. The most appropriate network is found through utility maximization [44][50][64].

As for the collaborative models, fuzzy logic controller ranks the candidate networks based on the user's selection criteria, network data rate and SNR. In objective function, user's RSS, network's queue delay and policy preferences such as cost are fed as input parameters, and the function provides the allocation of services to APs and terminals. Lastly, in TOPSIS, the best path for flow distribution on muti-homed end-hosts is computed. Also, network's QoS (delay, jitter, and BER), user's traffic class and most importantly QoE are also considered [38][72][8].

Other options could be to harness impatient or patient algorithms which are based on user-centric solutions. The Impatient algorithm uses a very simple policy: use 3G whenever WiFi is unavailable; else use WiFi. The Patient waits and sends data on WiFi until the delay tolerance threshold, and only switches to 3G if all of the data are not sent on WiFi before the delay tolerance threshold [7]. This are the standard approaches investigated in WiFi-based offloading.

7.2 Performance Metrics Influencing Data Offloading and System model

As for decision making functionality, UE or Mobile Network Operator (MNO) selects the access network by considering probabilistic demands. Network related, terminal related, user related and application related metrics need to be considered pertaining to handover decisions. However, the paramount elements amongst them are the user-related ones as QoE is at the very heart of contemporary mobile business performance expectations. Related parameters include throughput, energy consumption of the terminal, security etc. It is interesting to note that an adult's preferences along these dimensions would potentially differ from that of a young person. For instance, security-wise an adult might not prefer to watch videos through WEP or WPA on WiFi networks but EAP-SIM on 3GPP network. Maybe this choice could be trivial for a young person and actually he would prefer a free communication band, but considering recently emerging security challenges, operators need to pay importance to the subjects of security and privacy pertinent to each and every user they serve [32].

Handoff decision criteria can be categorized as below.

- **Network-related** considering coverage, bandwidth, latency, link quality (RSS (Received Signal Strength), BER (Bit Error Rate), cost, security level.
- **Terminal-related** considering, e.g., velocity or battery power
- **User-related** considering user profiles and preferences
- **Service-related** considering service capabilities, QoS, QoE, security level [32].

The Quality of Service (QoS) and Quality of Experience (QoE), mobility and network architecture are important factors during decision making or network selection phase. In general, the following QoS and QoE metrics are important to be checked while offloading the data traffic due to the nature of real-time applications and any subset can be used as the attribute set for MADM algorithms.

(a) **End to end delay (s):** This includes processing, queuing in both ingress and egress, and propagation delay. The end-to-end delay of a video signal is the time taken for the packets to enter the transmitter at one end, be encoded into a digital signal, travel through the network, and be regenerated by the receiver at the other end.

(b) **Received Signal Strength (dBm):** This attribute is calculated based on the received signal power at the user terminal.

(c) **Packet Loss (%):** This is calculated based on the dropped packets due to either network problems or some queuing problems.

(d) **Throughput (Kbps):** this is the total traffic where packets are successfully received by the destination excluding packets for other destinations.

(e) **MOS Value (Mean Opinion Score):** This corresponds to a numerical value, ranging between 1(worst) and 5(best) expressing the quality perceived by user. It is also used as a QoE metric.

(f**) Jitter (s):** In IP networks, jitter is the variation in the time-of-arrival of consecutive packets. Jitter results from a momentary condition where more packets are trying to get on a particular link than the link can carry away [45].

Considering these performance metrics as reference, we assume the following notation and model to represent the multiple user multiple attribute decision making problem:

- The total users set (i.e., all users in stadium) in the system is denoted as $U = \{u_1, u_2, u_3, ..., u_k\}$ where k (k>=2) denotes number of users.
- The multiple users' set (i.e., MOTO users) involved in the decision making process are denoted as $U' = \{u'_1, u'_2, u'_3, ..., u'_{k'}\}$ where k' (k'<=k) denotes number of users under multiple coverage.

- The multiple attribute set (i.e., values of factors (a) to (g)) is denoted as S = {s_1 ,s_2 , s_3, ...,s_m } where m (m>=2) denotes number of possible attributes.
- The multiple decision point set (i.e., Wi-Fi and eNB) is denoted as E = {e_1, e_2, e_3, …., e_P} where there are p (p≥2) possible decision points.
- The weight set is denoted as w ={w_1, w_2, w_3, …., w_m}, where each weight w_i is the weight assigned to attribute s_i i ∈{1,2,…,m}.

7.3 Multiuser offloading algorithms for heterogeneous networks

In this section, we first recall the main features of the standard TOPSIS framework, and then describe how we adapt it to our specific cases.

7.3.1 TOPSIS

TOPSIS (Technique for Order Preference by Similarity to Ideal Solution) [46][26], due to its easy implementation, is a suitable candidate to select the optimal target network for a given a set of given observed attributes for a user.

In the first step of TOPSIS algorithm a decision matrix **A** is created as

$$\mathbf{A} = [a_{ij}] = \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1m} \\ a_{21} & a_{22} & \dots & a_{2m} \\ . & . & \dots & . \\ . & . & \dots & . \\ . & . & \dots & . \\ a_{p1} & a_{p2} & \dots & a_{pm} \end{bmatrix} (i = 1, \dots, p; j = 1, \dots, m)$$

where m denotes to size of the multiple attribute set and p refers to size of the multiple decision points. Note that that all the attributes are transformed to have positive impact if necessary.

In second step, a normalized decision matrix is formed by using the following equation:

$$r_{ij} = \frac{a_{ij}}{\sqrt{\sum_{k=1}^{p} a_{kj}^2}}$$

Then the normalized matrix R is obtained as:

$$\mathbf{R} = [r_{ij}] = \begin{bmatrix} r_{11} & r_{12} & \dots & r_{1m} \\ r_{21} & r_{22} & \dots & r_{2m} \\ . & & & . \\ . & & & . \\ . & & & . \\ r_{p1} & r_{p2} & \dots & r_{pm} \end{bmatrix}$$

In third step, a weighted normalized decision matrix is created by multiplying each column of the matrix by corresponding weight w_i by,

$$\mathbf{v_i} = w_i * \mathbf{r_i}, \qquad \mathbf{r_i} = [r_{1i}, \dots, r_{pi}]^T, \qquad i = \{1, 2, \dots, m\}$$

where $\sum_{i=1}^{m} w_i = 1$.

In fourth step, the positive (A^+) and negative (A^-) solutions are formed by using the following formulas:

$$A^+ = \left\{ \left(\max_i v_{ij} \,|\, j \in \{1,2,..m\} \right) \right\}$$

$$A^- = \left\{ \left(\min_i v_{ij} \,|\, j \in \{1,2,...m\} \right) \right\}$$

At the end of fourth step, we end up with the following sets: $A^+ = \{v_1^+, v_2^+, \dots, v_m^+\}$ and $A^- = \{v_1^-, v_2^-, \dots, v_m^-\}$.

