Storage Capacity of Opportunistic Information Dissemination Systems

Gianluca Rizzo, Noelia Pérez Palma, Marco Ajmone Marsan and Vincenzo Mancuso

Abstract—Floating Content (FC) is a paradigm for localized infrastructure-less content dissemination, that aims at sharing information among nodes within a restricted geographical area by relying only on opportunistic content exchanges. FC provides the basis for the probabilistic spatial storage of shared information in a completely decentralized fashion, usually without support from dedicated infrastructure. One of the key open issues in FC is the characterization of its performance limits as functions of the system parameters, accounting for its reliance on volatile wireless exchanges and on limited user resources. This paper takes a first step towards tackling this issue, by elaborating a model for the storage capacity of FC, i.e., for the maximum amount of information that can be stored through the FC paradigm. The storage capacity of FC, and of similar probabilistic content dissemination systems, is evaluated with a powerful information theoretical approach, based on a mean field model of opportunistic information exchange. In addition, an extremely simple explicit approximate expression for storage capacity is derived. The numerical results generated by our analytical models are compared to the predictions of realistic simulations under different setups, proving the accuracy of our analytical approaches, and characterizing the properties of the FC storage capacity.

Index Terms—Floating content, Distributed caching, Opportunistic communications, Proactive caching.

———————————— ✦ ————————————

1 INTRODUCTION

Networking of mobile users through opportunistic device-to-device (D2D) communications has recently received considerable attention from the research community [1], [2], [3], although technology to support such infrastructure-less scenario initially was barely available. In recent years, with the development of more reliable versions of Wi-Fi Direct [4] and Bluetooth Low Energy (BLE) [5], and with the standardization of Sidelink in LTE and 5G [6], [7], [8], opportunistic mobile networks have become a viable possibility. Their application domains, traditionally including scenarios in which infrastructure is not available, such as disaster areas or battlefields [9], have now spread to pandemic-driven warnings and spontaneous protests. Indeed, several countries have recently invested in the development of proximity-based applications based on D2D communications to help contact tracing, thus stimulating technology development in this area. In addition, smartphone apps like FireChat[1] and Bridgefy[2], exploit D2D communications to enable information exchanges either when Internet access is unavailable, or when localized distribution is desired, or where infrastructure-based communications are not trusted. For instance, FireChat was the communication medium of choice in several civil protests.

The key questions about the usefulness of infrastructure-less communications are about how fast and reliable they are, and how much information they can store. While several studies have shown that localized infrastructure-less content dissemination schemes such as *Floating Content* (FC) [10], [11], [12], Locus [13] or Hovering information [14], [15] can be effective in the above-mentioned contexts, little has been done so far to quantify their

storage capacity, which is the focus of this work.

In this paper we focus on FC, which aims at disseminating information over a defined geographic area (called replication zone or RZ), based solely on direct D2D connectivity [16]. By so doing, FC stores information spatially in a probabilistic fashion, despite the mobility of user devices (UEs) and the unreliability of information exchanges, and with no need for centralized servers. Thus, FC can deliver the stored content proactively to users which are expected to traverse a specific region (the Zone of Interest, or ZOI), before they reach it. Hence, the main performance metric in such systems is the *success ratio*, i.e., the average fraction of nodes that enter the ZOI with content.

Clearly, guaranteeing (probabilistically) content persistence and a given target performance in such a volatile setting, without the support of a centralized static infrastructure, comes at a cost. The main additional cost as compared to traditional centralized infrastructure-based solutions (e.g., with respect to commonly studied distributed information storage schemes in which mobile UEs cache popular content items [17], [18], [19], [20]), is in a drastic increase both in content redundancy across the user population, and in the volume of communications required to reach the target population of users.

A strong point of FC is that, by enabling direct D2D content transfer and sharing without routing through access points (APs) or base stations (BSs), it offers a parsimonious approach to distributed edge storage because it can achieve higher energy efficiency while decreasing the utilization of BS resources.

A major open issue for the practical viability of FC as a distributed edge storage system is the characterization of its scalability, i.e., of the amount of information that FC is able to store for a given set of system parameters. This is the problem we address in this paper.

G. Rizzo is with HES SO Valais, Switzerland, and UNIFG, Italy. N. Perez Palma is with IMDEA Networks Institute and University Carlos III of Madrid, Spain. M. Ajmone Marsan is with Politecnico di Torino, Italy and IMDEA Networks Institute, Spain. V. Mancuso is with IMDEA Networks Institute, Spain.

