

Resource allocation for network-controlled device-to-device communications in LTE-Advanced

G. Nardini¹ · G. Stea¹ · A. Viridis¹ · D. Sabella² · M. Caretti²

Published online: 12 January 2016
© Springer Science+Business Media New York 2016

Abstract Network-controlled device-to-device (D2D) communication allows cellular users to communicate directly, i.e., without passing through the eNodeB, while the latter retains control over resource allocation. This allows the same time–frequency resources to be allocated to spatially separated D2D flows simultaneously, thus increasing the cell throughput. This paper presents a framework for: (1) selecting which communications should use the D2D mode, and when, and (2) allocating resources to D2D and non-D2D users, exploiting reuse for the former. We show that the two problems, although apparently similar, should be kept separate and solved at different timescales in order to avoid problems, such as excessive packet loss. We model both as optimization problems, and propose a heuristic solution to the second, which must be solved at millisecond timescales. Simulation results show that our framework is practically viable, it avoids the problem of packet losses, increases throughput and reduces delays.

Keywords LTE-Advanced · Device-to-device · Resource allocation · Optimization

1 Introduction

Network-controlled device-to-device (D2D) communication [1] is gaining rapid acceptance among telecommunication operators as a means to achieve very low latencies and to favor spectrum reuse within LTE-Advanced networks [2]. In a traditional cellular network, every communication is between a User Equipment (UE) and the eNodeB (eNB), thus two UEs can only communicate via a two-hop path, even though they are within hearing range of each other. Since resource allocation is always performed by the eNB—in both the uplink and the downlink—and UEs are simply told where to transmit/listen to their data, the eNB may *simultaneously* target both endpoints, having one transmit on the same Resource Blocks (RBs) where the other is told to receive. This brings several benefits: on one hand, a *single-hop* path is used, hence the delay is reduced and the radio resources required to support the same flows are (roughly) halved. On the other hand, and more importantly, since a UE's transmitting power is small, *spatial reuse* is made possible: the same RBs can be allocated simultaneously to flows whose receivers are sufficiently far apart to tolerate the interference of each other's transmitter. In the presence of highly localized transmissions, such as vehicular ad hoc networks, gaming, peer-to-peer file exchange, there is in fact room for a considerable increase in the cell capacity, with little investment from the operator.

Enabling network-controlled D2D requires solving several problems. One is *discovering* which pairs of UEs are eligible for D2D communication, due to proximity

✉ G. Stea
giovanni.stea@unipi.it

G. Nardini
g.nardini@ing.unipi.it

A. Viridis
a.viridis@iet.unipi.it

D. Sabella
dario.sabella@telecomitalia.it

M. Caretti
marco.caretti@telecomitalia.it

¹ Dipartimento di Ingegneria dell'Informazione, University of Pisa, Pisa, Italy

² Telecom Italia Lab, Turin, Italy

and—possibly—other enablers (such as terminal class or subscription profile). Another, orthogonal to the former, is *resource allocation*. This problem has two facets: *mode selection*, i.e., determining which D2D-eligible flows will actually use the direct path, and MAC-level *packet scheduling*, i.e., determining which RBs go to which flows, in the presence of both D2D and non-D2D users. The latter problem presents peculiar challenges with respect to traditional LTE scheduling, due to spatial reuse among D2D UEs and to the need of allocating resources to D2D and non-D2D UEs simultaneously. Several works have appeared lately on resource allocation for D2D systems (see, e.g., survey [3], and [4–11]). Some only tackle mode selection (e.g., [4, 5]). Others advocate solving both problems at the same time ([7, 8, 10]). Others, finally, consider both problems as disjoint, and solve them imposing restrictions on the allocated resources ([6, 9, 11]).

In this paper, we explore resource allocation for network-controlled D2D systems. We first show that the two problems of mode selection and packet scheduling are fundamentally different, and should be solved at different timescales: the former at *coarse* timescales (hundreds of ms or more) and the latter at the Transmission Time Interval (TTI) timescale, i.e., 1 ms. Faster mode selection, in fact, brings *route flapping*, which increases packet losses; on the other hand, slower packet scheduling fails to exploit channel adaptation, thus making resource allocation ineffective. We then design a practical framework where the two problems are solved at the relevant timescales, and information is fed back from one to the other. We show that it is practically feasible to select modes by solving a Mixed Integer-Linear Problem *at optimality*, and for packet scheduling we propose both an optimal formulation and a low-complexity heuristic that consider both D2D and non-D2D UEs simultaneously. Our framework is *general*, in that it easily accommodates different objectives, for instance maximizing the overall throughput, minimizing the energy consumed by the eNB, and others. Our results, obtained in a simulation environment that includes all the LTE data-plane protocols, show that the mode selection period influences packet loss. Moreover, selecting modes *carefully*, by taking into account global metrics (such as the available resources in both the uplink and the downlink) brings significant benefits, such as higher cell throughput and reduced delays.

The rest of the paper is organized as follows: Sect. 2 reports background on LTE-A. We detail the system model hypotheses in Sect. 3, and present our resource allocation framework in Sect. 4. Performance results are reported in Sect. 5, whereas Sect. 6 reviews the related work. Finally, Sect. 7 concludes the paper.

2 Background on LTE-A

We first provide background regarding LTE-A layering and resource allocation, and then move to introducing device-to-device communications.

User data coming down from the IP layer traverse the Packet Data Convergence Protocol (PDCP), and from there arrive at the Radio Link Control (RLC) in the form of RLC SDUs. As shown in Fig. 1, the RLC performs segmentation and concatenation of RLC SDUs into RLC PDUs, whose size is decided by the MAC. More in detail the RLC layer can be configured to work in three modes, namely *transparent* (TM), *unacknowledged* (UM) or *acknowledged* (AM). The first one performs a one-to-one mapping between RLC SDU and PDUs. The second one performs *segmentation/concatenation* of SDUs on transmission, and *reassembly*, duplicate detection and reordering of PDUs on reception. The third one adds reliable PDU transmission capabilities on top of that. RLC PDUs are then encapsulated into MAC PDUs: each of them contains the information being sent over the physical layer, and is called Transmission Block (TB) [12]. RLC PDUs are physically buffered at the RLC.

MAC-layer transmissions are arranged in subframes and paced at Transmission Time Intervals (TTIs) of 1 ms. In the downlink (DL), the eNB allocates a vector of *Resource Blocks* (RBs) to the UEs associated to it on each TTI. Each RB carries a fixed number of symbols, but a *different* number of bits depending on the Channel Quality Indicator (CQI) reported by the UE. The CQI depends on the measured Signal to Interference and Noise Ratio (SINR), and it determines the modulation that the eNB will use (hence, indirectly, the number of bytes per RB). Error recovery at the MAC level is provided by a Hybrid ARQ (H-ARQ) scheme, which allows a configurable number of retransmissions. Downlink H-ARQ processes are *asynchronous*, i.e., a retransmission can be scheduled by the eNB at any future TTI.

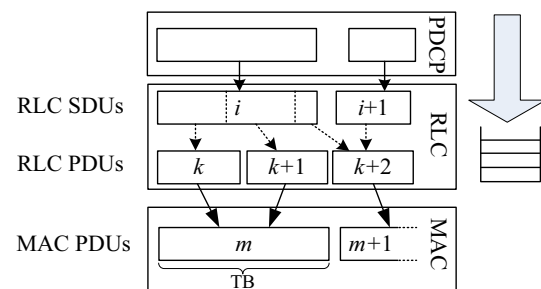


Fig. 1 Example of user data flow through the LTE protocol layers

In the uplink (UL), the eNB issues *transmission grants* for each UE, specifying which RBs they can use, using what transmission format. However, UL buffers reside at the UEs, hence the latter must notify their buffer status to the eNB. This is done via quantized *Buffer Status Reports (BSRs)*, transmitted (either alone or trailing a data transmission) *in band*, i.e. together with the data. Thus, they can only be sent i) when the UE is scheduled, and ii) if there is enough space to do so (a BSR can take up to 24 bits). Therefore, an *out-of-band* Random Access Procedure (RAC) is used to allow a UE to signal the presence of backlog. RAC requests may collide, and collisions are resolved through a backoff procedure. RAC requests are responded in-band, by scheduling the UE in a future TTI, and are re-iterated if unanswered. The standard handshake for UL transmissions, shown in Fig. 2, takes five messages: first the UE sends a RAC request; the eNB responds with a short grant, large enough for a BSR; the UE sends its BSR; the eNB sends a larger grant according to its scheduling policy, and the UE transmits its data. The middle two interactions can be dispensed with if the UE is enabled to send one or more PDUs together with the BSR, a technique known as *bandwidth stealing*, which saves UL resources and reduces latency. Semi-Persistent Scheduling (SPS) can also be used to transmit periodic traffic, e.g., VoIP. It consists in the eNB issuing a long-term, periodic grant to a UE, which can then transmit in the pre-assigned RBs without signaling. However an SPS grant sets—once and for all—the *format* of the uplink transmissions based on the channel conditions at the moment of decision, thus preventing link adaptation. The unavoidable mismatch between the *presumed* and *actual* UE channel quality results in either a higher Block Error Rate (BLER) or an overdimensioned grant, thus reducing cell capacity. Uplink H-ARQ is *synchronous* i.e. a retransmission is scheduled exactly eight TTI after the previous attempt.

2.1 Network-controlled device-to-device communications

Device-to-device (D2D) or *direct* communications are being standardized in the 3GPP community [1]. Assume UE *a* wants to transmit to UE *b*: this would normally require a




Fig. 2 Standard uplink scheduling (*left*) and bandwidth stealing (*right*)

two-hop transmission with the eNB as a relay. Instead, network-controlled D2D consists in having *a* use one or more RBs for transmission and *b* use the same RBs for reception. The link where this communication occurs is often referred to as *sidelink (SL)*, to distinguish it from the UL and DL.

