

Not All Content is Created Equal: Effect of Popularity and Availability for Content-Centric Opportunistic Networking

Pavlos Sermpezis
Mobile Communications Department
EURECOM, France
pavlos.sermpezis@eurecom.fr

Thrasyvoulos Spyropoulos
Mobile Communications Department
EURECOM, France
thrasyvoulos.spyropoulos@eurecom.fr

ABSTRACT

Mobile users are envisioned to exploit direct communication opportunities between their portable devices, in order to enrich the set of services they can access through cellular or WiFi networks. Sharing contents of common interest or providing access to resources or services between peers can enhance a mobile node's capabilities, offload the cellular network, and disseminate information to nodes without internet access. Interest patterns, i.e. how many nodes are interested in each content or service (popularity), as well as how many users can provide a content or service (availability) impact the performance and feasibility of envisioned applications. In this paper, we establish an analytical framework to study the effects of these factors on the delay and success probability of a content/service access request through opportunistic communication. We also apply our framework to the data offloading problem and provide insights for its optimization.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Store and forward networks, Wireless communication; C.4 [Performance of Systems]: Modeling techniques

Keywords

Performance analysis; Opportunistic networks; Mobile data offloading

1. INTRODUCTION

Opportunistic or Delay Tolerant Networks (DTNs) consist of mobile devices (e.g. smartphones, laptops) that can exchange data using direct communication (e.g. Bluetooth, WiFi Direct) when they are within transmission range. While initially proposed for communication in extreme environments, the proliferation of "smart" mobile devices has led researchers to consider opportunistic networks as a way to support existing infrastructure and/or novel applications, like file sharing [1, 2], crowd sensing [3, 4], collaborative computing [5, 6], offloading of cellular networks [7, 8, 9], etc.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
MobiHoc'14, August 11–14, 2014, Philadelphia, PA, USA.
Copyright 2014 ACM 978-1-4503-2620-9/14/08 ...\$15.00.
<http://dx.doi.org/10.1145/2632951.2632976>.

This trend is also shifting the focus from end-to-end to *content-centric* communications. Some content-centric applications for which opportunistic networking has been considered are: (i) *content sharing* [1, 10, 11]: the source(s) of a "content" (e.g. multimedia file, web page) might want to distribute it (e.g. user generated content) or is willing to share it with other nodes (e.g. content downloaded earlier); (ii) *service or resource access* [5, 6]: nodes offer access to resources (e.g. Internet access) or services (e.g. computing resources); (iii) *mobile data offloading* [7, 8, 9]: the cellular network provider, instead of serving separately each node requesting a given "content" (e.g. a popular video, or software update), distributes a few copies of the "content" in some relay nodes (or *holders*) and they can further forward it to any other node that makes a request for it.

The performance of these mechanisms highly depends on *who* is interested, in *what*, and *where* it can be found (i.e. which other nodes have it). While the effect of node mobility has been extensively considered (e.g. [1, 10, 12]) content popularity has been mainly considered from an algorithmic perspective (e.g. [9, 11]), and in the context of a specific application. Despite the inherent interest of these studies, some questions remain: Would a given allocation policy work well in a different network setting? Are there interest patterns that would make a scheme generally better than others? Key factors like content popularity and content availability might impact the performance or even decide the feasibility of a given application altogether. In this paper, we try to provide some initial insight into these questions, by contributing along the following key directions:

- We propose a simple analytical framework that is applicable to a range of mobility and content popularity patterns seen in real networks; to our best knowledge, this is the first application-independent effort in this direction (Section 2).
- We provide closed form expressions for important metrics that require few statistics about the aggregate node mobility and content popularity; these results facilitate *online* performance prediction and protocol tuning, compared to approaches requiring detailed per node statistics, as e.g. [9] (Section 3).
- While a detailed application-specific optimization is beyond the scope of this paper, we demonstrate how our analysis can be applied to an example application, mobile data offloading, and can help optimize its performance in a generic setting (Section 4).

Finally, we discuss related work in Section 5, and conclude our paper in Section 6.

2. NETWORK MODEL

2.1 Mobility Model

We consider a network \mathcal{N} , where N nodes move in an area, much larger than their transmission range. Data packet exchanges between a pair of nodes can take place only when they are in proximity (*in contact*). Hence, the time points, when the contact events take place, and the nodes involved, determine the dissemination of a message.

We assume that the sequence of the contact events between nodes i and j is given by a random point process with rate λ_{ij} ¹. Analyses of real-world traces suggest that the times between consecutive contacts for a given pair can often be approximated (completely or in the tail) as either exponentially [13, 14] or power-law (e.g. pareto) distributed [15]. Both distributions can be described with a main parameter λ_{ij} (the *contact rate*), and our analysis will be applied to both.

Hence, we can describe the network \mathcal{N} with the contact (or meeting) rates matrix $\Lambda = \{\lambda_{ij}\}$. Depending on the underlying mobility process, there might be large differences between the different λ_{ij} values in this matrix. Furthermore, it is often quite difficult, in a DTN context, to know Λ exactly, or estimates might be rather noisy. For these reasons, we consider the following simple model for Λ :

Assumption 1. *The contact rates λ_{ij} are drawn from an arbitrary distribution with probability density function $f_\lambda(\lambda)$ with known mean μ_λ and variance σ_λ^2 ($CV_\lambda = \frac{\sigma_\lambda}{\mu_\lambda}$).*

By choosing the right function f_λ the above model can capture heterogeneity in the pairwise contact rates, or noise in the estimates. In practice, one would fit the empirical distribution observed in a given measurement trace with an \hat{f}_λ and use it in the analysis.

2.2 Content Traffic Model

We assume that each node might be *interested* in one or more “contents”. A content of interest might refer to (i) a single piece of data (e.g. a multimedia file, a google map) [7], (ii) all messages/data belonging to a category of interests (e.g. local events, financial news) [2, 16], (iii) updates and feeds (e.g. weather forecast, latest news) [17], etc.

A number of content-sharing applications and mechanisms have been proposed in previous literature, from publish-subscribe mechanisms to “channel”-based sharing and device-to-device offloading, etc., (e.g. [2, 3, 4, 17]). To proceed with our analysis we need to setup a simple model of content/service access that can yet capture different (but of course not all) content-centric applications and approaches.

The main notation we use in our model and analysis is summarized in Table 1.

2.2.1 Content Popularity

We assume that when a node is interested in a content or service, it queries other nodes it *directly* encounters for it. We denote the event that a node $i \in \mathcal{N}$ is *interested* in a content \mathcal{M} (or, equivalently, i requests \mathcal{M}) as: $i \rightarrow \mathcal{M}$. We further denote the set of all the contents that nodes are interested in, as: $\mathbf{M} = \{\mathcal{M} : \exists i \in \mathcal{N}, i \rightarrow \mathcal{M}\}$. $|\mathbf{M}| = M$, where $|\cdot|$ denotes the cardinality of a set.

¹We ignore the contact duration and assume infinite bandwidth; assumptions that are common (e.g. [1, 9]) and orthogonal to the problem we consider here.

Definition 1 (Content Popularity). *We define the popularity of a content \mathcal{M} as the number of nodes $N_p^{(\mathcal{M})}$ that are interested in it²:*

$$N_p^{(\mathcal{M})} = |\mathcal{C}_p^{(\mathcal{M})}|, \text{ where } \mathcal{C}_p^{(\mathcal{M})} = \{i \in \mathcal{N} : i \rightarrow \mathcal{M}\} \quad (1)$$

We further denote the percentage of contents with a given popularity value n as

$$P_p(n) = \frac{1}{M} \sum_{\mathcal{M} \in \mathbf{M}} \mathcal{I}_{N_p^{(\mathcal{M})}=n}, \quad n \in [0, M] \quad (2)$$

where $\mathcal{I}_{N_p^{(\mathcal{M})}=n} = 1$ when $N_p^{(\mathcal{M})} = n$ and 0 otherwise.