1. https://apps.apple.com/us/app/firechat/id719829352
2. https://bridgefy.me/

1.1 Contribution

In this paper, we propose a simple analytical model of the storage capacity of probabilistic distributed edge storage systems such as FC, based on a mean field model of the opportunistic information exchange, which allows for a first order characterization of the scaling laws of the storage capacity of these systems. Specifically, the contributions of this paper are as follows:

- We develop an analytical model of FC performance, based on a mean field model of the dynamics of the population of users storing the information items, and of the population of users which are in the process of exchanging (sending or receiving) such items;
- We derive analytical expressions of the FC storage capacity, as a function of node mobility and of the geometry of the replication zone;
- We formulate an optimization problem for the derivation of the maximum amount of information which can be stored with FC, showing that it can be solved efficiently. To the best of our knowledge, this is the first work to characterize analytically the storage capacity of probabilistic distributed storage schemes such as FC;
- We evaluate numerically our results, validating our assumptions against simulation, under different mobility models, showing the accuracy of our mean field approach, and characterizing the properties of FC storage capacity as a function of the main system parameters.
- Leveraging the insights provided by our model on the FC behavior as an opportunistic information storage system, we derive an extremely simple and intuitive closed form approximate expression for FC storage capacity.

This paper builds upon the conference publication [21]. With respect to the conference version, this paper includes the following new items:

1) We propose a generalized version of our model, which applies to system configurations in which some of the UEs are static, e.g., in order to account for the presence of parked vehicles that participate in the opportunistic information exchange and storage. We discuss the impact of the fraction of static users, revealing interesting tradeoffs.
2) We evaluate the impact of non-uniform mobility models on the accuracy of our approach, by assessing the quality of the model performance predictions with both synthetic and measurement-based mobility traces. We study by simulation the Random Waypoint and Random Direction mobility models in addition to mobility trajectories from the city of Luxembourg. In this way, we show that our analytical model is robust with respect to a relaxation of the uniformity assumption.
3) We derive a simple explicit approximate expression of the storage capacity of a FC system, thus enabling an immediate estimation of system performance, and hence a dimensioning of the main system parameters that allow meeting specified goals.

The paper is organized as follows. In Section 2, we discuss some of the previous literature in the field. In Section 3, we introduce the main system characteristics, and in Section 4, we present our mean field approach to performance analysis of FC. In Section 5, we characterize the storage capacity of FC. In Section 6, we numerically assess the accuracy of our results, and evaluate the impact of the main system parameters on FC storage capacity. Finally, Section 7 concludes the paper.

2 RELATED WORK

Our analysis of the FC paradigm is a contribution to a more general research effort aimed at optimizing the utilization of resources (bandwidth and user storage) in gossip-based information diffusion paradigms. Since the epidemic spreading of information may easily saturate network resources (e.g., as a result of broadcast storms), a large array of techniques and approaches has been proposed (see, e.g., [22], [23] and references therein) in order to obtain an efficient use of resources in content diffusion processes.

These techniques mainly depend on the specific goal of the content diffusion process. In closed systems, in which nodes cannot leave a given area, the aim of gossip-based schemes is to achieve completeness, i.e., to deliver content to all users in the given area, in the most resource-efficient way [24].

The present work however, like the majority of realistic applications, focuses on scenarios in which nodes may join and leave [25] the area. These applications adopt different approaches to identify the set of users to which a given content should be delivered. An example is given by schemes for popularity-based peer-to-peer opportunistic content replication, e.g., for cooperative in-network content caching [26], [27], [28], [29], [30], where content must be delivered to a subset of nodes: those that requested such content. In applications in which the delivery delay plays a key role, (e.g., delivery of data related to unexpected events, or to potential safety hazards), the target population coincides with (a large fraction of) all nodes present in the given area at the time in which the hazardous or unexpected event takes place [22], [23].

In all such systems, control over content availability and persistence is achieved by involving in the scheme also nodes which are not among the set of requesters. Thus, a key performance tradeoff is between the amount of involved nodes (and thus resource utilization) on the one side, and content availability and likelihood of content persistence on the other. Several techniques for limiting content replication and/or the amount of involved nodes have been proposed, based on, e.g., a maximum lifetime or hop count of the content, among others [31]. In this respect, the specificity of FC schemes lies in associating content with a given spatial context, and in controlling the amount of nodes involved in the scheme by means of location-based criteria, such as the definition of the RZ borders.

