Should I seed or should I not: On the remuneration of seeders in D2D offloading

Filippo Rebecchi, Marcelo Dias de Amorim, Vania Conan

Should I seed or should I not: On the remuneration of seeders in D2D offloading

Filippo Rebecchi⋆, Marcelo Dias de Amorim⋄, and Vania Conan⋆

⋆Thales Communications & Security
{filippo.rebecchi,vania.conan}@thalesgroup.com

⋄LIP6 / CNRS UPMC Sorbonne Universités
marcelo.amorim@lip6.fr

***Abstract*—Traffic offloading using opportunistic device-to-device (D2D) communications is a new and exciting opportunity for cellular operators to cope with the unprecedented mobile data growth. A limitation of existing proposals is that they assume that all terminals are, by default, involved in the D2D forwarding process. In particular, they do not capture the need to *reward* seed users. For this reason, we include a rewarding cost in the design of the opportunistic offloading strategy. In our solution, we make the difference between nodes that receive content through the cellular channel only (*leechers*) and nodes that take part in the forwarding process (*seeders*). The key point for an operator is to design a global strategy to select which nodes act as seeders and which ones as leechers, in order to reduce the total dissemination cost. We formulate this question as a stochastic control problem that we solve using an application of Pontryagin's Maximum Principle. We provide a mathematical framework to devise the optimal strategy for opportunistic offloading under a generic cost model. First, we show that an optimal solution exists; then, from this policy, we extract some insights to develop heuristics. Finally, we discuss the advantages of the proposed model compared to the classic seeder-only model. We demonstrate that separating seeders/leechers leads to better incentive strategies in the most demanding cases of content with a large span of delivery delays.**

***Index Terms*—Data offloading, optimal control, device-to-device communications, epidemics.**

I. Introduction

Device-to-device (D2D) communications are a well-timed strategy for operators to face the ever-increasing mobile data demand by *offloading* part of the traffic from their cellular infrastructure. Motivated by the delay-tolerance and redundancy of some types of content, operators may send data only to a subset of requesting users (*seeders*), which act as opportunistic forwarders to help propagate content using D2D communications. The combination of two complementary channels (cellular and D2D) provides extra capacity, helping reduce the impact of redundant traffic.

Several past works demonstrated that opportunistic D2D offloading does help relieve congested cellular networks [1], [2], [3]. Nevertheless, the effectiveness of this strategy revolves around the willingness of users to cooperate as forwarders. Incentives are then needed to stimulate participation, and to reward users that act as data relays. Existing proposals assume that all users are, by default, also seeders in the D2D domain (i.e., potential forwarders). Such an assumption may lead to suboptimal results when seeders are rewarded for transmitting content on behalf of the infrastructure. In those scenarios,

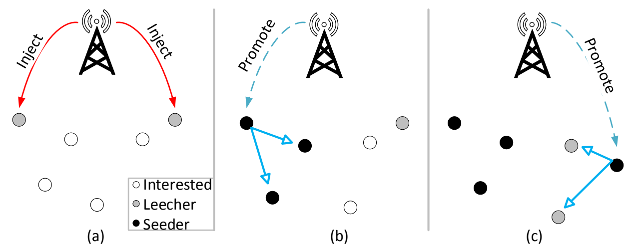

Fig. 1. Offloading process: the infrastructure selects two nodes as content initial leechers (Fig. 1(a)), deciding that one of the leechers should be promoted as seeder (Fig. 1(b)). Later on, the infrastructure estimates that it is worth promoting another node because the D2D transmissions are not enough to guarantee sufficient dissemination (Fig. 1(c)).

uncontrolled D2D communications may generate additional costs without necessarily bringing dissemination gains.

By borrowing the concept from the peer-to-peer jargon, we introduce *a clear separation between leechers and seeders*, as shown in Fig. 1. Leechers receive content via the cellular infrastructure, but only nodes promoted to the seeder state are allowed to forward the content opportunistically. The separation between leechers and seeders provides operators with an additional degree of freedom. The balance between instantaneous cost and future benefits of *leech* and *seed* decisions is strategic to data dissemination, given that the available resources (cellular bandwidth and rewards) are limited. Following the terminology introduced in [4], data dissemination follows a ”push-pull” model, where receivers can be both consumers and disseminators of data.

We investigate the following problem: *which fraction of leechers should be promoted as seeders and when should this happen*? We provide an answer to this question using a mathematical support that leads to an optimal solution. Content diffusion in opportunistic networks is comparable to the spreading of a disease in a population. We model the dissemination process using a variant of the classic Susceptible-Infected-Recovered (SIR) epidemic model from Kermack-McKendrick [5]. Since operators strive to optimize the distribution cost, we translate the possible decisions they can take into a cost function. We apply Pontryagin's Maximum Principle to minimize the cost function subject to the state-equations that govern content distribution. How to select the best seeders has been considered in the literature from the spatial re-use, throughput, and interference points of view, rather than from an incentive/economic perspective. As far as we know, no existing work on D2D offloading contemplates the difference between leechers and seeders, failing thus to quantify the trade-off that exists between performance and cost in more realistic systems. Performance evaluation of truly opportunistic forwarding has been treated extensively by means of ordinary differential equations (ODEs) [6] or Markov chains [7]. Instead, we propose an extended model that couples opportunistic dissemination with infrastructure control. Indeed, a central offloading coordinator controls the cellular injections and the promotion of users to the seeder state to reach optimal dissemination. The centralized scheme avoids the need for users to exchange periodically contextual information, such as meeting times, which instead can be directly extracted from the cellular infrastructure. Moreover, users are not required to continuously compute their seeding status, thus economizing battery and computation power.

In summary, the main contributions of this paper are:

- **Leecher-seeder model.** We propose a D2D offloading model where the network operator controls content injections and the promotion of users to the seeder state.
- **Optimal control.** We formulate D2D data offloading as an optimal control problem to minimize the dissemination cost in a hybrid scenario. A cost function trades off monetary and network resources consumed to reach a certain dissemination level and user satisfaction.
- **Cost functions.** We prove how to solve the offloading problem for different cost functions. We show that under plausible cost functions, the control of the injection is continuous, while promotions have an on-off behavior.
- **Evaluation.** We evaluate the sensitivity of the optimal controls to different contact rates and delay tolerance values. We evaluate the optimality of our strategy against other heuristics. Finally, we confirm the benefit of the proposed seeder-leecher model compared to the simplified seeder-only model currently employed in literature.

The remainder of this paper is structured as follows: Section II introduces the leecher-seeder model. Cost-related issues are discussed in Section III. The optimal control problem is formulated and solved in Section IV. Section V presents numerical results. We postpone related work to Section VI. Finally, Section VII draws the conclusion and perspectives.