In the fifth step, the Euclidean distance S_i^+ of each multiple decision point from the positive point A^+ and S_i^- of each multiple decision point from the negative point A^- are calculated by,

$$S_i^+ = \sqrt{\sum_{j=1}^{p} (v_{ij} - v_j^+)^2}, \quad i = \{1,..., p\}$$

$$S_i^- = \sqrt{\sum_{j=1}^{p} (v_{ij} - v_j^-)^2}, \quad i = \{1,..., p\}$$

In the final step, the relative similarity of the alternatives from the positive and negative point is calculated as

$$C_i = \frac{S_i^-}{S_i^- + S_i^+}, \quad i = \{1,..., p\}$$

where

$$0 \le C_i \le 1$$

.

The final solution is selected by

$$e^* = e_{i^*}$$

where

$$i^* = \arg\max_i C_i, \quad i = \{1,..., p\}$$

.

7.3.2 Standard TOPSIS (ST) method

With standard TOPSIS method, user's individual's benefits are considered. The method details are explained in the following steps.

Input: Set of technologies (multiple decision sets) E, and the TOPSIS matrix of user $v_{i'} \in V$ denoted by

$$\mathbf{A}^{i'} = \left[a_{ij} \right] = \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1m} \\ a_{21} & a_{22} & \dots & a_{2m} \\ . & . & \dots & . \\ . & . & \dots & . \\ . & . & \dots & . \\ a_{p1} & a_{p2} & \dots & a_{pm} \end{bmatrix} (i = 1,...., p; j = 1,...., m)$$

Output: Standard TOPSIS channel utilization vector $\mathbf{CU^e} = [CU_1^e, CU_2^e, \dots, CU_{k'}^e]$, $e \in E$.

Step1: Run TOPSIS algorithm using all $A^{i'}$ simultaneously and select the optimal decision points $e^* = e_n \in E$ for all users.

Step2: Update the channel utilization vector **CU**^{e} by summing the channel utilization demands of each user on the selected access point.

7.3.3 Capacity aware multi-user iterative TOPSIS (CAT) algorithm

In order to obtain certain benefits for access channel selection and resource allocation problem between multiple users, we propose Capacity aware iterative multi-user TOPSIS algorithm. In the capacity aware approach, a new network level attribute $\hat{u}_t$ is inserted to the TOPSIS matrix of all users and the decision is calculated sequentially in a centralized platform as users arrive, thus optimizing the total system benefit as well. The network level attribute $\hat{u}_t$ denotes the remaining available capacity in the target access technology.

Input: Set of technologies E, and the TOPSIS matrix of user $v_{i'} \in V$ denoted by

$$\mathbf{A}^{i'} = \left[a_{ij}\right] = \begin{bmatrix} a_{11} & a_{12} & \dots & a_{1m} & \hat{u}_1 \\ a_{21} & a_{22} & \dots & a_{2m} & \hat{u}_2 \\ . & . & \dots & . & . \\ . & . & \dots & . & . \\ . & . & \dots & . & . \\ a_{p1} & a_{p2} & \dots & a_{pm} & \hat{u}_p \end{bmatrix} (i = 1, \dots, p; j = 1, \dots, m)$$

Output: Capacity-aware channel utilization vector **CU**^{e} = [CU_1^e, CU_2^e, …. , $CU_{k'}^e$], e $\in$ E.

Step1: Set **CU**^{e}=[0] and j=0 (j ≤ k' is the (MOTO) user number)

Step2: Put j= j+1, as user $v_{i'}$ arrives.

Step3: Run TOPSIS algorithm using $A^{i'}$ of user $v_{i'}$ and select the optimal decision point e* =$e_n \in$ E.

Step4: Update the channel utilization vector **CU**^{e} by

$CU_{i'}^e$ = Channel utilization demand of user $v_{i'}$

Step 5:

IF (j = k') then **stop**

ELSE goto **Step2**

7.4 Performance Results

7.4.1 Simulation Scenario

Opportunistic offloading platforms such as MOTO platform are designed to overcome capacity problems in dense heterogeneous networks. In this context, we chose a 10K capacity stadium scenario with several Wi-Fi and eNB access points deployed. Out of the 10K people in the stadium 50, 100 or 200 terminals are selected randomly to represent the MOTO users. A sample user distribution for 100 MOTO users in a 10K capacity stadium is provided in Figure 47. The assumed locations of 8 Wi-Fi APs and 2 eNB APs are also depicted in Figure 47. The locations of APs are chosen heuristically and can be modified based on requirements.

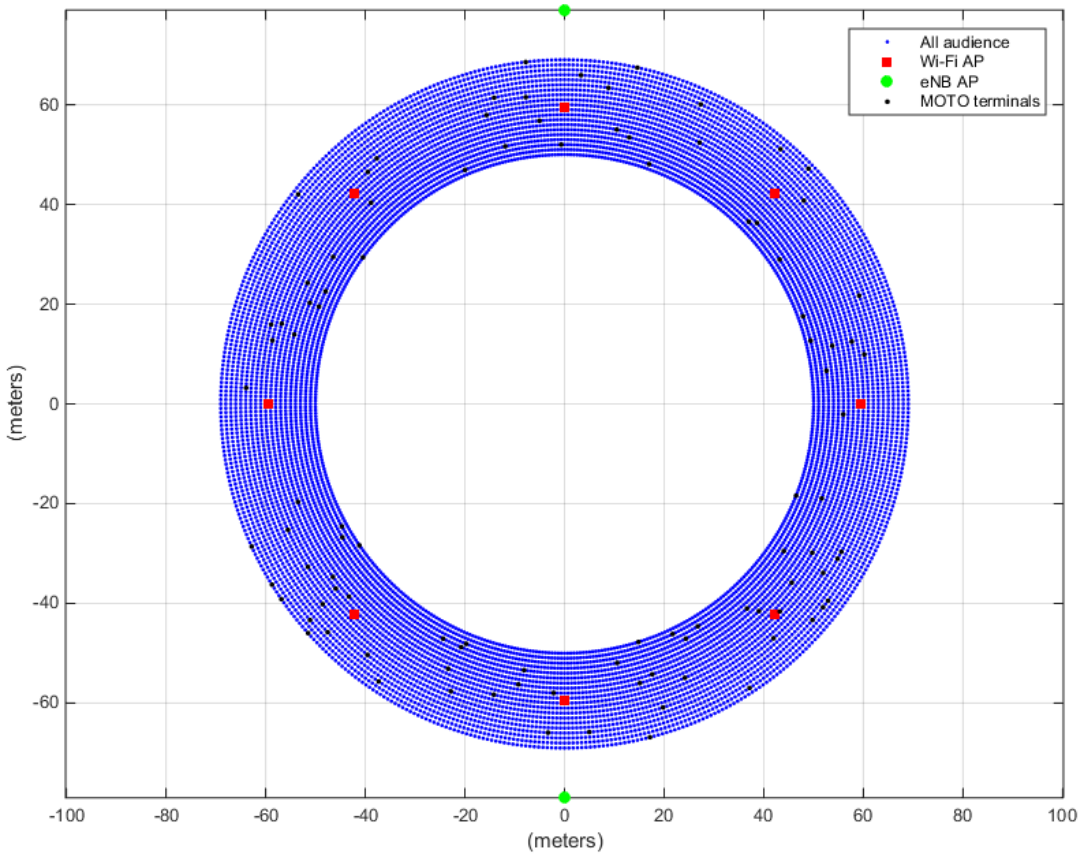

Figure 47. A sample of MOTO users' distribution and locations of APs

In the scenario, the MOTO users can download streaming video from 360p to 1080p with bit rates adapted from those of the worldwide video streaming service, YouTube [81]. MOTO users can download the content either from the closest Wi-Fi AP or the closest eNB AP or via D2D.

