Optimal duty cycling in mobile opportunistic networks with end-to-end delay guarantees

Elisabetta Biondi, Chiara Boldrini, Andrea Passarella, and Marco Conti
IIT-CNR, Via G. Moruzzi 1, 56124 Pisa, Italy
Email: first.last@iit.cnr.it

***Abstract*—Opportunistic communications have been recently proposed as a key strategy for offloading traffic from 3G/4G cellular networks, which is particularly beneficial in case of crowded areas where many users are interested in similar contents. To conserve energy, duty cycling schemes are typically applied, and therefore contacts between nodes may become intermittent and sporadic also in dense networks. It is thus of paramount importance to accurately tune the duty cycling policy in order to meet energy requirements without compromising the quality of communications. In this paper, building upon a model of duty cycling in opportunistic networks that we have validated in a previous work, we study how to optimise the value of the duty cycle in order to provide probabilistic guarantees on the delay experienced by messages. More specifically, for a broad range of end-to-end delay distributions, we provide closed-form approximated solutions for deriving the optimal duty cycle such that the probability that the delay is smaller than a target value z is greater than or equal to a configurable probability p.**

I. Introduction

Opportunistic networks have been conceived at the intersection between Mobile Ad hoc NETworks (MANET) and the Delay Tolerant (DTN) paradigm. In the conventional model, they exploit the movements of the nodes of the network (people with their smart, handheld devices like tablets and smartphones) in order to deliver messages to their destinations according to the store-carry-and-forward paradigm: nodes hold messages while they move and forward them to other nodes that are in radio contact, until messages reach their final destination. Opportunistic communications were initially seen as a standalone solution for those scenarios in which the nodes of the network were sparse and the infrastructure unavailable (disaster/emergency scenarios, developing countries, etc.). Recently, however, they have become one of the key strategies for mobile data offloading [1], whose main goal is to offload the traffic from cellular networks to other types of networks (e.g., WiFi infrastructured or MANET) in a synergic way, in order to address the overloading of the 3G/4G infrastructure.

In case of crowded environments (and thus dense networks) overloading may be even more critical, and opportunistic networking techniques can be usefully applied, as follows. Due to the typical Zipf-like shape of content interest, it is likely that large fractions of users in the crowd are interested in few, very popular contents (e.g., those mostly related to the area where the crowd gathers). Multicast can be a solution to reduce the traffic load only when content requests can be synchronised. When requests are generated dynamically by users, exploiting communications between users' devices is a more flexible solution, as content can be sent through the cellular network only to a few of them, exploiting opportunistic communications for the rest. The D2D technology addresses this goal to some extent, and is currently proposed in latest LTE releases. In this paper we focus on offloading through ad-hoc WiFi or Bluetooth technologies, as this approach permits to exploit additional portions of the spectrum (and, therefore, additional bandwidth) with respect to that allocated to cellular networks. A possible roadblock in this scenario is the fact that direct communications consume significant energy. To address this, nodes are typically operated in duty cycling mode, by letting their WiFi (or Bluetooth) interfaces ON only for a fraction of time. The joint effect of duty cycling and mobility is that, even if the network is dense, the resulting patterns in terms of communication opportunities is similar to that of conventional opportunistic networks, as devices are able to directly communicate with each other only when they come in one-hop radio range *and* both interfaces are ON.

The net effect of implementing a duty cycling scheme is thus the fact that some contacts between nodes are missed because the nodes are in power saving mode. Hence, detected intercontact times, defined as the time between two consecutive contact events during which a communication can take place for a pair of nodes, are longer than intercontact times determined only by mobility, when a duty-cycling policy is in place. This heavily affects the delay experienced by messages, since the main contribution to message delay is in fact due to the intercontact times. In our previous work [2], we have focused on exponentially distributed intercontact times and we have studied how these are modified by duty cycling, obtaining that intercontact times remain exponentially distributed but their rate is scaled by the inverse of the duty cycle (see Proposition 1, Section III). Building upon this result, we have then investigated how the first moments of the end to end delay vary with the duty cycle for a number of opportunistic forwarding schemes. In addition, we have found that energy saving and end-to-end delay both scale linearly with the duty cycling. Therefore, for a single message delivery, the same energy saved through duty cycling is spent because the network must stay alive longer. Thus, the main advantage of duty cycling is enabling the network to carry more messages by being alive longer (rather than improving the energy spent for each single delivery).