In other words, $P_p(n)$ defines a probability distribution over the different contents and associated popularities. In practice, it can be chosen according to common practices (e.g. skewed, pareto) [1, 9, 11], or be fitted to real data, if available.

2.2.2 Content Availability

We assume that a request for a content or service is completed, when (and if) a node that holds (a copy of) the requested content is *directly* encountered. We denote the event that a node i holds (a copy of) a content \mathcal{M} as $i \leftarrow \mathcal{M}$, and we define the availability $N_a^{(\mathcal{M})}$ of a content \mathcal{M} as

Definition 2 (Content Availability). *The availability of a content message \mathcal{M} is defined as the number of nodes $N_a^{(\mathcal{M})}$ that hold a copy of it.*

$$N_a^{(\mathcal{M})} = |\mathcal{C}_a^{(\mathcal{M})}|, \text{ where } \mathcal{C}_a^{(\mathcal{M})} = \{i \in \mathcal{N} : i \leftarrow \mathcal{M}\} \quad (3)$$

The availability of a given content might often (although not always) be correlated with the popularity of that content. A cellular network provider might *allocate* more holders for popular contents [9]. In a content-sharing setting, where some nodes might be more willing than others to maintain and share (“seed”) a content after they’ve downloaded and “consumed” it, popular content will end up being shared by more nodes. We will model such correlations in a *probabilistic* way, as follows.

Definition 3 (Availability vs. Popularity). *The availability of any content message \mathcal{M} is related to its popularity through the relation*

$$P\{N_a^{(\mathcal{M})} = m | N_p^{(\mathcal{M})} = n\} = g(m|n) \quad (4)$$

The above conditional probabilities can describe a wide range of cases where availability depends on popularity, and some additional randomness might be present due to factors like: natural churn in the nodes sharing the content, content-dependent differences in the sharing policies applied by nodes, estimation noise, etc. Some special cases of this model include: (i) *uncorrelated availability*, where $g(m|n) \equiv g(m)$; (ii) *deterministic availability*, where:

$$N_a^{(\mathcal{M})} = \rho(N_p^{(\mathcal{M})}) \Leftrightarrow g(m|n) = \begin{cases} 1, & m = \rho(n) \\ 0, & \text{otherwise} \end{cases}$$

where $\rho(n) : [1, N] \rightarrow [0, N]$ can be an arbitrary function. One such example could be a deterministic approximation of $g(m|n)$ with its average value, namely $\rho(n) = \bar{g}(n) \equiv \sum_m m \cdot g(m|n)$.

3. ANALYSIS OF CONTENT REQUESTS

We will now analyze how different popularity, availability, and mobility patterns (possibly arising from different applications, policies, and network settings) affect key metrics like: (i) the delay to

²This could be an average, calculated over some time window.

Table 1: Important Notation

MOBILITY (Section 2.1)		
λ_{ij}	Contact rate between nodes i and j	
$f_\lambda(\lambda)$	Contact rates distribution	
$\mu_\lambda, \sigma_\lambda^2$	Mean value/ variance of contact rates, $CV_\lambda = \frac{\sigma_\lambda}{\mu_\lambda}$	
CONTENT TRAFFIC (Section 2.2)		
$i \rightarrow \mathcal{M}$	Node i is interested / requests content \mathcal{M}	
\mathbf{M}	Set of contents in the network, $ \mathbf{M} = M$.	
$N_p^{(\mathcal{M})}$	Popularity of content \mathcal{M}	Def. 1
$\mathcal{C}_p^{(\mathcal{M})}$	Set of nodes interested in content \mathcal{M}	Def. 1
$P_p(n)$	Probability distribution of content popularity	Eq. (2)
$i \leftarrow \mathcal{M}$	Node i holds a copy of content \mathcal{M}	
$N_a^{(\mathcal{M})}$	Availability of content \mathcal{M}	Def. 2
$\mathcal{C}_a^{(\mathcal{M})}$	Set of nodes that hold a copy of content \mathcal{M}	Def. 2
$g(m n)$	Availability - Popularity relation	Def. 3
$\rho(n)$	Deterministic case for $g(m n)$	
$\bar{g}(n)$	The average value of $g(\cdot n)$	
ANALYSIS (Section 3.1)		
$P_p^{req.}(n)$	Popularity distribution of a random request	Lemma 3.1
$P_a^{req.}(n)$	Availability distribution of a random request	Lemma 3.2
T_{ij}	Time of next meeting between nodes i and j	
$T_{\mathcal{M}}$	Content access time	
$X_{\mathcal{M}}$	Sum of meeting rates of j and nodes $\in \mathcal{C}_a^{(\mathcal{M})}$	Eq. (6)

access a content of interest, (ii) the probability to retrieve a content before a deadline. A key parameter for these metrics is the number of holders for the requested content (availability). The higher this number, the sooner a requesting node will encounter one of them.

While content availability might sometimes be time dependent [11], or the content holders might be chosen based on their mobility properties [9], we first make two additional assumptions that allow us to derive simple, useful expressions. In Section 3.3, we relax both these assumptions.

Assumption 2. The popularity $N_p^{(\mathcal{M})}$ and availability $N_a^{(\mathcal{M})}$ of a content \mathcal{M} do not change over time.

Assumption 3. The set of requesters $\mathcal{C}_p^{(\mathcal{M})}$ and holders $\mathcal{C}_a^{(\mathcal{M})}$ of a content \mathcal{M} are independent of node mobility.

Assumption 2 is valid (or a good approximation), for example, when the number of holders is chosen by the cellular operator [8, 9] or content provider, and other nodes cannot act as holders or do not have incentives to do so. It is also valid when a given service (e.g. Internet access, or specific sensor) is offered only by a certain number of devices [6], or the “content” refers to a *channel* or *category* and not a particular file [17]. Nevertheless, if a content is disseminating and new nodes are willing to share it [7], then its availability might change over time.

Assumption 3 is a reasonable approximation when a *mobility oblivious* allocation policy is considered (e.g. [11], or the homogeneous algorithm of [9]) or when there is no knowledge of the interests-mobility correlation, if any. Nevertheless, there exist scenarios where who holds what content might depend on the contact rates with other nodes [10, 9].

3.1 Preliminary Analysis

Assume we observe the network for a long time, during which a large number of requests have been made. Assume further that we pick one such request randomly, which happens to be for content \mathcal{M} , and we want to predict its performance. We need to answer the following two questions:

Q.1 What is the popularity of \mathcal{M} ?

Q.2 How fast does a requesting node meet \mathcal{M} 's holders?

Q.1 is needed to predict the availability for \mathcal{M} , which according to Assumption 2 does not depend on the exact time of the request. Given the availability of \mathcal{M} , Q.2 will estimate the (sum of) contact rates between the requesting node and the holders, according to Assumptions 1 and 3.

Answering Q.1

It is easy to see that the popularity of \mathcal{M} should be proportional to $P_p(n)$: the higher the number of different contents with a popularity value n , the higher the chance that \mathcal{M} will be of popularity n . However, the higher the popularity of a content, the more the requests made for it. Hence, a first important observation is that the popularity of the content of such a *random request* is *not* distributed as $P_p(n)$ but is also proportional to the popularity value n .

Consider a stylized example, where only two contents exist in the network, content A with popularity value 10 and content B with popularity value 1. Hence, “half” the contents are of high popularity (10), and “half” of low (1), or in other words $P_p(10) = P_p(1) = \frac{1}{2}$. However, if we observe the network for a long time, 10 times more requests will be made, on average, for content A. Consequently, if we select a request randomly, there is a $10\times$ higher chance that it will be for content A, that is, for the content of popularity 10. Normalizing to have a proper probability distribution gives us the following lemma.