The number of MOTO users that will download the content from other MOTO users via D2D is assumed to be given. Given the number such D2D users, the D2D users are selected as the MOTO users that are farther away from the Wi-Fi and eNB APs. If a MOTO user is selected to use D2D, the range is assumed to be 10m, similar to that of Bluetooth. The D2D users are allowed to download the content only via surrounding MOTO users that are assigned to Wi-Fi or eNB APs.

The attributes for the TOPSIS algorithms are chosen as the ones that are most related with video transmission. First attribute is the received signal strength which effects the maximum bit rate that the user can download content. For the simulation the received signal strength is calculated by free space propagation principles based the distance between the transmitter (AP) and receiver (MOTO terminal). Second attribute is the video bitrate (i.e., throughput) calculated based on a mapping table using the received signal strength. Another attribute is the latency which is assumed to be constant for a given access technology. The last attribute is the remaining capacity in the AP that effects only the CAT algorithm and is ignored in ST. For all simulations the attribute weights given in Table 7 are used in order to obtain a balanced offloading outcome. The latency attribute is assigned the lowest weight due to low sensitivity of delay in the target scenario.

Table 7: Attribute weights for simulations

Attribute	Weight
Received Signal Strength	0.3
Throughput	0.3
Latency	0.1
Remaining Capacity	0.3

The simulation results demonstrate the average of 500 simulations. The simulations will provide the best Wi-Fi/eNB offloading solution for a given number of D2D users.

7.4.2 Performance Metrics

The simulation results are presented mainly based on two different performance metrics. First one is the total capacity demand (utilization demand) of the MOTO users from the existing APs based on their individual video bitrates. If the total capacity demand is above 100% in the results, this directly implies that the network is congested and some users cannot receive the video properly. The other performance metric is the subjective user satisfaction metric. User satisfaction is defined different for users based on their access technology. Wi-Fi/eNB users are assumed to be satisfied if they can download the content with the bit rate determined based on the received signal strength. Due to higher demand compared to available capacity in the APs, some users will become unsatisfied. On the other hand, D2D users are assumed to be satisfied if there exists at least one satisfied Wi-Fi/eNB user within their D2D communication range.

7.4.3 Results

In this part, we present the results based on two performance metrics as described in Section 7.4.2. First results are given in Table 8 which depicts the user distribution and resulting total capacity demands when ST and CAT algorithms are utilized for varying number of users. It is assumed that capacity of 75Mbps per each eNB AP and 36 Mbps per each Wi-Fi AP is reserved (available) for MOTO users. The rest of the capacity is assumed to be utilized by other mobile users in the stadium. In the results of Table 8, all users are assigned to a Wi-Fi or eNB AP when the number of D2D users is zero. As the number of MOTO users increase, the network (with the available capacity stated above) becomes more congested and user satisfaction degrades as expected. For all three different congestion levels, the CAT method yields improved balance between user distributions and capacity demand, as well as better user satisfaction. The results indicate the superiority of CAT algorithm owing to the utilization of network level information.

Table 8: User distribution, capacity demands and user satisfaction for CAT and ST at zero D2D density

# of Users	Type of MADM	USERS' DISTRIBUTION (avg.%)		TOTAL CAPACITY DEMAND (%)		MOTO User Satisfaction (%)
		3GPP	WLAN	3GPP	WLAN	
50	CAT	37.34%	62.66%	76.13%	53.25%	99.80%
	ST	80.82%	19.18%	158.40%	21.01%	67.50%
100	CAT	36.87%	63.13%	151.50%	107.80%	73.31%
	ST	80.33%	19.67%	316.90%	41.92%	40.47%
200	CAT	36.52%	63.48%	303.10%	217.10%	34.29%
	ST	80.69%	19.40%	631.80%	84.61%	23.08%

Next results are presented in Figure 48 and Figure 49. In these results, the number of D2D users increase from zero to number of MOTO users (#MOTO:50, 100, 200) to present the impact of D2D users' density on the user satisfaction as described in Section 7.4.2. In Figure 48, for each simulation result corresponding to different number of MOTO users, a fixed capacity of 75Mbps per each eNB AP and 36 Mbps per each Wi-Fi AP is reserved (available) for MOTO users. As the number of users increase the network becomes more congested. When the number of MOTO users is 50 (#MOTO:50), the total demand is lower than the available capacity. Thus, it is possible to satisfy all users if users are properly assigned (offloaded) to the access points. As observed from Table 8 and Figure 48, CAT algorithm satisfies all users (when the number of D2D users is zero) thanks to the utilization of network level information. On the other hand, ST algorithm, which only considers individual scores, is unsuccessful at satisfying all users despite the available capacity. As the number of MOTO users increases (#MOTO:100, 200), the demand of users exceeds available capacity leading to high congestion levels where 100% user satisfaction via Wi-Fi and eNB APs is impossible. In this case, it is clear that D2D offloading shows positive impact on the user satisfaction, such that, up to 13% and 67% increase in user satisfaction can be achieved when D2D density is 24% and 68% for CAT algorithm, for 100 and 200 MOTO users, respectively. The utilization rate of APs on the optimal D2D density points are provided in Table 9, which correspond to the best load-balance points with maximum user satisfaction.