Our work in [2] assumed that the value of the duty cycle was given and studied its effects on important performance metrics such as the delay, the network lifetime, and the number of messages successfully delivered to their destination. More in general, the duty cycling can be seen as a parameter that can be configured, typically, based on some target performance metrics. To this aim, the main contribution of this paper is a mathematical model that allows us to tune the duty cycle in order to meet a given target performance, expressed as a probabilistic guarantee (denoted as p) on the delay experienced by messages. Considering probabilistic, instead of hard, guarantees, allows us to cover a very broad range of application scenarios also beyond best-effort cases – all but those requiring real-time streaming. Specifically, we study the case of exponential, hyper-exponential and hypo-exponential delays (please recall that any distribution falls into one of these three cases, at least approximately [3]), deriving the optimal duty cycle for each of them. For the simple case of exponential delays we are able to provide an exact solution. For the other two cases, we derive an approximated solution and the conditions under which this approximation introduces a small fixed error ε (which is always below 0.14) on the target probability p. Specifically, in the worst case, the approximated duty cycle introduces an error on the target probability p of about 0.1 (hyper-exponential case) and 0.14 (hypo-exponential case), while in the other cases the error is well below these thresholds.

The paper is organised as follows. In Section II we overview the literature on duty cycle optimisation for opportunistic networks. After having introduced the network and duty cycle model that we consider in this work (Section III) we derive in Section IV the optimal duty cycles for the case of exponential, hyper-exponential, and hypo-exponential delays. Then, in Section V, given a target performance for the delay, we discuss how the optimal duty cycle affects the volume of messages delivered during the network lifetime and we highlight that in the case of hyper-exponential delays it is possible to achieve a lower duty cycle than hypo-exponential delays for a given target performance. Finally, Section VI concludes the paper.

II. Related work

There are not many contributions in the DTN literature studying the optimisation of the duty cycling policy. In [4], using a fixed duty cycle scheme, Wang et al. study the relationship between the probability of missing a contact and the associated energy consumption (considered inversely proportional to the contact probing interval). Building upon these results, [4] provides some heuristic algorithms to achieve an optimal contact probing. Differently from this work, in this paper we mathematically define the optimisation problem and we provide an analytical, closed form, result.

In [5], Gao and Li focus on the design of an adaptive duty cycle that minimises wakeups during intercontact times (which are useless, from a contact probing standpoint). Differently from [5], we have chosen to optimise the duty cycle directly, based on the performance goal that we want to achieve. While it is true that an optimisation based on intercontact times impacts directly on the delay performance, it is not straightforward how to control the one based on the other. With our model, instead, we can directly go from the requirements in term of probability of staying below a fixed delay threshold to the corresponding duty cycle value. In addition, differently from [5], we focus on a fixed duty cycle, similar to [6] [7] [4]. It is still an open research point which duty cycling strategy is to be preferred. However, preliminary results in [4] show that, under some assumptions, fixed duty cycle is the optimal strategy.

Another contribution focused on duty cycle optimisation is [8], in which Altman and Azad study the optimisation of node activation in DTN relying on a fluid approximation of the system dynamics. However, the problem analysed is different from the one studied in this paper, since in [8] nodes, once activated, remains active. In addition, this model is based on the assumption of i.i.d. intercontact times, while it has been shown that realistic intercontact times are intrinsically heterogeneous. For this reason, here we focus on heterogenous (but still independent) intercontact times.

III. Preliminaries

We assume that user mobile devices alternate between ON and OFF states, whose duration is fixed. We denote as duty cycle Δ the ratio between the duration of the ON and OFF states, and as T their sum. We assume that when a node is in the ON state it is able to detect contacts with other nodes. Please refer to [2] for a discussion on how to apply this model to popular technologies such as Bluetooth and WiFi Direct. For the sake of simplicity, coarse synchronisation (e.g., controlled by the cellular infrastructure in the case of mobile data offloading) can be used to guarantee that ON intervals overlap between any pair of nodes, such that they can communicate during a contact if this overlaps with their ON phases. Under this assumption, in [2] we have investigated the effect of duty cycling on the detection of encounters between pairs of nodes. As discussed in Section I, this problem is extremely relevant to opportunistic networks, in which messages are delivered by means of consecutive exchanges between encountering nodes. In fact, the net effect of a duty cycling policy is to reduce the number of contacts that can be exploited for exchanging messages. More specifically, we have shown that, when intermeeting times follow an exponential distribution[1], the contact rate between a tagged node pair is approximately decreased by a factor Δ. We summarise this result below.

Proposition 1: Considering a tagged pair of nodes i and j with exponential intercontact time of rate λ_{ij}, the detected intercontact time, i.e., the effective intercontact time when a duty cycling policy is in place, features approximately an exponential distribution with rate $\Delta\lambda_{ij}$, as long as $\lambda_{ij}T \ll 1$, where T is the duty cycling period.

In [2] we have shown that the condition $\lambda_{ij}T \ll 1$ holds for the majority of contact traces available in the literature. Please note also that the above result has been obtained assuming that the duration of a contact is negligible with respect to the duration of the OFF period, which is reasonable (for example, results in [11] show that in absence of duty cycling the median contact duration is below 48s, while the period of typical duty cycling policies is in the order of several minutes).