The standard is mainly concerned with *one-to-many* D2D communications. In the latter, the UE transmits to proximate neighbors, either acting as a relay to increase the cell coverage, or sending data of its own. Two modes are foreseen for network-controlled resource allocation. In the first mode, the eNB schedules RBs for *a*'s transmission and informs *b* that it should receive in the same RBs where the grant is being issued (slots 5 and 3 of Fig. 2, respectively). Alternatively, UE *a* selects RBs from a resource pool on its own, where the pool may be configured statically or semi-statically by the eNB. Clearly, the first mode allows more flexibility at the cost of higher control overhead. In a Frequency Division Duplexing (FDD) deployment, D2D communications can take place using physical resources taken from either or both the UL and DL subframes. The standard recommends using the RLC UM for D2D communications, and ACKing of PDUs is not provided [1].

On the other hand, some literature—e.g. [13]—(although no standard document) exists on *one-to-one* D2D communication. In the latter, the endpoints are identified before communication takes place, through means that are outside the scope of this paper, with the eNB acting as a resource allocator using either of the two above modes.

Finally, D2D communications can also happen *without* the assistance from the network. For example, D2D can be used in out-of-coverage scenarios, especially for public safety applications. In this case, UEs cannot rely on the eNB support for scheduling, synchronization and other control functions. The rest of the paper focuses only on network-controlled, one-to-one D2D communications.

3 System model

We focus on a single cell served by an eNB. Each UE in the cell is an endpoint to at least one unidirectional unicast flow. Call **F** the set of flows. Let **D** and $\bar{\mathbf{D}} = \mathbf{F} \setminus \mathbf{D}$ be the sets of *D2D-eligible* and *D2D-ineligible* flows (see Fig. 3). The former are those whose endpoints are both in the cell, (electromagnetically) near enough for direct communication to take place. The latter are flows whose endpoints are outside the reach of each other, possibly in different cells, or are legacy devices, which can only communicate through the eNB. The means through which a flow (and, accordingly, its transmitter and receiver, there being no ambiguity) is classified as D2D-eligible (e.g. endpoints' geographical proximity) is outside the scope of this paper, and we refer the interested reader to [14] for an example of

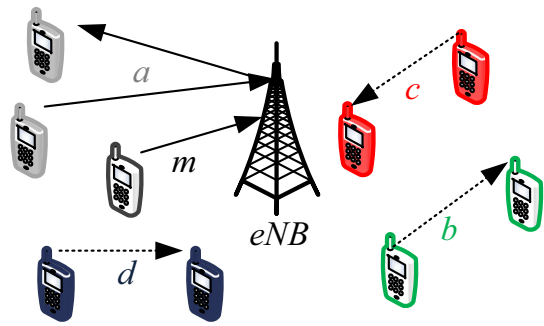


Fig. 3 DM and IM flows: a, b, c, d are in \mathbf{D} , m is in $\bar{\mathbf{D}}$. b, c, d are in DM whereas a is in IM

how to achieve this. Endpoints of flows in \mathbf{D} can either communicate directly, i.e., in *D2D mode (DM)*, or using the eNB as a relay, i.e., in *Infrastructure mode (IM)*. Each flow i has a *rate request* A_i , which the system should try to meet if possible at all. This can be inferred from signaling information at flow setup, e.g. the application QoS Class Indicator (QCI). For flows in \mathbf{D} , we assume that the eNB is aware of: (1) the CQIs in the UL and the DL of an IM communication, and (2) the CQI of the DM link.¹

We consider an FDD system, where UL and DL subframes have M RBs each. We assume that DM transmissions take place in the UL subframe, which is less loaded than the DL one (due to the well-known traffic asymmetry) and allows better overall SINR, especially when the D2D users are far from the eNB [15]. Accordingly, UEs are equipped with a Single Carrier-Frequency Division Multiple Access (SC-FDMA) receiver [13]. As far as H-ARQ is concerned, [13] argues that the receiver feedback (i.e., ACK or NACK) may be either *direct* or *indirect*, as shown in Fig. 4a, b. In both cases, the ultimate destination of the feedback is the transmitter. Instead, we posit that the receiver feedback should reach *both* the transmitter and the eNB (see Fig. 4c), since the former needs to know if retransmission is required (in eight TTI, it being UL), while the latter has to allocate RBs for it to take place.

Several DM flows may reuse the same RBs, hence their transmissions may interfere with each other. We model interference using a *binary* model: two DM flows are either *interfering*, if the power they receive from each other is above a given threshold, hence they must occupy disjoint RBs, or *non-interfering*, in which case they may share the same RBs. Interference conditions can thus be modeled via a *conflict graph* (CG) [16]. Vertices of the CG are flows in \mathbf{D} and an edge exists between two interfering (i.e.,

¹ The means by which this is achieved are outside the scope of this paper. However, for the sake of concreteness, we mention that the eNB may request D2D-eligible receivers to measure and report certain signals produced by the related transmitter, e.g. as proposed in [14].




Fig. 4 H-ARQ feedback reporting

conflicting) flows. Let \mathbf{E} be the set of edges in the CG. Figure 5 represents the CG for the scenario of Fig. 3, whereas Fig. 6 shows a possible resource allocation which exploits spatial reuse under the CG constraints. Reuse between one IM flow and one or more DM ones is not considered: we will discuss later on that enabling it would make solving the MS problem at optimality prohibitive.

Allocating resources in this system requires two decisions:

1. which flows in \mathbf{D} should be scheduled as DM or IM, a decision which we call *mode selection (MS)*;
2. how to allocate RBs to *all* flows in \mathbf{F} on each TTI, once the above decision has been made, which we call *packet scheduling (PS)*.

Hereafter, we present a general framework to solve the above problems.

4 Resource allocation framework

We first argue that MS and PS should occur at very different timescales: the MS at a coarse one (e.g., hundreds or thousands of TTIs) and the PS once per TTI. Intuitively, the MS should—in theory—be run as frequently as possible, so as to ride channel peaks and select the best mode for each flow. However, when a flow is switched (e.g.) from DM to IM, a burst of losses may occur. In fact, a sender uses different radio bearers to peer with different receivers. As stated in [17], different radio bearers come with different PDCP and RLC entities, which in turn have independent sequence numbering. Thus, *PDU sequence number cannot be assumed to be synchronized* on different RLC connections. With reference to Fig. 7, consider a DM flow having UEs a and b as endpoints. At time t_1 , b 's RLC expects a PDU whose sequence number is 101. However, if the mode is switched to IM at t_1 , then b will peer with the eNB's RLC entity, which will expect a different next sequence number (251, in this example). Note, further, that the eNB relays PDUs from a to b employing *two* separate RLC entities in any case, peering with a and b respectively, and that UL RLC PDUs coming from a will travel the

Fig. 5 Example of conflict graph including D2D-eligible flows





Fig. 6 Example of UL subframe allocation with spatial reuse, compatible with the mode selection of Fig. 3

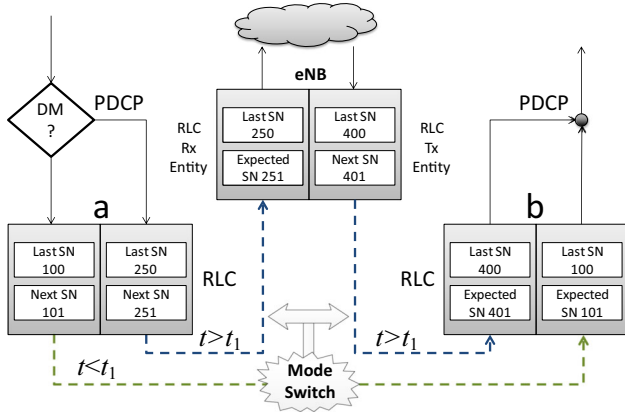


Fig. 7 Effects of mode-switch on RLC communication

protocol stack up to at least the IP layer before climbing it down again for the DL leg. RLC PDUs in the DL leg will thus have yet another, unrelated numbering, as well as a different fragmentation. Having *separate* RLC entities implies that, after a mode switch, all PDUs buffered at the “old” RLC entity will need to be discarded, since their numbering is unsuitable for the new RLC connection. Moreover, the flow may not switch back to the old one anytime soon (or ever again), hence keeping them buffered for possible later use makes no sense. Depending on how fast the PDCP sends down SDUs to the RLC, and how slowly the MAC requests RLC PDUs to the RLC, these losses may be significant. Moreover, *fragments* of RLC SDUs that have already made it to *b* at the time of the mode switch will be discarded as well (when *b*’s RLC reassembly timeout expires), since their missing counterparts, buffered at *a*’s “old” RLC entity, will be discarded. The same problem occurs, obviously, also when switching from IM to DM, where it is exacerbated by the fact that losses can occur independently on the UL and DL leg of the IM path. We also observe that using a different RLC—notably AM, despite recommendation to the contrary in the standard [1]—would not solve the problem. In fact, although AM allows a sender to know which RLC PDUs have/have not




Fig. 8 High-level block diagram of the resource allocation framework

made it to the peer entity, when switching from IM to DM there is no way for *a* to know what has got to *b*, since *b*’s RLC—*b* being two hops away—is not peering with *a*’s at that moment.

Now, losses due to frequent mode switching impair the throughput of the flow and of the cell as well, as we show later on in Sect. 5. Minimizing the impact of these losses is in fact one of the main reasons why MS *must* be performed at *coarse* timescales, the other one being that its overhead is non negligible. On the other hand, it is self-evident that PS should occur on a per-TTI basis, since the backlog and channel state of UEs change at a comparable timescale. Long-term resource booking, as discussed in Sect. 2, is in fact inefficient. Finally, we underline that the sets of flows on which MS and PS decisions are to be made are different: MS is for flows in **D**, PS is for *all* flows in **F**.

Figure 8 reports a high-level view of our solution. The MS module makes decisions at *periods* of *T* TTIs (in the order of hundreds or thousands of TTIs). It selects between DM and IM for each flow in **D**, taking into account rate requests, *long-term* channel quality statistics, availability of resources in both the UL and DL subframes, and simultaneous presence of D2D-ineligible flows. The PS module actually allocates resources on each TTI, taking MS decisions as inputs, and based on the *instantaneous* channel quality and UE buffer status. An *aggregator* is required to average TTI-level information over a period and provide input to the MS module, i.e., achievable rates and subframe occupancy.