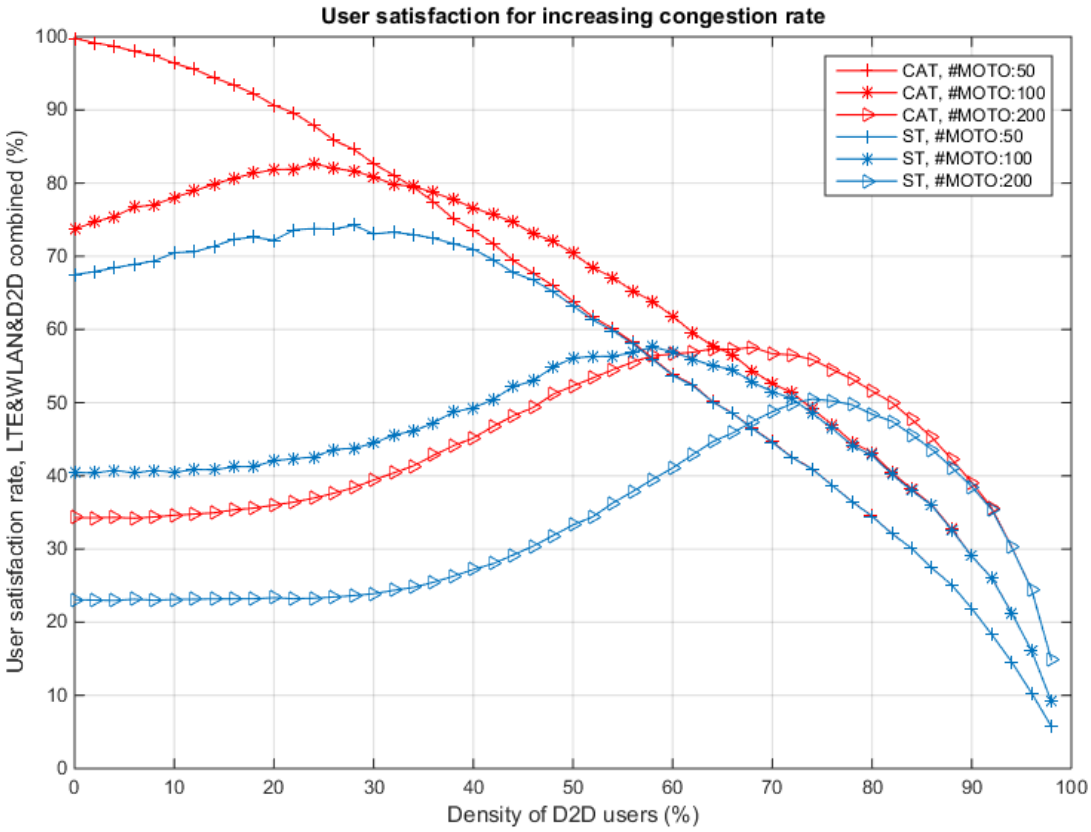

Figure 48: User satisfaction curves for increasing congestion rates

Table 9: User distribution, capacity demands and user satisfaction for CAT and ST at optimal D2D density for increasing congestion rate

# of Users	Type of MADM	Opt. D2D Density (%)	USERS' DISTRIBUTION (avg.%)		TOTAL CAPACITY DEMAND (%)		MOTO User Satisfaction (%)
			3GPP	WLAN	3GPP	WLAN	
50	CAT	0%	37.34%	62.66%	76.13%	53.25%	99.80%
	ST	28%	52.82%	19.18%	158.40%	21.01%	74.26%
100	CAT	24%	23.62%	52.38%	97.46%	97.68%	82.67%
	ST	58%	23.73%	18.27%	226.30%	41.45%	57.62%
200	CAT	68%	8.73%	23.27%	73.83%	96.88%	57.56%
	ST	74%	11.43%	14.57%	102.10%	67.12%	50.49%

In Figure 49, simulation results for fixed congestion rate (i.e., when "number of MOTO users / available capacity" is same) for different number of MOTO users (#MOTO:50, 100, 200) is provided. The rationale behind this simulation is to demonstrate the benefit of D2D offloading as the density of MOTO users increase in a heterogeneous network. The results clearly demonstrate that for same congestion level the user satisfaction increases with increasing number of MOTO users. This is mainly caused by the increase in the probability of the existence of LTE/WLAN users in the range of D2D users. The utilization rate of APs on the optimal D2D density points are provided in Table 10, which correspond to the best load-balance points with maximum user satisfaction.

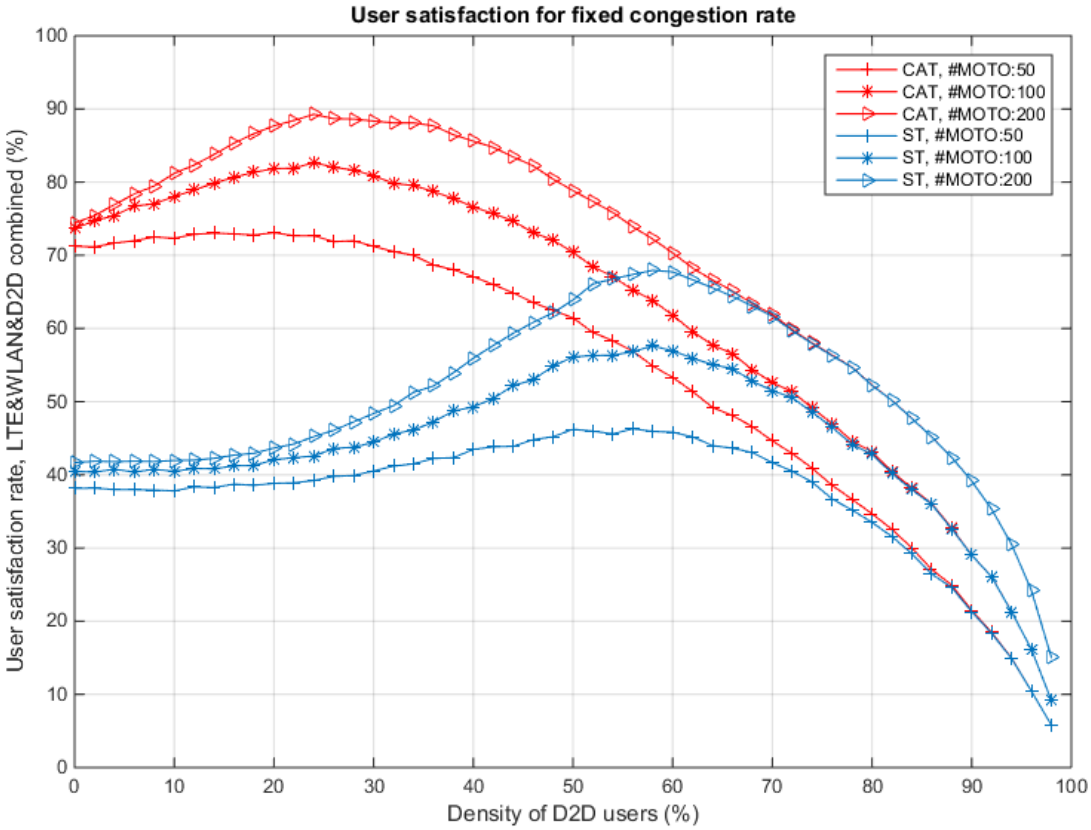

Figure 49: User satisfaction curves for fixed congestion rate

Table 10: User distribution, capacity demands and user satisfaction for CAT and ST at optimal D2D density for fixed congestion rate

# of Users	Type of MADM	Opt. D2D Density (%)	USERS' DISTRIBUTION (avg.%)		TOTAL CAPACITY DEMAND (%)		MOTO User Satisfaction (%)
			3GPP	WLAN	3GPP	WLAN	
50	CAT	14%	30.62%	55.38%	123.60%	100.30%	73.03%
	ST	56%	25.54%	18.46%	106.70%	40.40%	46.30%
100	CAT	24%	23.62%	52.38%	97.46%	97.68%	82.67%
	ST	58%	23.73%	18.27%	99.27%	40.30%	57.62%
200	CAT	24%	23.11%	52.89%	95.87%	99.38%	89.27%
	ST	58%	23.68%	18.32%	98.79%	40.62%	67.96%

7.5 Modified Fully-Centralized TOPSIS Algorithm including D2D

Device-to-device (D2D) communications in cellular networks can be defined as direct communication between two mobile users without passing through the Base Station (BS). The D2D communication is generally not transparent to the cellular network and it can occur either on the cellular spectrum (i.e., inband) or unlicensed spectrum (i.e., outband). It is important to note that in a conventional cellular network, all communications go through the BS even if both communicating users are close to each other for the D2D communications. This conventional architecture suits well for low data rate mobile applications and services, such as voice call and text messaging, in which users are usually not close enough to have direct communication. However, mobile users in today's cellular networks use high data rate services (e.g., video sharing, gaming, location-based services, proximity-aware social networking) and they might be close enough to each other for the D2D communications. Therefore, the D2D communications in such scenarios can highly improve the spectral efficiency of the network. In addition to improving spectral efficiency, the D2D communications can potentially improve throughput, energy efficiency, and reduce communication delay depending on the application scenarios.