Lemma 3.1. The probability that a random request is for a content of popularity equal to n is given by

$$P_p^{req.}(n) = \frac{n}{E_p[n]} \cdot P_p(n)$$

where $E_p[n] = \sum_n n \cdot P_p(n)$ is the average content popularity³.

Answering Q.2

The answer to question Q.2 consists of two separate steps: (i) we calculate the number of holders of \mathcal{M} , and then (ii) we calculate how fast the requesting node can meet these holders. Towards answering (i), Lemma 3.2 maps the popularity of the content involved in a random request (derived in Lemma 3.1) to the number of holders for this content. This number is a random variable dependent both on the popularity distribution $P_p(n)$, and on the availability function $g(m|n)$.

Lemma 3.2. The probability that a random request is for a content of availability equal to m is given by

$$P_a^{req.}(m) = \frac{E_p[n \cdot g(m|n)]}{E_p[n]}$$

Proof. For a random request for content \mathcal{M} , using the property of conditional expectation, we can write [18]:

$$P_a^{req.}(m) = \sum_n P\{N_a^{(\mathcal{M})} = m | N_p^{(\mathcal{M})} = n\} \cdot P_p^{req.}(n)$$

where $P_p^{req.}(n)$ is defined in Lemma 3.1. Substituting, from Def. 3 and Lemma 3.1, the above terms, we successively get

$$\begin{aligned} P_a^{req.}(m) &= \sum_n g(m|n) \cdot \frac{n}{E_p[n]} \cdot P_p(n) \\ &= \frac{\sum_n g(m|n) \cdot n \cdot P_p(n)}{E_p[n]} = \frac{E_p[n \cdot g(m|n)]}{E_p[n]} \end{aligned}$$

which proves the Lemma. \square

³We use subscript p to denote an expectation over the popularity distribution $P_p(n)$, and n denotes the random popularity values.

Table 2: Performance Metrics when $f_\lambda \sim \text{Gamma}$ with μ_λ, CV_λ and $P_p(n) \sim \text{Pareto}(n_0, \alpha = 2)$.

$\rho(n) = c \cdot n$	$E[T_{\mathcal{M}}] = \frac{1}{\mu_\lambda \cdot CV_\lambda^2} \left[\frac{c \cdot n_0}{CV_\lambda^2} \cdot \ln \left(\frac{1}{1 - \frac{CV_\lambda^2}{c \cdot n_0}} \right) - 1 \right]$
$\rho(n) = c \cdot \ln(n)$	$P\{T_{\mathcal{M}} \leq TTL\} = 1 - \frac{1}{(1 + \ln(\gamma)) \cdot \gamma^{\ln(n_0)}}$ where $\gamma = (1 + \mu_\lambda \cdot CV_\lambda^2 \cdot TTL) \frac{c}{CV_\lambda^2}$

Having computed the statistics for the content availability, we can now calculate how fast the requesting node, say j , meets any of the holders i (i.e. nodes $i \in \mathcal{C}_a^{(\mathcal{M})}$). As discussed in Section 2.1, the inter-contact intervals are shown to be either exponentially or pareto distributed:

Exponential Inter-Contact Times. Let T_{ij} denote the inter-contact times between node j and a node $i \in \mathcal{C}_a^{(\mathcal{M})}$, and let T_{ij} be exponentially distributed with rate λ_{ij} . If we denote with $T_{\mathcal{M}}$ the first time until j meets any of the nodes $i \in \mathcal{C}_a^{(\mathcal{M})}$ (and, thus, accesses the content), then: $T_{\mathcal{M}} = \min_{i \in \mathcal{C}_a^{(\mathcal{M})}} \{T_{ij}\}$, i.e. $T_{\mathcal{M}}$ is distributed as a minimum of exponential random variables, and it holds that [18]:

$$T_{\mathcal{M}} \sim \text{exp}(X_{\mathcal{M}}) \Leftrightarrow P\{T_{\mathcal{M}} > t\} = e^{-X_{\mathcal{M}} \cdot t} \quad (5)$$

where

$$X_{\mathcal{M}} = \sum_{i \in \mathcal{C}_a^{(\mathcal{M})}} \lambda_{ij} \quad (6)$$

Pareto Inter-Contact Times. Inter-contact times between node j and a node $i \in \mathcal{C}_a^{(\mathcal{M})}$ are pareto distributed with *shape* and *scale* parameters α_{ij} and t_0 , respectively:

$$T_{ij} \sim \text{pareto}(\alpha_{ij}) \Leftrightarrow P\{T_{ij} > t\} = \left(\frac{t_0}{t}\right)^{\alpha_{ij}} \quad (7)$$

Then, it can be shown for $T_{\mathcal{M}} = \min_{i \in \mathcal{C}_a^{(\mathcal{M})}} \{T_{ij}\}$ that

$$T_{\mathcal{M}} \sim \text{pareto}(A_{\mathcal{M}}) \Leftrightarrow P\{T_{\mathcal{M}} > t\} = \left(\frac{t_0}{t}\right)^{A_{\mathcal{M}}} \quad (8)$$

where $A_{\mathcal{M}} = \sum_{i \in \mathcal{C}_a^{(\mathcal{M})}} \alpha_{ij}$.

Remark: In this case the contact rates will be $\lambda_{ij} = \frac{1}{E[T_{ij}]} = \frac{1}{t_0} \cdot \left(1 - \frac{1}{\alpha_{ij}}\right)$, $\alpha_{ij} > 1$. However, for simplicity, we can use the parameters α_{ij} instead of the rates λ_{ij} , and, correspondingly, a distribution $f_\alpha(\alpha)$, instead of $f_\lambda(\lambda)$.

Clearly, knowing $X_{\mathcal{M}}$ (resp. $A_{\mathcal{M}}$) is needed to proceed with the desired metric derivation. Based on the preceding discussion, $X_{\mathcal{M}}$ (resp. $A_{\mathcal{M}}$) is a random variable that depends on: (i) the number of content holders m (i.e. the cardinality of set $\mathcal{C}_a^{(\mathcal{M})}$ in Eq.(6)), and (ii) the meeting rates with the holders. Applying Assumption 3, it holds that, conditioning on m , $X_{\mathcal{M}}$ (Eq. (6)) is a sum of m i.i.d. random variables $\lambda_{ij} \sim f_\lambda(\lambda)$, i.e

$$X_{\mathcal{M}} \sim f_{m\lambda}(x) = (f_\lambda * f_\lambda \cdots * f_\lambda)_m, \quad (9)$$

where $*$ denotes convolution, and mean value [18]:

$$E[X_{\mathcal{M}} | N_a^{(\mathcal{M})} = m] = E_{m\lambda}[x] = m \cdot \mu_\lambda \quad (10)$$

Similarly, for Pareto intervals ($f_\alpha(\alpha)$, μ_α):

$$A_{\mathcal{M}} \sim f_{m\alpha}(x) = (f_\alpha * \cdots * f_\alpha)_m, \quad E_{m\alpha}[x] = m \cdot \mu_\alpha$$

Due to space limitations, the remaining analysis will refer to the case of exponential inter-contact times. The analysis for the Pareto case is similar; when necessary, we might highlight if there are any

major differences. We refer the interested reader to [19] for the detailed results of the Pareto case.

3.2 Performance Metrics

We consider two main performance metrics: the *average delay* and *delivery probability*. Based on the analysis of Section 3.1, we derive results under generic content traffic (i.e. $P_p(n)$ and $g(m|n)$) and mobility (i.e. $f_\lambda(\lambda)$) patterns.

