Modelling D2D Communications in Cellular Access Networks via Coupled Processors

Christian Vitale
IMDEA Networks Institute, Madrid, Spain
and University Carlos III of Madrid, Spain
christian.vitale@imdea.org

Vincenzo Mancuso
IMDEA Networks Institute, Madrid, Spain
vincenzo.mancuso@imdea.org

Gianluca Rizzo
HES SO Valais, Switzerland
gianluca.rizzo@hevs.ch

***Abstract*—This paper presents a first fully analytical approach to performance evaluation of D2D communication systems, which does not assume the system to be in saturated conditions. We adopt a *Coupled Processors* model to describe a cellular scenario with D2D users sharing radio resources with cellular users, i.e., adopting *in-band underlay* D2D schemes. We derive sufficient conditions for stability of such system, characterizing the effect of D2D transmissions on cellular user performance. Moreover, we present a computationally feasible method for the determination of a proportionally fair allocation of resources. We show that, in non-saturated networks, such an allocation sensibly improves the one derived under the assumption of saturation. Our results show the importance of accurately modelling the interdependence in users performance in the design and evaluation of a D2D cellular system.**

I. Introduction

The increasing penetration of interconnected devices (smartphones, sensors, home appliances, etc.) is giving progressively rise to a wealth of new services within the context of the "Internet of Things". In cellular access networks, this has lead to the emergence of Device-to-Device (D2D) communications, a direct communication mode between two mobile users in which the exchanged information does not traverse the Base Station (BS) or the core network. Initially proposed for multihop relays in cellular networks, D2D communications are finding an increasing number of use cases in vehicular networks, in content distribution and cellular offloading. In such applications, D2D has the potential to lower packet delays, increase energy efficiency, fairness and throughput while improving the spectral efficiency of dense cellular networks (e.g., adopting new flexible paradigms like in LTE-A).

Different types of D2D schemes have been proposed [1]. *In-band* schemes either allow D2D transmissions to occur over dedicated cellular resources (*in-band overlay* schemes or over the same resources used by legacy cellular users (*in-band underlay* schemes. Here we focus on the latter, though our approach can be easily extended to in-band overlay schemes.

Performance analysis of in-band underlay schemes is challenging, due to the complexity of D2D systems and of their interaction with cellular system operations. Indeed, as D2D and cellular users share the same resources, their performances are strongly correlated through interference and through the scheduling algorithm. Available results are mainly based on simulations [9], and they assume either the system to be saturated, or that only a single D2D pair can be scheduled in the same Resource Block (RB) as a cellular user [5]. This leads to pessimistic, overly conservative results, particularly in non-saturated settings. Moreover, such an approach makes it difficult to characterize the main performance trade-offs in such D2D systems, in order to design efficient scheduling and rate allocation algorithms.

978-1-4244-8953-4/11/$26.00 © 2015 IEEE

We propose an analytical approach to performance characterization of D2D schemes, which captures the dependencies between interfering transmissions and achievable rates. Our approach is based on the Coupled Processors (CP) model [2]. Such model naturally applies to D2D systems, as it explicitly accounts for the achievable transmission rates when the correlation between the service rates of multiple queues is known. More specifically, we adopt the method of analysis of CP systems recently proposed in [13] and based on Network Calculus, enabling a fully analytical study of CP systems.

The main contributions of this paper are as follows. We present a novel analytical approach to the study of D2D schemes in LTE-like cellular networks, which applies to scenarios with one or more D2D pairs scheduled in the same RB of a cellular user, or when little knowledge of traffic distribution is available. We derive new sufficient conditions for stability of transmission queues in a D2D system, and we show how to evaluate the effects of D2D transmissions over cellular user performance. Moreover, we present a computationally feasible method for the determination of a proportionally fair allocation of resources, which allows trading the amount of fairness of the solution for computational complexity. Finally, we validate our results through simulations, assessing numerically the quality of the bounds and of the optimal allocations derived with our approach. Our results show the importance, in the performance study of D2D systems, of accurately modelling the mutual correlations in performance among users in a D2D cellular system.

II. System Model

We consider a base station (or evolved Node B, eNB, in LTE) belonging to an LTE access network, serving U cellular users. Under the coverage area of the considered eNB, there are also D D2D transmitter-receiver pairs. We consider a static scenario, in which D2D pairs do not change over time and in which users do not move. We assume D2D transmissions happen on the uplink channel. The main reason for this choice is that on the uplink the eNB is the receiver of all cellular transmissions. Hence, it has a complete knowledge of the interference, and it can implement interference-aware scheduling strategies for D2D transmissions.

