The Cost of Being Altruistic: Optimal D2D Offloading under Rewarding Conditions

Filippo Rebecchi, Marcelo Dias de Amorim, Vania Conan

The Cost of Being Altruistic: Optimal D2D Offloading under Rewarding Conditions

Filippo Rebecchi[1,2], Marcelo Dias de Amorim[2], Vania Conan[1]

[1] *Thales Communications & Security, 4 Av. des Louvresses, 92230 Gennevilliers, France*
[2] *LIP6/CNRS – UPMC Sorbonne Universités, 4 Place Jussieu, 75005 Paris, France*

Existing device-to-device (D2D) offloading techniques assume that all nodes storing data are potential forwarders. This leads to suboptimal results whenever the system has to reward forwarders. How to design a global strategy that keeps the number of seed users low (to save cellular bandwidth) and selects the appropriate set of forwarders (to know which ones to reward) remains an open issue. We formulate this question as a stochastic control problem that we solve using an application of Pontryagin's Maximum Principle (PMP). We provide a framework that works under a generic cost model. We show analytically that an optimal solution exists and compute when operators benefit from this policy.

Keywords: Mobile data offloading, optimal control, information epidemics, Pontryagin's Maximum Principle.

1 Introduction

Device-to-device (D2D) communications are a well-timed strategy for operators to face the ever-increasing demand for mobile data by *offloading* part of the traffic from their cellular infrastructure. Motivated by the delay-tolerance of some types of content, operators may send data to a subset of users (seeders) and let them propagate it by means of opportunistic D2D transmissions.

Existing proposals in the literature assume that all seeders are, by default, also opportunistic forwarders [HHK+12, RdAC14]. Such an assumption leads to suboptimal results when forwarders must be rewarded for helping relieve the load on the cellular channel. The problem is that, if performed in a uncontrolled fashion, opportunistic diffusion may generate additional costs without necessarily bringing gains to data dissemination. Moreover, a well-planned rewarding strategy is essential to encourage mobile users to participate in the offloading process.

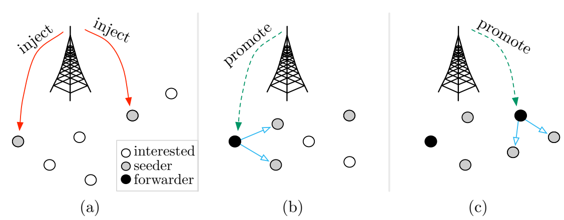

Figure 1: The infrastructure injects two copies of the content (Fig. 1(a)) and decides that one of the nodes should be promoted as forwarder (Fig. 1(b)). Later, it decides to promote another seeder as opportunistic forwarding seems not to guarantee sufficient dissemination (Fig. 1(c)).

The balance between instantaneous cost and future benefits of the *injection* and *forwarding* decisions is strategic to the dissemination process, given that available resources (bandwidth and rewards) are limited. Fig. 1 illustrates the offloading process with a central coordinator (at the infrastructure side) that controls the cellular injections and the promotion of users to the forwarding state. We investigate the following problem: *which fraction of seeders should be promoted as forwarders and when should this happen*? We translate the possible decisions operators can take (injection and forwarding) into a cost function. We apply Pontryagin's Maximum Principle (PMP) [BGP60] to minimize the cost function subject to the state-equations governing the network evolution. Numerical results provide us with insights into the interactions between seeders, forwarders, and the evolution of data dissemination. We also reveal that, under rewarding conditions, the forwarding decision is just as critical as the choice of seeders.

2 System description

We model the evolution of the content dissemination using a variant of the classic SIR model.[†] Some users request data and are referred to as *interested*. Initially, all the nodes are in the *interested* state. At this stage, the operator can only use cellular transmissions to reach these users. Nodes that receive the content through the cellular channel enter the *seeder* state, (not yet playing any active role in data distribution). At this point, the coordinator can promote a fraction of them to the *forwarder* state to trigger the opportunistic diffusion of the content.

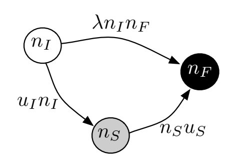

Figure 2: State transition rates.