II. System description and scope

Overview. We formulate D2D offloading as an optimal control problem, modeling data dissemination using a variant of the classic SIR model for epidemics. In our model, some users request data and are referred to as *interested*. Initially, all the nodes are in the *interested* state waiting for content. At this stage, the operator can only use cellular transmissions to reach a subset of the *interested* users. Nodes that receive the content enter the *leecher* state, still not playing any active role in data distribution. At this point, the coordinator can promote a fraction of them to the *seeder* state to diffuse the content.[1]

TABLE I
List of parameters.

Parameter	Definition
$n_I(t)$	fraction of interested nodes at t
$n_L(t)$	fraction of leecher nodes at t
$n_S(t)$	fraction of seeder nodes at t
$\lambda(t)$	contact rate at t
$u_I(t)$	direct injection rate at t
$u_P(t)$	promotion rate at t
$I_{max}(t)$	maximum injection rate at t
T	content lifetime
$\Phi(.)$	final payoff function for interested nodes at T
$f(.)$	instantaneous cost function for injection
$g(.)$	instantaneous cost function for promotion

Network model. The system consists of N mobile nodes and one content to be distributed by the infrastructure to all nodes within the lifetime T. Intermediary nodes can be used as opportunistic relays. Following the notation introduced above, nodes can be in the *interested*, *leecher*, or *seeder* states. Their respective fractions are $n_I(t)$, $n_L(t)$, $n_S(t)$, and $n_I(t)+n_L(t)+n_S(t) = 1\ \forall t \in [0, T]$. Therefore, we can always represent the system using only two states. An interested node can receive the content whenever in contact with a seeder, but not with a leecher. For the sake of clarity, Table I provides a summary of the parameters used along the paper (some of them will be explained in the following sections).

Cost. Incentives to reward user participation in data dissemination can be offered by using virtual credit schemes or discounts. For now, we do not consider any additional cost related to overhead, signalization, or maintenance of promotions, which are left for future work. Hence, promotion to the seeder state does not represent a cost in itself for operators, but allows users to be rewarded for their contribution in data dissemination. On the other hand, injections through the cellular infrastructure entail direct costs, in general related to resource availability in the access network. We will further elaborate the reflection on cost-related aspects in Section III.

Encounters and communication opportunities. In the real world, the system under observation can be described with discrete values (e.g., the number of users and the number of cellular transmissions performed). For the sake of modeling, we consider instead continuous values for the state and the control values. We assume that N is large and that encounters are homogeneous. Consistently with the literature, we use a *mean field* model that is accurate for a large population. D2D dissemination can be regarded as the spread of infective disease – not surprisingly, *epidemic* routing is a conventional forwarding strategy in opportunistic networks. As with a disease contagion in a population, in our model content spreads from *seeders* to *interested* nodes when such a pair enters in physical

[1] Note that we do not align to the traditional SIR nomenclature. Instead, we use a nomenclature borrowed from peer-to-peer networks: “susceptible” users in the SIR model are analogous to *interested* in our model. Similarly, “infective” and “recovered” nodes are named respectively *seeders* and *leechers*.

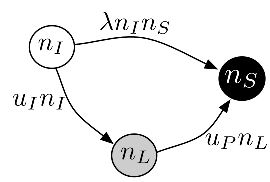

Fig. 2. State transition rates for the leecher-seeder model.

proximity.[2] We describe the state evolution with a system of ODEs along with a set of initial and terminal constraints. The contact rate $\lambda(t)$ rules the encounter of any two nodes at time t. We assume that $\lambda(t)$ includes also the uncertainties introduced by the wireless channel and the movement of nodes. At time t, we have $n_S(t)$ seeders capable of meeting $n_I(t)$ interested nodes. Consequently, interested nodes become seeders with rate $\lambda(t)n_I(t)n_S(t)$ (the transitions are illustrated in Fig. 2).

Injections and promotions. In line with the existing literature on data offloading [8], we consider a central coordinator that manages the cellular injections, adding to its duties also the responsibility to decide the promotion of nodes. It turns out that cellular injections increase the rate at which nodes become leechers (namely, nodes that have the content but do not distribute it). The intensity at which injections are performed is governed by the signal $u_I(t)$, a bounded Lebesgue integrable function with $0 \leq u_I(t) \leq 1 \ \forall t \in [0, T]$. The value of the injection rate is restricted in the interval $[0, I_{max}(t)]$, which measures the instantaneous available load on the cellular network. Consequently, $u_I(t)n_I(t) \leq I_{max}(t)$ describes the rate of injected copies.

Leechers carry the content but need to be promoted in order to contribute to data dissemination. Operators can promote only the necessary fraction of seeders. This is done via a control channel binding users to the central coordinator. As a result, leechers shift to the seeder state with intensity regulated by the signal $u_P(t)$, a bounded Lebesgue integrable function with $0 \leq u_P(t) \leq 1 \ \forall t \in [0, T]$. This increases the fraction of nodes in the forwarder state by a rate $u_P(t)n_L(t)$. Therefore, the following system of ODEs controls the evolution of the interested, leecher, and seeder nodes in the system:

$$\frac{dn_I}{dt} = -\lambda(t)n_I(t)n_S(t) - u_I(t)n_I(t), \quad (1a)$$

$$\frac{dn_L}{dt} = u_I(t)n_I(t) - u_P(t)n_L(t), \quad (1b)$$

$$\frac{dn_S}{dt} = \lambda(t)n_I(t)n_S(t) + u_P(t)n_L(t), \quad (1c)$$

with initial states $n_I(0) = i_0$, $n_L(0) = l_0$, and $n_S(0) = 1 - i_0 - l_0$. For the offloading problem we consider, we have $i_0 = 1$ and $l_0 = 0$, since we consider all users to be in the interested state at the beginning of content diffusion.

The equations above describe how the states n_I, n_L, n_S change at time t as a reaction to the control signals u_I, u_L. Note that $\frac{\partial n_I}{\partial t} + \frac{\partial n_L}{\partial t} + \frac{\partial n_S}{\partial t} = 0$; therefore, the model can always be expressed using only two out of the three equations.

[2]As already mentioned, our system shares also several similarities with peer-to-peer (P2P) networks, which features a large number of nodes downloading data from seeder nodes.