In the literature, two types of D2D communication, i.e., inband D2D and outband D2D communications, have been proposed. Inband D2D communication is utilized to control and to optimize cellular spectrum more efficiently. There are two types of inband D2D technologies: underlay and overlay. Inband underlay D2D aims at sharing same radio resources between cellular and D2D communications links. On the other hand, inband overlay type requires a dedicated cellular communication channel. The problem with these approaches is the interference caused by D2D links are affecting overall cellular spectrum performance.

Whereas, outband D2D communication is utilized to harness the unlicensed spectrum band; therefore, the interference issue between D2D and cellular link is eliminated. However, that kind of approach requires an additional wireless network interface at user side; for example, WLAN Direct or Bluetooth. There are two types of outband D2D technologies, namely, controlled and autonomous. Controlled outband technology provides the second interface's control to the cellular network whereas autonomous suggests to give the control of the second interface to the users.

The key topic related to this activity is mainly related to how to include also D2D technologies into TOPSIS. TOPSIS is a centralized scheme, and therefore the most suitable type of D2D technology that we considered is inspired by LTE-D2D, i.e. where the operator controls the fine details of direct communications between mobile devices.

We hereafter provide a brief sketch of how we extended the algorithms presented before in this sense. Assume that we list all the combinations for p different technologies as the decision point in decision matrix as

$$D = \begin{array}{c} \\ sce_1 \\ sce_2 \\ sce_3 \\ \dots \\ sce_l \end{array} \begin{array}{c} \begin{array}{cccccc} T_1 & T_2 & T_3 & T_4 & \dots & T_m \end{array} \\ \begin{bmatrix} . & . & . & . & . & . \\ . & . & . & . & . & . \\ . & . & . & . & . & . \\ \dots & \dots & \dots & \dots & . & . \\ . & . & . & . & . & . \end{bmatrix} \end{array}$$

Where sce_i is the i-th scenario, $l = \widehat{\mathbf{CU}}^{e*} = p^{k'}$ and k' denotes the users under multiple coverage.

For example, if we have three technologies, LTE, WLAN and D2D and two users, U1, U2 are under multiple coverage, then there are 9 possibilities on the rows of D matrix.

sce_1= [LTE, LTE] sce_4= [WLAN, LTE] sce_7 = [D2D, LTE]

sce_2= [LTE, WLAN] sce_5 = [WLAN, D2D] sce_8 = [D2D, WLAN]

sce_3= [LTE, D2D] sce_6 = [WLAN, WLAN] sce_9 = [D2D, D2D]

Then on the columns of D matrix, there are attributes such as:

T_1) Total QoS= $\sum QoS_i$

T_2) Total Delay = $\sum \tau_i$

T_3) Total Energy = $\sum E_i$

T_4) Total Throughput = $\sum Th_i$

T_5) Total MOS values = $\sum MOS_i$

…

T_m) …

After constructing the decision matrix, we can run TOPSIS to decide for the appropriate possibility of scenarios (sce_1, …, sce_9) for users that are under multiple coverage. Clearly, careful evaluation of this extension is needed, as the D2D alternatives are typically much more dynamic. Even more importantly, the number of possible combinations increases exponentially with the number of technologies, and this needs to be taken into consideration when the number of users increases. If we assume 100 users and 4 attributes, the decision matrix will be of size 3^{100} x 4. Thus, the modified algorithm has a complexity of $O(p^{Nusers} \text{ x } m)$, which may be computationally impractical to implement in many cases.

Note, however, that the methodology we have presented in this section provides a practical operation way of allocating traffic between LTE, WiFi and D2D. Specifically, given a certain layout of users, the tools developed in this part of the analysis allow us to find the optimal set of users to be served via D2D, such that the capacity of LTE and WiFi are not over-used, and the maximum number of users is satisfied with respect to the traffic constraints. This configuration can be found through the simulation tools developed, which are lightweight to implement, and therefore provide a reasonable first-level approximation for the real performance of the system.

8 Conclusions

At the end of the activities of WP3 we can draw the main conclusions of the results obtained in this work package. As already mentioned in Section 1, work has progressed mostly along the lines planned at the beginning of the project. When some restructuring of activities has been needed (in T3.3), we have been able to refocus them on synergic topics with respect to the objectives of the WP and of the project overall.

The main conclusions related to activities on T3.1 have been drawn already in D3.2. It is important to note that these results have been instrumental for some of the key achievements presented in this deliverable, such as the analytic tools to configure the trade-off between energy efficiency and capacity (see Section 2.2).

In summary the main outputs of the activities carried out in the rest of the project can be presented as follows.

1. **We have derived a set of practical tools to configure an offload network, based on strong mathematical formulations and analysis**. Specifically, notable results in this class include:
    1.1. A set of **conditions** to be tested, given a representation of the mobility patterns of the users, to **select appropriate routing protocols in opportunistic networks**. Opportunistic network routing protocols are one of the fundamental components of an offload network, as they define the capacity of the D2D communications, and thus the amount of traffic that can be offloaded from an LTE network. The tools we have derived allows an operator to (i) tune the opportunistic protocols to guarantee delivery of contents (i.e., to avoid packet loss or indefinitely long delays), and (ii) select the most appropriate protocols to avoid such problems (or identify protocols that cannot be used because they are prone to such problems).
    1.2. A set of **closed form expressions** to estimate the **end-to-end delay** of opportunistic networking protocols, from a given description of the users mobility patterns. These expressions can be used to **estimate the offloading gain** of the offload network. In fact, given a certain requirement in terms of end-to-end delay, it is possible to estimate the probability of not meeting that deadline through D2D communications, and therefore the share of traffic that will have to be sent via LTE. From a complementary standpoint, it is possible to use the same tool to set the maximum delay requirements, for a given target in terms of offloading gain.
    1.3. A set of **closed form expressions** to tune the **trade-off between energy saving** at mobile devices and the **additional capacity** (or QoS delay constraints) that can be obtained from the opportunistic network. Using these expressions, a network operator can, again, decide what is the offloading gain that can be obtained for a certain energy saving rate, or set the optimal energy saving configuration given a requirement in terms of offloading gain (or maximum end-to-end delay).
2. **We have analysed in depth several key features of the various components of an offload network, characterising its behaviour from a number of standpoint.** In particular:
    2.1. We have derived an **LTE throughput model** that considers **at the same time PHY and MAC-layer mechanisms**. This model allowed us to understand with a compact analytical tool key performance drops of LTE, and conditions where offloading can greatly help. Such a model is quite original with respect to the existing literature, where performance are typically analysed for each individual layer separately.
    2.2. We have characterised the **performance of a completely integrated network considering an LTE infrastructure, a WiFi infrastructure, and D2D communications**. In particular, we have characterised, using real mobility traces, the performance gain of using D2D with respect to more conventional WiFi offloading solutions, as a function of key systems and environment parameters.