We consider a log-distance path loss model. Specifically, path loss (in dB) between a transmitter and a receiver at a distance r is given by $L(r) = L_0 + 10\eta \log_{10}\left(\frac{r}{r_0}\right) + X_S$, where L_0 is the path loss at a reference distance r_0, η is the path loss exponent, and X_S is a Gaussian random variable with standard deviation σ_S, modelling the effects of shadowing [12]. As medium access technology, we assume that Orthogonal Frequency Division Multiple Access (OFDMA) is used. Transmission time is split into slots of fixed duration, while multiple and independent sub-carriers are obtained over a wide channel bandwidth. A RB is the smallest resource that can be assigned to a particular user. It is composed by a specific time slot and by a given set of subcarriers.[1]

We model capacity through the Shannon formula. Given a transmitter i (either cellular or D2D) and its receiver j at a distance $r_{i,j}$, the amount of bits transmitted per RB is:

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

where B is the bandwidth of the RB, τ is the duration of a time slot, $P_T(i)$ is the per-RB transmission power of i, N_0 is the noise spectral density, I_C is the inter-cell interference, while b_M is the maximum amount of bits that can be transmitted in a RB when the best modulation and coding scheme is used. The summation of the interference at the denominator goes over all the active transmitters in the RB, both cellular and D2D. We assume that the transmission power used by the devices is fixed over the RBs.We assume that uplink power control is in place both at the cellular and at the D2D transmitters, with a specific per-RB SNR target, compensating for path loss and shadowing. Due to the short distances covered by D2D communications, the transmission power used by D2D communications is typically small, and the amount of interference it generates becomes relevant mainly in particularly dense scenarios. But as such dense networks are gradually becoming a reality in present day networks, we included in our model the effects of interference caused by D2D transmissions to other concurrent D2D and cellular transmissions.

A. A D2D in-band underlay scheme for LTE

In in-band underlay D2D transmissions, a D2D pair can be scheduled by the eNB on RBs already assigned to cellular transmitters or to other D2D pairs. Fig. 1 shows an example of radio resource utilization with in-band underlay D2D.

The scheduling policy we consider is a variation of the one proposed in FlashLinQ [14]. FlashLinQ is a state-of-the-art PHY-MAC architecture for D2D that allows the scheduling of different transmitters (D2D or cellular) in the same time and frequency resource, through an OFDMA-like access selection mechanism. The scheduling of the transmitters is performed at RB level. Furthermore, FlashLinQ does not distinguish between cellular and D2D users. We believe indeed that the D2D users have to be considered as primary users of the cellular access network, and that D2D transmissions have to be scheduled while guaranteeing that the given cellular user performance target is met. We introduce therefore a two-tier scheduling policy, in which cellular transmitters are scheduled first, as follows:

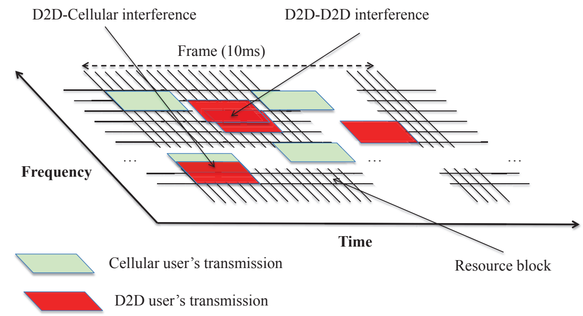

Fig. 1: Resource allocation with in-band underlay D2D.

- **Cellular UEs transmitters**: The scheduling policy at the eNB is Equal Time [8], which is interference-unaware.[2] In order to have a full understanding of the effect of the D2D transmissions on the cellular ones, we assume that cellular users are in saturation, i.e., they always have a packet ready to send. As a result, the eNB assigns RBs to each cellular user with the same probability.
- **D2D transmitters (FlashLinQ policy)**: For each RB, all the D2D transmitters having at least one packet to send are considered, one at at time, in a random order. Each D2D transmitter is scheduled for transmission if (i) its interference on the cellular transmitter scheduled in the same RB is below a given limit, and if (ii) its Signal to Interference Ratio (SIR) is above a given threshold, considering the cellular user scheduled in the particular RB under analysis and all the D2D transmitters already scheduled. This ensures that the impact of D2D transmissions on cellular ones is kept below a given threshold and, furthermore, that D2D users achieve a target minimum throughput.

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

III. A CP Model for D2D Schemes

In this section we describe our approach to the analysis of in-band underlay D2D schemes. First, we introduce a model that is able to characterize the main features of such D2D schemes. Then we present our main analytical results.

To study the performance of the cellular and D2D systems, we use a CP model [2]. A CP system (CPS) is a set of parallel

[1]Note that our approach can be easily extended to Single Carrier - FDMA.

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

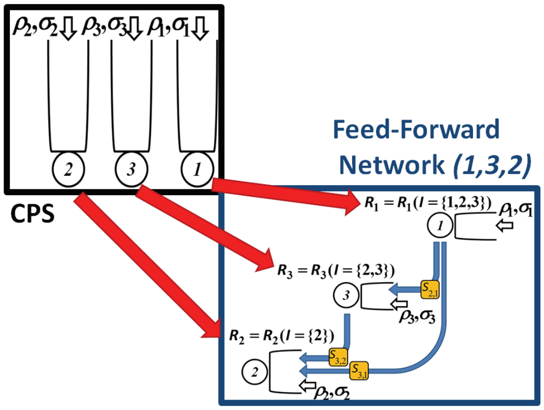

Fig. 2: A three-nodes CPS, and one of its upper bounding networks, associated to the sorting $\{1, 3, 2\}$.

queues, whose service rates at any time t are determined by the set of active queues at that time. More formally, if at time t the set of active queues is $I(t)$, the service rate of queue $i \in I(t)$ is $R_i(t) = R_i(I(t))$. Therefore, at each time t, the set of active queues $I(t)$ univocally determines the state of the system. Note that the set of system states $\mathcal{I}$ is finite, and represents all the possible subsets of active queues of the CPS.