III. Distributing content is costly

The optimal offloading strategy consists in minimizing the amount of nodes still in the interested state at the end of content lifetime, while implementing a cost-savvy injection/promotion campaign. If cellular operators had no capacity limitations, then the optimal strategy would consist in injecting the maximum amount of data via the cellular channel. Instead, when capacity is limited, operators may seek to exploit D2D communication capabilities of their customers. Nevertheless, rewarding users can be costly, as both operational and budgetary constraints should be taken into account when serving content. Based on these aspects, Eq. 2 considers a cost function J that is general enough to grasp various types of cost incurred by operators:

$$J = \underbrace{\Phi[n_I(T)]}_{\text{payoff}} + \int_0^T \underbrace{f[u_I(t)\, n_I(t)]}_{\text{injection}} + \underbrace{g[\lambda\, n_I(t)\, n_S(t)]}_{\text{reward}} \, dt. \quad (2)$$

In Eq. 2, $\Phi[n_I(T)]$ represents the final payoff, i.e., the cost that the operator has to pay for having failed to satisfy the fraction $n_I(T)$ of users by the deadline. This may lead to loss of earnings due to missed deliveries or to extra costs in terms of final injections [2], [9], [10]. Word of mouth among angry consumers may also boost the commercial fallout of missed or late deliveries making this term highly non-linear. $f[u_I(t)\, n_I(t)]$ captures the cost, in terms of network resources, of injections over the cellular channel. Despite advanced cellular technologies (e.g., LTE) increase the overall system capacity with multi-user diversity, each additional user to be served reduces the rate of improvement [11]. In other words, the radio resources allocated to transmit data do not follow a linear trend with the number of users. In addition, one must consider that overhead is larger when many users share the same bandwidth, rather than a single user benefiting from the whole bandwidth [12]. For these reasons, the injection term is likely another non-linear cost for operators. Finally, seeders are rewarded with $g[\lambda\, n_I(t)\, n_S(t)]$ (instantiated as discounts or virtual credits), accorded each time they make an opportunistic transmission. Examples of possible incentive strategies in D2D networks are provided in [4], [13].

The integral in Eq. 2 portrays the growing cost over time of these two latter terms. Note that promotion control u_L does not appear inside the cost function. As outlined in Section II, promoting a node to the forwarder state does not directly generate a cost. However, seeders will be able to transmit data opportunistically, possibly increasing the rewarding cost for the operator. For physical reasons, $\Phi(.)$, $f(.)$, and $g(.)$ should be monotonically increasing functions, with $\Phi(0) = f(0) = g(0) = 0$ (the cost for doing nothing is zero).

Relationship between control and cost. At any time, the offloading controller must choose the values of the control signals u_I and u_L. The decision is taken by assessing the fraction of nodes in each compartment (n_I, n_L, and n_S), the time remaining before the deadline, and the contact rate between nodes. Applied controls lead to two consequences: (i) direct effect, which generates the instantaneous costs $f(.)$ and $g(.)$ for the operator, and (ii) indirect effect, represented by the future change in states formalized by Eq. 1. The optimal offloading strategy requires the coordinator to plan its injection and promotion strategies by *minimizing the cost for the operator while maximizing the rate of change of state variables*.

Extreme cases. We aim to further clarify this concept by focusing on the injection control u_I. We consider two extreme strategies that do not consider the evolution of dissemination. The first myopic strategy always injects the maximum amount of copies through the cellular channel. In this case, the offloading controller does not take into consideration the D2D capabilities of nodes. It turns out that this strategy is strongly suboptimal. The other extreme is a minimum injection strategy, where the coordinator injects only once. It is intuitive that an optimal decision avoids such extremes.

IV. Optimal solution

We now have the sufficient background to derive the optimal offloading solution for our problem. By applying $n_S(t) = 1-n_I(t)-n_L(t)$, we can formulate the optimal control problem considering only the two state variables n_I and n_L. The system is controlled by the tuple $< u_I, u_P >$, which belongs to the set of all the admissible controls $U = \{u_I, u_P\}$, where u_I, u_P are Lebesgue integrable with $u_I, u_P \in [0, 1]$. The goal is to characterize the optimal controls $< u_I^*, u_P^* >$ that minimize the cost function J, subject to the constraints defined in Eq. 1:

$$\min_{u_I(t),u_P(t)\in U} J, \text{ subject to:} \tag{3a}$$

$$\frac{dn_I}{dt} = -\lambda(t)n_I(t)(1 - n_I(t) - n_L(t)) - u_I(t)n_I(t), \tag{3b}$$

$$\frac{dn_L}{dt} = u_I(t)n_I(t) - u_P(t)n_L(t), \tag{3c}$$

$$\begin{aligned} &n_I(t) \geq 0, n_L(t) \geq 0, n_S(t) \geq 0,\\ &n_I(t) + n_L(t) + n_S(t) = 1, \end{aligned} \tag{3d}$$

where the initial states are $i_0 = 1$ and $l_0 = 0$.

A. Existence of an optimal control

The existence of an optimal solution can be determined by applying the Filippov-Cesari theorem (Theorem 4.1 in [14]). To prove this theorem, the following conditions should be satisfied: (i) the functions inside the integral in Eq. 2 are continuous, bounded, and convex in controls, with bounded derivatives, (ii) the control signals $u_I(t)$ and $u_P(t)$ take values in a closed set, and (iii) Eqs. 3b and 3c are linear in the controls. These conditions guarantee the existence of an optimal solution.

B. General solution – Pontryagin's Maximum Principle

Since an optimal control exists, we apply Pontryagin's Maximum Principle [15] to derive necessary conditions on the optimal control (Theorem 3.4 in [16]). The conditions of Pontryagin's maximum principle reduce the computation of an optimal strategy to the solution of a boundary value problem for a system of differential equations. Let the tuple $(n_I^*(.), n_L^*(.), u_I^*(.), u_P^*(.))$ be an optimal solution to the problem formalized in Eq. 3.[3] There exist continuous and piecewise continuously differentiable adjoint functions $p_I^*(t)$ and $p_L^*(t)$ that maximize the present-value Hamiltonian function H.

For the sake of ease of mathematical manipulation, we transform the problem into a *maximization problem* by multiplying the Hamiltonian by -1. We also remove the dependence from time whenever possible, in order to make reading easier:

$$\begin{aligned} H(n_{I,L}, u_{I,P}, p_{I,L}, t) = &-f[u_I\, n_I] \\ &- g[\lambda\, n_I\,(1 - n_I - n_L)] \\ &+ p_I[-\lambda n_I(1 - n_I - n_L) - u_I n_I] \\ &+ p_L[u_I n_I - u_P n_L]. \end{aligned} \tag{4}$$

The Hamiltonian function, in analogy with the corresponding concept occurring in traditional mechanics, balances the rate of change of states and the cost incurred by operators. Indeed, the Hamiltonian is a generalized profit rate that includes both direct and indirect effects, and has to be maximized at each instant. The *weights* for the state variables are given by the adjoint functions p_I and p_L, which represent the marginal increase of H due to an increment in the state. Consequently, the adjoint equations p_I and p_L evaluated at the optimum are:

$$\begin{aligned} \frac{dp_I^*}{dt} &= -\frac{\partial H(.)}{\partial n_I}\bigg|_{n^*_{I,L},u^*_{I,P},p^*_{I,L}} = \\ &= \frac{\partial f(.)}{\partial n_I} + \frac{\partial g(.)}{\partial n_I} - p_I\left[\lambda\,(2n_I - 1 + n_L) - u_I\right] - p_L\, u_I, \end{aligned} \tag{5a}$$

$$\begin{aligned} \frac{dp_L^*}{dt} &= -\frac{\partial H(.)}{\partial n_L}\bigg|_{n^*_{I,L},u^*_{I,P},p^*_{I,L}} = \\ &= \frac{\partial g(.)}{\partial n_L} - p_I\lambda n_I + p_L u_P. \end{aligned} \tag{5b}$$

with transversality (terminal) conditions that describe what must be satisfied at the end of the time horizon T:

$$p_I(T) = \frac{\partial\Phi(n_I^*(T), T)}{\partial n_I}, p_L(T) = \frac{\partial\Phi(n_I^*(T), T)}{\partial n_L} = 0. \tag{6}$$