From a network operator perspective, the objective of resource allocation could be one among meeting rate requests, minimizing energy consumption, maximizing fairness among UEs or some combination of the above. The framework proposed in this paper is general enough to account for each of the above objectives. For ease of explanation, we will detail the solution in the hypothesis that sources are constrained by a traffic profile specifying a maximum long-term sending rate, and we will then show that little modifications are required to pursue alternative objectives.

4.1 Mode selection

The MS module runs a decision algorithm that states on which path a communication should take place, whether DM or IM. While flows in $\bar{\mathbf{D}}$ are not involved in MS, they

do occupy resources in both the UL and DL subframes. Let M^{UL} and M^{DL} be the average, taken over the last period, of the RBs left unused by flows in $\bar{\mathbf{D}}$ in the UL and DL subframes, respectively. Although we assume that DM communications insist on the UL subframe, when a flow is scheduled as IM its data must travel to the eNB, then it must come down in the DL. However, if the DL subframe does not have enough resources for scheduling that flow, communication may incur excessive delay. Thus, occupancy of the DL subframe should also be considered in the decision process of the MS module. For each flow $i \in \mathbf{D}$, let R_i^l be the average MAC-level rate achievable on link $l \in \{SL, UL, DL\}$, the latter two being the legs of the IM communication. These values are computed by the aggregator by averaging the rates of the last period. The output of the MS can be obtained as the optimum of an optimization problem. The latter can be formulated in two different ways, based on whether or not spatial reuse is allowed among DM flows. The solutions state which mode a flow will use to communicate during the next period, and give an indication of what percentage of the UL subframe it is likely to occupy. We discuss MS without spatial reuse first, and then add spatial reuse.

4.1.1 Mode selection without spatial reuse

The problem is as follows:

$$\begin{aligned}
 & \max \sum_{i \in \mathbf{D}} [(x_i^{SL} \cdot R_i^{SL} + x_i^{UL} \cdot R_i^{UL}) - (s_i + t_i) \cdot Q_i/T] \\
 & \text{s.t.} \\
 & \sum_{i \in \mathbf{D}} x_i^{UL} \cdot R_i^{UL} / R_i^{DL} \leq M^{DL} & (i) \\
 & \sum_{i \in \mathbf{D}} (x_i^{SL} + x_i^{UL}) \leq M^{UL} & (ii) \\
 & x_i^{SL} \cdot R_i^{SL} + x_i^{UL} \cdot R_i^{UL} \leq A_i & \forall i \in \mathbf{D} \quad (iii) \\
 & x_i^{SL} \leq \min\{M^{UL}, A_i/R_i^{SL}\} \cdot d_i & \forall i \in \mathbf{D} \quad (iv) \\
 & x_i^{UL} \leq \min\{M^{UL}, A_i/R_i^{UL}\} \cdot (1 - d_i) & \forall i \in \mathbf{D} \quad (v) \\
 & s_i \geq d_i - D_i^{old} & \forall i \in \mathbf{D} \quad (vi) \\
 & t_i \geq D_i^{old} - d_i & \forall i \in \mathbf{D} \quad (vii) \\
 & d_i, s_i, t_i \in \{0, 1\} \forall i \in \mathbf{D}, \quad x_i^{SL}, x_i^{UL} \in \mathbb{R}^+ \forall i \in \mathbf{D} & (viii)
 \end{aligned} \tag{1}$$

The objective function (to be maximized) is the sum of the throughput of each flow, which consists of two parts. The first one is the average rate that the flow is expected to achieve in the coming period: x_i^{SL} , x_i^{UL} are the average number of RBs virtually occupied by a flow i over the SL and UL, respectively. Note that we consider these to be real-valued, since we are concerned with long-term average rates, hence a continuous relaxation (i.e., a *fluid* approach) is justified. Since a flow i cannot be in both DM and IM

simultaneously in a period, either x_i^{SL} or x_i^{UL} will be null. The second addendum estimates the throughput *lost* due to a mode switch: when a flow is switched, buffered RLC PDUs (Q_i) will be discarded, hence a throughput loss equal to Q_i/T needs be considered, T being the MS period. s_i , t_i are two binary variables, either of which is set when the mode is switched. The rationale is to prevent a flow to continuously change the communication mode, unless the throughput gain is higher than the reduction due to discarded PDUs. Constraint (i) states that we cannot send in UL more traffic than we can carry in DL: the total UL throughput is in fact $x_i^{UL} \cdot R_i^{UL}$, and this must be equal to the DL one (otherwise either queueing at the eNB occurs or resources are wasted). The resources required in the DL subframe to achieve this goal are thus $x_i^{UL} \cdot R_i^{UL} / R_i^{DL}$. In practice, constraint (i) implies that—if the DL subframe is full—then it is wise to set as many flows as possible to DM. Constraint (ii) states that the allocation for all flows in \mathbf{F} must fit in the UL subframe. Constraint (iii) specifies that the flow throughput cannot exceed its maximum sending rate. Since this is a maximization problem, the solver will push the left hand side of (iii) towards the upper bound A_i in any case. If achieving *all* the A_i is impossible, then the MS will still select modes trying to achieve the highest cell throughput. Note that it is the PS that actually allocates RBs, and it does so using actual backlogs and CQIs, available on each TTI. Variable d_i is binary, and is set

when flow i is in DM. Hence, constraints (iv–v) ensure that either x_i^{DM} or x_i^{IM} is null. Constraints (vi–vii) state that binary variables s_i and t_i are set if there is a mode switch, i.e. if d_i differs from its previous value D_i^{old} . Finally, constraint (viii) specifies the problem variables and their domain. Note that *only* variable d_i will be used later on by the PS. The other variables, notably x_i^{SL} , x_i^{UL} , s_i and t_i , are only helper variables, since resource allocation is performed by the PS.

The above is a Mixed Integer-Linear problem (MIP), with $O(D)$ variables and constraints. MIPs are NP-hard in general. However, they can be solved quite fast, even at optimality, if their size is not too large.

4.1.2 Mode selection with spatial reuse


Spatial reuse among DM flows improves the overall throughput, as more UEs are served simultaneously in the same RBs. On the other hand, interference comes into play and must be managed. Conflict Graph (CG) serves this purpose. DM flows for which an edge exists in the CG must be serialized, i.e., must occupy different RBs. In the UL, Single Carrier-Frequency Division Multiple Access (SC-FDMA) imposes the restriction that RBs allocated to a UEs should be adjacent, hence we need a variable π_i to denote the RB index where flow i starts its transmissions. $\pi_i + x_i^{SL}$ determines the RB index flow i ends its transmissions. The left side of Fig. 9 shows an example of a valid allocation for two conflicting flows i and j , whereas the right side of Fig. 9 shows an invalid allocation, due to the overlapping between RBs allocated to i and j .

The MS problem with spatial reuse is the following.

$$\begin{aligned}
 & \max \sum_{i \in \mathbf{D}} [(x_i^{SL} \cdot R_i^{SL} + x_i^{UL} \cdot R_i^{UL}) - (s_i + t_i) \cdot Q_i/T] \\
 & s.t. \\
 & \sum_{i \in \mathbf{D}} x_i^{UL} \cdot R_i^{UL} / R_i^{DL} \leq M^{DL} \tag{i} \\
 & n^{SL} + \sum_{i \in \mathbf{D}} x_i^{UL} \leq M^{UL} \tag{ii} \\
 & x_i^{SL} \cdot R_i^{SL} + x_i^{UL} \cdot R_i^{UL} \leq A \quad \forall i \in \mathbf{D} \tag{iii} \\
 & x_i^{SL} \leq \min\{M^{UL}, A_i/R_i^{SL}\} \cdot d_i \quad \forall i \in \mathbf{D} \tag{iv} \\
 & x_i^{UL} \leq \min\{M^{UL}, A_i/R_i^{UL}\} \cdot (1 - d_i) \quad \forall i \in \mathbf{D} \tag{v} \\
 & \pi_i + x_i^{SL} \leq \pi_j + M^{UL} \cdot [o_{ij} + (1 - d_i) + (1 - d_j)] \quad \forall (i, j) \in \mathbf{E} \tag{vi} \\
 & \pi_j + x_j^{SL} \leq \pi_i + M^{UL} \cdot [(1 - o_{ij}) + (1 - d_i) + (1 - d_j)] \quad \forall (i, j) \in \mathbf{E} \tag{vii} \\
 & \pi_i + x_i^{SL} \leq n^{SL} + M^{UL} \cdot (1 - d_i) \quad \forall i \in \mathbf{D} \tag{viii} \\
 & s_i \geq d_i - D_i^{old} \quad \forall i \in \mathbf{D} \tag{ix} \\
 & t_i \geq D_i^{old} - d_i \quad \forall i \in \mathbf{D} \tag{x} \\
 & d_i, s_i, t_i \in \{0, 1\} \quad \forall i \in \mathbf{D}, \quad o_{ij} \in \{0, 1\} \quad \forall (i, j) \in CG \tag{xi} \\
 & x_i^{SL}, x_i^{UL}, \pi_i \in \mathbb{R}^+ \quad \forall i \in \mathbf{D}, \quad n^{SL} \in \mathbb{R}^+ \tag{xii}
 \end{aligned}$$