3.2.1 Content Access Delay

Result 1. *The expected content access delay can be computed with the expression*

$$E[T_{\mathcal{M}}] = \frac{1}{E_p[n]} \cdot E_p \left[n \cdot \sum_m E_{m\lambda} \left[\frac{1}{x} \right] \cdot g(m|n) \right]$$

Proof. The time $T_{\mathcal{M}}$ a node j needs to access a content \mathcal{M} is exponentially distributed with rate $X_{\mathcal{M}}$. However, $X_{\mathcal{M}}$ is a random variable itself, distributed with $f_{m\lambda}(x)$ (Eq. (9)). Thus, we can write for the expected content access delay:

$$\begin{aligned} E[T_{\mathcal{M}}] &= \sum_m E[T_{\mathcal{M}} | N_a^{(\mathcal{M})} = m] \cdot P_a^{req.}(m) \\ &= \sum_m \int E[T_{\mathcal{M}} | X_{\mathcal{M}} = x, N_a^{(\mathcal{M})} = m] \cdot f_{m\lambda}(x) dx \cdot P_a^{req.}(m) \\ &= \sum_m \int \frac{1}{x} \cdot f_{m\lambda}(x) dx \cdot P_a^{req.}(m) \end{aligned} \quad (11)$$

The last equality follows from the fact that the expectation of an exponential⁴ random variable with rate x is $\frac{1}{x}$.

Expressing the integral in Eq. (11) as an expectation over the $f_{m\lambda}(x)$ and substituting $P_a^{req.}(m)$ from Lemma 3.2, gives

$$\begin{aligned} E[T_{\mathcal{M}}] &= \sum_m E_{m\lambda} \left[\frac{1}{x} \right] \cdot \frac{E_p[n \cdot g(m|n)]}{E_p[n]} \\ &= \frac{1}{E_p[n]} \cdot \sum_m E_{m\lambda} \left[\frac{1}{x} \right] \cdot E_p[n \cdot g(m|n)] \end{aligned} \quad (12)$$

Rearranging the expectations and summation in Eq. (12) we get the expression of Result 1. \square

If the functions $f_\lambda(\lambda)$, $g(m|n)$ and $P_p(n)$ are known, the expected delay $E[T_{\mathcal{M}}]$ can be computed directly from Result 1, as shown in the following example.

Example Scenario: The contact rates (f_λ) follow a *gamma distribution*, as suggested in [20], with μ_λ and CV_λ . Content popularity $P_p(n)$ is Pareto distributed, as observed in [21], with *scale* and *shape* parameters n_0 and $\alpha = 2$, respectively. Finally, we consider a (deterministic) allocation of holders, $\rho(n) = c \cdot n$ (see Section 2.2.2). Then a closed form expression for $E[T_{\mathcal{M}}]$ is given in the first row of Table 2.

However, in a real implementation, it might not be always possible to know the *exact* distributions of the contact rates (f_λ) and/or the availabilities ($g(m|n)$), needed to compute the expression of Result 1. In the following theorem, we derive an expression for $E[T_{\mathcal{M}}]$ that requires only the *average statistics* (which are much easier to estimate or measure in a real scenario), namely (i) the mean value of the contact rates, μ_λ , and (ii) the average availability for contents of a given popularity, $\bar{g}(n)$.

⁴Considering Pareto intervals, one only needs to change the integral in Eq. (11) as: $\int \frac{x \cdot t_0}{x-1} \cdot f_{m\alpha}(x) dx$.

Theorem 3.3. A lower bound for the expected content access delay is given by

$$E\{T_{\mathcal{M}}\} \geq \frac{1}{\mu_{\lambda} \cdot E_p[n]} \cdot E_p \left[\frac{n}{\bar{g}(n)} \right]$$

Proof. In Result 1 we can express $E_{m\lambda} \left[\frac{1}{x} \right]$ as $E_{m\lambda}[h(x)]$, where $h(x) = \frac{1}{x}$. Since $h(x)$ is a convex function, applying *Jensen's inequality*, i.e. $h(E[x]) \leq E[h(x)]$, gives

$$E_{m\lambda} \left[\frac{1}{x} \right] \geq \frac{1}{E_{m\lambda}[x]} = \frac{1}{m \cdot \mu_{\lambda}} \quad (13)$$

where, in the equality, we used Eq. (10).

Substituting Eq. (13) in the expression of Result 1, gives

$$E\{T_{\mathcal{M}}\} \geq \frac{1}{\mu_{\lambda} \cdot E_p[n]} \cdot E_p \left[n \cdot \sum_m \frac{1}{m} \cdot g(m|n) \right] \quad (14)$$

The sum in Eq. (14) is the expectation over $g(\cdot|n)$, i.e.

$$\sum_m \frac{1}{m} \cdot g(m|n) = E_g \left[\frac{1}{m} \right] \quad (15)$$

Applying, as before, *Jensen's inequality*, we get

$$\sum_m \frac{1}{m} \cdot g(m|n) = E_g \left[\frac{1}{m} \right] \geq \frac{1}{E_g[m]} = \frac{1}{\bar{g}(n)} \quad (16)$$

where we used for $E_g[m]$ the notation $\bar{g}(n)$.

Combining Eq. (16) and Eq. (14), the expression of the theorem follows directly. \square

3.2.2 Content Access Probability

One often needs to also know the probability that a node can access a content by some deadline, i.e. $P\{T_{\mathcal{M}} \leq TTL\}$. E.g, a node might lose its interest in a content (e.g. news) after some time, or in an offloading scenario a node might decide to access a content directly to the base station.

Result 2. The probability a content to be accessed before a time *TTL* can be computed with the expression

$$P\{T_{\mathcal{M}} \leq TTL\} = 1 - \frac{E_p \left[n \cdot \sum_m E_{m\lambda} \left[e^{-x \cdot TTL} \right] g(m|n) \right]}{E_p[n]}$$

Proof. Conditioning on the values of $N_a^{(\mathcal{M})}$ and $X_{\mathcal{M}}$, as in Eq. (11), we can write:

$$\begin{aligned} P\{T_{\mathcal{M}} \leq TTL\} &= \\ &= \sum_m \int P\{T_{\mathcal{M}} \leq TTL | X_{\mathcal{M}} = x, N_a^{(\mathcal{M})} = m\} \cdot f_{m\lambda}(x) dx \cdot P_a^{req.}(m) \\ &= 1 - \sum_m \int e^{-x \cdot TTL} \cdot f_{m\lambda}(x) dx \cdot P_a^{req.}(m) \end{aligned} \quad (17)$$

where the last equality follows because $T_{\mathcal{M}}$ is exponentially⁵ distributed with rate $X_{\mathcal{M}} = x$. After some similar steps as in Theorem 3.3, the final result follows. \square

The expression of Result 2 for the example scenario of Section 3.2.1, with a different allocation function $\rho(n) = c \cdot \ln(n)$, is given in the second row of Table 2.

⁵In the Pareto case, the integral in Eq. (17) changes as: $\int \left(\frac{t_0}{TTL} \right)^x \cdot f_{m\alpha}(x) dx$, for $TTL \geq t_0$.



Figure 1: Markov Chain for the dissemination of a content with initial popularity and availability n and m , respectively.

Theorem 3.4. An upper bound for the probability to access a content by a time *TTL* is given by

$$P\{T_{\mathcal{M}} \leq TTL\} \leq 1 - \frac{1}{E_p[n]} \cdot E_p \left[n \cdot e^{-\bar{g}(n) \cdot \mu_{\lambda} \cdot TTL} \right]$$

Proof. The bound follows easily by observing that $h(x) = e^{-x \cdot TTL}$ is a convex function, and applying *Jensen's inequality* and the methodology of Theorem 3.3. \square

3.3 Extensions

In this section, we study how the results of Section 3.2 can be modified, when we remove the Assumptions 2 and 3. Due to space limitations, we state only the main findings; the detailed proofs can be found in [19].