2.3. We have analysed the **performance of D2D offloading when content requests are not synchronised** (which completely rules out the possibility of using LTE multicast). Also in this case, we have shown that offloading can bring very significant capacity gains, by exploiting limited temporary storage at mobile devices.

3. We have designed **advanced scheduling solutions** to split traffic across various network technologies in offload networks. As opposed to the typical approach in the literature, we have focused on scheduling on a time grain much coarser than packet level, complementing scheduling algorithms implemented, for example, inside LTE. In particular:

   3.1. We have defined how to **schedule traffic inside an LTE network between the multicast and the D2D communication modes**. Through this, we have also demonstrated the advantage of using offloading in presence of multicast traffic, and presented (and evaluated) a learning algorithm to automatically tune the split of traffic to be sent of the two communication modes.

   3.2. We have investigated how to **schedule traffic inside the LTE core network**, in such a way to **reduce the energy consumption of LTE access devices** (picocells). Specifically, we have shown that it is possible to concentrate users on fewer picocells, and switch-off the remaining ones, without compromising the throughput performance achieved by the users.

   3.3. We have investigated how to **optimally split traffic between LTE, WiFi and D2D communications** in order to optimise the overall users' satisfaction. To this end, we have developed a centralised optimisation algorithm (amenable to be used by a network operator), and a methodology to identify the optimal fraction of users to be served through D2D, WiFi and LTE, given a certain configuration of users in the physical space.

4. **Last but not least, these results have had a significant impact on the rest of the project activities**. For example:

   4.1. Specific mechanisms have been defined in the **MOTO architecture to incorporate the control settings defined by the analytical models**. Corresponding APIs have been defined to implement the resulting control decisions (see D2.2.2).

   4.2. The end-to-end delay analytical tools have been used to **set the QoS achievable in an offload network, and its trade-off with energy efficiency**. This additional control point is supported by appropriate interfaces at the architectural level (see D2.2.2).

   4.3. The **tools** derived to characterise the **convergence and end-to-end delay** of opportunistic networks have also been used to **expose to the users the trade-off between privacy and performance** achievable through the network. This allowed us to make the users' aware of this trade-off, and select where they are ready to operate in this space (from highest-performance/minimal privacy to lowest-performance/maximal privacy), see D4.3.

   4.4. The **mechanisms defined in WP3** have been **extensively tested in simulation**, to better characterise their performance in the MOTO reference integrated simulation environment (see D5.2).

In conclusion, these results have shed significant light on the expected performance of offloading in future generation mobile networks, and have provided practical tools to control the various trade-offs involved, and the behaviour of a heterogeneous network based on the concept of offloading. We believe that the building blocks provided by WP3 are therefore an important step in the use of offloading techniques in the framework of 5G networks.

References

[1] The MOTO consortium, Deliverable D3.1 "Initial results on offloading foundations and enablers", available at http://www.fp7-moto.eu/wp-content/uploads/2013/10/moto_D3.1_v1.0_PU1.pdf

[2] The MOTO consortium, Deliverable D3.2 "Spatiotemporal characterization of contact patterns in dynamic networks" available at http://www.fp7-moto.eu/wp-content/uploads/2014/12/moto-D3.2.pdf

[3] The MOTO consortium, Deliverable D3.3.1 "Design and evaluation of enabling techniques for mobile data traffic offloading (release a)" available at http://www.fp7-moto.eu/wp-content/uploads/2014/12/D3.3.1.compressed.pdf

[4] The MOTO consortium, Deliverable D2.2.1 "General Architecture of the Mobile Offoading System (Release a)", available at http://www.fp7-moto.eu/wp-content/uploads/2014/01/Deliverable_D21_1.01.pdf

[5] The MOTO Consortium, "D5.1.1 – Description and development of MOTO simulation tool environment – Release a", available at http://www.fp7-moto.eu/wp-content/uploads/2014/01/moto_D5.1.1v1.01.pdf

[6] R. Akl, S. Valentin, G. Wunder, and S. Stanczak, "Compensating for CQI Aging By Channel Prediction: The LTE Downlink," in *Proc. of IEEE GLOBECOM'12*, 2012, pp. 4821–4827.

[7] A. Balasubramanian, R. Mahajan, and A. Venkataramani. Augmenting mobile 3G using WiFi. In Proceedings of the 8th international conference on Mobile systems, applications, and services, pages 209-222. ACM, 2010

[8] A. Ben Nacef, N. Montavont, A generic end-host mechanism for path selection and flow distribution, in: IEEE 19th International Symposium on Personal, Indoor and Mobile Radio Communications PIMRC, 2008, pp. 1–5

[9] Elisabetta Biondi, Chiara Boldrini, Andrea Passarella, and Marco Conti, "Optimal duty cycling in mobile opportunistic networks with end-to-end delay guarantees", *European Wireless*, Barcelona, Spain, 14-16 May 2014

[10] Elisabetta Biondi, Chiara Boldrini, Marco Conti, Andrea Passarella, "Duty Cycling in Opportunistic Networks: the Effect on Intercontact Times",*The 17th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (ACM MSWiM 2014)*, Montreal, Canada, September 21-26 2014.

[11] Y. Blankenship, P. Sartori, B. Classon, V. Desai, and K. Baum, "Link error prediction methods for multicarrier systems," in *Proc. of IEE VTC- Fall'04*, vol. 6, 2004, pp. 4175–4179.

[12] Chiara Boldrini, Marco Conti, Andrea Passarella, "The stability region of the delay in Pareto opportunistic networks", *IEEE Transactions on Mobile Computing*, vol.14, no.1, pp.180, 193, Jan. 1, 2015: http://dx.doi.org/10.1109/TMC.2014.2316506

[13] C. Boldrini, M. Conti, and A. Passarella, "Performance modelling of opportunistic forwarding under heterogenous mobility," *Computer Communications*, pp. 1–17, 2014.

[14] K. Brueninghaus, D. Astely, T. Salzer, S. Visuri, A. Alexiou, S. Karger, and G.-A. Seraji, "Link performance models for system level simulations of broadband radio access systems," in Proc. of IEEE PIMRC'05, 2005.

[15] R. Bruno, A. Masaracchia, A. Passarella, "Robust Adaptive Modulation and Coding (AMC) Selection in LTE Systems using Reinforcement Learning", Proc. of IEEE VTC2014-Fall, 14–17 September 2014, Vancouver, Canada.