Given the infrastructure-less and probabilistic nature of such distributed storage paradigms, a crucial issue is determining the conditions under which content persists for a significant amount of time in the given area. [17] proposes a model for the interval of the time during which the content floats. [10] introduces the criticality condition, i.e., a sufficient condition for the content to float indefinitely with very high probability, under various mobility models. [32] demonstrates the feasibility of FC (in terms of ability to sustain content persistence within the RZ for a given period of time) even in setups with sparse node distributions. [33] proposes a model for content persistence for outdoor pedestrian mobility over large open spaces, such as city squares. [34], [35] characterize the mean time to information loss in several scenarios, based on synthetic mobility and on measurement-based vehicular mobility traces. [36], [37] propose a modeling approach based on mean field theory and a stochastic susceptible-infected-recovered (SIR) epidemic model to evaluate content lifetime, and to determine

sufficient conditions for content persistence. Other works (e.g., [38], [39]) focus on how to engineer the replication and storage strategies in realistic settings in order to efficiently guarantee a given success probability within a given time range. Despite focusing on resource efficiency, none of these works investigates those issues arising when several different contents are exchanged among a same set of nodes, and thus the impact of their approach on the amount of content which can be supported in a same area by the FC scheme.

A few works propose techniques for coping with resource contention among different contents floating in a same set of nodes. [40] proposes a content-centric dissemination scheme. Its solution is based on a policy which sets the order of content exchanges on a contact between two nodes, and the probability for a node to drop a content in a way which tries to maximize the total delivery rate over a set of contents of different popularity. The paper [41] proposes a scheme in which, in scenarios with several different floating contents and overlapping RZs, users prioritize the contents to exchange based on measures such as RZ size, or total amount of users with content in each RZ. In [42], authors assume that there are multiple different content items present in the network and that replicating all of them using FC may lead to overloading the wireless network and/or the storage capacity of nodes. In order to avoid these issues, they propose a strategy which controls the number of copies of a particular content item within a RZ. None of these works however characterizes the limit performance of FC in terms of maximum number of contents which can be supported with a given minimum performance, as a function of system parameters, like we do in this paper.

An important application of gossip-based information diffusion protocols is distributed caching [43], [44], [45]. In VANETs, self-organized storage or cloudlets adopt vehicle-to-vehicle communications in order to retain information in a given area for a range of time, particularly in those conditions in which infrastructure is not available [46], [47], [48]. However, these works focus on such issues as the relationship between cache hit ratio and the amount of content redundancy, on communication overhead, on reliability, leaving open the key issue of the relationship between the performance of such caching systems and the limit capacity of the distributed storage system on which they rely. [34], [49] characterize the mean time during which information persists in a FC-based storage system. These works are based on simulations in various settings (such as highways [35], or city centers), and on an informal definition of FC-based storage. Most importantly, they still neglect the characterization of the capacity of such a distributed storage system. Similar to the present paper, [50], [51] consider scenarios where a few nodes are static, and act as "repositories" in order to improve content retrieval and persistence.

Storage capacity quantifies the amount of information that can be actually stored, as a function of the number and size of contents to be made available, and hence it is a fundamental descriptor of the usefulness of such systems. To the best of our knowledge, the only work in the literature which tackles a problem similar to the one we adderss in the present paper is [52]. Its authors consider a FC-based file storage scheme for vehicular nodes over a two-lane highway, in which the file is split into blocks, and erasure coding is used to enable recovery of the whole file. For such a specific setup, and for the time period in which all blocks keep floating in the RZ, the authors propose a simple upper bound for storage capacity, which corresponds to the right member of the inequality

TABLE 1: Main notation used in the paper.

Notation	Parameter	Unit
M	Data storage capacity of each node	$bits$
D	Total node density	m^{-2}
ψ	Fraction of moving nodes	m^{-2}
$\aleph$	Ratio of static over moving nodes	
γ	Intensity (mean number of floating elements per unit area)	m^{-2}
C_0	Mean capacity of the link between two nodes in contact	$bits/s$
L	Content size	$bits$
R	Replication Zone (RZ) radius	m
$i \in \{s, d\}$	Node label (s for static, d for moving)	
g_i	Mean contact rate per unit area involving static (resp. moving) nodes	$s^{-1}m^{-2}$
α	Linear flow	$s^{-1}m^{-1}$
τ_0	Transfer setup time	s
$\tau_s(\tau_d)$	Contact time static-moving (resp. moving-moving) nodes	s
$f_i(\tau_i)$	PDF of τ_i	

in our Equation (17) in Section 5. However, they completely omit the issue of characterizing the tightness of the bound as a function of system parameters. In addition, their approach is based on an ad-hoc approach which is only suitable for the considered road geometry and for the patterns of contact and mobility that the system geometry produces. As a result, their approach cannot be generalized to other settings, such as, e.g., the urban vehicular scenarios that we consider in this paper.