Constraints (i–v), (ix–x) are the same as in (1), the only difference being in (ii), where instead of the *sum* of the RBs over all DM flows, we use the *overall amount* n^{SL} of the RBs occupied by DM flows, thus accounting for reuse ($n^{SL} \leq \sum_{i \in \mathbf{D}} x_i^{SL}$). Constraints (vi–vii) define the sequencing of conflicting DM flows i and j in the subframe, i.e. the direction of the edges in the CG. Note that, due to the term $(1 - d_i) + (1 - d_j)$, these three constraints are automatically verified unless i and j are both DM flows. In this last case, whether (vi) or (vii) are active is determined by binary variable o_{ij} , which is 0 if i precedes j and 1 otherwise. Finally, constraint (viii) defines the right extreme n^{SL} of the portion of subframe allocated to DM flows. This MIP has $O(D + |\mathbf{E}|)$ variables and constraints, \mathbf{E} being the set of edges in the CG. Therefore, we can expect spatial-reuse MS to be harder than the one with mutually exclusive resources, all else being equal. We will show later on in Sect. 5 that the cost of solving (2) at optimality is however compatible with the timescale at which we want it to work. Note that allowing reuse between one IM flow and one or more DM ones would entail including D2D-ineligible flows into the CG as well. These in turn form a \bar{D} -clique of conflicts, thus adding \bar{D}^2 edges to \mathbf{E} . While the formulation changes only slightly, the time involved in

Fig. 9 Valid (left) and invalid (right) allocation for two conflicting flows



solving it at optimality becomes prohibitive, so that heuristics are needed. Finally, we observe that, when spatial reuse is allowed, values R_i^{DM} do account for the interference generated by all the (non-conflicting) D2D flows which may have shared RBs with flow i . Interference decreases the SINR, which in turn decreases R_i^{DM} . We will show in Sect. 5 that this

Under spatial reuse, we need to devise different schedulers, that exploit sharing RBs among non-conflicting DM flows. We call \mathbf{B}_{DM} and \mathbf{B}_{IM} the set of *backlogged* DM and IM flows, respectively. Let Q_i denote the backlog of flow i . The max-throughput schedule for both DM and IM flows is the optimum of the following problem:

$$\begin{aligned}
 & \max \sum_{i \in \mathbf{B}_{DM}} (b_i \cdot CQI_i^{SL} - p_i) + \sum_{i \in \mathbf{B}_{IM}} (b_i \cdot CQI_i^{UL} - p_i) \\
 & \text{s.t.} \\
 & n^{SL} + \sum_{i \in \mathbf{B}_{IM}} b_i \leq M \quad (i) \\
 & b_i \cdot CQI_i \leq Q_i + p_i \quad \forall i \in \mathbf{B}_{DM} \cup \mathbf{B}_{IM} \quad (ii) \\
 & p_i \leq CQI_i - 1 \quad \forall i \in \mathbf{B}_{DM} \cup \mathbf{B}_{IM} \quad (iii) \\
 & \pi_i + b_i \leq \pi_j + M \cdot o_{ij} \quad \forall (i, j) \in \mathbf{E} \quad (iv) \\
 & \pi_j + b_j \leq \pi_i + M \cdot (1 - o_{ij}) \quad \forall (i, j) \in \mathbf{E} \quad (v) \\
 & \pi_i + b_i \leq n^{SL} \quad \forall i \in \mathbf{B}_{DM} \quad (vi) \\
 & \pi_i \in \mathbb{Z}^+ \quad \forall i \in \mathbf{B}_{DM} \quad (vii) \\
 & b_i, p_i \in \mathbb{Z}^+ \quad \forall i \in \mathbf{B}_{DM} \cup \mathbf{B}_{IM} \quad (viii) \\
 & n^{SL} \in \mathbb{Z}^+ \quad (ix) \\
 & o_{ij} \in \{0, 1\} \quad \forall (i, j) \in \mathbf{E} \quad (x)
 \end{aligned} \tag{3}$$

makes the MS problem robust to misconfigurations of the conflict threshold.

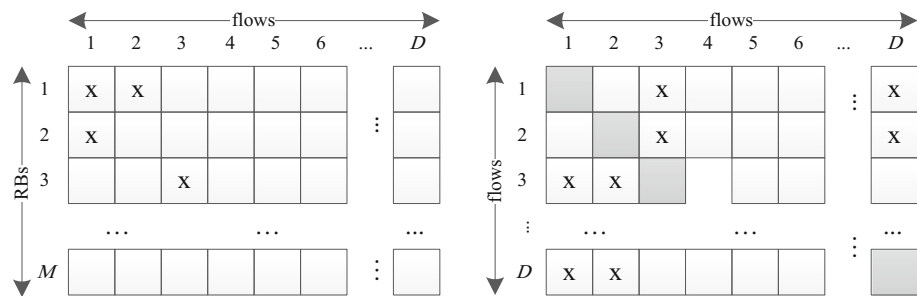
4.2 Packet scheduling

Once modes are selected, we need a packet scheduler that actually allocates the RBs of the UL subframe on each TTI, to all flows in \mathbf{F} .² In fact, values x_i^{SL} , x_i^{UL} computed by the MS are based on *long-term* average rates, inferred from those achieved in the past period, and maximum rates A_i . PS instead exploits fresh CQIs and actual instantaneous backlogs to make an efficient TTI-level resource allocation. Note that, at this stage, DM flows are a *subset* of \mathbf{D} , whereas IM flows are a set that *includes* $\overline{\mathbf{D}}$. When spatial reuse is not allowed, any traditional scheduling algorithm suitable for the uplink (e.g., MaxC/I, Proportional Fair, Round Robin, etc.) can be used, since DM and IM flows need not be distinguished in principle. Priority-aware versions of the above schedulers can also be used, if necessary, e.g. to prioritize DM flows over IM ones. Strictly speaking, the PS need not even have the same objective as the MS: for instance, both MaxC/I and PF scheduling can coexist with MS (1). Using one or the other will yield *different* long-term rates R_i^l (likely, lower for PF than for MaxC/I), and MS will adapt accordingly.

Problem (3) aims to maximize the amount of backlog cleared (minus the padding p_i), where b_i is the number of RBs allocated to flow i . CQI_i is the number of bits per RB for flow i in its selected mode. Constraint (i) states that all the schedule must fit in the UL subframe. Constraint (ii) states that the number of RBs given to a flow is upper bounded by its backlog, plus some allowance for padding, which is however limited to less than one RB's worth by constraint (iii). Constraints (iv–vi) enforce conflict graph precedence among the allocations for DM flows. Whether (iv) or (v) are active is determined by binary variable o_{ij} , which is 0 if i precedes j and 1 otherwise. These constraints are similar to (vi–vii) of Problem (2). However, recall that variables o_{ij} in (2) are only helper variables and are not conveyed to the PS. In fact, PS is free to decide to allocate either i before j in the subframe or vice versa. Mutual exclusion among IM flows, implicitly defined by the LTE standard, is taken into account in constraint (i). Note that b_i and π_i are integer variables now, since we are now assigning RBs in a single subframe instead of making a long-term resource planning. Problem (3) is a MIP with $O(|\mathbf{B}_{DM}| + |\mathbf{B}_{IM}|)$ variables. Solving it in a TTI time may be infeasible if many flows are backlogged simultaneously, hence we present a low-complexity heuristic.

Assume that the allocation is recorded into a binary $M \times D$ RB matrix (shown in Fig. 10, left). Element (i, j) of the RB matrix is equal to one if RB i is allocated to flow j . The CG can also be represented by a $D \times D$ binary

² Note that PS on the DL subframe is unaffected by our framework, hence is not mentioned further in the paper.

Fig. 10 RB matrix (*left*) and Conflict Graph matrix (*right*)**Fig. 11** Pseudo-code for the best-fit heuristic

```

1  IM_bottom = M // last RB unoccupied by IM
2  DM_top = 0 // topmost RB occupied by DM
3  RB_matrix = all 0s
4  IM_allocable = TRUE
5  sort flows by decreasing CQI
6  while (flow list not empty)
7    extract top flow i
8    B = ceiling(Qi/CQIi)
9    if (i is IM)
10   | if (IM_allocable)
11   |   B = min{B, IM_bottom-DM_top}
12   |   allocate B RBs to flow i
13   |   IM_bottom -= B
14   |   IM_allocable = (IM_bottom > DM_top)
15   | end if
16   else // DM flow
17   | j = 1
18   | best = try = 0
19   | best_RBs = try_RBs = {}
20   | while (j <= IM_bottom)
21   |   if (RB_matrix[j] AND CG[i]) == 0
22   |   | try++
23   |   | try_RB = try_RB U {j}
24   |   else
25   |   | if (B <= try < best) OR (best < try < B)
26   |   |   best = try
27   |   |   best_RBs = try_RBs
28   |   | end if
29   |   | try_RB = {}
30   |   | try = 0
31   |   end if
32   |   j++
33   | end while
34   | while (B > 0) AND (best_RBs not empty)
35   |   extract top k from best_RBs
36   |   set bit (k,i) in RB_matrix
37   |   B--
38   end if
39 end while

```

matrix (shown in Fig. 10, right). Element (i,j) of the CG matrix is equal to one if arc (i,j) is in the CG. With reference to the pseudocode of Fig. 11, the heuristic works as follows: DM and IM flows are sorted in a single list by decreasing CQI (line 5). Cyclically, flows are extracted from the top of the list and served until their backlog B is cleared or no more RBs are available. However, IM flows are allocated *backwards* starting from the *end* of the UL subframe (lines 9–15), whereas DM flows are allocated

forward starting from the *beginning* of the subframe, using a *best fit* algorithm to regulate the conflicts: flow i is allocated a set of contiguous RBs not occupied by conflicting DM flows (lines 17–37) until the part of subframe occupied by IM flows is met. This allows us to do without enforcing a fixed resource partitioning between the two groups of flows. Instead, the boundary changes dynamically, based on the backlog and conflict state of the flows. For each DM flow i , requesting b_i RBs, the algorithm




Fig. 12 Detailed information flow for the framework

searches for the available groups of contiguous RBs and allocates either (line 25):

- b_i RBs in the *smallest* group larger than or equal to b_i , if there is one, or,
- the size of the *largest* group smaller than b_i , otherwise.

Checking if RB j can be allocated to DM flow i only requires to AND row j of the RB matrix and row i of the CG matrix (line 19). The complexity of the heuristic is $O(F \log F + D \cdot M)$: the first term is due to flow sorting (and is common to schedulers without spatial reuse, such as MaxC/I or PF), and the second is due to the inner while loop that finds available RBs for a DM flow, assuming that the bitwise AND of line 21 can be done in constant time. Such assumption is reasonable with large machine words. This complexity is clearly affordable in a TTI, given that M can be up to 100 in practice.