According to the maximum principle, there exist optimal controls, a tuple $< u_I^*, u_P^* >\in U$ of continuous and piecewise continuously differentiable functions, and their corresponding solutions n_I^*, n_L^* that maximize the Hamiltonian H satisfying Eqs. 5 and 6:

$$u^*_{I,P}(t) \in \arg\max_{u_{I,P}\in U} H(n_{I,L}, u_{I,P}, p_{I,L}, t). \tag{7}$$

The canonical system, composed of four coupled ODEs (Eqs. 3b, 3c, 5a, 5b) and the transversality conditions (Eqs. 6), determines a boundary value problem (BVP).

[3]Throughout the paper, variables with the star superscript (e.g., $n_I^*(t)$) represent the optimum value.

C. The case of data offloading

Let us now focus on the case of the seeder-leecher model by solving the optimization problem for a class of cost functions $\Phi(.)$, $f(.)$, $g(.)$ in Eq. 2. The choice of these functions follows the discussion in Section III. However, thanks to the flexibility of the model, they can be replaced at will, to take into account the specificities and operating costs of certain networks. Besides the existence constraints discussed in Section IV-A, by physical reasons, Φ, f, and g should be monotonically increasing functions starting at zero.

We consider an exponential function for the final payoff $\Phi(x) = e^x - 1$, a power-law function for the cellular injections $f(x) = bx^\alpha$, with $\alpha \geq 2$, and a linear function $g(x) = cx$ to reward seeders that distribute content. The final payoff function $\Phi(x)$ starts at zero and then increases exponentially to model the cost for missing the delivery by the deadline T. $f(x)$ represents the cost for injecting data on the cellular channel during the content lifetime. The power-law accounts for the cost of simultaneous data transmissions. Indeed, the more simultaneous cellular data transmissions, the more costly they become (in terms of consumed radio resources) at the base station. The power-law coefficient α depends on the considered network and on its overall congestion conditions. Since the domain of $f \in [0, 1]$, lower values of α give the steeper curves. The cost function $g(x)$ is linear, as we consider that the reward offered to forwarders for each opportunistic transmission they perform is fixed. By substituting the cost functions $\Phi(.)$, $f(.)$, and $g(.)$ in Eqs. 4, 5, and 6, the Hamiltonian, the adjoint functions, and the transversality conditions become:

$$\begin{aligned} H = &-b(n_I u_I)^\alpha - c(\lambda\, n_I\,(1 - n_I - n_L))+ \\ &+ p_I[-\lambda n_I(1 - n_I - n_L) - u_I n_I]+ \\ &+ p_L[u_I n_I - u_P n_L], \end{aligned} \tag{8a}$$

$$\begin{aligned} \frac{dp_I^*}{dt} = &\, b\,\alpha\, n_I^{\alpha-1} u_I^\alpha + c\,\lambda\,(1 - 2n_I - n_L) \\ &- p_I\,[2\lambda n_I - \lambda + \lambda n_L - u_I] - p_L\, u_I, \end{aligned} \tag{8b}$$

$$\frac{dp_L^*}{dt} = -c\,\lambda n_I - p_I \lambda n_I + p_L u_P, \tag{8c}$$

$$p_I(T) = e^{n_I^*(T)}, \tag{8d}$$

$$p_L(T) = 0. \tag{8e}$$

Injections. Given that Eq. 8a is strictly concave in the control variable u_I, we extract $u_I(t)$ using the Hamiltonian maximization condition ($\frac{\partial H}{\partial u_I} = 0$ evaluated at the optimum), along with the restriction on the maximum injection rate ($n_I(t) u_I(t) \leq I_{max}(t)\ \forall t \in [0, T]$). By defining the function $\psi(t) = \sqrt[\alpha-1]{\frac{p_I^* - p_L^*}{-\alpha b}}$, we can write the optimal solution as:

$$u_I^*(t) = \begin{cases} 0, & \text{if } \psi(t) < 0, \\ \frac{\psi(t)}{n_I(t)}, & \text{if } 0 \leq \psi(t) \leq I_{max}(t), \\ \frac{I_{max}(t)}{n_I(t)}, & \text{if } \psi(t) \geq I_{max}(t). \end{cases} \tag{9}$$

Equivalently, we have that $u_I^*(t) = \frac{\min[\max[\psi(t), 0], I_{max}]}{n_I(t)}$.

Promotions. In the case of promotions, since the Hamiltonian is linear in the control variable u_P, the maximization condition $\frac{\partial H}{\partial u_P} = 0$ is trivially satisfied and independent of u_P. The control in this case is called *singular* (Definition 3.40 in [16]) with a *bang-bang* solution, i.e., a control that switches discontinuously between one extreme to the other.

Since the maximization condition cannot help determine the optimal control, it is possible to rewrite Eq. 8a as:

$$\begin{aligned} H = &- b(n_I u_I)^\alpha - c(\lambda\, n_I\,(1 - n_I - n_L))+ \\ &+ p_I[-\lambda n_I(1 - n_I - n_L) - u_I n_I]+ \\ &+ p_L[u_I n_I] - u_P \underbrace{[p_L\, n_L]}_{\sigma}. \end{aligned} \tag{10}$$

By defining a switching function $\sigma = (p_L n_L)$, to maximize H the control should take its maximum (minimum) value when $\sigma < 0$ ($\sigma > 0$). By construction, $u_P \in [0, 1]$, and:

$$u_P^*(t) = \begin{cases} 0, & \text{if } \sigma > 0, \\ 1, & \text{if } \sigma < 0. \end{cases} \tag{11}$$

In order to be able to retrieve the evolution of the state and adjoint variables, we have to solve a system of coupled ODEs (respectively Eqs. 3b, 3c, 8b, 8c, 8d, and 8e), with a mix of initial and final conditions (boundary values). We solved it numerically by using the shooting method from the **R** package *bvpSolve* to compute the evolution of the state and adjoint variables as well as the optimal control [17].

V. Numerical results

To identify the optimal offloading strategy, we conduct numerical analyses using the software **R**. Firstly, we execute a sensitivity analysis on the values of λ and T in order to understand their implications in the offloading strategy. Performance depends strongly on these two parameters, so do the injection and promotion strategies. Then, we explore under which circumstances the seeder-leecher model brings advantages over the classic seeder-only model. Finally, we address some implementation issues by comparing the optimal strategy with several heuristic strategies, investigating under which conditions and limits they can be adopted.