Exploiting the result in Proposition 1, in our previous work [2] we have evaluated how intercontact times modified by the duty cycling policy affect the first two moments of the pairwise end-to-end delay for a set of representative (both social-oblivious and social-aware) opportunistic forwarding

[1]Exponential intercontact times are a popular assumption in the related literature [9] [10], even if a general consensus on the best probability distribution to approximate the realistic intercontact process has not been reached yet.

strategies. Specifically, we have derived the following properties, which we will use extensively throughout the paper:

P1 The dependence of the coefficient of variation c of the delay from Δ is negligible.

P2 The expected delay when a duty cycling policy is in place (denoted as $E[D_\Delta]$) is approximately equal to the expected delay $E[D]$ with no duty cycle scaled by a factor $\frac{1}{\Delta}$, i.e, $E[D_\Delta] = \frac{E[D]}{\Delta}$.

P3 The second moment of the delay when a duty cycling policy is in place (denoted as $E[D^2_\Delta]$) is approximately equal to the second moment of the delay $E[D^2]$ with no duty cycle scaled by a factor $\frac{1}{\Delta^2}$, i.e, $E[D^2_\Delta] = \frac{E[D^2]}{\Delta^2}$.

IV. Setting the duty cycle for achieving a probabilistic guarantee on the delay

In this section we discuss how to derive the optimal duty cycle Δ_{opt} such that the delay of a tagged message remains, with a certain probability p, under a target fixed threshold z or, in mathematical notation, $\Delta_{opt} = \min\{\Delta : P\{D_\Delta < z\} \geq p\}$. Since the delay increases with Δ, the latter is equivalent to finding the solution to the following[2]:

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

Please note that in the following we will denote the CDF of D_Δ as $F_\Delta(x)$. In order to find the solution to Equation 1, the distribution of the delay D_Δ should be known. Although it is in general unfeasible to obtain an exact closed form for the distribution of D_Δ (except for some trivial cases, such as when the source node can only deliver the message to the destination directly), it is often possible to compute its moments, either exactly or approximately, under different distributions for intercontact times, as shown, e.g., in [9][12]. When the first two moments of the delay can be derived, it is possible to approximate its distribution with either a hypo-exponential or hyper-exponential random variable, using the moment matching approximation technique [3]. So, assuming that we have derived the first moment $E[D_\Delta]$ and the second moment $E[D^2_\Delta]$ of the delay using, e.g., the models in [9] [12], exploiting property P1, we can compute the coefficient of variation c as $\sqrt{\frac{E[D^2]}{E[D]^2} - 1}$. Then, when c is greater than one, D_Δ can be approximated using a 2-stages hyper-exponential distribution with the same moments of D_Δ, as stated in the following Lemma.

Lemma 1 (Hyper-exponential approximation): The two moments matching approximation of D_Δ with coefficient of variation $c \geq 1$ is a 2-stages hyper-exponential distribution with parameters $(\lambda_1, p_1), (\lambda_2, p_2)$ given by the following:

$$\begin{cases} p_1 = \frac{1}{2}\left(1 + \sqrt{\frac{c^2-1}{c^2+1}}\right) \\ \lambda_1 = \frac{2p_1}{E[D_\Delta]} \end{cases} \qquad \begin{cases} p_2 = 1 - p_1 \\ \lambda_2 = \frac{2p_2}{E[D_\Delta]} \end{cases} \quad (2)$$

Vice versa, when the coefficient of variation of the delay is smaller than 1 (but greater than $\frac{1}{\sqrt{2}}$ [13]), D_Δ can be approximated with an hypo-exponential distribution with CDF $F_X(x) = 1 - \frac{\mu_2}{\mu_2-\mu_1}e^{-\mu_1 x} + \frac{\mu_1}{\mu_2-\mu_1}e^{-\mu_2 x}$, for all $x \geq 0$, according to the following lemma.

[2]In the rest of the paper, for convenience of notation, we will drop subscript *opt* since all Δ we derive are the optimal ones.

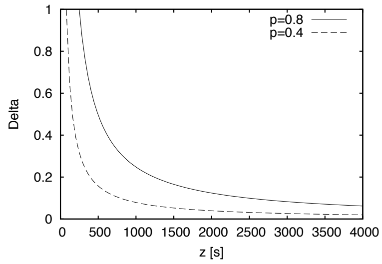

Fig. 1. Δ optimum for the exponential delay.

Lemma 2 (Hypo-exponential approximation): The two moments matching approximation of D_Δ with a coefficient of variation $c \in (\frac{1}{\sqrt{2}}, 1)$ is an hypo-exponential distribution with rates μ_1, μ_2 given by the following:

$$\begin{cases} \mu_1 = \frac{2}{E[D_\Delta]} \cdot \frac{1}{1+\sqrt{1+2(c^2-1)}} \\ \mu_2 = \frac{2}{E[D_\Delta]} \cdot \frac{1}{1-\sqrt{1+2(c^2-1)}} \end{cases} \quad (3)$$

In the rest of the section, we will analyse the optimisation problem in Equation 1 assuming that D_Δ features an exponential (Section IV-A), hyper-exponential (Section IV-B) or hypo-exponential distribution (Section IV-C). Please note that all three cases are possible starting from exponential intercontact times.