We are now in a position to recap the details of our framework, which we do in Fig. 12. The latter shows in detail the flow of information within our framework. In particular, we highlight two aspects. First, MS only indicates which flows are DM and which ones are IM, thus the PS is free to obtain the TTI-level resource allocation without using the long-term allocation performed by the MS. Second, PS provides the MS (via the aggregator) the information about the amount of available RBs for both UL and DL subframes, i.e. the RBs left unused by flows in $\bar{\mathbf{D}}$.

4.3 Generalizations

We now show how the two main blocks of our framework, namely MS and PS can be adapted to pursue different high-level objectives.

As far as MS is concerned, problems (1) and (2) can be readily transformed into others, simply by modifying the objective function and/or few constraints. For instance, proportional fairness can be simply enforced by using the following objective function in (1) or (2):

$$\max \sum_{i \in \mathbf{D}} \left[\left(\frac{x_i^{SL} \cdot R_i^{SL} + x_i^{UL} \cdot R_i^{UL}}{r_i} \right) - (s_i + t_i) \cdot Q_i/T \right],$$

where r_i is the long-term rate of the flow in the previous period. Moreover, an objective that is of interest for future

5G networks is maximizing the energy efficiency of the system, computed as the spectral efficiency divided by the energy consumed by the whole system [18–21]. According to well-established eNB energy consumption models [22], the consumed power is proportional to the number of RBs transmitted by the eNB in the DL subframe. On the other hand, the eNB consumes little energy for DM flows, i.e., the one required to respond RAC handshakes, which is negligible, and the one for transmitting signaling and control information. Since we assume a constant transmit power for the eNB, maximizing an energy efficiency function like the one proposed in [18–21] results in minimizing the energy consumed by the eNB. This can be achieved by limiting the transmissions having a DL leg, i.e., by sending as many of \mathbf{D} 's flows in DM as possible. Thus, a *min-energy* (mE) MS can be obtained by substituting objective function $\min \sum_{i \in \mathbf{D}} x_i^{UL} \cdot R_i^{UL} / R_i^{DL}$ in (1–2). The latter sets modes so as to minimize the portion of the DL subframe booked for flows in \mathbf{D} . Constraint (i) then becomes redundant, and the inequality in constraint (iii) must be reversed, since we are now facing a minimization problem and we need to satisfy the requested rate of flows in \mathbf{D} (otherwise the solver would just nullify each UL assignment). The alert reader will observe that other objectives (e.g., max–min fairness) can be accommodated via similar modifications.

Regarding PS, the same generalizations as above can be applied to problem (3), *mutatis mutandis*. For instance, an optimal PF with spatial reuse can be formulated simply by dividing each term in the objective function by the long-term rate of their flow. More interestingly, the heuristic of Fig. 11 can be reused as well: all it takes is to sort flows *differently* (e.g., by decreasing PF score) to enforce the required criterion. Moreover, an orthogonal choice can even be made on whether to schedule DM and IM flows jointly, sorting both in the same list—as we do in our framework—or by prioritizing either set, by using *two* sorted lists and scanning one before the other.

5 Performance evaluation

In this section we evaluate the performance of the framework presented in Sect. 4. Our evaluation is carried out using SimuLTE [23, 24], a C++ system-level simulator developed for the OMNeT++ framework [25]. SimuLTE simulates the data plane of the LTE/LTE-A radio access network, including all the protocol stack from the PDCP to the physical layer. It also simulates resource allocation at the eNB in the DL and UL subframes. In order to support network-controlled D2D communications, we had to implement several protocol-related functions besides all the framework described in Sect. 4, and, while doing so,

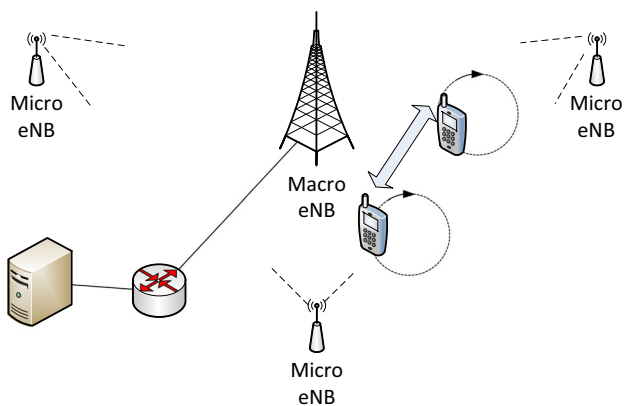


Fig. 13 Reference scenario for the simulations

Table 1 Simulation parameters

Parameter	Value
Carrier frequency	2 GHz
Bandwidth	5 MHz
Number of RBs	25
Mobility model	Circle (OMNeT++ model)
Path loss model	ITU urban macro
Fading model	Jakes
eNB Tx power	40 dBm (macro), 20 dBm (micro)
eNB antenna gain	18 dB
Noise figure	5 dB
Cable loss	2 dB
Simulation time	30 s
No. independent replicas	5

we had to make decisions on a non-negligible practical issue, which we have not observed in any related work so far: MS interferes with H-ARQ. In fact, a mode switch might occur when some H-ARQ process is in retransmission. At least two strategies may be envisaged to deal with this: one is to only switch *idle* H-ARQ processes, allowing busy ones to complete their cycle of retransmission before switching them too, the other one is that all pending retransmissions are aborted and started anew in the new mode. Solving this issue is clearly outside the scope of this paper, thus we stick to the first solution henceforth.

We simulate an LTE-A cell managed by one eNB endowed with an omnidirectional antenna (see Fig. 13). Inter-cell interference is generated by three micro-eNBs located at 1000 m, at a 120° angle. Micro-eNBs are anisotropic, pointing toward the macro-eNB, and the attenuation is $A(\theta) = \min \{12 \cdot (\theta/70^\circ)^2, 25\}$, θ being the angle between the micro and the receiver. eNBs use the same transmit power over the whole spectrum. The interference generated by micro-eNBs reduces CQIs in the DL so as to simulate a more realistic scenario. We assume 5 MHz

bandwidth, resulting in 25 RBs per subframe, in order to reach saturation with a reasonable number of users. Path loss, shadowing and fading models are taken from [26]. System parameters are summarized in Table 1. We simulate both D2D-eligible and D2D-ineligible flows. The latter can be either UL or DL flows, which means that these flows have either the transmitter or the receiver within the cell, and communicate with a server which sits beyond a router. We assume instead that the packets of IM flows having *both* endpoints in the cell reverse their direction *at the eNB*, without going any further (e.g., all the way up and down some dedicated network element in the Evolved Packet Core network). Such assumption, which is common to all the literature, allows us to discount non-radio access delays and bandwidth and processing bottlenecks. Unless specified otherwise, UEs follows circular trajectories whose center is at a maximum distance of 100 m from the eNB and whose radius ranges from 20 to 50 m, at a maximum speed of 10 m/s. The centers of the trajectories of two peering UEs are at a maximum distance of 25 m. This is done to allow communicating UEs to spend amounts of time where *both* modes are actually feasible, albeit with different CQIs. The UM RLC is used at all nodes.

Our evaluation is aimed at highlighting which benefits are due to which design choice. Thus, we start with evaluating the MS period and overhead; we then compare our MS with others available in the literature. Moreover, we show that the computation cost of the optimal PS is infeasible, and that our best-fit heuristic performs close to the optimum. Finally, we assess the impact of spatial reuse.

5.1 Mode selection period and overhead

We first investigate the effects of varying the MS period on the loss rate. Then, we provide figures regarding the overhead of solving MS at optimality.

As previously said, mode switching implies discarding of RLC buffers. Figure 14 shows the packet loss *at the RLC level* as a function of the MS period T . It is clear that periods below 500 ms are unadvisable, since the mode for D2D-eligible flows is switched too frequently. Moreover, note that we allow pending H-ARQ retransmissions to complete before switching, and losses could be even higher otherwise. From now on, we set the period at $T = 1$ s.

We then simulate 25, 30 and 35 D2D-eligible flows, 10 D2D-ineligible UL flows and 50 D2D-ineligible DL flows, transmitting CBR traffic (a 100-byte packet every 10 ms). Figure 15 shows the cumulative distribution function (CDF) of the time needed to solve the MS problem at optimality, using the ILOG CPLEX solver [27] on an Intel(R) Core(TM) i7 CPU at 2.80 GHz, with 8 GB of RAM and a Linux Ubuntu 12.04 operating system. As the figure shows, these times are mostly lower than 100 ms,

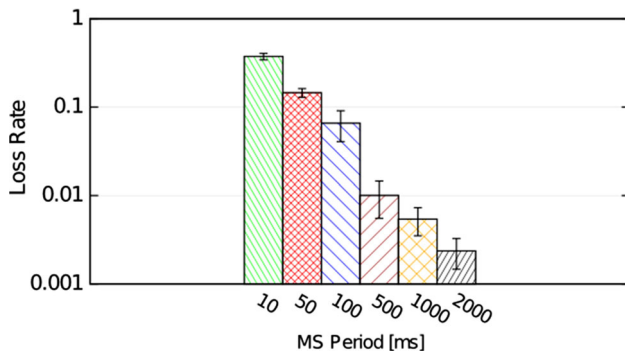


Fig. 14 Average RLC packet loss ratio with different MS periods

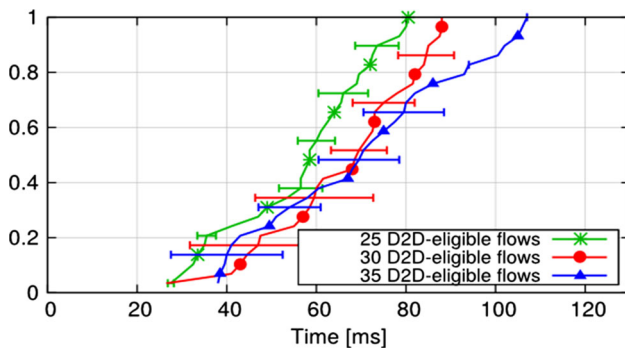


Fig. 15 CDF of the execution time of the MS problem

thus they are fully compatible with the period at which it is safe to run MS, e.g. 1 s. Note that no attempt at an optimized solution *algorithm* for problem (2) has been made (e.g., one exploiting the structure of the problem), and that the times recorded in the figure include context switches and the time spent exchanging the problem instance and solution via memory-mapped files between SimuLTE and CPLEX. We are then justified in believing that these times can be abated even further if needed and that the optimal solution can be found in reasonable time also with a higher number of D2D-eligible flows. Finally, we remark that the size of problem (2) does not depend on the number of RBs, thus its complexity is affordable also with higher bandwidth.