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

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

$$R_d(I) = \frac{\phi}{U} \sum_u \frac{\sum_{I_O} b_d^u(I_O)}{|I|!}, \tag{2}$$

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

In such a CP model of our D2D system, there are a few aspects which could be a source of discrepancies between the behavior of the D2D system and its CP model. First, the instantaneous service rate of the CPS queues has been set equal to the average rate of the corresponding D2D transmitters in a given system state. Moreover, despite in our D2D system the system state may change only when a new RB is scheduled (as the cellular transmissions are scheduled sequentially, on a per-RB basis), the equivalent CP model works at the "fluid" limit, i.e., each queue serves its traffic as if it were infinitely divisible, and it could change state at any time t. In Section VI we validate numerically our model, evaluating the impact of these modeling choices, and showing that such a CP system models accurately the dynamics of the original D2D system.

A. Analysis of the D2D Model

The analysis of the D2D model we perform in this paper is based on the study of its CPS characterization through the approach in [13]. Such an approach is convenient because it is purely analytic and because it enables the derivation of heuristics that suitably trade-off accuracy for computational complexity. This approach consists in the derivation, from a given CPS, of a set of feed-forward networks of interconnected GPS nodes (as depicted in Fig. 2), whose stability implies the stability of the CPS.

Given a particular CPS with D parallel queues, [13] derives one network for each possible sorting d_O of $\mathcal{D}$. We name w_O the generic feed-forward network built from the sorting $d_O = [d_1, ..., d_D]$. w_O is composed by D GPS nodes, each of which is in a one-to-one mapping with one of the D queues of the CPS model. This means that the arrivals at the CPS queues are exactly the same at their corresponding GPS nodes in the introduced feed-forward networks, at any time t. The GPS nodes are connected according to the sorted list d_O, in a feed-forward configuration: the first node receives the traffic that in the CPS corresponds to the first queue in d_O, i.e., d_1; the second node receives the traffic corresponding to d_2, plus a *rescaled* version of the output of the first GPS node; in general, the ith GPS node receives the traffic of the d_ith CPS queue plus a rescaled version of the traffic served by GPS nodes $j \in \{1, ..., i-1\}$. Fig. 2 shows an example with $D = 3$.

Rescaled traffic flows represent the feed-forward coupling between CPS queues. When w_O is analysed, the service rate of the ith GPS node is the same service rate of the d_ith CPS queue when the most "coupled" configuration possible of CPS queues $\{d_i, d_{i+1}, ..., d_D\}$ is active. In practice, the service rate of the i-th GPS node in w_O, that we name from now on $R_{d_i}^{up}$, is the worst with respect to the activity of all the possible subsets of CPS queues $\{d_{i+1}, ..., d_D\}$. Note that the service rate of a GPS node is constant in w_O, so that the feed-forward traffic flow absorbs part of the fixed capacity of a node. As shown in [13], if carefully tuned, such feed-forward *coupling* traffic guarantees that the stability of a GPS node in w_O implies the stability of the corresponding queue at the CPS. Obviously, each d_O yields a different feed-forward network. Indeed, depending on d_O, a different service rate is used at the GPS nodes and a different feed-forward *coupling* traffic enters the GPS nodes. All in all, a queue in the CPS is stable if at least one of the GPS nodes onto which it has been mapped is stable [13]. We achieve the tuning of the feed-forward *coupling* traffic at the GPS nodes through the use of so-called *scalers* [6], which are system components that allow to formally reduce or amplify the volume of traffic moving from a queue to another in the network. As shown in [13], the scaling factor $S_{i,k}$ that modifies the traffic going from the

k-th to the i-th GPS node of the generic network w_O is:

$$S_{i,k} = \frac{R_{d_i}^{up} - R_{d_i}^{k,up}}{R_{d_k}^{up}}, \tag{3}$$

where $R_{d_i}^{k,up}$ is the achieved rate of the D2D transmitter d_i when the most "coupled" configuration of the transmitters $\{d_k, ..., d_D\}$ is active. If so, *at any time* t the real traffic of the transmitter d_i is always served at an equal or lower rate than in the CPS model, ensuring that, if the backlog at the GPS node does not grow indefinitely, the corresponding CPS backlog does not grow indefinitely as well. Such scalers also guarantee that the traffic at each of the GPS nodes i in w_O is such that the traffic mapping the activity of the D2D pairs is at least served at the minimum possible rate it can be served at the equivalent CPS queue. We call such rate $R_{d_i}^{min}$.