A. The exponential case

The simplest case is when the delay features a coefficient of variation c equal to one. In this hypothesis, the distribution of the delay is exponential with parameter $\lambda_\Delta = E[D_\Delta]^{-1}$. Then, it is straightforward to derive Theorem 1.

Theorem 1: The optimal duty cycle when D_Δ features an exponential distribution is given by the following:

$$\Delta = -\frac{\log(1-p)}{\lambda z}, \quad (4)$$

where we indicate with λ the parameter of the exponential distribution obtained with $\Delta = 1$, i.e., $\lambda = E[D]^{-1}$.

Proof: We know that $\lambda_\Delta = \frac{1}{E[D_\Delta]}$, hence, since $E[D_\Delta] \sim \frac{E[D]}{\Delta}$ (Property P2), we have that $\lambda_\Delta = \lambda\Delta$. Thus, we can rewrite Equation 1 as $1 - e^{-\lambda\Delta z} = p$, from which Δ can be easily obtained. ■

For the sake of example, in Figure 1 we plot Δ obtained from Theorem 1 setting $p = 0.8$. $E[D]$ is set to $154s$, which is the average expected delay obtained in [2] for a simple social-aware policy that selects the next relay of a message based on its contact rate with the destination and assuming the average contact rate equal to $4.07 \cdot 10^{-3}s^{-1}$ (the average contact rate measured in the RollerNet contact dataset [14]). Figure 1 shows that, as expected, when the target delay threshold is too small, it is impossible to achieve it with a probabilistic guarantee p, regardless of the value of the duty cycle. Instead, starting from $z = -\frac{\log(1-p)}{\lambda}$, Δ is inversely proportional to z.

Varying the parameters z and p, Equation 1 describes a surface in $\mathbb{R}^3$, and more precisely the surface K given by:

$$K = \{(z, p, \Delta) \in \mathbb{R} \times [0,1] \times [0,1] : P\{D_\Delta < z\} = p\}. \quad (5)$$

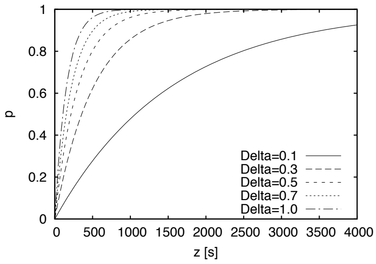

Fig. 2. Level set K_Δ for different values of Δ for the exponential delay

Given a certain duty cycle $\Delta \in (0,1]$, we can thus describe K as the union of its level sets K_Δ or, in other terms, $K = \bigcup_{\Delta\in(0,1]} K_\Delta$ where:

$$K_\Delta = \{(z,p) \in \mathbb{R} \times [0,1] : P\{D_\Delta < z\} = p\}\,. \qquad (6)$$

K_Δ is thus the set of pairs (z,p) that can be obtained with a given duty cycling Δ. It can be useful to plot K_Δ for different Δ in order to study whether it is possible to slightly compromise on the target performance in order to achieve a lower duty cycle. Assuming that we want $z = 250s$, in the exponential case (Figure 2) we can achieve it with a probability 0.8 with $\Delta = 1$ or with 0.68 with $\Delta = 0.7$, thus saving battery lifetime. Similarly, if we want to guarantee a target probability $p = 0.8$, with $\Delta = 1$ we obtain approximately $z = 250s$. If we are more flexible in terms of z, we can choose level set $K_{0.7}$ which gives $z = 350s$. This kind of analysis can be performed also for the hyper-exponential and hypo-exponential delays, with similar results.

B. The hyper-exponential case

When the coefficient of variation of the delay is greater than one, the delay can be approximated with an hyper-exponential distribution as stated in Lemma 1. This means that Equation 1 becomes $1 - p_1 e^{-\lambda_1 z} - p_2 e^{-\lambda_2 z} = p$, where parameters $(\lambda_1, p_1), (\lambda_2, p_2)$ are given by Equation 2. From Equation 2, λ_1 and λ_2 depend on Δ (while p_1 and p_2 do not), thus, denoting with λ_1^0 and λ_2^0 the rates when $\Delta = 1$ and exploiting property P2, we can write Equation 1 as follows:

$$1 - p_1 e^{-\lambda_1^0 \Delta z} - p_2 e^{-\lambda_2^0 \Delta z} = p. \qquad (7)$$

The exact solution Δ to this equation cannot be found analytically because Equation 7 cannot be inverted. However, in Theorem 2 below, we show how to obtain an approximated solution Δ_a that introduces a small error at most equal to ε.