5.2 Comparison with other MS schemes

We compare our MS scheme against two static and two dynamic schemes. The static ones select all flows in \mathbf{D} to be DM and IM, respectively (the second choice corresponding to traditional cellular communications). The two dynamic schemes are a *Random* one and the one in [4], which we call *Max-CQI* for brevity. The Random scheme selects the *same* percentage of flows as our scheme to be DM or IM, but assigns them to either mode *randomly*. This allows us to ascertain what the added value

is of selecting not only the right *number*, but the right *set* of flows. The Max-CQI one assumes instead that the eNB knows *three* SINR values for each flow in \mathbf{D} , one for IM and *two* for DM, namely *with* and *without* spatial reuse. Then, the expected throughput is estimated as a function of the SINR and the available resources for each flow. The latter are obtained based on the *number* of UEs, instead of on their backlog. Each flow in \mathbf{D} will use the mode with the highest expected throughput, without taking into account any global metric. Note that [4] does not mention how to do PS: thus, to make the comparison more challenging, we endow both the Random and the Max-CQI schemes with the *same* conflict resolution as ours wherever reuse is allowed, run them at the same period, and employ our very same PS heuristic at the TTI timescale. Thus, any performance difference is *only* due to the MS scheme employed.

We first simulate 25 D2D-eligible flows, sending 100-byte packets each 10 ms, and 50 D2D-ineligible DL flows, that send packets with the same period and constant size. The latter varies from 200 to 600 bytes, so as to congest the downlink subframe progressively. All UEs are in a 100 m radius from the eNB. Figure 16 shows the sum of the application-level throughputs at the destinations of flows in \mathbf{F} , whereas Fig. 17 shows the percentage of D2D-eligible flows sent to DM, both as a function of the overall offered

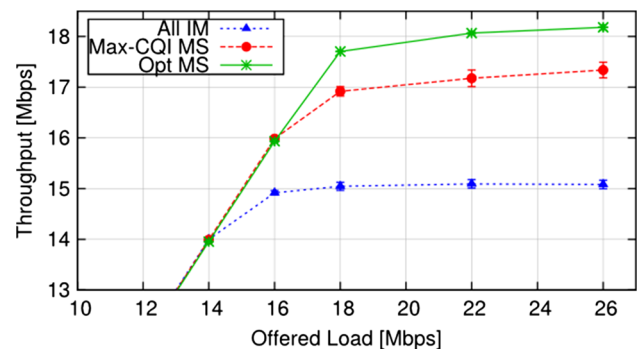


Fig. 16 Application-layer cell throughput

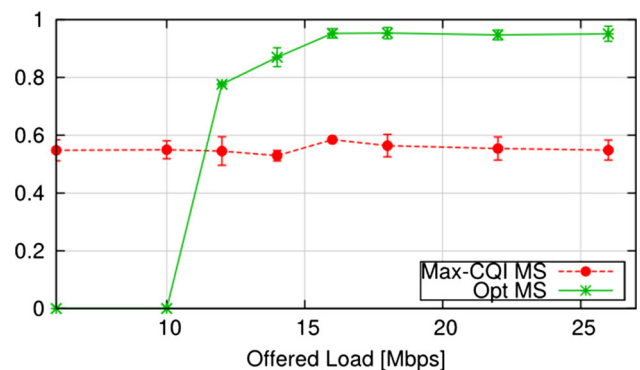


Fig. 17 Average fraction of DM flows

load. When all flows are statically assigned to IM (*All-IM* scheme), there soon comes a point (at about 16 Mbps of offered load) where the cell cannot accommodate all the DL flows. In fact, packets sent to the eNB have to traverse the DL leg to reach their destination, further increasing the congestion of the DL subframe. The Max-CQI scheme selects DM for some flows and exploits spatial reuse in the UL subframe, but ends up sending too much in IM in any case (45 % of D2D-eligible flows, as shown in Fig. 17), since it is oblivious of the increasing congestion in the DL. Our scheme, instead, outperforms both because it avoids sending in the UL what cannot come down in the DL. In fact, when the cell approaches saturation, our MS scheme offloads the DL subframe by sending more than 90 % of D2D-eligible flows in DM. This shows that making the MS aware of the available resources in both UL and DL significantly improve the cell throughput. Note that, in this scenario, the Random and *All-DM* schemes perform exactly like ours. This is no wonder, since all these schemes assign all D2D-eligible flows to DM when the cell saturates. Our scheme also reduces the communication latency. Figure 18 shows the CDF of the application-level delay for flows in **D**, at an offered load of 18 Mbps. The Max-CQI and the All-IM schemes saturate the cell and data experience queueing at the eNB, thus their delay is higher than our scheme's. Although the Random MS achieves the same throughput as our MS, this result is paid for with higher delay, since selecting DM for flows having bad channel conditions, e.g. with endpoints too far apart, increases their latency.

The advantages of our scheme with respect to the All-DM and Random ones are best appreciated in a denser scenario, with 50 D2D-eligible flows. UEs are static and the distance between flow endpoints is at most 25 m. Sources are CBR, sending 500-byte packets every 10 ms. Figure 19 shows the application-layer cell throughput against the conflict threshold. Low threshold values make the conflict graph highly connected, which minimizes interference, but also hampers spatial reuse and, as a




Fig. 18 CDF of application-layer delay, offered load = 18 Mbps




Fig. 19 Application-layer cell throughput

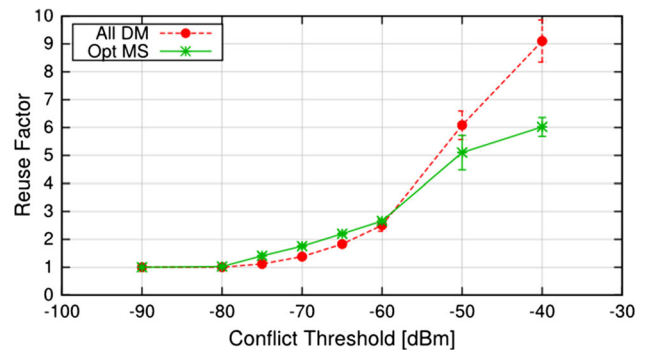


Fig. 20 Reuse factor

consequence, the overall throughput. As the threshold increases, conflicts disappear progressively, allowing spatial reuse, which results in a remarkable increase of the cell throughput. However, increasing the threshold further reveals another, negative effect: cumulative interference prevails, reducing the SINR of DM flows, hence ultimately the cell throughput. This is in fact what happens with the All-DM curve. Our MS scheme shows little dependency on the conflict threshold, unless of course the latter is so low as to prevent reuse at all (e.g., < -70 db). In fact, the *rates* actually employed in our MS are fed back from the MAC, and they do include the interference made by flows marked as non-conflicting in the CG. Thus, with a high conflict threshold, some D2D-eligible flows are still assigned to IM to protect them from an otherwise intolerable interference. This is made evident in Fig. 20, where we plot the average reuse factor per RB (i.e., the average number of flows using that RB). When the threshold is set to -40db, our proposed scheme limits the reuse factor to 6, against 9 of the All-DM scheme. The reasons behind the poor performance of the Random scheme are excessive packet loss, due to poor choice of the flows sent to DM and uncorrelated MS in adjacent periods, i.e., some flows swap modes at each MS period, hence their RLC buffers are repeatedly discarded.

5.3 Packet scheduling optimality and time cost

We now discuss the optimality of the best-fit heuristic. For the optimal scheme, we run simulations where Problem (3) has been solved at each TTI, using CPLEX with an optimality gap of 5 %, hence results of the heuristic are rescaled by 0.95 to err on the safe side. We simulate an increasing number of D2D-eligible flows, ranging from 10 to 30, and 10 D2D-ineligible UL flows. All sources send 100-byte packets every 10 ms. In order to evaluate only the impact of PS, we disable the MS part of our framework and assume that all D2D-eligible flows are DM flows. Figure 21 shows the sum of the application-layer throughput for all flows in \mathbf{F} using the best-fit heuristic and the optimal PS (scale on the left vertical axis), together with the corresponding optimality ratio (the dash-dotted line, whose scale is on the right vertical axis). With 20–25 flows, the UL subframe is not completely utilized and the heuristic achieves practically the same throughput as the optimal solution. As more flows need to be scheduled, the optimality ratio decreases slightly, but stays above 0.85 in any case. Unfortunately, it is impractical to solve scenarios with higher load at optimality, due to the complexity of Problem (3), which increases with the number of flows. Figure 22 reports the average time needed to solve one




Fig. 21 Application-layer cell throughput and optimality ratio, optimal PS versus best fit PS




Fig. 22 Average solving time for optimal PS

instance of the optimization problem. As expected, the optimal solution is not affordable within the TTI limit, not even at low loads. In particular, computation times are one to two orders of magnitude above the time limit already with 20–40 flows. With 50 flows, the average time is about 14 s, and is not reported in Fig. 22 for ease of reading. Although it is likely that these times can be abated through careful reformulation of the problem and/or exploiting more powerful hardware, a TTI-bounded solution could be achieved only with low loads (e.g., 20–30 flows in \mathbf{F}), i.e., exactly where the best-fit heuristic is near-optimal.