Next we characterize stability and throughout in the system.

IV. Analytical Results

In the following we assume that arrivals are upper bounded by leaky bucket arrival curves [3]. If d is the D2D pair under analysis, ρ_d is the leaky bucket rate and σ_d the burstiness. This assumption does not limit the applicability of our analysis. Indeed, in practical settings almost any source can be described by some leaky bucket arrival curve, possibly by means of some conservative assumptions on the statistics of the traffic (e.g., burstiness of the flow).

A. Stability Region of the System

For a generic D2D system with D pairs and U cellular users, we can obtain sufficient conditions for stability exploiting the model introduced in Section III. For sufficient conditions for stability we mean a conservative bound for the actual stability region, described in terms of arrival rates. The following theorem introduces a set of inequalities ensuring that the arrival traffic leaky bucket descriptors $(\rho_d, \sigma_d), d \in \mathcal{D}$, yield a stable behavior.

Theorem 4.1 (Stability Region): Consider an in-band underlay D2D system with D pairs having traffic demands upper bounded by leaky bucket arrival curves with parameters $(\rho_d, \sigma_d), d \in \mathcal{D} = \{1, ..., D\}$, and the generic sorting $d_O = \{d_1, ..., d_D\}$ of $\mathcal{D}$. For each d_O, the following inequalities yield a conservative estimate of the stability region:

$$\rho_{d_i} \leq \max\left(R_{d_i}^{min}, R_{d_i}^{up} - \sum_{k=1}^{i-1} \frac{R_{d_i}^{up} - R_{d_i}^{k,up}}{R_{d_k}^{up}} \rho_{d_k}\right), \forall d_i \in d_O. \tag{4}$$

The union of all conservative estimates of the stability region, obtained under all possible permutations d_O, lays within the stability region of the D2D system.

Considering that each possible sorting d_O represents one of the possible feed-forward networks we can gather from Section III-A, Theorem 4.1 just applies well-known Network Calculus results. For the proof see [13]. If stability can be proved for any of the possible feed-forward networks we gather from the CPS characterization of the D2D system we started with, then it is clear that also the D2D system is stable. In order to get some insights, (4) is the maximum arrival rate for the traffic of transmitters d_i that can be handled by its corresponding GPS node while considering as always active the queues $\{d_{i+1}, ..., d_D\}$ and the remaining queues as much active as their actual demands allow.

B. Saturation Throughput of Cellular Users

The model proposed in Section III can be also used to obtain a lower bound on the saturation throughput of each cellular user u in the scenario, for the stable long term rates of D2D users $\{\rho_1, ..., \rho_D\}$. To achieve this goal, we build a new CPS model with $D + U$ queues, of which D are in one-to-one mapping with the D2D transmitters and U model the activity of cellular users. Exactly as in Section III, we aim at studying such CPS through a set of feed-forward networks such that, *at any time* t, the traffic of the cellular transmitters at the corresponding GPS nodes is always served at most as fast as in the CPS. Then, if ρ_u is the maximum arrival rate of queue u for which we can prove stability, ρ_u is also a valid lower bound for the saturation throughput of u. Indeed, when the transmission queue of u is stable, all the traffic that enters the queue also leaves it.

Given a particular set of long term rates $\{\rho_1, ..., \rho_D\}$ for the D2D users, which we are able to prove to be stable, we build the feed-forward networks following the same mechanism presented in Section III. In each of the feed-forward networks, the jth GPS node represents either a cellular or a D2D user, and its capacity $R_{d_j}^{up}, d_j \in d_O$ and d_O sorting of $\mathcal{D} \cup \mathcal{U}$, is computed as the minimum of the service rates for d_j considering as active transmitters all possible subsets of transmitters mapped onto GPS nodes $\{j+1, ..., D+U\}$. Please note that the cellular transmitters mapped onto one of the GPS nodes $\{j+1, ..., D+U\}$ have to be considered in all possible subsets, as they are in saturation. For each resulting system state $S = \{d_j\} \cup A$, $A \subseteq \{d_{j+1}, ..., d_{D+U}\}$, the service rate of d_j is computed by averaging over the possible scheduling permutations of interfering D2D:

$$R_{d_j}(S) = \begin{cases} \phi \frac{\sum_{S_O} b_{d_j}(S_O)}{D_S!}, & \text{if } d_j \text{ is D2D, } U_S = 0, \\ \frac{\phi}{U_S} \sum_{i \in \mathcal{U}_S} \frac{\sum_{S_O} b_{d_j}^i(S_O)}{D_S!}, & \text{if } d_j \text{ is D2D, } U_S \geq 1, \\ \frac{\phi}{U_S+1} \frac{\sum_{S_O} b_{d_j}(S_O)}{D_S!}, & \text{if } d_j \text{ is cellular,} \end{cases} \tag{5}$$

where U_S and D_S are the number of cellular and D2D transmitters in S, respectively, $\mathcal{U}_S$ is the set of cellular transmitters in S, S_O is a sorted permutation of the D2D transmitters in S, and $b_{d_j}(S_O)$ is the amount of bits per RB transmitted by d_j when the S_O list is chosen by the scheduling algorithm. $b_{d_j}(S_O)$ is computed as in (1).