Network model. The system consists of N mobile nodes and one content to be distributed by the infrastructure to all the nodes within the lifetime T. Following the notation introduced above, nodes can be in the *interested*, *seeder*, or *forwarder* states. Their respective fractions are $n_I(t)$, $n_S(t)$, $n_F(t)$.

Communication opportunities. We use a *mean field* model that is accurate for a large population. As with a disease contagion in a population, content spreads from *forwarder* to *interested* nodes when in physical proximity. State evolution can be described by a system of ODEs and a set of initial constraints. The contact rate $\lambda(t)$ rules the encounter of any two nodes. We assume, for the purposes of this paper, that encounters are homogeneous. As shown in Fig. 2, interested nodes become forwarders with rate $\lambda(t)n_I(t)n_F(t)$.

Injections and Promotions. The central offloading coordinator manages the cellular injections and the promotion of seeders to the forwarder state. Injections increase the rate at which nodes switch from the interested to the seeder state. The intensity of injections is denoted by $u_I(t)$. Consequently, $u_I(t)n_I(t) \leq I_{max}(t)$ describes the rate of injected copies. The injection rate is bounded by $I_{max}(t)$, which measures the maximum available load on the cellular network. Seeders carry the content but have to be promoted in order to contribute to data dissemination. As a result, nodes shift to the forwarder state with intensity $u_S(t)$. This increases the fraction of nodes promoted to the forwarder state by a rate $u_S(t)n_S(t)$.

Cost. We consider a general cost function J as defined in Eq. 1:

$$J(T) = \underbrace{\Phi[n_I(T)]}_{\text{payoff}} + \int_0^T \underbrace{f[u_I(t)\,n_I(t)]}_{\text{injection}} + \underbrace{g[\lambda n_I(t)\,n_F(t)]}_{\text{reward}}\,dt, \tag{1}$$

where $\Phi[n_I(T)]$ is the final cost incurred by the operator for not having satisfied the fraction $n_I(T)$ of users by the deadline T. This can be seen as the loss of earnings due to missed deliveries, or the extra costs paid due to final injections [RdAC14]. $f[u_I(t)\,n_I(t)]$ captures the instantaneous cost in terms of network resources for injections over the cellular channel. Forwarders are rewarded with $g[\lambda n_I(t)\,n_F(t)]$, which represents reductions or virtual credits accorded to users each time they make an opportunistic transmission. The integral portrays the growing cost over time of these two latter terms.

3 Formulation

We express the optimal control problem considering only two state variables n_I and n_S. This is possible because we always have $n_F(t) = 1 - n_I(t) - n_S(t)$. The system is controlled by the tuple $< u_I, u_S >$ belonging to the set of all the admissible controls $U = \{u_I, u_S\}$, where $u_I, u_S \in [0,1]$ are Lebesgue integrable. The idea is to minimize the cost function J subject to a series of state evolution constraints:

$$\min_{u_I(t),\, u_S(t)\, \in U} J, \text{ subject to: } \begin{cases} \frac{\partial n_I}{\partial t} = -\lambda(t)n_I(t)(1 - n_I(t) - n_S(t)) - u_I(t)n_I(t), \\ \frac{\partial n_S}{\partial t} = u_I(t)n_I(t) - u_S(t)n_S(t), \end{cases} \tag{2}$$

where $n_I(t) \geq 0$, $n_S(t) \geq 0$, $n_I(t) + n_S(t) + n_F(t) = 1$, and $n_I(0) = n_f(0) = 0$, $n_S(0) = 1$.

[†] We adopted a slight variation of the traditional nomenclature of the SIR model: "susceptible" users in the original SIR model are analogous to *interested* of our model. Similarly, "infective" and "recovered" nodes are named *forwarders* and *seeders*, respectively.