5.4 Spatial reuse

Spatial reuse allows different flows to be allocated in the same RBs, so it can improve the system performance. On the other hand, interference must be managed. Figure 23 shows the distribution of the SINR that a receiver would perceive from its peering UE located in the center of the area—i.e., the triangle at coordinates (1000,1000)—in different resource sharing situations. SINR increases going from red to green. Circles and squares represent, respectively, active and inactive transmitters of other DM flows (corresponding receivers are not shown). If uncontrolled spatial reuse is employed, all UEs may be allocated in the same RBs as the central UE, thus producing high interference. This results in low SINR for the receiving UE, unless the latter is within a 15–20 m radius around the transmitter (see the green area in the leftmost part of Fig. 23). On the contrary, if spatial reuse is prevented, the SINR distribution is the one depicted in the central part of Fig. 23. Here, flows are allocated in different RBs and the green area becomes much wider, allowing the UEs to use higher CQIs. The latter represents the interference situation when a reuse-unaware scheduler is used, e.g. MaxC/I. However, such approach is oblivious to the fact that different DM flows may be far enough for their mutual interference to be negligible. In this case, enabling spatial reuse would be preferable, so as to reduce the resources required in the UL subframe. The rightmost part of Fig. 23 shows the same scenario as above, where *some* of the neighboring UEs can reuse the same RBs as the central UE (the circles in the figure), whereas some others cannot (the squares in the figure). The green area is obviously smaller than that in the central figure, but the SINR is still acceptable within a radius of 90–100 m around the transmitter. Note that this result is obtained by allowing resource sharing only among UEs located far from the central one, which confirms that a careful selection of which flows can reuse the same resources is required.

We assess whether the advantages of spatial reuse overtake the disadvantages caused by the interference in a dense scenario. To this aim, we compare our MS with spatial reuse to that without spatial reuse. The former uses




Fig. 23 SINR distribution: with spatial reuse (*left*), without spatial reuse (*center*), with partial spatial reuse (*right*)

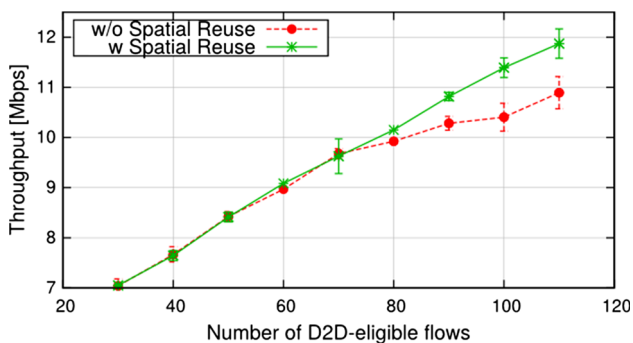


Fig. 24 Application-layer cell throughput with and without spatial reuse

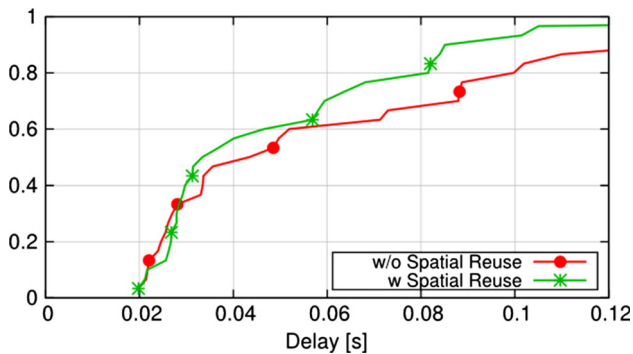


Fig. 25 CDF of application-layer delay w./w.o. spatial reuse, 30 flows in **D**

the best-fit heuristic as a PS algorithm, whereas the latter uses MaxC/I.

We simulate an increasing number of D2D-eligible flows, with 10 D2D-ineligible UL flows and 50 D2D-ineligible DL flows, all sending 100-byte packets every 10 ms. The center of the circular trajectories followed by UEs is at a maximum distance of 300 m from the eNB. Figure 24 reports the sum of the application-layer throughput for flows in **F**. The benefits of spatial reuse are clearly visible and begin at moderate loads (80 flows in **D**), where the scheme without spatial reuse

begins to saturate the UL subframe. Moreover, as shown in Fig. 25 for 30 flows in **D**, even at *low* traffic loads (when MaxC/I achieves the same throughput) delays are smaller with spatial reuse, since UEs’ data experience less buffering delay at the sender side. Such results are obtained thanks to the combined use of the conflict graph and the MS. In fact, besides the restrictions imposed by the former, MS can decide to assign some D2D-eligible flows to IM (i.e., those receiving high cumulative interference *if* they are assigned to DM), which do not share resources with DM flows.

5.5 Discussion

The above sections have highlighted the following major results on how to deal with mode selection and packet scheduling for D2D communications. First, MS cannot be performed too fast. However, since packet scheduling need to be run every TTI, it turns out that every attempt to jointly perform MS and PS may lead to inefficient management of the resources. Nevertheless, MS and PS cannot be treated as independent problems: results have shown that higher cell throughput is achieved when MS is aware of how much resources need to be allocated to D2D-ineligible flows, in both the UL and the DL subframes. In addition to this, not only the right *percentage* of DM flows should be selected, but also the choice of *which* flows use DM should be taken carefully, lest UEs with bad SINR on the DM link are selected for direct transmissions.

Second, a binary interference model is an efficient way to deal with the inter-UE interference, despite its simplicity. Acting on the conflict threshold allows one to obtain different level of protection and, consequently, different level of spatial reuse. Cumulative interference is accounted for in any case, since the MS exploits channel state reports from UEs. For instance, whenever the CQI (i.e., the available rate) of a flow in DM is low because of cumulative interference, the MS module may decide to assign that flow to IM.

Third, D2D communications require specific PS algorithms that are aware of the presence of both D2D-eligible

and D2D-ineligible flows, spatial reuse and interference constraints. Solving at optimality the above problem is not feasible at TTI timescale. However, it is possible to devise efficient heuristics that can be run at each scheduling interval. Since D2D communications occur preferably in the UL subframe, the packet scheduler must take into account the constraints of the SC-FDMA scheme, i.e. RBs allocated to a UE must be adjacent.

Finally, the negative effects of spatial reuse on the SINR, shown in Fig. 23, can be contained through the combined use of a conflict graph in both MS and PS. This way, spatial reuse brings benefits in terms of both throughput and latency.

6 Related work

Available works about D2D can be categorized in *inband* and *outband* [3], depending on whether D2D communications take place in licensed or unlicensed spectrum. Since our work falls into the first category, we only review literature about inband D2D. The main challenges concern the selection of the communication path (direct or eNB-relayed) for a potential D2D pair of UEs, namely mode selection, and allocation of resources, so as to mitigate the interference arising from spatial reuse among cellular and D2D UEs.

Besides the algorithm proposed in [4] and discussed in Sect. 5, [5] also addresses the mode selection problem. The work in [5] is similar to [4], except that the former also proposes a method for selecting the transmission power of a D2D transmitter. As in [4], no global metric is taken into account and packet scheduling is not discussed.

Packet scheduling is the focus of [9, 11, 6]. [9] models allocation as a Mixed Integer Non-Linear Problem (MINLP). Due to its NP-hardness, authors propose a graph-based heuristic scheme, according to which each link can be allocated *at most* one RB per scheduling interval, which is clearly inadequate. The resource allocation scheme in [11] first assigns RBs to cellular UEs, then tries to allocate RBs to as many D2D pairs as possible, minimizing the interference produced on cellular UEs. In this scheme, spatial reuse among DM flows is not allowed, and flows are allocated only *one* RB, similarly to [9]. [6] is a two-stage scheduling mechanism, where the first and the second steps are run at large (i.e., several *tens* of TTIs) and short (i.e., one TTI) time scales, respectively. First, the eNB prepares a RB allocation, deciding which set of flows in \mathbf{F} can share which RBs, so as to keep the mutual interference below a predefined threshold. Second, the eNB schedules cellular UEs according to some scheduling policy (e.g., MaxC/I, PF etc.), whereas D2D pairs are allocated the RBs granted in the first phase. RB allocation is modeled as an optimization problem, which aims to maximize the minimum number of RBs allocated to each cellular UE while

satisfying the required number of RBs allocated to D2D pairs. However, the amount of resources requested by D2D pairs can be very different from one TTI to the next, thus long-term allocation can become largely suboptimal after few TTIs.

Algorithms that tackle mode selection and packet scheduling jointly have also been studied in [7, 8] and [10]. They formulate optimization problems under different objective functions (e.g., maximize the overall throughput, minimize the power consumption) and different types of constraints. In any case, the problems are found to be NP-hard, and heuristics are presented. We have already discussed at length that MS cannot be fast and PS cannot be slow, if the effectiveness of resource management is a concern. Thus, a joint approach must deal with the trade-off between tolerable packet losses and optimality of PS.

Works [28, 29] discuss packet scheduling taking into account the RB adjacency imposed by SC-FDMA, which is employed in the uplink subframe in LTE-A networks. In [28], the eNB schedules cellular UEs first, then pairs one D2D flow with one cellular UE. To this aim, authors present a MINLP problem and the corresponding greedy heuristic. Work [29] assumes that RBs have already been allocated to cellular UEs, then tries to optimize both power and RB allocation for D2D flows. In particular, RB allocation consists in partitioning the available RBs into N_D disjoint contiguous subsets, one per D2D flow. Since the number of possible allocations is $N_D!$, dynamic programming is used to reduce the complexity of the problem. In both works, one-to-one pairing between one cellular and one D2D flow is advocated. This is constraining, since a D2D flow can only use as many resources as its paired cellular UE, hence may result in suboptimal decisions. For instance, if one greedy cellular flow gets all the UL RBs, only one D2D flow will be scheduled in that TTI. Moreover, no spatial reuse among D2D flows is possible. On the other hand, our framework allocates resources simultaneously for cellular and D2D flows and exploits spatial reuse judiciously.