A. Injections and promotions

We investigate "when" the infrastructure should inject copies of the content and promote nodes as seeders. Figs. 3 and 4 display the time evolution of both states and control variables for two different deadlines and different contact rates.

By comparing Figs. 3 and 4, we discover that there are more promotions for short deadlines. We conclude that content lifetime strongly influences promotions. Short lifetimes (e.g., $T = 5$) are not sufficient to yield complete dissemination under the provided cost-function. Instead, for $T = 10$ we obtain complete data delivery (at least for the best contact rates). When content dissemination is incomplete by the deadline, the final payoff $\Phi(n_I(T))$ takes a large part of the cost function J. This behavior confirms a well-known phenomenon in opportunistic networking: increased delivery delays improve the fraction of nodes that receive content through opportunistic communications. The added dissemination time allows fewer injections and promotions, thus lowering costs for cellular

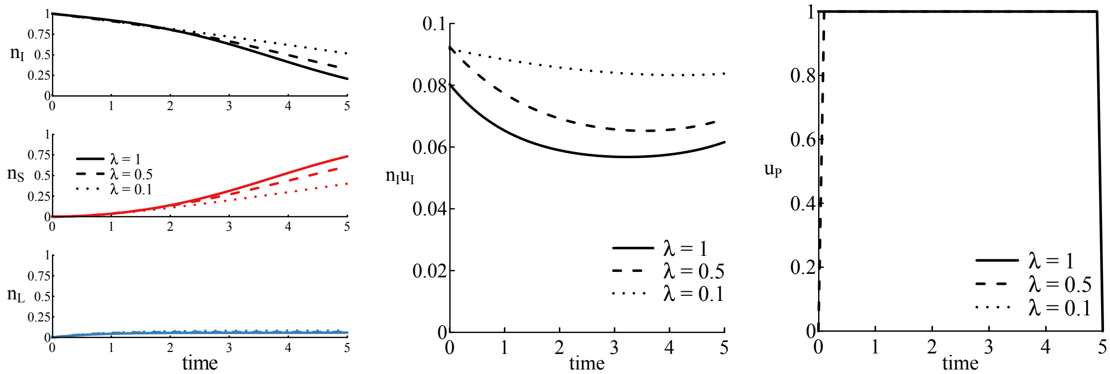

(a) state evolution (b) injections (c) promotions

Fig. 3. Optimal offloading for different contact rates λ. $T = 5s$ Other parameters: $I_{max} = 0.1$, $\alpha = 2$, $b = 10$, $c = 1$, $i_0 = 1$, $s_0 = 0$.

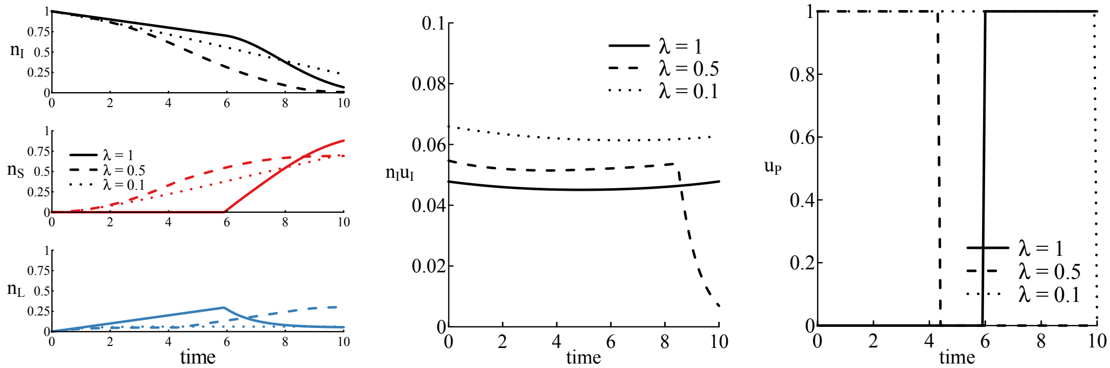

(a) state evolution (b) injections (c) promotions

Fig. 4. Optimal offloading for different contact rates λ. $T = 10s$. Other parameters: $I_{max} = 0.1$, $\alpha = 2$, $b = 10$, $c = 1$, $i_0 = 1$, $s_0 = 0$.

operators. On the other hand, increased delivery times could hurt user satisfaction.

Besides the deadlines, low contact rates also lead to stronger injection and promotion controls. Focusing on injections (Figs. 3b and 4b), we observe that the control is stronger in the beginning and in the end of the dissemination, following a nearly symmetrical pattern. This pattern depends on the interactions between n_I and n_S in Eq. 1a. Therefore, the injection rate is higher when a few nodes are in the seeder or the interested states – respectively at the beginning and at the end of the dissemination. Injections help overcome the slow start and the convergence time of opportunistic dissemination. Wang et al. previously pointed out the symmetric trend of the injection control, although for a simplified model [18].

Promotions (visible in Fig. 3c and 4c) follow a completely different pattern. For $T = 5$, the control is always at its maximum. The shorter deadline is responsible for the poor dissemination. Indeed, even with an extreme promotion strategy (always on), there are still nodes in the interested state at the deadline. Longer deadlines, instead, allow devising more effective promotion strategies. For $T = 10$, $u_P(t)$ shows three different patterns depending on the value of the contact rate. For $\lambda = 0.1$, the control is always at its maximum for the entire dissemination. $\lambda = 0.5$ presents an on-off behavior, with promotions that stop when the amount of seeders reaches significant levels (in order to self-sustain without costing too much to operators). Finally, for $\lambda = 1$, promotions are activated only after half of the dissemination period. This collides with the desire to attain the widest possible dissemination of the content. Although at first sight this might seem counter-intuitive, we must not forget that operators pay a small fee for each opportunistic transmission performed by users. Under higher contact rates, opportunistic dissemination has to be limited in order to save monetary resources.

As anticipated in Section IV-C, u_S takes the form of a *bang-bang* control with exactly one on-off switch. Injections performed when $u_S(t) = 0$ serve only to satisfy the fraction of users that will likely not receive the content by the deadline, without further improving the dissemination (because these nodes are not promoted, remaining leechers). Moreover, we point out that the optimal strategy does consider moments where no additional seeders are needed ($u_S(t) = 0$). This strengthens the idea that separating seeders and leechers is beneficial from a cost-benefit point of view.