3.3.1 Popularity / Availability Time Dependence

Let us assume a scenario where, initially, some nodes hold some content items (e.g. data files), in which some other nodes are interested. This can be, for example, a content sharing scenario with contents being, e.g., some google maps. When a node interested in a content item, meets a holder and gets the content, it can hold it in its memory and act as a holder too. Specifically, we describe such scenarios as:

Definition 4.

- I. When a requester accesses a content, acts as a holder for it.
- II. The initial content popularity and availability patterns are given by $P_p(n)$ and $g(m|n)$.

Result 3. Under time changing availability / popularity (Def. 4), the expected content access delay is approximately given by

$$E\{T_{\mathcal{M}}\} = \frac{1}{\mu_{\lambda} \cdot E_p[n]} \cdot E_p \left[\ln \left(1 + \frac{n}{\bar{g}(n)} \right) \right]$$

Sketch of proof: Let us consider a content \mathcal{M} of initial popularity $N_p^{(\mathcal{M})}(0) = n$ and availability $N_a^{(\mathcal{M})}(0) = m$. When the first requester accesses the content, the number of holders will increase to $m + 1$ and the remaining requesters will be $n - 1$. Building a Markov Chain as in Fig. 1, where each state denotes the number of holders, it can be shown for the expected delay of moving from state $m+k$ to state $m+k+1$, $k \in [0, 1]$, that it holds $E\{T_{k,k+1}\} \approx \frac{1}{(m+k) \cdot (n-k) \cdot \mu_{\lambda}}$. Computing the times $E\{T_{k,k+1}\}$ and averaging over all the contents, gives the expected delay.

The model of Def. 4 can be further extended, e.g. for scenarios where nodes *might* act as holders (with some probability) or holders can also drop some contents. We defer such a study as a part of a future work.

3.3.2 Mobility Dependent Allocation

Definition 5 (Mobility Dependent Allocation). The probability π_{ij} a node i to be a holder for a content in which a node j is interested, is related to their contact rate λ_{ij} such that $\pi_{ij} = \pi(\lambda_{ij})$, where $\pi(\cdot)$ is a function from \mathbb{R}^+ to $[0, 1]$.

Result 4. Under Def. 5, Theorems 3.3 and 3.4 and Result 3 hold if we replace μ_λ with $\mu_\lambda^{(\pi)}$, where

$$\mu_\lambda^{(\pi)} = \frac{E_\lambda[\lambda \cdot \pi(\lambda)]}{E_\lambda[\pi(\lambda)]}$$

where $E_\lambda[\cdot]$ denotes an expectation taken over the contact rates distribution $f_\lambda(\lambda)$ (Assumption 1).

Sketch of proof: Since the requesters-holders contact rates are mobility dependent, the contact rates between them are not distributed with the contact rates distribution $f_\lambda(\lambda)$, but with a modified version of it, i.e. with a distribution:

$$f_\pi(\lambda) = \frac{1}{E_\lambda[\pi(\lambda)]} \cdot \pi(\lambda) \cdot f_\lambda(\lambda)$$

Hence, Eq. (9) and Eq. (10) need to be modified as:

$$X_{\mathcal{M}} \sim f_{m\pi}(x) = (f_\pi * f_\pi \cdots * f_\pi)_m$$

$$E[X_{\mathcal{M}} | N_a^{(\mathcal{M})} = m] = E_{m\pi}[x] = m \cdot \frac{E_\lambda[\lambda \cdot \pi(\lambda)]}{E_\lambda[\pi(\lambda)]} = m \cdot \mu_\lambda^{(\pi)}$$

Example Scenario. The holders of a content \mathcal{M} are selected taking into account their contact rates with the requesters, as following: Each node i (candidate to be a holder) is assigned a weight $w_i = \prod_{j \in c_p^{(\mathcal{M})}} \lambda_{ij}$. Using such weights, the selection of holders that rarely meet the requesters is avoided. Then, each node is selected to be one of the $N_a^{(\mathcal{M})}$ holders with probability $p_i = \frac{w_i}{\sum_i w_i}$. With respect to Def. 5, it turns out that this mechanism is (approximately) described by $\pi(\lambda) = c \cdot \lambda$. Substituting $\pi(\lambda)$ in Result 4, gives

$$\mu_\lambda^{(\pi)} = \frac{E_\lambda[\lambda \cdot \pi(\lambda)]}{E_\lambda[\pi(\lambda)]} = \frac{E_\lambda[\lambda^2]}{E_\lambda[\lambda]} = \mu_\lambda \cdot (1 + CV_\lambda^2) \quad (18)$$

3.4 Model Validation

As a first validation step, we compare our theoretical predictions to synthetic simulation scenarios conforming to the models of Section 2, in order to consider (a) various mobility and content traffic patterns, and (b) large networks.

Simulation Scenarios: We assign to each pair $\{i, j\}$ a contact rate λ_{ij} , which we draw randomly from a distribution $f_\lambda(\lambda)$, and create a sequence of contact events (Poisson process with rate λ_{ij}). Then, we create M contents and assign to each of them a popularity value ($N_p^{(\mathcal{M})}$), drawn from the distribution $P_p(n)$. According to the given function $g(m|n)$, we assign the availability values ($N_a^{(\mathcal{M})}$). Finally, for each content \mathcal{M} , we randomly choose the $N_p^{(\mathcal{M})}$ nodes that are interested in it and its $N_a^{(\mathcal{M})}$ holders.

Mobility / Popularity patterns: In most of the scenarios we present, we use the gamma distribution for the contact rates (i.e. $f_\lambda(\lambda)$), since it has been shown to match well characteristics of real contact patterns [20]. Also, content popularity in mobile social networks has been shown to follow a power-law distribution, e.g. [21]. Therefore, we select $P_p(n)$ to follow Discrete (Bounded) Pareto or Zipf distributions, similarly to the majority of related works [11, 9, 1].

In Fig. 2 we present the simulation results, along with our theoretical predictions, in scenarios of $N = 10000$ nodes with varying mobility and content popularity patterns. The mean contact rate is $\mu_\lambda = 1$ and content popularity follows a Bounded Pareto distribution with shape parameter (i.e. exponent) α and $n \in [50, 1000]$. The availability function is $\rho(n) = 0.2 \cdot n$ (i.e. deterministic). An almost perfect match between simulation results (markers) and

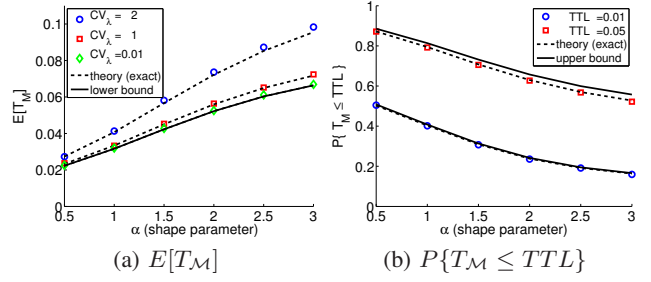


Figure 2: (a) $E[T_{\mathcal{M}}]$ and (b) $P\{T_{\mathcal{M}} \leq TTL\}$ in scenarios with varying content popularity (α : shape parameter) and $\rho(n) = 0.2 \cdot n$.