3 SYSTEM DESCRIPTION

In this section we briefly discuss the basic FC operations, and we define the key performance indicators of a FC system. In addition, we state the main modeling assumptions that we use in the analysis of the FC system storage capacity, and we introduce our notation.

3.1 Floating Content basic operation

We consider a plane over which two populations of nodes are present. The two populations comprise moving (or dynamic – we use the two terms as synonyms in this paper) and static nodes, respectively. Each node is equipped with a wireless transceiver and a data storage. Every node knows its position in space.

We say two nodes are *in contact* when they are able to directly exchange information via wireless communications.

At a time t_0 that defines the start of the FC system operation, a node (the *seeder*) generates a piece of content (e.g., a text message, or a picture)[3].

Within a region of the plane called *replication zone* (RZ) that contains the location of the seeder at t_0, whenever a node with content comes in contact with a node without it, content is exchanged. When a node moves out of the RZ, content is discarded.

When two nodes come in contact, a setup time τ_0 is required before content transfer, because nodes need to detect the presence of neighbors, verify the availability of content, and set up the transfer. A content transfer is successfully completed when the contact time and the channel capacity between the two nodes are such that content can be transferred in full, and the receiving node has free storage space for the content.

---

3. It is also possible to consider the case of multiple seeders that simultaneously receive a piece of content, e.g., through the cellular infrastructure.

Every node can exchange content (i.e., send or receive) with one node at a time. Moreover, nodes do not interrupt a content exchange with one node in order to start another exchange with a different node, unless the former exchange is completed, or the two nodes are not in contact any more. However, other modes of communications, e.g., broadcast, can be easily accounted for in our approach.

From this description, it emerges clearly that in FC, when proper conditions (in terms of user density and mobility, and of size of the RZ) are met, content *floats* (i.e., it persists probabilistically) in the RZ, even after the seeder has left it. In practice, content never floats forever, unless, e.g., one or more static nodes are present in the considered area, or it is possible to actively influence node mobility (such as with UAVs) to avoid content disappearance from the RZ.

At each contact, each content which is owned by only one of the two nodes, and for which there is enough memory to store it at the other node, has the same probability to be chosen for the transfer, regardless of which of the two nodes it resides on. Content exchanges between two nodes are unidirectional.

Limited modifications in the system operations can be easily accounted for in our approach, which can easily be expanded to include, e.g., bidirectional content exchanges, the effects of content broadcasting, or any specific priority scheme for contents to be exchanged, for example based on content relevance or popularity.

The system corresponding to one content item floating in its RZ according to the scheme described above is called a *Floating Element* (FE). In this work, we focus on systems composed by several floating elements.

With respect to the relative position of the RZ of each FE, we consider two types of systems. In *distributed floating systems*, the centers of the RZs of each FE are distributed on the plane so that RZs can partially overlap. This is typical of setups where each application or end user defines its own ZOI location. In *localized floating systems* instead, all FEs share a same RZ and ZOI. This second scenario allows ruling out the impact of randomness in RZ overlapping on the storage capacity of the system, and hence it allows investigating the maximum amount of information which can be stored in a given location (the RZ) via the FC paradigm.

3.2 Floating Content performance indicators

The goal of the FC paradigm is to ensure, through opportunistic replications, that the content item is delivered to a given target population of users. Hence, one of the main performance metrics for FC is content *availability* at time $t \geq t_0$, i.e., the mean fraction of nodes with content, at time t in the RZ. High values of availability in the RZ typically correlate with low likelihood of quick content disappearance, and with high probability of transferring content to nodes entering the RZ.

The specific definition of the probability of successful delivery (i.e., the *success probability* P_{succ}) varies according to the performance objectives induced by the application, and by the way in which the population of target users is identified.

In general therefore, the expression of P_{succ} is not only a function of content availability, of the geometry of the ZOI, and of user mobility patterns, but also of the specific application supported by the FC scheme.
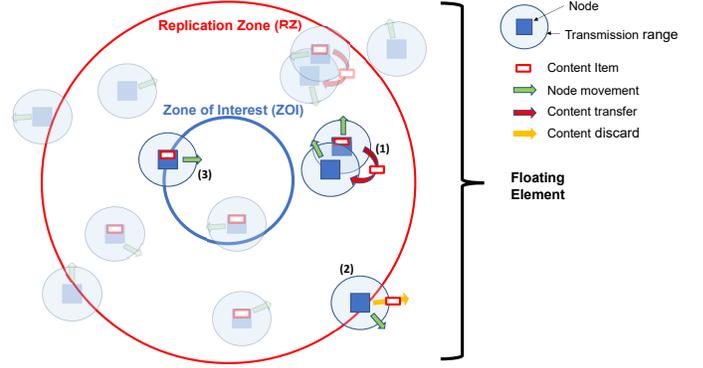