B. Controlling seeders: Giving operators an edge

We examine in which cases it is worth separating leechers and seeders from the operator's point of view. Including the

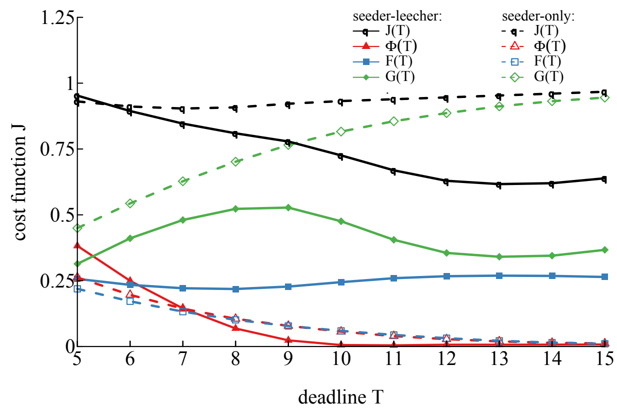

Fig. 5. Cost functional J and its main components for the optimal strategy using two offloading models (seeder-seeder and two-state), varying the deadline T . Other parameters: $\lambda = 0.5$ $I_{max} = 0.1$, $\alpha = 2$, $b = 10$, $c = 1$, $i_0 = 1$, $s_0 = 0$.

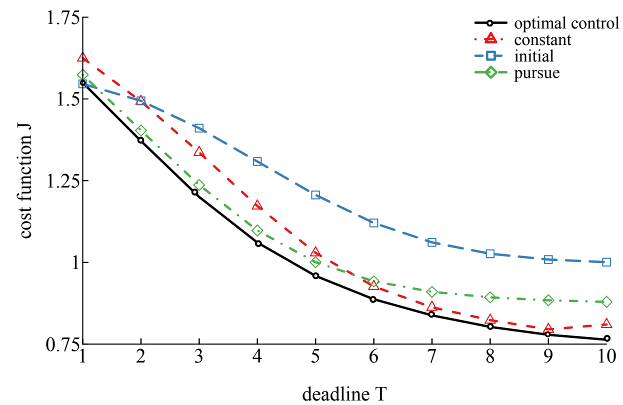

Fig. 6. Cost functional J for different control strategies varying the deadline T. Other parameters: $\lambda = 0.5$ $I_{max} = 0.05$, $\alpha = 2$, $b = 10$, $c = 1$, $i_0 = 1$, $s_0 = 0$.

leecher state in the picture is motivated by the fact that not all users carrying the content may be required to forward data. These scenarios hinge on a combination of factors, such as the contact pattern or the delay tolerance of the content. Eventually, adding an intermediate state between *interested* and *seeder* becomes cardinal when operators want to reward user participation. As many works suggest, offering some kind of incentive (i.e., discounts or virtual credits) motivates user participation [19]. However, current models in the literature consider only two states (interested and seeder).

The optimal solution for the seeder-only model will serve for comparison, and is detailed in Appendix A. In this case, the only control is the injection rate u_I. We plot in Fig. 5 the evolution of the cost function J divided by its three main components $\Phi(T)$, $F(T) = \int_0^T f[u_I(t)\, n_I(t)]dt$, and $G(T) = \int_0^T g[\lambda\, n_I(t)\, n_S(t)]dt$. $\Phi(T)$ is the final payoff value, due to nodes that have not received the content by the deadline, $F(T)$ is the total cost of injections, and $G(T)$ is the total cost of rewarding.

Separating seeders and leechers is advantageous for operators if compared to the classic seeder-only model. In Fig. 5, we can observe that, as the deadline increases, the leecher-seeder model improves its aggregate performance compared to seeder-only. From our analysis, for $T > 5$, the number of uninfected nodes at the deadline decreases steadily, reducing the relative weight of the term $\Phi(T)$ on the overall cost. Rewarding cost $G(T)$ takes the larger part of J, accounting for the cost to reward seeders. Seeding costs increase linearly for the seeder-only model as the deadline stretches, making up nearly the entirety of J. This confirms that an uncontrolled number of seeders can interfere with the will of operators to cut operational expenses. Instead, *a separation between leechers and seeders offers improved flexibility in the control of the offloading evolution, bringing clear advantages in terms of distribution costs*.

Note that for short deadlines (for $T \leq 5$ in the example), all connection opportunities should be used as captured by the seeder-only model. Indeed, the cost-functional is dominated by the final payoff $\Phi(T)$, whose value depends on missed data deliveries. The leecher-seeder model introduces an additional transition delay from the interested to the seeder state (the ODE formulation requires a non-null time to transit to a state); thus, its benefits come into play in the non-trivial situations of content with larger deadline requirements.

C. Implementation considerations

We first investigate how intuitive heuristics not requiring any optimization framework perform compared to the optimal strategy, and what lessons can be learned having the knowledge of the optimal injection and promotion controls. Finally, we consider the impact of realistic values of the contact rate.

Heuristics. The optimal solution requires to solve numerically the optimization problem, which is a complex task. This may not be possible, given the size of the network and time constraints, when the available processing power is limited. For this reason, we analyze the performance of simple heuristics. We compare the optimal strategy against three other heuristics. The first heuristic, named *initial control*, mimics an operator wanting to rely only on an initial subset of seeders. These seeders are the only way to distribute content by the deadline. This strategy relies on an initial injection at the rate $I_{\max}$, without performing any further injections. The second strategy, named *constant control*, steadily injects at a fixed rate of $\frac{I_{\max}}{2}$. These two strategies are static and do not require any knowledge of how the dissemination evolves. In both cases, the promotion control $u_R(t)$ is fixed at 1 for all the dissemination delay. Finally, we consider a more dynamic strategy, named *pursue control*, where both injection and promotion controls follow the evolution of the interested nodes $n_I(t)$. In this case, the control is strong at the beginning of the dissemination, gradually descending as the time passes by, following $n_I(t)$. The rationale behind this choice is that copies with a large tolerance to dissemination time are more effective in content dissemination. We compare these strategies in terms of the cost function introduced in Eq. 2, varying the deadline T and the contact rate λ.

In Fig. 6, we plot the cost functional J by varying the deadline. As expected, the functional J for the optimal control is always smaller than all the other heuristic strategies. However, for shorter deadlines ($T < 6$), the *pursue* strategy is very close

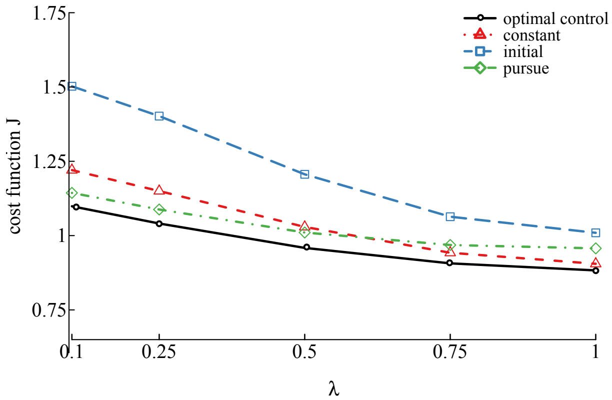

Fig. 7. Cost functional J for different control strategies varying the contact rate λ. Other parameters: $T = 10s$, $I_{max} = 0.05$, $\alpha = 2$, $b = 10$, $c = 1$, $i_0 = 1$, $s_0 = 0$.

to the optimal. This is the first lesson we can draw: with shorter deadlines, a control that follows the rate of interested nodes comes close to the optimal. As the deadline increases, the efficiency of the *pursue* strategy decreases. On the other hand, we note that the *constant* strategy approaches the optimum for larger deadlines. Indeed, the steady injection profile at rate $\frac{I_{max}}{2}$ is very similar to the one in the optimal strategy (depicted in Fig. 4b). Promotions are not adapted, concurring to increase the overall cost. Lastly, relying only on an initial set of seeders, without any additional injection during the dissemination duration (such as in the *initial* strategy), brings considerable efficiency drops. This happens because the set of initial leechers (which is inherently limited by $I_{\max}$) cannot cover the entire network in T. Fig. 7 shows the trend of the cost functional J varying the contact rate λ. In general, the relative performance of the heuristics is the same as in the previous case.