Theorem 2: Let us λ^0 denote $E[D]^{-1}$ and λ_1^0, λ_2^0 the rates of the hyper-exponential delay (Equation 2) for $\Delta = 1$. When delay D_Δ has coefficient of variation greater than one, given a threshold z of the delay and a target probability p, for every fixed $\varepsilon \geq \min\{\varepsilon_1, \varepsilon_2\}$ (whose definition is provided in the proof below), the duty cycle defined by:

$$\Delta_a = \begin{cases} \frac{1}{z}\left[-\frac{1-p-p_2}{\lambda_2^0 p_2} + \right. \\ \left. + \frac{1}{\lambda_1^0} W\left(\frac{p_1^2}{p_2^2} e^{\frac{\lambda_1^0(1-p-p_2)}{\lambda_2^0 p_2}} \right) \right] & \text{if } \varepsilon_1 < \varepsilon_2 \\ -\frac{\log 1-p}{\lambda^0 z} & \text{if } \varepsilon_1 \geq \varepsilon_2, \end{cases} \qquad (8)$$

where W is the Lambert function[3], verifies that $|F_{\Delta_a}(z) - p| \leq \varepsilon$ and so it is a good approximation of the solution to Equation 7.

Proof: We will provide below an intuitive sketch of the proof whose detailed version can be found in [15]. The idea for finding an approximate solution to Equation 7 is to identify an approximation $\tilde{F}(z)$ that is close to $F_\Delta(z)$ under some conditions. So, we build a function $\tilde{F}$ for which it is possible to solve Equation 7 and for which $\min\{\varepsilon_1, \varepsilon_2\}$ is the error introduced (we will clarify this point below). Specifically, we have identified the following function:

$$\tilde{F}(z) = \begin{cases} 1 - p_1 e^{-\lambda_1^0 \Delta z} - p_2(1 - \lambda_2^0 \Delta z) & \text{if } \varepsilon_1 < \varepsilon_2 \\ 1 - e^{-\lambda^0 \Delta z} & \text{if } \varepsilon_1 \geq \varepsilon_2 \end{cases} \qquad (9)$$

Let us denote with $\tilde{F}_1(z)$ and $\tilde{F}_2(z)$ the two parts of $\tilde{F}(z)$ in the above equation. In $\tilde{F}_1(z)$, we have approximated the third term on the left hand side of Equation 7 using the Taylor expansion, after noting that this term contributes to $F_\Delta(z)$ less and less as the coefficient of variation c increases. Vice versa, the pure exponential behaviour ($\tilde{F}_2(z)$) dominates when c is close to 1. Both $\tilde{F}_1(z)$ and $\tilde{F}_2(z)$ can be solved to find Δ, from which Equation 8 follows.

The quality of these two approximations depends on the desired tolerance to the error that we inevitably introduce when we approximate $F_\Delta(z)$. If we tolerate a large error, either approximation can be chosen. Instead, if we want to achieve the smallest error, depending on the coefficient of variation of D_Δ we might have to prefer the one or the other. In the following we briefly discuss how to identify the minimum error introduced by $\tilde{F}_1(z)$ and $\tilde{F}_2(z)$, which we denote with ε_1 and ε_2 respectively. Let us start with $\tilde{F}_1(z)$. We want to find the region for which $|F_{\Delta_a}(z) - p| \leq \varepsilon$ or, equivalently, $|F_{\Delta_a}(z) - \tilde{F}_1^{(\Delta_a)}(z)| \leq \varepsilon$, where we denote with superscript (Δ_a) the fact that the CDF is computed using the approximated solution for Δ. Solving the above inequality, we find that it holds for all $p < p_{max}$, where p_{max} is a function of c and ε (due to lack of space, we do not report its formula here, please refer to [15] for details). Specifically, p_{max} monotonically increases with ε. So, if we want to derive the minimum error for which inequality $|F_{\Delta_a}(z) - p| \leq \varepsilon$ holds for all p, we have to solve equation $p_{max}(c, \varepsilon) = 1$. We obtain the following:

$$\varepsilon_1 = \frac{(a-1)\left(-(a-1)W\left(\frac{(a+1)^2 e^{\frac{a+1}{a-1}}}{(a-1)^2}\right) + a + 1\right)^2}{4(a+1)^2}, \qquad (10)$$

where again $W(x)$ denotes the Lambert function and a is defined as $\sqrt{1 + 2(c^2 - 1)}$.