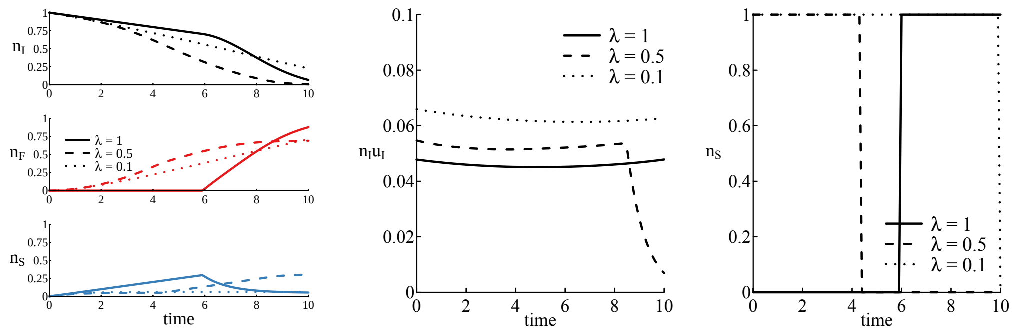

(a) state evolution (b) injections (c) promotions

Figure 3: Optimal offloading evaluated for different contact rates λ. $T = 10s$. $I_{max} = 0.1$, $\alpha = 2$, $b = 10$, $c = 1$.

The existence of an optimal solution can be proved using the Filippov-Cesari theorems [FR75]. We apply PMP to solve the above problem and derive the optimal control (Theorem 3.4 in [GCF$^+$08]). Let the tuple $< n_I^*(.), n_S^*(.), u_I^*(.), u_S^*(.) >$ be an optimal solution of Eq. 2.[‡] Then, there exist continuous and piecewise continuously differentiable adjoint functions $p_i^*(t)$ and $p_s^*(t)$ that maximize the *Hamiltonian* function:

$$H(n_{I,S}, u_{I,S}, p_{i,s}, t) = -f[u_I n_I] - g[\lambda n_I(1 - n_I - n_S)] + p_i[-\lambda n_I(1 - n_I - n_S) - u_I n_I] + p_s[u_I n_I - u_S n_S]. \quad (3)$$

The optimal adjoint equations are $\dot{p}_i^*(t) = -\frac{\partial H(.)}{\partial n_I}\Big|_{n^*_{I,S}, u^*_{I,S}, p^*_{i,s}}$ and $\dot{p}_s^*(t) = -\frac{\partial H(.)}{\partial n_S}\Big|_{n^*_{I,S}, u^*_{I,S}, p^*_{i,s}}$.

In the following, we consider an exponential function for the final payoff $\Phi(x) = e^x - 1$, a power-law function for the direct injections $f(x) = bx^\alpha$ ($\alpha \geq 2$), and a linear function $g(x) = cx$ to reward forwarders. Due to the limited space available, we skip the derivation of the generic framework, which follows a somewhat standard pattern, and we provide the solution directly.

Injections. Given that $f(x)$ is strictly convex, we can 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:

$$u_I^*(t) = \frac{\min[\max[\psi(t), 0], I_{max}]}{n_I(t)}, \qquad \psi(t) = \sqrt[\alpha-1]{\frac{p_i(t)^* - p_s(t)^*}{-\alpha\, b}}. \quad (4)$$

Promotions. Since Eq. 3 is linear in the control variable u_S, the maximization condition is trivially satisfied and independent of u_S. The control in this case is called *singular* (Definition 3.40 in [GCF$^+$08]) with a *bang-bang* solution, i.e., a control that switches discontinuously between one extreme to the other. We define the switching function $\sigma = (p_s n_S)$. By construction $u_S \in [0, 1]$, then it follows that $u_S^*(t) = 1(-\sigma)$.

4 Numerical results

To solve the system of coupled differential equations, we adopt the shooting method from the **R** package *bvpSolve* to compute the evolution of the state variables as well as the optimal control [SCM10]. Fig. 3 gives an example of the state and control variables for different values of the contact rate λ. In general, injections are stronger at the beginning and at the end of the dissemination period (Fig. 3b).