[16] Raffaele Bruno, Antonino Masaracchia, Andrea Passarella, and Stefano Mangione, "Analysis of MAC-level Throughput in LTE Systems with Link Rate Adaptation and HARQ Protocols", Proc. of IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (IEEE WoWMoM), Boston, USA, 14-17 June 2015.

[17] Raffaele Bruno, Antonino Masaracchia, and Andrea Passarella, "Offloading through Opportunistic Networks with Dynamic Content Requests", The IEEE Workshop on CellulAR Traffic Offloading to Opportunistic Networks (IEEE CARTOON 2014), Philadelphia, Pennsylvania, USA, October 27, 2014

[18] Raffaele Bruno, Antonino Masaracchia, and Andrea Passarella, "Analysis of MAC-layer Throughput in LTE Systems with Link Rate Adaptation and HARQ Protocols", IIT-CNR Technical Report, 2014.

[19] V. Buenestado, J. Ruiz-Aviles, M. Toril, S. Luna-Ramirez, and A. Mendo, "Analysis of Throughput Performance Statistics for Bench- marking LTE Networks," IEEE Communications Letters, vol. 18, no. 9, pp. 1607–1610, September 2014.

[20] H. Cai and D. Eun, "Crossing over the bounded domain: From exponential to power-law intermeeting time in mobile ad hoc networks," *IEEE/ACM Trans. on Netw.*, vol. 17, no. 5, pp. 1578– 1591, 2009.

[21] F. Capozzi, G. Piro, L. Grieco, G. Boggia, and P. Camarda, "Downlink Packet Scheduling in LTE Cellular Networks: Key Design Issues and a Survey," IEEE Communications Surveys & Tutorials, vol. 15, no. 2, pp. 678–700, 2013.

[22] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of human mobility on opportunistic forwarding algo- rithms," *IEEE Trans. Mobile Comput.*, pp. 606–620, 2007.

[23] V. Conan, J. Leguay, and T. Friedman, "Characterizing pair- wise inter-contact patterns in delay tolerant networks," in *Autonomics'07*, 2007.

[24] COST Action 231, "Digital mobile radio future generation systems," Final Report - EUR 18957, 1999. "ns-3 Model Library – Release ns-3.17", May 14, 2013.

[25] S. Donthi and N. Mehta, "An Accurate Model for EESM and its Application to Analysis of CQI Feedback Schemes and Scheduling in LTE," IEEE Transactions on Wireless Communications, vol. 10, no. 10, pp. 3436–3448, October 2011.

[26] Dutta, A., et al.: Seamless Handover across Heterogeneous Networks - An IEEE802.21 Centric Approach. In: IEEE WPMC (2006).

[27] A. Elnashar and M. El-Saidny, "Looking at LTE in Practice: A Perfor- mance Analysis of the LTE System Based on Field Test Results," IEEE Vehicular Technology Magazine, vol. 8, no. 3, pp. 81–92, September 2013.

[28] J. Francis and N. Mehta, "EESM-Based Link Adaptation in Point-to- Point and Multi-Cell OFDM Systems: Modeling and Analysis," *IEEE Transactions on Wireless Communications*, vol. 13, no. 1, pp. 407–417, January 2014.

[29] S. Gamboa A. Pelov, P. Maillé X. Lagrange, N. Montavont, "*Energy efficient cellular networks in the presence of delay tolerant users*" Global Communications Conference (GLOBECOM), 2013 IEEE, Dec. 2013

[30] Weisi Guo, Siyi Wang , T O'Farrell,S. Fletcher, "*Energy Consumption of 4G Cellular Networks: A London Case Study*", Vehicular Technology Conference (VTC Spring), 2013 IEEE 77th, June 2013

[31] Z. Haas and T. Small, "A new networking model for biological applications of ad hoc sensor networks," *IEEE/ACM Trans. on Netw.*, vol. 14, no. 1, pp. 27–40, 2006.

[32] Hadiji, F.; Zarai, F.; Kamoun, A., "Architecture of mobile node in heterogeneous networks," Communications and Information Technology (ICCIT), 2012 International Conference on , vol., no., pp.260,264, 26-28 June 2012

[33] Z. He and F. Zhao, "Performance of HARQ With AMC Schemes In LTE Downlink," in Proc. of IEEE CMC'10, 2010, pp. 250–254.

[34] K. Kant, *Introduction to computer system performance evaluation*. McGraw-Hill, 1992.

[35] H. Kobayashi, *Modeling and analysis: ain introduction to system performance evaluation methodology*. Addison-Wesley, 1981.

[36] M.A. Imran, E. Katranaras, "Energy Efficiency analysis of the reference system, areas of improvements and targets breakdown," Tech.Rep., Dec.2010

[37] T. Karagiannis, J.-Y. Le Boudec, and M. Vojnovic and, "Power law and exponential decay of intercontact times between mobile devices," *IEEE Trans. Mobile Comput.*, vol. 9, no. 10, pp. 1377 –1390, 2010.

[38] G. Koundourakis, D. Axiotis, M. Theologou, Network-based access selection in composite radio environments, in: IEEE Wireless Communications and Networking Conference, 2007. WCNC'2007, 11–15 March 2007, pp. 3877–3883

[39] A. Kuhne and A. Klein, "Throughput analysis of multi-user ofdma- systems using imperfect cqi feedback and diversity techniques," *IEEE Journal on Selected Areas in Communications*, vol. 26, no. 8, pp. 1440– 1450, October 2008.

[40] J. Ikuno, S. Pendl, M. Simko, and M. Rupp, "Accurate SINR estimation model for system level simulation of LTE networks," in *Proc. of IEEE ICC'12*, 2012, pp. 1471–1475.

[41] C. Lee, "Heterogeneity in contact dynamics: helpful or harmful to forwarding algorithms in DTNs?" in *WiOPT'09*. IEEE, 2009, pp. 1–10.

[42] J. Leinonen, J. Hamalainen, and M. Juntti, "Capacity Analysis of Down- link MIMO-OFDMA Resource Allocation with Limited Feedback," IEEE Transactions on Communications, vol. 61, no. 1, pp. 120–130, January 2013.

[43] Z. Lin, P. Xiao, and B. Vucetic, "SINR distribution for LTE downlink multiuser MIMO systems," in Proc. of IEEE ICASSP'09, April 2009, pp. 2833–2836.

[44] X. Liu, V.O.K. Li, P. Zhang, Joint radio resource management through vertical handoffs in 4G networks, in: IEEE Global Telecommunications Conference, 2006. GLOBECOM'06, November 2006, pp. 1–5

[45] M. Logothetis, K. Tsagkaris, P. Demestichas, Application and mobility aware integration of opportunistic networks with wireless infrastructures, Computers & Electrical Engineering, Available online 24 August 2012, ISSN 0045-7906, 10.1016/j.compeleceng.2012.07.014

[46] Z. Markovic, (2010). Modification of TOPSIS method for solving of multi-criteria tasks. YugoslavJournal of Operations Research. 20, p.117-143

[47] M. Mezzavilla, M. Miozzo, M. Rossi, N. Baldo, and M. Zorzi, "A Lightweight and Accurate Link Abstraction Model for the Simulation of LTE Networks in Ns-3," in Proc. of ACM MSWiM '12, 2012, pp. 55–60.