Let us now consider $\tilde{F}_2(z)$. We are able to prove that function $|F_{\Delta_a}(z) - p|$ has a maximum in p^*. We derive p^* by finding the p in which the derivative of $|F_{\Delta_a}(z) - p|$ becomes zero. Then, ε_2 can be computed as $\varepsilon_2 = |F_{\Delta_a}(z) - p^*|$, obtaining the following:

$$\varepsilon_2 = \frac{1}{2}(a+1)^{-2/a}\left(a\sqrt{2-a^2}+1\right)^{1/a} \cdot \\ \cdot \left(\frac{(a-1)(a+1)^2}{a\sqrt{2-a^2}+1} - \frac{a\sqrt{2-a^2}+1}{a+1} + 2 \right), \qquad (11)$$

[3] The Lambert function is defined as $W(x)e^{W(x)} = x$, for all $x \geq -\frac{1}{e}$

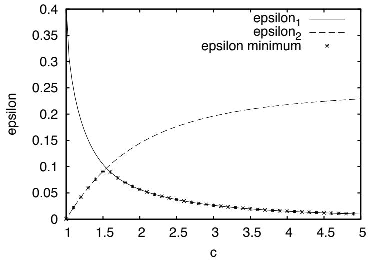

Fig. 3. Error introduced by $\tilde{F}_1(z)$ and $\tilde{F}_2(z)$, varying c.

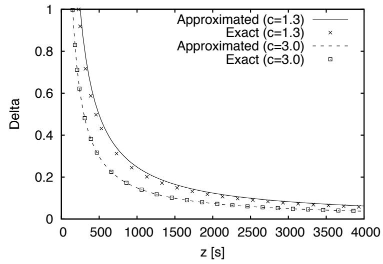

Fig. 4. Δ optimum (approximated vs exact) for the hyper-exponential delay with target probability $p = 0.8$ and varying c.

where again $a = \sqrt{1 + 2(c^2 - 1)}$. Thus, for both ε_1 and ε_2 we have derived a closed-form expression that tells us that the error that we make with our approximation is fixed for a given c. ■

In Figure 3, we show how ε_1 and ε_2 vary with respect to the coefficient of variation c. As expected, for small c (recall that we are in the hyper-exponential case, so $c > 1$ by definition) the exponential assumption $\tilde{F}_2$ allows us to achieve smaller errors. The opposite is true for large c. The worst case is reached for $c \sim 1.5$, when the minimum error is around 0.1, which is still low. In Figure 4 we plot how the optimal duty cycle varies with z, setting the target probability to $p = 0.8$, for two values of coefficient of variation ($c = 1.3$ and $c = 3$). In both cases the approximation is good (the exact value is computed with standard numerical techniques to solve Equation 7). Specifically, when $c = 1.3$ the minimum error that can be achieved is 0.06 and is provided by $\tilde{F}_2(z)$, hence confirming the predominance of the exponential behaviour for c close to 1. Vice versa, when $c = 3$ the minimum error is 0.026 and is provided by $\tilde{F}_1(z)$. It is also interesting to notice that smaller duty cycles can be achieved when c increases, i.e., when the variability of the delay is higher. The importance of this result will be further discussed in Section V.

C. The hypo-exponential case

When the coefficient of variation c of the delay D_Δ is smaller than one, following Lemma 2, it is possible to approximate the delay with a hypo-exponential distribution. In particular, using property P2, if we denote with μ_1^0 and μ_2^0 the parameters obtained when $\Delta = 1$ in Equation 3, we can rewrite Equation 1 making explicit the dependence on Δ:

$$1 - \frac{\mu_2^0}{\mu_2^0 - \mu_1^0} e^{-\mu_1^0 \Delta z} + \frac{\mu_1^0}{\mu_2^0 - \mu_1^0} e^{-\mu_2^0 \Delta z} = p. \qquad (12)$$

As in the hyper-exponential case, this equation can not be directly inverted for finding Δ, but it is possible to derive an approximate solution for which a small fixed (for a given c) error is introduced.

Theorem 3: Let μ_1^0 and μ_2^0 be the parameters given by Equation 3 with $\Delta = 1$. When the delay D_Δ has coefficient of variation smaller than one, the duty cycle defined by:

$$\Delta_a = \begin{cases} -\frac{1}{\mu_1^0 z} \log\left[(1-p) \cdot \frac{\mu_2^0 - \mu_1^0}{\mu_2^0}\right] & \text{if } \varepsilon_1 < \varepsilon_2 \\ -\frac{\log 1-p}{\lambda^0 z} & \text{if } \varepsilon_1 \geq \varepsilon_2, \end{cases} \qquad (13)$$

verifies that $|F_{\Delta_a}(z) - p| \leq \varepsilon$ (with $\varepsilon \geq \min\{\varepsilon_1, \varepsilon_2\}$, see the proof in [15]), and so it is a good approximation of the solution to Equation 12.

Due to lack of space and since the rationale follows that of the proof for Theorem 2, we omit the proof of the above theorem, which can however be found in [15].

In Figure 5 we plot ε_1 and ε_2 varying c. When c is close to one, both approximations are very good. For values of c roughly in the interval $(0.83, 0.97)$, $\tilde{F}_1(z)$ provides better results, while, for low values of c, $\tilde{F}_2(z)$ is to be preferred. In Figures 6(a) and 6(b) we show how the optimal duty cycle varies with z, setting the target probability to $p = 0.8$, for two values of coefficient of variation ($c = 0.75$ and $c = 0.9$, respectively). In both cases the approximation and the exact value are very close. In Figure 6(a) the minimum error that can be achieved is 0.13 and is provided by $\tilde{F}_2(z)$, while in Figure 6(b) the minimum error is 0.05 and is provided by $\tilde{F}_1(z)$.