Contact rate. Fig. 7 outlines the importance of the contact rate in the performance of the offloading strategy. λ is at the base of the opportunistic diffusion between mobile users. However, the contact rate depends on the mobility pattern of users, and can vary in time. λ may also include the uncertainties induced by the wireless channel and the movement of nodes. Operators should estimate the value of λ in order to adapt the optimal solution to current network conditions. In this context, offloading architectures that employ a feedback mechanism can prove very useful [10], [2].

To give the reader an idea of the values at stake, the meeting time between nodes is nearly exponentially distributed when nodes move in a bounded region (of area A) according to a common mobility model [7]. Under these assumptions, for the contact rate the following formula holds: $\frac{\lambda}{N} \approx \frac{8\omega r v}{\pi A}$, where $\omega = 1.3683$ for the random waypoint, $\omega = 1$ for the random direction model, r is the transmission range, and v is the speed of the users. This means that, considering an area of 500×500 m^2, 100 mobile users equipped with a transmission technology that allows direct transmissions up to 50 m and mobility of 25 km/h following the random waypoint model, we find a contact rate parameter $\lambda = 0.488$. If we consider, instead, real mobility traces from the Infocom dataset [20], the average contact rate is $\lambda = 0.14167$.

VI. Related Work

D2D offloading is an interesting method to reduce congestion in cellular networks [8]. Unicast and multicast D2D dissemination strategies are studied in [21], [22], [23]. By adding the infrastructure in the picture, the focus shifts towards the selection of an optimal subset of users to kick-start dissemination [1], [24]. However, these works consider only the optimal selection of the initial seed nodes without controlling the evolution of the dissemination. More related to our work, heuristics are presented in [2], [9]. A learning approach can help identify the best data carriers [10] or broadcast modulation [25]. The optimization techniques presented in this work were initially developed as a solution of aeronautic problems, where engineers wanted to control a system by minimizing a certain performance index. Such techniques have not been sufficiently considered in the case of cellular data offloading. Instead, the works that are the closest to ours come from the literature on opportunistic diffusion [7], [6] and peer-to-peer networks [26], [27]. Optimal control has been employed to allocate resources (e.g., vaccines or security patches) to prevent the diffusion of epidemic diseases [28], [29] or computer viruses [30]. The main difference to our work is that they aim at preventing the infection rather than encouraging it. Instead, applications of optimal control to boost epidemics exist in the area of marketing [31], and opportunistic networks [32].

VII. Conclusion and Outlook

We proposed a novel analytical framework for opportunistic offloading capturing the differences between leechers and seeders. Following this approach, mobile operators are able to finely control the dissemination evolution through external controls such as cellular injections and promotions. We have shown the existence of solutions for the model. We applied Pontryagin's Maximum Principle to devise an optimal offloading strategy that minimizes the distribution costs. One of the main strengths of the model is that every parameter can be easily tuned. We analyzed the sensitivity to different values of contact rate and delay-tolerance, and evaluated the advantages of the proposed model over a simple two-state model. We provided evidence that when we have enough time flexibility, introducing a separation between leechers and seeders is strongly beneficial for the cellular operator.

Besides evaluating the optimal strategy in case of time-varying parameters, we believe that future developments can leverage the techniques presented in this paper to handle a more general case of stochastic diffusion processes following a Markov decision model. In this case, the evolution of the diffusion evolves following stochastic values, and the applied control depends on the observation of the system. The epidemic diffusion model can be also extended taking into account seeders that stop sharing content due to battery or storage constraints. Finally, a birth-death process can be included to represent the arrival and departure of users in the interest area.

REFERENCES

[1] B. Han, P. Hui, V. S. A. Kumar, M. V. Marathe, J. Shao, and A. Srinivasan, "Mobile data offloading through opportunistic communications and social participation," *IEEE Transactions on Mobile Computing*, vol. 11, no. 5, pp. 821–834, May 2012.

[2] J. Whitbeck, Y. Lopez, J. Leguay, V. Conan, and M. D. de Amorim, "Push-and-track: Saving infrastructure bandwidth through opportunistic forwarding," *Pervasive and Mobile Computing*, vol. 8, no. 5, pp. 682–697, Oct. 2012.

[3] S. Andreev, A. Pyattaev, K. Johnsson, O. Galinina, and Y. Koucheryavy, "Cellular traffic offloading onto network-assisted device-to-device connections," *IEEE Communications Magazine*, vol. 52, no. 4, pp. 20–31, April 2014.

[4] T. Ning, Z. Yang, H. Wu, and Z. Han, "Self-interest-driven incentives for ad dissemination in autonomous mobile social networks," in *IEEE INFOCOM*, Turin, Italy, April 2013, pp. 2310–2318.

[5] W. Kermack and A. McKendrick, "Contributions to the mathematical theory of epidemics – I," *Bulletin of mathematical biology*, vol. 53, no. 1, pp. 33–55, 1991.

[6] X. Zhang, G. Neglia, J. Kurose, and D. Towsley, "Performance modeling of epidemic routing," *Computer Networks*, vol. 51, no. 10, pp. 2867–2891, Jul. 2007.

[7] R. Groenevelt, P. Nain, and G. Koole, "The message delay in mobile ad hoc networks," *Performance Evaluation*, vol. 62, no. 1–4, pp. 210 – 228, 2005.

[8] F. Rebecchi, M. Dias de Amorim, V. Conan, A. Passarella, R. Bruno, and M. Conti, "Data offloading techniques in cellular networks: A survey," *IEEE Communications Surveys Tutorials*, vol. 17, no. 2, pp. 580–603, Secondquarter 2015.

[9] F. Rebecchi, M. D. de Amorim, and V. Conan, "DROiD: Adapting to individual mobility pays off in mobile data offloading," in *IFIP Networking*, Trondheim, Norway, Jun. 2014.