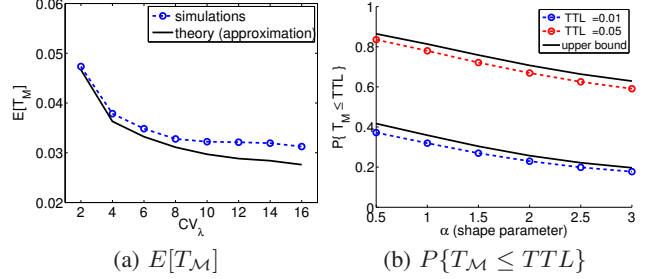


Figure 3: (a) $E[T_{\mathcal{M}}]$ in scenarios under Def. 4 and (b) $P\{T_{\mathcal{M}} \leq TTL\}$ in scenarios under Def. 5. $\rho(n) = 0.2 \cdot n$.

the theoretical predictions (dashed lines) of Results 1 and 2 can be observed. In Fig. 2(a), the lower bound (continuous line) of Theorem 3.3 is very tight for low mobility (i.e. small CV_λ) and/or content popularity (i.e. small α) heterogeneity. For the delivery probability $P\{T_{\mathcal{M}} \leq TTL\}$ (Fig. 2(b)), we present the results for two different values of TTL in scenarios with $CV_\lambda = 2$. Here, the upper bound (continuous line) of Theorem 3.4 is very close to the simulation results, despite the very heterogeneous mobility.

In Fig. 3(a) we compare Result 3 with simulations on scenarios conforming to Def. 4: $P_p(n)$ is a *Bounded Pareto* distribution with $\alpha = 2$, and $f_\lambda(\lambda) \sim \text{Pareto}$. It can be seen that our theoretical prediction (approximation) achieves good accuracy even in these very heterogeneous mobility scenarios.

Results for scenarios with mobility-dependent availability (Def. 5) are presented in Fig. 3(b). $P_p(n)$ is selected as before and $f_\lambda(\lambda) \sim \text{Gamma}$ with $\mu_\lambda = 1, CV_\lambda = 0.5$. The allocation of holders is made as in the example in Section 3.3.2. The upper bounds of Result 4 are tight in all scenarios, similarly to the case without mobility dependence (Fig. 2(b)).

Finally, we need to mention that we have also performed a large number of other scenarios, with similar conclusions.

4. CASE STUDY: MOBILE DATA OFFLOADING

The results of Section 3 can be used to predict the performance of a given content allocation policy or content-sharing scheme. In this section, we show how these results could be also used to design / optimize policies. We focus on an application that has recently attracted attention, that of *mobile data offloading* using opportunistic networking [7, 8, 9]. Nevertheless, the same methodology applies for a range of other applications where the number of content/service providers must be chosen.

In a mobile data offloading scenario, the cellular network provider distributes content copies only to some of the nodes interested in this content (holders), in order to reduce the cellular traffic (possibly offering some incentives to the holder nodes). The remaining (interested) nodes must then retrieve the content from the designated holders during direct encounters. A tradeoff is involved between the amount of traffic offloaded and the average delay for non-holders. In some cases, an additional *QoS* constraint might exist: if the delay to access a content exceeds a *TTL*, a requesting node will download it from the infrastructure [7, 8, 9].

Hence, the objective in offloading optimization problems is how the cellular network provider should *choose the set of holders for each content in order to optimize a performance metric*, under a given constraint (e.g. energy or buffer size) and a given popularity distribution $P_p(n)$.

We study cases with and without *QoS* constraints in Sections 4.1 and 4.2, respectively. For simplicity, we use the expressions of Theorems 3.3 and 3.4 as approximations for $E[T_{\mathcal{M}}]$ and $P\{T_{\mathcal{M}} \leq TTL\}$. Since, these expressions imply that (a) the exact mobility patterns are not known (i.e. only μ_λ is needed) and (b) contents with the same popularity are equivalent, our goal is to select the number of holders for each content with a given popularity. In other words, we try to find the optimal allocation function $g(m|n)$.

4.1 Case 1: no QoS constraints

When no *QoS* constraints exist, the cellular operator decides the maximum amount of traffic that it wishes to serve directly over the infrastructure. Under this constraint, which can be translated as a constraint on the number of holders for different contents, the objective is to minimize the expected delay $E[T_{\mathcal{M}}]$. The following result, formalizes this optimization problem and provides with the optimal solution for $g(m|n)$.

Result 5. *The minimum expected content access delay, under the constraint of an average number of $c_{\mathcal{M}}$ copies per content, i.e.:*

$$\min\{E[T_{\mathcal{M}}]\} \quad s.t. \quad \sum_{\mathcal{M}} N_a^{(\mathcal{M})} = M \cdot c_{\mathcal{M}}, \quad N_a^{(\mathcal{M})} \geq 0$$

can be achieved when the allocation function, $g(m|n)$, is deterministic and equal to

$$\rho^*(n) = \frac{c_{\mathcal{M}}}{E_p[\sqrt{n}]} \cdot \sqrt{n}$$

Proof. Using as an approximation for $E[T_{\mathcal{M}}]$ the expression of Theorem 3.3, we can write

$$E[T_{\mathcal{M}}] = \frac{1}{\mu_\lambda \cdot E_p[n]} \cdot E_p \left[\frac{n}{\bar{g}(n)} \right]$$

Jensen's inequality used in Eq. (16), becomes equality when $g(m|n)$ is deterministic. This suggests that among all the functions $g(m|n)$ with the same average value $\bar{g}(n)$, the minimum delay can be achieved in the case: $\rho(n) = \bar{g}(n)$. Thus, the $E[T_{\mathcal{M}}]$ minimization problem becomes equivalent to

$$\min\{E_p \left[\frac{n}{\rho(n)} \right]\} = \sum_n \frac{n}{\rho(n)} \cdot P_p(n) = \sum_n \frac{n}{\rho_n} \cdot P_p(n) \quad (19)$$

where we expressed the expectation as a sum and denoted $\rho_n = \rho(n)$.

Moreover, we can express the content copies constraint as

$$c_{\mathcal{M}} = \frac{\sum_{\mathcal{M}} N_a^{(\mathcal{M})}}{M} = E_p[\rho(n)] = \sum_n \rho_n \cdot P_p(n) \quad (20)$$

Using Eq. (19) and Eq. (20), the optimization problem becomes

$$\min_{\bar{\rho}} \left\{ \sum_n \frac{n}{\rho_n} \cdot P_p(n) \right\} \quad s.t. \quad \sum_n \rho_n \cdot P_p(n) = c_{\mathcal{M}} \quad (21)$$

where $\bar{\rho}$ denotes the vector with components ρ_n .

The optimization problem of Eq. (21) is *convex* and, thus, it can be solved with the method of Lagrange multipliers [22]. Hence, we need to find the values of $\bar{\rho}$ for which it holds that

$$\nabla \left(\sum_n \frac{n}{\rho_n} \cdot P_p(n) \right) + \nabla \lambda_0 \left(\sum_n \rho_n \cdot P_p(n) - c_{\mathcal{M}} \right) = 0$$

where λ_0 is the langrangian multiplier. Here, the constraint $\rho_n \geq 0$ needs also to be taken into account. However, it is proved to be an inactive constraint (the solution satisfies it) and thus we omit it at this step for simplicity. Similarly, we assume a large enough network, i.e. always holds $\rho_n \leq N$.

The differentiation over ρ_n gives

$$\rho_n = \frac{1}{\sqrt{\lambda_0}} \cdot \sqrt{n} \quad (22)$$

Substituting Eq. (22) in the constraint expression $\sum_n \rho_n \cdot P_p(n) = c_{\mathcal{M}}$ (Eq. (21)), we can easily get

$$\sqrt{\lambda_0} = \frac{\sum_n \sqrt{n} \cdot P_p(n)}{c_{\mathcal{M}}} = \frac{E_p[\sqrt{n}]}{c_{\mathcal{M}}} \quad (23)$$

Then, substituting Eq. (23) in Eq. (22), gives

$$\rho(n) = \rho_n = \frac{c_{\mathcal{M}}}{E_p[\sqrt{n}]} \cdot \sqrt{n} \quad (24)$$

Finally, the values of Eq. (24) satisfy the *Karush-Kuhn-Tucker* conditions, which means that the solution of Eq. (24) is a global minimum [22]. \square

Result 5 is a generic result, since it holds under *any* content popularity pattern. We also note that an allocation policy of $\rho(n) \propto \sqrt{n}$ has also been shown to achieve optimal results in (conventional) peer-to-peer networks [23]. This is an interesting finding, given the inherent differences between the two settings (e.g. node mobility).