Promotions (Fig. 3c) display three different patterns. For $\lambda = 0.1$, the control is always at its maximum. In low contact-rate scenarios, considering a separate forwarder state brings no improvements. In the two latter cases, instead, the optimal strategy does not contemplate an indiscriminate transition toward the forwarder state. For instance, $\lambda = 0.5$ presents an on-off behavior, with promotions that stop only when the amount of forwarders reaches significant levels. Finally, when $\lambda = 1$, the promotion is switched on only after half of the dissemination period. Although at first sight this might seem counter-intuitive, we must not forget that, in the model, operators have to reward each D2D transmission performed by users. We draw the lesson that *under high contact rates, opportunistic dissemination has to be limited in order to save monetary resources.*

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

Similarly, we investigate under which values of T it is worth considering a separation between seeders and forwarders. We compare our model to a classic two-state model, where all the seeders are also forwarders. Fig. 4 shows the evolution of the cost function J divided by its three main components $\Phi(T)$, $F(T)$, and $G(T)$.[§] With shorter deadlines, when nodes have few contact opportunities, the two-state model benefits from a small advantage in terms of cost, as J is dominated by the final payoff $\Phi(T)$. In this scenario, there is no need to consider an additional state – seeder in our model – because it tends to slow down content diffusion (nodes have to transit from the seeder state before being promoted to the forwarder state). Conversely, for longer deadlines, the three-state model significantly improves the cost functional J. For $T > 5$, the number of uninfected nodes at the deadline decreases, reducing the weight of $\Phi(T)$ on the overall cost. The largest part of J is due to the reward of opportunistic forwarders (portrayed by $G(T)$). The cost for rewarding users increases linearly for the two-state model as the deadline increases. An uncontrolled number of forwarders interferes with the will of operators to cut operational costs. Instead, *a separation between seeders and forwarders offers improved flexibility in the control of the offloading evolution, allowing the implementation of cost-savvy strategies*.

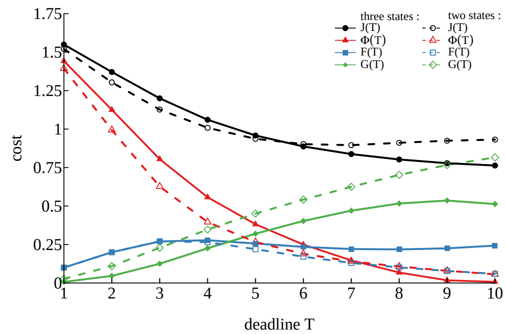

Figure 4: Cost functional J for the optimal strategy using two offloading models (seeder-forwarder and two-state), varying the deadline T, $\lambda = 0.5$, $I_{max} = 0.1$, $\alpha = 2$, $b = 10$, $c = 1$.

5 Conclusion and outlook

We proposed a novel analytical framework for opportunistic offloading that captures the differences between seeders and forwarders. Mobile operators can finely control the dissemination evolution through infrastructure injections and forwarders' promotion. After formalizing the diffusion and the cost model, we applied the Pontryagin's Maximum Principle to devise the optimal strategy that minimizes the aggregate cost for the operator. The solution is then evaluated numerically for a sample cost-function, and compared with state-of-the-art offloading model. Future developments will consider a more general case of stochastic diffusion processes following a Markov decision model. The dissemination model can also be extended by taking into account forwarders that stop sharing content due to battery or storage constraints.

References

[BGP60] V. G. Boltyanskii, R. V. Gamkrelidze, and L. S. Pontryagin. The theory of optimal processes. I. The maximum principle. Technical report, DTIC Document, 1960.

[FR75] W. Fleming and R. Rishel. *Deterministic and stochastic optimal control*. Springer, 1975.

[GCF+08] 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.

[HHK+12] 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 Trans. on Mobile Computing*, 11(5):821–834, May 2012.

[RdAC14] F. Rebecchi, M. Dias de Amorim, and V. Conan. DROiD: Adapting to individual mobility pays off in mobile data offloading. In *IFIP Networking*, Trondheim, Norway, June 2014.

[SCM10] K. Soetaert, J. Cash, and F. Mazzia. Package bvpSolve, solving boundary value problems in R. *Journal of Statistical Software*, 33(9):1–25, 2010.

[§] In this case, $F(T)$ and $G(T)$ represents the total cost over the distribution period T due to injection and rewards respectively.