For each GPS node, the scalers that modify the incoming "coupling" traffic are set as in (3), where $R_{d_j}^{k,up}$ is computed as the minimum of the service rates for d_j obtained by considering all the possible subsets of transmitters mapped onto GPS nodes $\{k, ..., j-1, j+1, ..., D+U\}$ as active transmitters. Again, all the cellular transmitters mapped onto one of the GPS nodes $\{k, ..., j-1, j+1, ..., D+U\}$ have to be considered in all possible subsets, since in saturation. In addition, in order to consider the cellular transmitters in saturation, in each of the GPS nodes onto which the cellular transmitters have been mapped, we feed traffic having finite burstiness and infinite long term arrival rate.

When the particular ordering d_O of $\mathcal{D} \cup \mathcal{U}$ is analysed, a valid lower bound for the saturation throughput of cellular transmitter u, i.e., $\rho_u(d_O)$, mapped in d_O onto the jth GPS node, is the maximum achievable rate computed from (4), i.e.:

$$\rho_u(d_O) = \max\left(R_u^{min}, R_u^{up} - \sum_{k=1}^{j-1} \frac{R_u^{up} - R_u^{k,up}}{R_{d_k}^{up}} \rho_{d_k}\right). \quad (6)$$

Finally, the overall lower bound for the saturation throughput of u is $\rho_u = \max_{d_O} \rho_u(d_O)$.

V. Proportionally Fair Optimization

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

A. Problem Formulation

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

$$\begin{aligned}
&\underset{\rho^*}{\text{maximize}} \quad \max_{d_O} \sum_{i=1}^{D} w_{d_i} \log\left(\rho_{d_i}^*\right), \\
&\text{subject to:} \quad \forall i = 1, \ldots, D, \\
&\rho_{d_i}^* + \min\left(R_{d_i}^{up} - R_{d_i}^{min}, \sum_{k=1}^{i-1} \frac{R_{d_i}^{up} - R_{d_i}^{k,up}}{R_{d_k}^{up}} \rho_{d_k}^*\right) \leq R_{d_i}^{up}, \\
&\rho_{d_i}^* \leq \rho_{d_i},
\end{aligned} \quad (7)$$

where the maximization is performed over any possible ordering d_O of $\mathcal{D}$ tuning the long term rates ρ^* of the shapers we want to introduce in the network and where w_{d_i} is the weight assigned to the transmitter d_i in order to achieve proportional fairness. The first constraint ensures that the long term rate $\rho_{d_i}^*$ we pick leads to a stable transmission queue for the generic transmitter D2D d_i (directly from Theorem 4.1), while the second ensures that the long term rate we pick for the shapers is at most the one of the demand of the generic transmitter d_i.

When dealing with a particular sorting d_O, the presence of the $\min$ function in the constraints leads to a piece-wise linear feasibility region. To deal with the problem, we use the Big-M transformation, so that the resulting optimization problem has a convex objective function and belongs to the mixed-binary programming family. We achieve the global optimum of the optimization through *branch-and-bound* ($B\&B$) [4].

B. Derivation of a Heuristic

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

The second and most important approximation we introduce is in the set of networks that we evaluate. Instead of choosing all the sortings d_O, we try to identify the subset of feed-forward networks that most probably contains the stable set of long term rates ρ^* that maximise our proportional fair problem. We limit to those networks the evaluation of (7).

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

Obviously, the higher the weight of the D2D pair in the proposed optimization, the higher is the contribution to the utility function of the system. Therefore, the objective of the heuristic we propose is to maximise the achieved throughput of the D2D pairs associated to the largest weights. Accordingly, the small subset of networks we decide to evaluate is initialized by probabilistically mapping D2D transmitters onto GPS nodes sequentially, starting from the earliest stages of the feed forward network, with probabilities proportional to the weights w_i of D2D transmitters in the objective function of the optimization problem (7). We repeat the probabilistic mapping process to populate a small set of starting points for the heuristic. For each starting point network, we identify the set of $D-1$ "neighbour" networks each obtained by swapping two adjacent GPS nodes. Within this set of neighbours, the network with the higher proportional fairness for the system is taken as the new reference network, and the evaluation of neighbour feed-forward networks is repeated until a local maximum for the proportional fairness is found. Obviously, the trade-off among computation and accuracy of the solution is given by the number of starting points we choose.