Finally, our result is also consistent in scenarios with *mobility dependent holders allocation*. For example, after choosing the number of copies for a content (Result 5), the selection of holders can be made, taking into account mobility utility metrics, e.g. meeting frequency [10] or node centrality [1].

4.2 Case 2: QoS constraints

In cases where a maximum delay *TTL* is required, the objective is to minimize the traffic load served by the infrastructure. The metric used in related work, e.g. [9], is the data offloading ratio, R_{off} , which is defined as the percentage of content requests that are served by nodes. Since requests are served by the infrastructure only after the time *TTL* elapses, it follows that in our framework: $R_{off} = P\{T_{\mathcal{M}} \leq TTL\}$.

Hence the optimization problem is equivalent to

$$\max P\{T_{\mathcal{M}} \leq TTL\} \quad s.t. \quad \sum_{\mathcal{M}} N_a^{(\mathcal{M})} = M \cdot c_{\mathcal{M}}, \quad N_a^{(\mathcal{M})} \geq 0$$

Using the same notation and arguments as in the Section 4.1 and the expression of Theorem 3.4 as an approximation for $P\{T_{\mathcal{M}} \leq TTL\}$, the above optimization problem becomes:

$$\min_{\rho(n)} \{E_p \left[n \cdot e^{-\rho(n) \cdot \mu_\lambda \cdot TTL} \right]\} \quad s.t. \quad E_p[\rho(n)] = c_{\mathcal{M}} \quad (25)$$

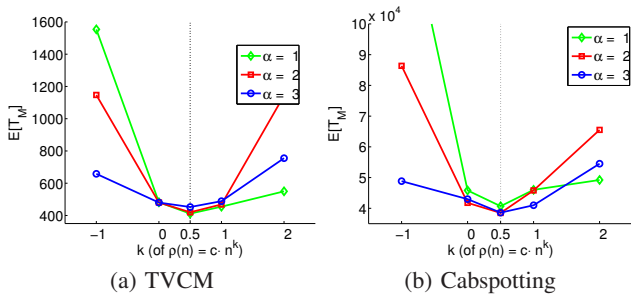


Figure 4: Content access delay $E[T_{\mathcal{M}}]$ of different allocation policies $\rho(n) = c_k \cdot n^k$, where $c_k = \frac{c_{\mathcal{M}}}{E_p[n^k]}$.

with $\rho(n) \geq 0$, or, equivalently:

$$\begin{aligned} \min_{\vec{\rho}} \quad & \left\{ \sum_n n \cdot e^{-\rho_n \cdot \mu_{\lambda} \cdot TTL} \cdot P_p(n) \right\} \\ \text{s.t.} \quad & \sum_n \rho_n \cdot P_p(n) = c_{\mathcal{M}}, \quad \rho_n \geq 0 \end{aligned} \quad (26)$$

The optimization problem of Eq. (26) is convex. Although a closed form solution, as in Result 5, cannot be derived⁶, it can be solved numerically, using well known methods.

4.3 Performance Evaluation

To investigate whether the policies suggested as optimal by our theory indeed perform better, we conducted simulations on various synthetic scenarios and on traces of real networks, where node mobility patterns usually involve much more complex characteristics than our model (Assumption 1).

The results in the majority of scenarios considered have been encouragingly consistent with our theoretical predictions. Hence, we only present here a small, representative sample, due to space limitations. Specifically, we consider the following traces coming from state-of-the-art mobility models or collected in experiments.

TVCM mobility model [24]: Scenario with 100 nodes divided in 4 communities of unequal size. Nodes move mainly inside their community and leave it for a few short periods.

SLAW mobility model [25]: Network with 200 nodes moving in a square area of 2000m (the other parameters are set as in the source code provided in [25]).

Cabspotting trace [26]: GPS coordinates from 536 taxi cabs collected over 30 days in San Francisco. A range of 100m is assumed.

Infocom trace [27]: Bluetooth sightings of 98 mobile and static nodes (iMotes) collected in an experiment during Infocom 2006.

4.3.1 Case 1: no QoS constraints

In each scenario, we compare different allocation functions $\rho(n) = c_k \cdot n^k$, where $c_k = \frac{c_{\mathcal{M}}}{E_p[n^k]}$ is a normalization factor such that the constraint $E_p[\rho(n)] = c_{\mathcal{M}}$ is satisfied.

In Fig. 4 we present simulation results in scenarios for the *TVCM* (Fig. 4(a)) and *Cabspotting* (Fig. 4(b)) traces. Content popularity ($P_p(n)$) follows a *Zipf* distribution with $n \leq 30$ and exponent $\alpha = \{1, 2, 3\}$. The availability constraint is set to $c_{\mathcal{M}} = 10$. It can be seen that the optimal delay $E[T_{\mathcal{M}}]$ is achieved for $k = 0.5$, as Result 5 predicts (despite the fact that we used the expression of the lower bound as an approximation for the expected delay $E[T_{\mathcal{M}}]$).

4.3.2 Case 2: QoS constraints

To evaluate the performance of the allocation function $\rho(n)$ that follows after solving Eq. (26) (i.e. *optimal* allocation), we compare

⁶The difference here is that the constraint $\rho_n \geq 0$ is active.

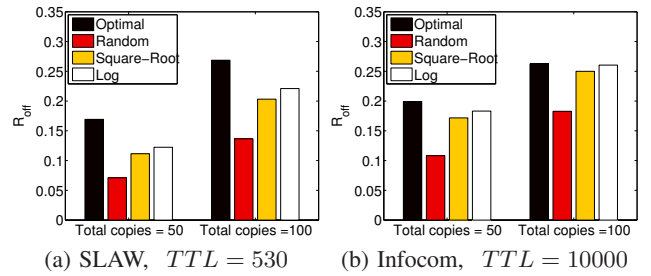


Figure 5: Offloading Ratio R_{off} of different allocation policies $\rho(n)$.

the *offloading ratio* R_{off} it achieves with the offloading ratios of the following policies:

Random: We randomly select a content and give a copy of it to a node. We repeat $M \cdot c_{\mathcal{M}}$ times.

Square Root: We select $\rho(n) \propto \sqrt{n}$ (i.e. the allocation that achieves the minimum expected delay $E[T_{\mathcal{M}}]$).

Log: We select $\rho(n) \propto \log n$.

Random policy has been used in related work as a baseline [9] and *square root* policy is the optimal policy when the metric of interest is the content access delay (Section 4.1). Finally, we observed that the *optimal* policy (Eq. (26)), in the scenarios considered, allocated copies only to the 10% – 20% highest popularity contents. The *log* policy allocates in a similar manner the copies (e.g. no copies to contents with low popularity).

Simulation results on the *SLAW* and *Infocom* scenarios are presented in Fig. 5(a) and 5(b), respectively. The parameters in these scenarios are: $M = 50$ messages, $P_p \sim Zipf$ with $n \in [1, 30]$ and $\alpha = 1$, total copies $M \cdot c_{\mathcal{M}} = \{50, 100\}$. As it can be seen our *optimal* policy (leftmost bar) achieves the highest offloading ratio R_{off} . The random policy is clearly inferior than the others. Between *square root* and *log* policies, it is the latter that achieves better performance. These results indicate that, to maximize R_{off} , it is better to allocate the available resources only for popular contents, and serve the non-popular exclusively through the infrastructure.