[48] M. Ni, X. Xu, and R. Mathar, "A channel feedback model with robust SINR prediction for LTE systems," in *Proc. of EuCAP'13*, 2013, pp. 1866–1870.

[49] D. Niyato, E. Hossain, A Cooperative game framework for bandwidth allocation in 4G heterogeneous wireless networks, in: IEEE International Conference on Communications, 2006. ICC'06, vol. 9, June 2006, pp. 4357–4362

[50] O. Ormond, P. Perry, J. Murphy, Network selection decision in wireless heterogeneous networks, in: IEEE 16th International Symposium on Personal, Indoor and Mobile Radio Communications, 2005. PIMRC 2005, vol. 4, 11–14 September 2005, pp. 2680–2684

[51] A. Passarella and M. Conti, "Analysis of individual pair and aggregate inter-contact times in heterogeneous opportunistic networks," *IEEE Trans. Mobile Comput.*, vol. 12, no. 12, pp. 2483 – 2495, 2013.

[52] Andrea Passarella and Marco Conti, "Impact of contact processes on convergence of forwarding protocols in heterogeneous opportunistic networks", IIT-CNR Technical Report

[53] K. Piamrat, A. Ksentini, J.Bonnin, C. Viho, Radio resource management in emerging heterogeneous wireless networks, Computer Communications, Volume 34, Issue 9, 15 June 2011, Pages 1066-1076, ISSN 0140-3664, 10.1016/j.comcom.2010.02.015.

[54] A. Picu, T. Spyropoulos, and T. Hossmann, "An analysis of the information spreading delay in heterogeneous mobility dtns," in *IEEE WoWMoM*, 2012, pp. 1–10.

[55] Filippo Rebecchi, Marcelo Dias de Amorim, Vania Conan: DROid: Adapting to individual mobility pays off in mobile data offloading**.**Networking 2014: 1-9

[56] Filippo Rebecchi, Marcelo Dias de Amorim, and Vania Conan. 2014. Flooding data in a cell: is cellular multicast better than device-to-device communications?. In *Proceedings of the 9th ACM MobiCom workshop on Challenged networks* (CHANTS '14).

[57] Filippo Rebecchi, Lorenzo Valerio, Raffaele Bruno, Vania Conan, Marcelo Dias de Amorim, Andrea Passarella, "A Joint Multicast/D2D Learning-Based Approach to LTE Traffic Offloading", submitted to an international journal.

[58] M. Rinne and O. Tirkkonen, "LTE, the radio technology path towards 4G," Computer Communications, vol. 33, pp. 1894–1906, 2010.

[59] R. Shassberger, "On the waiting time in the queuing system GI/G/1," *The Annals of Mathematical Statistics*, vol. 41, no. 1, pp. 182–187, 1970.

[60] H. Seo and B. Lee, "Proportional-fair power allocation with CDF- based scheduling for fair and efficient multiuser OFDM systems," IEEE Transactions on Wireless Communications, vol. 5, no. 5, pp. 978–983, May 2006.

[61] P. Serra, A. Rodrigues, "*Picocell positioning in an LTE network*" Anancom, Nov 2013

[62] M. Sharif and B. Hassibi, "On the capacity of MIMO broadcast chan- nels with partial side information," IEEE Transactions on Information Theory, vol. 51, no. 2, pp. 506–522, February 2005.

[63] H. Song, R. Kwan, and J. Zhang, "General results on SNR statistics involving EESM-based frequency selective feedbacks," IEEE Transactions on Wireless Communications, vol. 9, no. 5, pp. 1790–1798, May 2010.

[64] Q. Song, A. Jamalipour, A network selection mechanism for next generation networks, in: IEEE International Conference on Communications, 2005. ICC'2005, vol. 2, 16–20 May 2005, pp. 1418–1422

[65] T. Spyropoulos, K. Psounis, and C. Raghavendra, "Efficient routing in intermittently connected mobile networks: The single copy case," *IEEE/ACM Trans. on Netw.*, vol. 16, no. 1, pp. 63–76, 2008.

[66] T. Spyropoulos, K. Psounis, and C. Raghavendra, "Efficient routing in intermittently connected mobile networks: The multiple- copy case," *IEEE/ACM Trans. on Netw.*, vol. 16, no. 1, pp. 77–90, 2008.

[67] T. Spyropoulos, T. Turletti, and K. Obraczka, "Routing in Delay- Tolerant Networks Comprising Heterogeneous Node Populations," *IEEE Trans. Mobile Comput.*, pp. 1132–1147, 2009.

[68] E. Stevens-Navarro and V. W. S. Wong, "Comparison between vertical handoff decision algorithms for heterogeneous wireless networks," in Vehicular Technology Conference, 2006. VTC 2006-Spring. IEEE 63rd, vol. 2, Melbourne, Vic., May 2006, pp. 947

[69] R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*. The MIT Press, March 1998.

[70] A. Taha, H. Hassanein, H. Mouftah, On robust allocation policies in wireless heterogeneous networks, in: First International Conference on Quality of Service in Heterogeneous Wired/Wireless Networks, 2004. QSHINE 2004, 18–20 October 2004, pp. 198–205.

[71] N. Varanese, J. Vicario, and U. Spagnolini, "On the Asymptotic Throughput of OFDMA Systems with Best-M CQI Feedback," IEEE Wireless Communications Letters, vol. 1, no. 3, pp. 145–148, June 2012.

[72] A. Wilson, A. Lenaghan, R. Malyan, Optimising wireless access network selection to maintain QoS in heterogeneous wireless environments, Wireless Personal Multimedia Communications 2005. WPMC'05, 18–22 September 2005

[73] X. Yang, J. Bigham, L. Cuthbert, Resource management for service providers in heterogeneous wireless networks, in: IEEE Wireless Communications and Networking Conference 2005, vol. 3, 13–17 March 2005, pp. 1305–1310.

[74] "Key Performance Indicators (KPI) for Evolved Universal Terrestrial Radio Access Network (E UTRAN): Definitions," TS 32.450, Version 9.1.0 Release 9, June 2010.

[75] 3GPP, "Conveying MCS and TB size via PDCCH," R1-081483, Septem- ber 2010.

[76] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA): Physical layer procedures (Release 9)," TS 36.213 V9.3.0, September 2010.

[77] IEEE 802.16 Broadband Wireless Access Working Group, "Evaluation Methodology for P802.16m-Advanced Air Interface," IEEE 802.16m- 08/004r2, 2008.

[78] WiMAX Forum, "WiMAX System Evaluation Methodology," V2.1, July 2008.

[79] "ns-3 Model Library – Release ns-3.17", May 14, 2013.

[80] “ns-3 Model Library – Release ns-3.18”, August 29, 2013

[81] Recommended encoding settings for YouTube, https://support.google.com/youtube/answer/1722171?hl=en