[10] L. Valerio, R. Bruno, and A. Passarella, "Adaptive data offloading in opportunistic networks through an actor-critic learning method," in *ACM CHANTS*, Maui, HI, 2014, pp. 31–36.

[11] R. Kwan, C. Leung, and J. Zhang, "Multiuser scheduling on the downlink of an lte cellular system," *Research Letters in Communications*, vol. 2008, p. 3, 2008.

[12] C. Ide, B. Dusza, M. Putzke, C. Müller, and C. Wietfeld, "Influence of m2m communication on the physical resource utilization of lte," in *Wireless Telecommunications Symposium (WTS), 2012*, 2012, pp. 1–6.

[13] S. Zhang, J. Wu, Z. Qian, and S. Lu, "Mobicache: Cellular traffic offloading leveraging cooperative caching in mobile social networks," *Computer Networks*, vol. 83, pp. 184–198, 2015.

[14] W. Fleming and R. Rishel, *Deterministic and stochastic optimal control*. Springer, 1975.

[15] V. G. Boltyanskii, R. V. Gamkrelidze, and L. S. Pontryagin, "The theory of optimal processes. I. The maximum principle," DTIC Document, Tech. Rep., 1960.

[16] D. Grass, J. P. Caulkins, G. Feichtinger, G. Tragler, and D. A. Behrens, *Optimal control of nonlinear processes: with applications in drugs, corruption, and terror*. Springer, 2008.

[17] K. Soetaert, J. Cash, and F. Mazzia, "Package bvpSolve, solving boundary value problems in R," *Journal of Statistical Software*, vol. 33, no. 9, pp. 1–25, 2010.

[18] X. Wang, M. Chen, Z. Han, T. T. Kwon, and Y. Choi, "Content dissemination by pushing and sharing in mobile cellular networks: An analytical study," in *IEEE MASS*, Las Vegas, NV, 2012, pp. 353–361.

[19] X. Zhuo, W. Gao, G. Cao, and S. Hua, "An incentive framework for cellular traffic offloading." *IEEE Transactions on Mobile Computing*, vol. 13, no. 3, pp. 541–555, 2014.

[20] J. Scott, R. Gass, J. Crowcroft, P. Hui, C. Diot, and A. Chaintreau, "CRAWDAD data set cambridge/haggle (v. 2006-01-31)," Downloaded from http://crawdad.org/cambridge/haggle/, Jan. 2006.

[21] K. C.-J. Lin, C.-W. Chen, and C.-F. Chou, "Preference-aware content dissemination in opportunistic mobile social networks," in *IEEE INFOCOM*, Orlando, FL, March 2012, pp. 1960–1968.

[22] C. Boldrini, M. Conti, and A. Passarella, "Contentplace: social-aware data dissemination in opportunistic networks," in *ACM MSWiM*, 2008, pp. 203–210.

[23] Y. Han, H. Wu, Z. Yang, and D. Li, "A new data transmission strategy in mobile d2d networks — deterministic, greedy, or planned opportunistic routing?" *IEEE Transactions on Vehicular Technology*, vol. PP, no. 99, 2016.

[24] M. V. Barbera, A. C. Viana, M. D. de Amorim, and J. Stefa, "Data offloading in social mobile networks through VIP delegation," *Ad Hoc Networks*, 2014.

[25] F. Rebecchi, L. Valerio, R. Bruno, V. Conan, M. D. de Amorim, and A. Passarella, "A joint multicast/D2D learning-based approach to LTE traffic offloading," *Computer Communications*, vol. 72, pp. 26–37, 2015.

[26] K. Leibnitz, T. Hoßfeld, N. Wakamiya, and M. Murata, "Modeling of epidemic diffusion in peer-to-peer file-sharing networks," in *Biologically Inspired Approaches to Advanced Information Technology*. Springer, 2006, pp. 322–329.

[27] B. Fan, D.-M. Chiu, and J. C. Lui, "Stochastic differential equation approach to model bittorrent-like p2p systems," in *IEEE ICC*, vol. 2, Istanbul, Turkey, 2006, pp. 915–920.

[28] H. Behncke, "Optimal control of deterministic epidemics," *Optimal control applications and methods*, vol. 21, no. 6, pp. 269–285, 2000.

[29] H. Gaff and E. Schaefer, "Optimal control applied to vaccination and treatment strategies for various epidemiological models." *Mathematical biosciences and engineering: MBE*, vol. 6, no. 3, pp. 469–492, 2009.

[30] M. Khouzani, S. Sarkar, and E. Altman, "Optimal dissemination of security patches in mobile wireless networks," *IEEE Transactions on Information Theory*, vol. 58, no. 7, pp. 4714–4732, July 2012.

[31] K. Kandhway and J. Kuri, "How to run a campaign: Optimal control of SIS and SIR information epidemics," *Applied Mathematics and Computation*, vol. 231, pp. 79–92, 2014.

[32] M. Khouzani, S. Sarkar, and E. Altman, "Optimal control of epidemic evolution," in *IEEE INFOCOM*, Shanghai, China, Apr. 2011, pp. 1683–1691.

APPENDIX
TWO-STATE MODEL SOLUTION

The evolution of the state equations for the two-state model can be described using a single ODE to obtain the control u_I that manages the cellular injections. Considering $n_I(t) = 1 - n_S(t)$, the minimization problem becomes:

$$\min_{u_I(t) \in U} J, \text{ subject to:} \tag{12a}$$

$$\frac{\partial n_I}{\partial t} = -\lambda(t) n_I(t)(1 - n_I(t)) - u_I(t) n_I(t), \tag{12b}$$

$$0 \le n_I(t) \le 1,\, n_I(t) + n_S(t) = 1,$$

$$n_I(0) = i_0, n_S(0) = 1 - i_0. \tag{12c}$$

We consider the same cost functions defined in Section IV-C. Employing Pontryagin's maximum principle, we find the Hamiltonian H, the adjoint function p_I, and the transversality condition as follows:

$$H = -\, b(n_I u_I)^\alpha - c(\lambda\, n_I\, (1 - n_I)) + p_I[-\lambda n_I (1 - n_I) - u_I n_I], \tag{13a}$$

$$\frac{dp_I^*(t)}{dt} = b\, \alpha\, n_I^{\alpha-1} u_I^\alpha + c\, \lambda\, (1 - 2n_I) - p_i\, [2\lambda n_I - \lambda - u_I]\,, \tag{13b}$$

$$p_I(T) = e^{n_I^*(T)}. \tag{13c}$$

By defining the function $\psi(t) = \sqrt[\alpha-1]{\frac{p_I^*}{-\alpha b}}$, the Hamiltonian maximization condition (Eq. 7) along with the restrictions on the number of simultaneous transmissible copies through the infrastructure gives that$u_I^*(t) = \frac{\min[\max[\psi(t),0], I_{max}]}{n_I(t)}$. The injection control takes a similar form to the leechers-seeders model, with $\psi(t)$ depending, this time, only on $p_I^*(t)$.