VI. Numerical Evaluation

In this section we validate the accuracy of our CPS-based model for in-band underlay D2D communications. First, in a small scenario, we simulate a real D2D system, in order to evaluate the main analytical results we presented in the paper, i.e., the conservative characterization of the stability region and the lower bound of the saturation throughput of the cellular users. We chose a small scenario for graphical purpose, although all the results we show can be easily replicated for a larger number of devices. Subsequently, we show the advantage achieved exploiting our analysis when a

TABLE I: Simulation Setup

LTE Carrier	2.45 GHz
UL Bandwidth	20 MHz
B	200 kHz
Subframe Duration τ	1 ms
b_M(64-QAM with coding rate ∼0.93)	468 $bits/RB$
Simulated area	300×300 m
Max distance between D2D TX/RX	40 m
r_0	3 m
η	2.5
σ_S	8 dB
$\mathcal{N}$	$3.98 * 10^{-18}$ W/Hz
P_{TX}^{MAX}	200 mW
packet length	12,000 bits

proportional fair assignment of the resources is the goal. For the whole set of proposed simulations, Table I summarizes the values of the parameters we use. We also use power control, aiming to achieve a per-RB SNR of 50.

A. Stability Region and Saturation Throughput

We pick a scenario where 3 cellular transmitters and 2 D2D pairs are present.

The position of cellular transmitters ($UE1-UE2-UE3$) and eNB is fixed as in Fig. 3. The distance among cellular transmitters and eNB is exactly the same for the 3 UEs (100 m). In different simulations, the D2D transmitters are moved symmetrically from position A to position B, although their distance from the eNB is also fixed to 100 m. By moving the D2D pairs closer and closer, we want to evaluate the impact of coupling among transmissions.

To evaluate the impact of distance for D2D transmissions, Fig. 3 also shows our conservative estimate of the stability region for D2D pairs and the rates achieved in the simulator at the two extreme positions we studied (position A, maximum distance; position B, minimum distance). Specifically, Fig. 3 includes the contour of the achievable rates computed via Monte Carlo experiments. As expected, the stability region in the two cases is extremely different. When widely distant, the D2D pairs do not influence each other and they substantially increase the spectral efficiency of the cell (both are able to transmit almost in each case using the maximum rate). When close, the D2D transmission heavily impact on each other, and a small increase of the demand of one of the two causes a sensible reduction of the throughput achieved by the other. Regarding the quality of the stability region achieved analytically, the difference of the areas shown in the picture is of the 0.387% and the 0.092%, when the D2D transmitters are in position A and B respectively. We also evaluated randomly generate topologies with 3 cellular users and 2 D2D users, and the biggest difference observed was as little as 2.18%.

To dig into the results, Fig. 4 represents the effect of the D2D pairs over the cellular transmitters when the D2D transmitters are in the most coupled position, i.e., in position B in Fig. 3. The figure shows both the conservative estimate achieved by means of the introduced model, and the saturation throughput achieved by simulation. Inspecting the results, it is clear that UE TX3 achieves always the same throughput, independently from the D2D demands. Indeed, due to the proximity of the D2D receivers to UE TX 3, the scheduling policy does not allow D2D pairs to transmit so to avoid poor

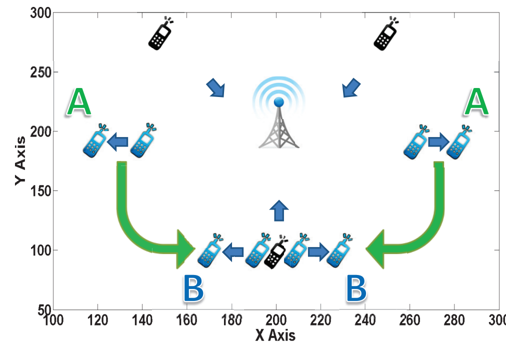

(a) Position of the devices.

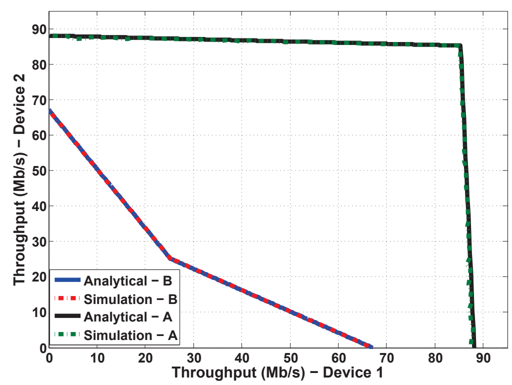

(b) Stability Region.

Fig. 3: Scenario under analysis.

rates. It is possible to evaluate from the particular chosen case the benefit that D2D transmission brings in throughput terms. Specifically, for each demand set of D2D transmitters, it is easy to compute the throughput reduction suffered by cellular users. However, even if the cellular transmitters loose just little achievable throughput in total, the sum of D2D throughputs is potentially very high (grater than 60 Mb/s). Hence, in-band underlay D2D ensures efficient utilization of resources.

Fig. 5 quantifies the difference among the conservative analytical estimation of the saturation throughput and the throughput achieved via simulations by cellular users (see Fig. 4). In the case under analysis, the worst per-user underestimation goes as far as 40%, although the average underestimation is as low as 11.18% and large underestimates only occur for low saturation throughputs. All other evaluated cases presented similar values for the underestimation of the saturation throughput of the cellular users, and the worst per-user average underestimation we observed was 13.01%.