7 Conclusions and future work

This paper presented a framework to allocate resources in D2D-enabled LTE-A cells. The framework relies on selecting UE modes at long timescales (e.g., 1 s) and scheduling resources at TTI timescales, allowing spatial reuse among D2D flows. Mode selection takes into account the availability of resources in *both* the uplink and downlink subframes *and* the mutual exclusion of interfering flows, leading to more effective decisions. We have shown that selecting modes *dynamically*, according to cell-wide measures (e.g., the occupancy of both the uplink and the downlink subframes) allows higher cell throughput to be achieved. Moreover, our framework is robust to (mis)configurations of the interference threshold.

As for packet scheduling, we propose a best-fit heuristic to manage the coexistence of DM and IM flows in the uplink subframe. Such heuristic is near-optimal and affordable.

Acknowledgments The subject matter of this paper includes description of results of a joint research project carried out by Telecom Italia and the University of Pisa. Telecom Italia reserves all proprietary rights in any process, procedure, algorithm, article of manufacture, or other result of said project herein described. The authors wish to thank Andrea Rossali, formerly MSc student at the University of Pisa, for his contributions to the code.

References

- 3GPP. (2014). Study on LTE device to device proximity services: Radio aspects (release 12). TS 36.843 v12.0.1. March 2014.
- 3GPP. (2015). Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); overall description; stage 2. TS 36.300 v12.5.0. March 2015.
- Asadi, A., Wang, Q., & Mancuso, V. (2014). A survey on device-to-device communication in cellular networks. *IEEE Communications Surveys and Tutorials*, 16(4), 1801–1819.
- Doppler, K., Yu, C. H., Ribeiro, C. B., & Janis, P. (2010). Mode selection for device-to-device communication underlying an LTE-Advanced network. In *Proceedings of WCNC*, (pp. 1–6). April 18–21, 2010.
- Gu, J., Bae, S. J., Choi, B. G., & Chung, M. Y. (2012). Mode selection scheme considering transmission power for improving performance of device-to-device communication in cellular networks. In *Proceedings of ICUIMC*. February 20–22, 2012, Kuala Lumpur, Malaysia.
- Lee, D. H., Choi, K. W., Jeon, W. S., & Jeong, D. G. (2014). Two-Stage Semi-Distributed Resource Management for Device-to-Device Communication in Cellular Networks. *IEEE Transactions on Wireless Communications*, 13(4), 1908–1920.
- Belleschi, M., Fodor, G., & Abrardo, A. (2011). Performance analysis of a distributed resource allocation scheme for D2D communications. In *IEEE GLOBECOM 2011 workshops*, (pp. 358–362). December 5–9, 2011.
- Xiao, X., Tao, X., & Lu, J. (2011). A QoS-aware power optimization scheme in OFDMA systems with integrated device-to-device (D2D) communications. In *Proceedings of VTC Fall*, (pp. 1–5). September 5–8, 2011.
- Zhang, R., Cheng, X., Yang, L., & Jiao, B. (2013). Interference-aware graph based resource sharing for device-to-device communications underlying cellular networks. In *Proceedings of WCNC*, (pp. 140–145). April 7–10, 2013.
- Wen, S., Zhu, X., Zhang, X., & Yang, D. (2013). QoS-aware mode selection and resource allocation scheme for Device-to-Device (D2D) communication in cellular networks. In *Proceedings of ICC*, (pp. 101–105). June 9–13, 2013.
- Xu, Y., Yin, R., Han, T., & Yu, G. (2012). Interference-aware channel allocation for Device-to-Device communication underlying cellular networks. In *Proceedings of IEEE ICC*, (pp. 422–427). August 15–17, 2012.
- Dahlman, E., Parkvall, S., & Skold, J. (2011). *4G: LTE/LTE-Advanced for mobile broadband*. Cambridge: Academic Press.
- Lin, X., Andrews, J., Ghosh, A., & Ratasuk, R. (2014). An overview of 3GPP device-to-device proximity services. *IEEE Communications Magazine*, 52(4), 40–48. doi:10.1109/MCOM.2014.6807945.
- Fodor, G., et al. (2012). Design aspects of network assisted device-to-device communications. *IEEE Communications Magazine*, 50(3), 170–177.
- Yu, C. H., Tirkkonen, O., Doppler, K., & Ribeiro, C. (2009). On the performance of device-to-device underlay communication with simple power control. In *Proceedings of IEEE VTC Spring*, (pp. 1–5). April 26–29, 2009.
- Zhou, X., Zhang, Z., Wang, G., Yu, X., Zhao, B. Y., & Zheng, H. (2015). Practical conflict graphs in the wild. *IEEE/ACM Transactions on Networking*, 23(3), 824–835. doi:10.1109/TNET.2014.2306416.
- 3GPP. (2015). Evolved Universal Terrestrial Radio Access (E-UTRA); Packet Data Convergence Protocol (PDCP) specification. TS 36.323 v12.3.0. March 2015.
- Li, C., Sun, F., Cioffi, J. M., & Yang, L. (2014). Energy efficient MIMO relay transmissions via joint power allocations. *IEEE Transactions on Circuits and Systems-II*, 61(7), 531–535.
- Li, C., Yang, H. J., Sun, F., Cioffi, J. M., & Yang, L. (2014). Approximate closed-form energy efficient PA for MIMO relaying systems in the high SNR regime. *IEEE Communications Letters*, 18(8), 1367–1370.
- Li, C., Cioffi, J. M., & Yang, L. (2014). Optimal energy efficient joint power allocation for two-hop single-antenna relaying systems. *European Transactions on Telecommunications*, 25(7), 745–751.
- Li, C., Zhang, S., Liu, P., Sun, F., Cioffi, J. M., & Yang, L. (2015). Overhearing protocol design exploiting inter-cell interference in cooperative green networks. *IEEE Transactions on Vehicular Technology*, 99, 1–7. doi:10.1109/TVT.2015.2389826.
- EARTH EU project website. <https://www.ict-earth.eu/>.
- SimuLTE webpage. <http://www.simulte.com>.
- Virdis, A., Stea, G., & Nardini, G. (2014). *SimuLTE: A modular system-level simulator for LTE/LTE-A networks based on OMNeT++*. SimuTech 2014, Vienna, AT, August 28–30, 2014.
- OMNeT++. <http://www.omnetpp.org>.
- 3GPP. (2010). Further advancements for E-UTRA physical layer aspects (release 9). TR 36.814 v9.0.0. March 2010.
- ILOG CPLEX Software. <http://www.ilog.com>.
- Zulhasnine, M., Changcheng, H., & Srinivasan, A. (2010). Efficient resource allocation for device-to-device communication underlying LTE network. In *IEEE 6th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, (pp. 368–375). October 11–13, 2010. doi:10.1109/WIMOB.2010.5645039.
- Wu, Y., Wang, S., Guo, W., Chu, X., & Zhang, J. (2014). Optimal resource management for device-to-device communications underlying SC-FDMA systems. In *9th International Symposium on Communication Systems, Networks and Digital Signal Processing (CSNDSP)*, (pp. 569–574). July 23–25, 2014. doi:10.1109/CSNDSP.2014.6923893.



Giovanni Nardini obtained his M.Sc. in Computer Systems Engineering (summa cum laude) at the University of Pisa, Italy, in 2013 and he is currently working towards his Ph.D. at the same University. His research interests are cellular networks, and specifically resource allocation in LTE-Advanced networks. In this fields, he has coauthored two patents and seven papers. He has been involved in EU-funded and industrial research projects.



Giovanni Stea is Associate Professor at the Department of Information Engineering of the University of Pisa, Italy, where he also got his Ph.D. in 2003. His current research interests include Quality of Service and resource allocation in wireline and wireless networks, performance evaluation through simulation and analytical models, traffic engineering. In these fields he has coauthored more than 70 peer-reviewed papers and fourteen patents. He has

been involved in national and European research projects, and he has led joint research projects with industrial partners. He is currently chairing University of Pisa's activities in EU-H2020 projects Flex5Gware and SAPIENT. He has served as a member of the technical and/or organization committees for several international conferences, including SIGCOMM, WoWMoM, and VALUE-TOOLS, and he is serving on the Editorial Board of the Springer Wireless Networks journal.



Antonio Virdis is a post-doc researcher at the University of Pisa, where he obtained his M.Sc. degree in Computer System Engineering in 2011, and his Ph.D. in Information Engineering in 2015. His research interests include Quality of Service, scheduling and resource allocation in LTE-Advanced wireless networks, network simulation and performance evaluation. He has been involved in national, EU-funded and industry-funded research

projects. He coauthored four patents and eleven peer-reviewed papers in the field of LTE-Advanced resource allocation.



Dario Sabella received his Dr. Eng. Degree in electronic engineering from Politecnico di Torino (Italy) in 2002, and post-degree specialization Master in Telecommunications in 2004. Since 2002 he has been working for Telecom Italia Lab, involved in several projects targeting future generation communications. He is currently working in the Wireless Innovation Department as responsible for some international projects for the mobile network. Moreover, Dario

participates or has participated in several European projects, including

EARTH, ARTIST4G, iJOIN as WP leader, 5GrEEn as project leader of Telecom Italia activities, Flex5Gware as WP leader. He has served as chairman or TPC member for several international conferences. He has co-authored several publications and patents in the field of wireless communications and energy efficiency. From 2012 he is member of AEIT association. From July 2015 he is vice-chair of ETSI MEC (Mobile Edge Computing) IEG.



Marco Caretti received a Laurea degree in Electronic Engineering from University of L'Aquila. He has been working with Telecom Italia Lab, the Research center of Telecom Italia group, since 2000 where he has been involved in several internal R&D projects on performance study of wireless access communication systems (EDGE, UMTS, WiMAX and LTE). His research interests include design and performance study of radio access technologies

(in particular LTE and LTE Advanced) and their evolution towards 5G systems. He coauthored several international patents and papers. Since 2005 he participated in several standardisation groups (WiMAX, 3GPP RAN1) and he is currently participating in 3GPP RAN4. Since 2013 he is in charge of an internal project focused on virtualized radio access solutions towards 5G systems.