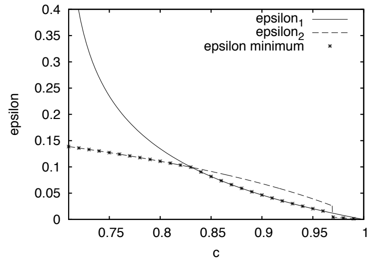

Fig. 5. Error introduced by $\tilde{F}_1(z)$ and $\tilde{F}_2(z)$, varying c.

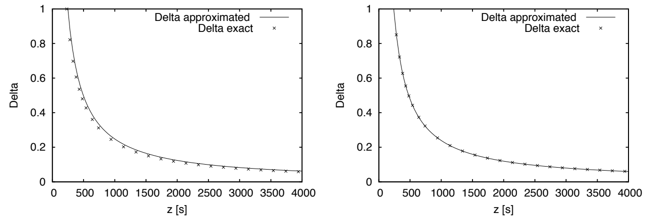

(a) $c = 0.75$ (b) $c = 0.9$

Fig. 6. Δ optimum (approximated vs exact) for the hypo-exponential delay with target probability $p = 0.8$ and varying c.

V. Optimal duty cycle and traffic gain

In this section we investigate how the choice of the optimal duty cycle affects the volume of traffic carried by the network. As already discussed, the advantage of implementing a duty cycle policy is that device batteries are preserved and, as a consequence, the lifetime of the network increases. Specifically, with a duty cycle Δ and a baseline network lifetime L (i.e., with $\Delta = 1$), the network lifetime when a duty cycling

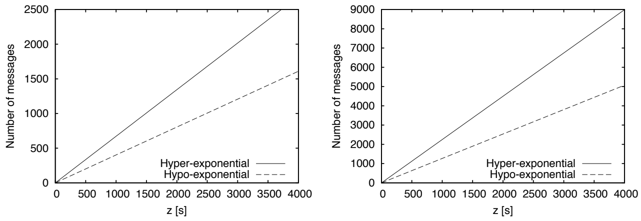

(a) $p = 0.4$

(b) $p = 0.8$

Fig. 7. $\mathcal{N}(\Delta)$ varying z for different p in the case of hyper-exponential ($c = 3$) and hypo-exponential ($c = 0.75$) delay.

policy is in place is given by $\frac{L}{\Delta}$. A longer network lifetime is very useful because it allows nodes to exchange messages for a longer time. If we assume, similarly to [2], that messages are generated according to a Poisson process with rate η, we can derive how the total number of messages $\mathcal{N}$ *delivered* by the nodes varies with Δ. Due to lack of space, in the following we only consider the hypo-exponential and hyper-exponential cases. First, in the theorem below we recall the main results for $\mathcal{N}$ derived in [2].

Theorem 4: If the delay D_Δ has coefficient of variation c greater than one, the volume $\mathcal{N}(\Delta)$ of messages delivered by the system under duty cycling Δ is given by:

$$\mathcal{N}(\Delta)=\frac{\eta L}{\Delta}-\eta E[D_\Delta]\left[1-\frac{1}{2}e^{\frac{-L}{E[D_\Delta]\Delta}}\left(e^{\left(1+\sqrt{\frac{c^2-1}{c^2+1}}\right)}+e^{\left(1-\sqrt{\frac{c^2-1}{c^2+1}}\right)}\right)\right]. \quad (14)$$

Instead, if the delay D_Δ has coefficient of variation c smaller than one, the volume $\mathcal{N}(\Delta)$ of messages delivered by the system under duty cycling Δ is given by:

$$\mathcal{N}(\Delta)=\frac{\eta L}{\Delta}-\eta E[D_\Delta]\cdot \left[1-\frac{1}{4\sqrt{1+2(c^2-1)}}\left(\left(1+\sqrt{1+2(c^2-1)}\right)^2 e^{-\left(\frac{2L}{\Delta E[D_\Delta]\left(1+\sqrt{1+2(c^2-1)}\right)}\right)} -\left(1-\sqrt{1+2(c^2-1)}\right)^2 e^{-\left(\frac{2L}{\Delta E[D_\Delta](1-\sqrt{1+2(c^2-1)})}\right)}\right)\right]. \quad (15)$$

If we substitute in the above equations the optimal Δ derived in the previous section, we obtain how $\mathcal{N}$ varies as a function of the target performance (z, p). In order to study this dependence, we set the network lifetime L to $60000s$ and we assume that each node generates one message every ten minutes ($\eta = \frac{1}{600}s^{-1}$). In Figure 7(a) we set p to the value 0.8 and we plot $\mathcal{N}$ varying z, while in Figure 7(b) we set $p = 0.4$. Besides the expected result that the less stringent the performance requirements (i.e., higher p) the higher the volume of traffic (because smaller duty cycles can be used), we observe an interesting difference between the two delay distributions. The traffic delivered under hyper-exponential delays is always higher than that exchanged under hypo-exponential delays. This is due to the fact that, as we have seen in Section IV-B, when c increases we can achieve smaller optimal duty cycle for a given target performance (z, p), hence saving more energy and increasing the lifetime of the network.