B. Optimization

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

In order to show the scalability of the approach we propose, here we choose larger scenarios. In Table II we present the

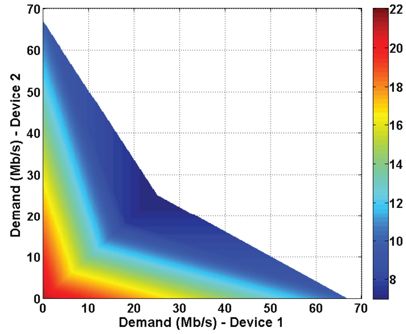

(a) UE TX1 (Model).

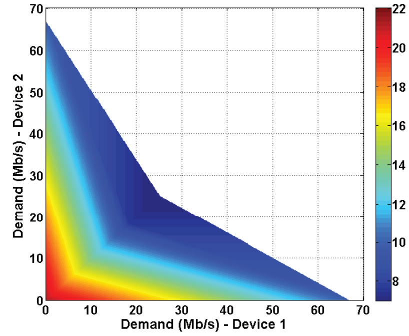

(b) UE TX2 (Model).

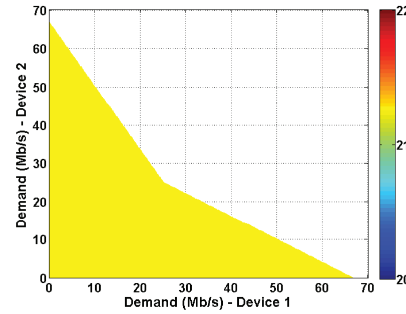

(c) UE TX3 (Model).

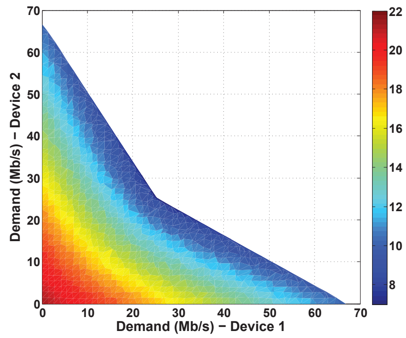

(d) UE TX1 (Simulation).

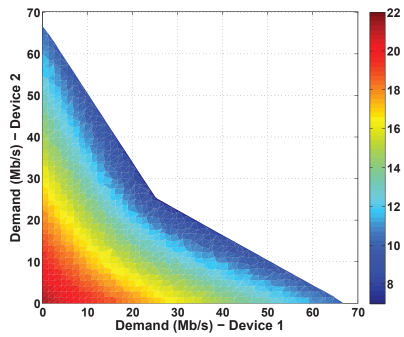

(e) UE TX2 (Simulation).

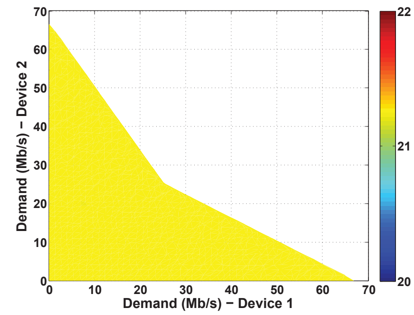

(f) UE TX3 (Simulation).

Fig. 4: Maximum Achievable Throughput, Cellular UEs. Analysis vs. Simulations.

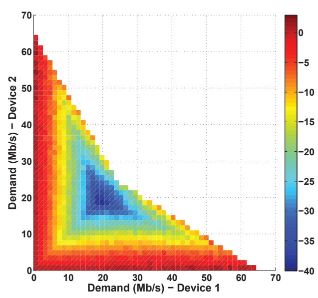

(a) UE TX1.

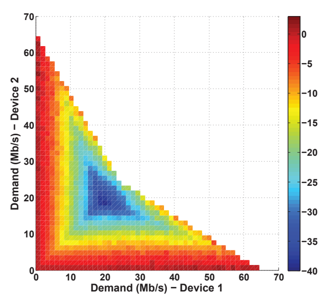

(b) UE TX2.

Fig. 5: Cellular saturation throughput differences between analysis and simulations (percentage w.r.t. simulations).

TABLE II: Heuristic vs. Brute force: Utility

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

results obtained when 20 cellular transmitters and 3, 4, 5 and 6 D2D transmitters are present. The choice of the number of transmitters is such that we can solve the optimization problem with the brute force approach.

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

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

TABLE III: Heuristic vs. Brute Force: Complexity

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

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

In order to show the gain that the knowledge of the conservative estimate of the stability region brings, we also simulate all the scenarios for which the optimization was performed. We simulate, for each of the cases, the operation of the in-band underlay D2D system, considering or not the presence of shapers at the D2D transmitters set as the output of the optimization problem we solved. In both cases we compute the throughput achieved from the D2D transmitters, and then we compare the corresponding log-utility (see (7)).

TABLE IV: Optimization vs. Saturation: Utility

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

As it easy to see from Table IV, even if we achieve just a conservative estimate of the whole stability region, the shapers improve sensibly the value of the utility achieved. In particular, in large scenarios, the utility improves up to 35.2% on average and 131.53% in the best case. Please note that the simulations with shapers ("Mean Opt." in Table IV) achieve almost the same utility as the one found with our heuristic (Table II).