5. RELATED WORK

Content-centric applications were introduced in opportunistic networking under the *publish - subscribe* paradigm [2, 17, 16, 10], for which several data dissemination techniques have been proposed. In [2], authors propose a mechanism that identifies social communities and the nodes-“hubs”, and builds an overlay network between them in order to efficiently disseminate data. SocialCast [16] based on information about nodes interests, social relationships and movement predictions, selects the set of holders. Similarly to the above approaches, ContentPlace [10] uses both community detection and nodes social relationships information, to improve the performance of the content distribution.

Under a different setting, [1, 11] study content sharing mechanisms with limited resources (e.g. buffer sizes, number of holders). In [1], authors analytically investigate the data dissemination cost-effectiveness tradeoffs, and propose techniques based on contact patterns (i.e. λ_{ij}) and nodes interests. Similarly, CEDO [11] aims at maximizing the total content delivery rate: by maintaining a utility per content, nodes make appropriate drop and scheduling decisions.

Recently, further novel content-centric application have been proposed, like location-based applications [3, 4] and mobile data of-

flooding [7, 8, 9]. The latter category, due to the rapid increase of mobile data demand, has attracted a lot of attention. In the setting of [7], content copies are initially distributed (through the infrastructure) to a subset of mobile nodes, which then start propagating the contents epidemically. Differently, in [8] the authors consider a limited number of holders, and study how to select the best holders-target-set for each message. In [9], the same problem is considered, and (centralized) optimization algorithms are proposed that take into account more information about the network: namely, size and lifetimes of different contents, and interests, privacy policies and buffer sizes of each node.

In the majority of previous studies, although node interests and content popularity are taken into account, the focus has been on the algorithms and the applications themselves. We believe that our study complements existing work, by providing a common analytical framework for a number of these approaches that can be used both for predicting the performance of proposed schemes, as well as proposing improved ones.

6. CONCLUSION

The increasing number of mobile devices and traffic demand, renders content-centric applications through opportunistic communication very promising. Hence, motivated by the lack of a common analytical framework, we modeled and analyzed the effects of content popularity / availability patterns in the performance of content-centric mechanisms.

As a part of future work we intend to study, in more detail, extensions of our model and to investigate further characteristics of content traffic patterns, like *traffic locality* in location based social networks, and their performance effects.

Acknowledgments

We would like to thank Chadi Barakat, Chiara Boldrini, Jörg Ott, and Andrea Passarella for their feedback while preparing this paper.

7. REFERENCES

- [1] W. Gao and G. Cao, "User-centric data dissemination in disruption tolerant networks," in *Proc. IEEE INFOCOM*, 2011.
- [2] E. Yoneki, P. Hui, S. Chan, and J. Crowcroft, "A socio-aware overlay for publish/subscribe communication in delay tolerant networks," in *Proc. ACM MSWiM*, 2007.
- [3] M. Satyanarayanan, "Mobile computing: The next decade," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 15, no. 2, pp. 2–10, 2011.
- [4] J. Ott and J. Kangasharju, "Opportunistic content sharing applications," in *Proc. ACM NoM Workshop*, 2012.
- [5] M. Conti, S. Giordano, M. May, and A. Passarella, "From opportunistic networks to opportunistic computing," *Communications Magazine, IEEE*, vol. 48, pp. 126–139, sept. 2010.
- [6] M. Pitkänen, T. Kärkkäinen, and et al., "SCAMPI: service platform for social aware mobile and pervasive computing," *ACM Comput. Commun. Rev.*, vol. 42, pp. 503–508, 2012.
- [7] J. Whitbeck, M. Amorim, Y. Lopez, J. Leguay, and V. Conan, "Relieving the wireless infrastructure: When opportunistic networks meet guaranteed delays," in *Proc. IEEE WoWMoM*, 2011.
- [8] B. Han, P. Hui, V. Kumar, M. Marathe, J. Shao, and A. Srinivasan, "Mobile data offloading through opportunistic communications and social participation," *IEEE Trans. on Mob. Comp.*, vol. 11, no. 5, pp. 821–834, 2012.
- [9] Y. Li, M. Qian, D. Jin, P. Hui, Z. Wang, and S. Chen, "Multiple mobile data offloading through disruption tolerant networks," *IEEE Transactions on Mobile Computing*, vol. PrePrints, 2013.
- [10] C. Boldrini, M. Conti, and A. Passarella, "Design and performance evaluation of contentplace, a social-aware data dissemination system for opportunistic networks," *Computer Networks*, vol. 54, no. 4, pp. 589–604, 2010.
- [11] F. Neves dos Santos, B. Ertl, C. Barakat, T. Spyropoulos, and T. Turlletti, "Cedo: Content-centric dissemination algorithm for delay-tolerant networks," in *Proc. ACM MSWiM*, 2013.
- [12] A. Picu, T. Spyropoulos, and T. Hossmann, "An analysis of the information spreading delay in heterogeneous mobility dtns," in *Proc. IEEE WoWMoM*, 2012.
- [13] W. Gao, Q. Li, B. Zhao, and G. Cao, "Multicasting in delay tolerant networks: a social network perspective," in *Proc. ACM MobiHoc*, 2009.
- [14] V. Conan, J. Leguay, and T. Friedman, "Characterizing pairwise inter-contact patterns in delay tolerant networks," in *Proc. ACM Autonomics*, 2007.
- [15] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of human mobility on opportunistic forwarding algorithms," *IEEE Trans. on Mobile Computing*, vol. 6, no. 6, pp. 606–620, 2007.
- [16] P. Costa, C. Mascolo, M. Musolesi, and G. Picco, "Socially-aware routing for publish-subscribe in delay-tolerant mobile ad hoc networks," *IEEE JSAC*, vol. 26, no. 5, pp. 748–760, 2008.
- [17] V. Lenders, G. Karlsson, and M. May, "Wireless ad hoc podcasting," in *Proc. IEEE SECON*, 2007.
- [18] S. M. Ross, *Introduction to Probability Models*. Academic Press, Elsevier, 9 ed., 2007.
- [19] P. Sermpezis and T. Spyropoulos, "Effects of content popularity in the performance of content-centric opportunistic networking: Analytical approach and applications," Tech. Rep. RR-14-292, Eurecom, 2014.
- [20] A. Passarella and M. Conti, "Analysis of individual pair and aggregate intercontact times in heterogeneous opportunistic networks," *IEEE Trans. on Mobile Computing*, vol. 12, no. 12, pp. 2483–2495, 2013.
- [21] Y. Li, M. Steiner, L. Wang, Z.-L. Zhang, and J. Bao, "Exploring venue popularity in foursquare," in *Proc. IEEE INFOCOM Workshops (NetSciCom)*, 2013.
- [22] A. Antoniou and W.-S. Lu, *Practical Optimization: Algorithms and Engineering Applications*. Springer, 2007.
- [23] E. Cohen and S. Shenker, "Replication strategies in unstructured peer-to-peer networks," in *Proc. ACM SIGCOMM*, 2002.
- [24] W.-J. Hsu, T. Spyropoulos, K. Psounis, and A. Helmy, "Modeling spatial and temporal dependencies of user mobility in wireless mobile networks," *IEEE/ACM Trans. on Networking*, vol. 17, no. 5, pp. 1564–1577, 2009.
- [25] K. Lee, S. Hong, S. J. Kim, I. Rhee, and S. Chong, "Slaw: A new mobility model for human walks," in *Proc. IEEE INFOCOM*, 2009.
- [26] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser, "CRAWDAD data set epfl/mobility (v. 2009-02-24)." <http://crawdada.cs.dartmouth.edu/epfl/mobility>.
- [27] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot, "Pocket switched networks and human mobility in conference environments," in *Proc. ACM WDTN 2005*.