VI. Conclusion

In this work we have studied how to optimise the duty cycle in order to guarantee, with probability p, that the delay of messages remains below a threshold z, assuming that inter-contact times are exponentially distributed. We have provided an exact solution for the case in which the delay follows an exponential distribution, and approximated solutions for the cases in which the coefficient of variation of the delay is greater than or smaller than 1. We have also demonstrated that the approximation of Δ introduces an error ε whose formula we have provided and that is small and fixed for a given coefficient of variation c of the delay. Finally we have focused on the volume of traffic delivered by the network when the optimal duty cycle is implemented, and we have discussed how the two parameters z and p impact on the number of messages delivered. Specifically, we have shown that the optimisation of the duty cycle is more efficient with hyper-exponential delays, as it achieves lower duty cycles and thus provides higher energy gains.

Acknowledgment

This work was partially funded by the European Commission under the EINS (FP7-FIRE 288021), MOTO (FP7 317959), and EIT ICT Labs MOSES (Business Plan 2014) projects.

References

[1] J. Whitbeck, Y. Lopez, J. Leguay, V. Conan, and M. D. De Amorim, "Push-and-track: Saving infrastructure bandwidth through opportunistic forwarding," *Perv. and Mob. Comp.*, vol. 8, no. 5, pp. 682–697, 2012.

[2] E. Biondi, C. Boldrini, A. Passarella, and M. Conti, "Duty cycling in opportunistic networks: intercontact times and energy-delay tradeoff," IIT-CNR 22-2013, Tech. Rep. 22/2013, http://cnd.iit.cnr.it/chiara/pub/techrep/biondi2013duty_tr.pdf.

[3] H. Tijms, *A First Course in Stochastic Models*. Wiley, 2003.

[4] W. Wang, V. Srinivasan, and M. Motani, "Adaptive contact probing mechanisms for delay tolerant applications," in *MobiCom*. ACM, 2007, pp. 230–241.

[5] W. Gao and Q. Li, "Wakeup scheduling for energy-efficient communication in opportunistic mobile networks," in *IEEE INFOCOM*, 2013.

[6] H. Zhou, J. Chen, H. Zhao, W. Gao, and P. Cheng, "On exploiting contact patterns for data forwarding in duty-cycle opportunistic mobile networks," *IEEE Trans. on Vehic. Tech.*, pp. 1–1, 2013.

[7] O. Trullols-Cruces, J. Morillo-Pozo, J. M. Barcelo-Ordinas, and J. Garcia-Vidal, "Power saving trade-offs in delay/disruptive tolerant networks," in *WoWMoM*. IEEE, 2011, pp. 1–9.

[8] E. Altman, A. Azad, T. Başar, and F. De Pellegrini, "Combined optimal control of activation and transmission in delay-tolerant networks," *IEEE/ACM Trans. on Netw.*, vol. 21, no. 2, pp. 482–494, 2013.

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

[10] W. Gao and G. Cao, "User-centric data dissemination in disruption tolerant networks," in *IEEE INFOCOM*, 2011, pp. 3119–3127.

[11] S. Gaito, E. Pagani, and G. P. Rossi, "Strangers help friends to communicate in opportunistic networks," *Computer Networks*, vol. 55, no. 2, pp. 374–385, 2011.

[12] C. Boldrini, M. Conti, and A. Passarella, "Performance modelling of opportunistic forwarding under heterogenous mobility," IIT-CNR, Tech. Rep. TR-12/2013, http://cnd.iit.cnr.it/chiara/pub/techrep/boldrini2013heterogenous_tr.pdf.

[13] G. Bolch, S. Greiner, H. de Meer, and K. Trivedi, *Queueing Networks and Markov Chains: Modeling and Performance Evaluation with Computer Science Applications*. Wiley, 2006.

[14] P.-U. Tournoux, J. Leguay, F. Benbadis, J. Whitbeck, V. Conan, and M. D. de Amorim, "Density-aware routing in highly dynamic dtns: The rollernet case," *IEEE Trans. on Mob. Comp.*, vol. 10, no. 12, pp. 1755–1768, 2011.

[15] E. Biondi, C. Boldrini, A. Passarella, and M. Conti, "Optimal duty cycling in mobile opportunistic networks with end-to-end delay guarantees," IIT-CNR 2014, Tech. Rep., http://cnd.iit.cnr.it/chiara/pub/techrep/biondi2014optimisation_tr.pdf.