VII. Related Work

To date, few works have tried to characterize the stability region of a D2D system. Indeed, the typical assumption is that the system is in saturation [16] or performance evaluation is performed via simulations [9]. The works that are closer to ours are [11] and [10]. [11] computes the Pareto boundary of the stability region without considering user traffic demands, but assuming that the transmission power of user devices can be tuned. Besides the fact that the analysis does not hold when the nodes are not in saturation, the paper applies to settings with few nodes. In our case, we assume that the power control mechanism acts on a different time scale than the scheduling of the resources. Nonetheless, our approach applies also to scenarios with large number of transmitters, and it remains accurate even when not all users are in saturation. The approach used by [10] exploits a CP modelling of D2D transmissions and could be used to achieve the stability region of the system. Two major assumptions are used though. The scheduling policy taken into account is the Full Frequency Reuse, i.e., D2D transmitters use the channel every time their transmission queue is not empty. Furthermore, the buffers at the transmitters have finite capacity. As a result, the model studies a simplified version of the D2D system, where a monotonic CP is derived from a Finite State Markov Chain characterization of the system. On the contrary, we generalize the study of D2D communications allowing buffers of infinite capacity and a complex, but realistic, scheduling policy. Such choice leads to a non-monotonic CP characterization of the system, that can be analysed, as far as we know, just through worst case analysis.

Several pieces of work try to optimize the mechanism which assign resources to the different users when the system is in saturation. Those optimizations also hold when the aim is fairness. For example [7] schedules resources in order to accomplish given QoS requirements, without taking into consideration the demands of the users while scheduling. Finally, [15] proposes a game theory approach for the scheduling of *each* RB, where the parameters are set in order for the transmitters to achieve a given fairness. Besides the complexity of the approach, the paper is designed to work when all the devices have traffic to transmit.

VIII. Conclusions

In this paper, we have introduced a novel analytical approach to D2D systems based on a Coupled Processors System model. We have shown how to use such characterization to optimize a D2D system under any given operational condition, i.e., also when saturation does not hold. Specifically, our method accurately estimates the stability region of D2D users and the saturation throughput of cellular users. In addition, we have shown how to apply our method to determine a proportionally fair allocation of throughput among D2D users.

IX. Acknowledgments

This work has been funded by the EU Seventh Framework Programme under grant agreement n° 318115 (CROWD).

The work has been partially supported by the Spanish MICINN grant TEC2011-29688-C02-01 and by program S2013/ICE-2919 of the Madrid Regional Government.

References

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

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

[3] J.-Y. L. Boudec and P. Thiran. *Network Calculus: A Theory of Deterministic Queuing Systems for the Internet*, volume 2050 of *LNCS*. Springer, 2001.

[4] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, USA, 2004.

[5] K. Doppler, M. Rinne, C. Wijting, C. B. Ribeiro, and K. Hugl. Device-to-device communication as an underlay to lte-advanced networks. *Communications Magazine, IEEE*, 47(12):42–49, 2009.

[6] M. Fidler and J. B. Schmitt. On the way to a distributed systems calculus: An end-to-end network calculus with data scaling. In *ACM SIGMETRICS Performance Evaluation Review*, volume 34, pages 287–298. ACM, 2006.

[7] M.-H. Han, B.-G. Kim, and J.-W. Lee. Subchannel and transmission mode scheduling for D2D communication in OFDMA networks. In *Proceedings of IEEE VTC-Fall*, pages 1–5, 2012.

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

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

[10] L. Lei, Y. Zhang, X. S. Shen, C. Lin, and Z. Zhong. Performance analysis of device-to-device communications with dynamic interference using stochastic petri nets. *IEEE Transactions on Wireless Communications*, 12(12):6121–6141, 2013.

[11] P. Phunchongharn, E. Hossain, and D. I. Kim. Resource allocation for device-to-device communications underlaying lte-advanced networks. *Wireless Communications, IEEE*, 20(4), 2013.

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

[13] C. Vitale, G. Rizzo, and B. Rengarajan. *Performance Bounds in Coupled Processor Systems*, 2013. Available at *http://publications.hevs.ch/index.php/attachments/single/665*.

[14] X. Wu, S. Tavildar, S. Shakkottai, T. Richardson, J. Li, R. Laroia, and A. Jovicic. Flashlinq: A synchronous distributed scheduler for peer-to-peer ad hoc networks. *IEEE/ACM Transactions on Networking (TON)*, 21(4):1215–1228, 2013.

[15] C. Xu, L. Song, Z. Han, Q. Zhao, X. Wang, and B. Jiao. Interference-aware resource allocation for device-to-device communications as an underlay using sequential second price auction. In *Proceedings of IEEE ICC*, pages 445–449, 2012.

[16] C.-H. Yu, K. Doppler, C. B. Ribeiro, and O. Tirkkonen. Resource sharing optimization for device-to-device communication underlaying cellular networks. *Wireless Communications, IEEE Transactions on*, 10(8):2752–2763, 2011.