Fig. 1: Basic operations of the Floating Content scheme within a Floating Element. When two nodes are in range within the RZ, content gets replicated (1), thus persisting probabilistically in the RZ of the Floating Element. When a node leaves the RZ, it discards the content (2). To every node entering the ZOI (3) corresponds a content request, e.g., on behalf of a user application, and thus a read operation.
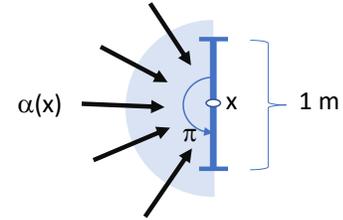

Fig. 2: Linear flow α is the rate at which nodes cross a segment of unit length centered at x, with a direction in the interval $(0, \pi)$.

3.3 Basic assumptions and notation

We assume that dynamic nodes move according to a stationary mobility model such that, at any time instant, nodes are uniformly distributed in space[4], with mean density ψD. We assume that static nodes are distributed over the plane according to a Poisson Point Process (PPP) with intensity $(1 - \psi)D$.

With τ_i, $i \in \{s, d\}$, we denote the node contact time, i.e., the duration of the time interval during which a static and a dynamic node (resp. two dynamic nodes) are in contact, and with $f_i(\tau_i)$ $i \in \{s, d\}$, we indicate the contact time PDF. With C_0 we denote the mean channel capacity between two nodes that are in contact.

Let g_i, $i \in \{s, d\}$, denote the mean rate of contacts per unit area between static and dynamic nodes, and among dynamic nodes, respectively.

We call *linear flow* the rate at which moving nodes traverse a segment of unit length in a same direction (e.g., from left to right of the segment), and we denote it by α. In order to derive an expression for α, let x denote a position in space and θ an angle, and let $\alpha(x, \theta)$ be the *angular node flow*, i.e., the rate of nodes moving in direction $(\theta, \theta + d\theta)$ across a small line segment of length ds centered at x and orthogonal to θ, divided by $ds \cdot d\theta$. The linear flow α is given by

$$\alpha = \int_0^\pi \alpha(x, \theta) d\theta. \tag{1}$$

In order to simplify notation, in what follows we assume that the mobility model is isotropic, hence, that α does not depend on the

---

4. We will discuss the impact of nonuniform user distributions later in the paper, through simulations.

position in space of the unit segment, nor on the specific direction of the node flow. However, note that our approach can be easily extended to more general, non-isotropic mobility models.

With $S_{i,d}$, $i \in \{s, d\}$ we denote the probability that a content transmission completes successfully during a contact with a static node, or between two moving nodes.

In what follows, without loss of generality, we assume that the RZ and the ZOI of a FE are circular and concentric, and that the radius of the ZOI is always strictly less than the radius of the RZ, R. For simplicity we assume that all RZs have a same RZ radius and all ZOIs have a same ZOI radius.

Moreover, all contents have the same size L bits, and M is the size in bits of the memory that can be used for FC in each node. Note however that our approach can be extended to contents of different size and to scenarios where nodes have different memory size, mainly at the cost of increasing the notation complexity.

In distributed floating systems, the centers of the RZ of each FE are taken to be distributed according to a PPP with intensity γ. In localized floating systems, the intensity γ is the ratio of the total number of FEs, over the RZ area.

4 A MEAN FIELD MODEL OF FLOATING CONTENT

In this section we introduce our mean field model of FC, present our main theorem, and discuss the validity of the mean field model over finite and infinite time horizons.

4.1 Content diffusion over finite time intervals

In FC, the set of nodes exchanging content within a RZ can be modeled as a system of interacting objects, in which interactions bring to changes in the state of the objects via content replication. When the number of these objects becomes large, the analytical performance study of such system becomes difficult, due to the exponential growth of the state space size.

Indeed, in order to model the dynamics of content diffusion and availability in our system, we focus on the temporal evolution of a set of parameters, related to different node populations in our system. If K is the total number of distinct contents in the system, for a RZ radius R at any time instant t we can associate to the system the state vectors $(N_s^1(t, R), ..., N_s^k(t, R), ...N_s^K(t, R))$, and $(N_d^1(t, R), ..., N_d^k(t, R), ...N_d^K(t, R))$ such that $N_s^k(t, R)$, $(N_d^k(t, R))$ is the number of static (resp. dynamic) nodes in the $k - th$ RZ with content at time t. Such a system can be assumed to evolve according to Markovian dynamics, and it can therefore be modelled as a Continuous Time Markov Chain (CTMC).[5] In order to derive meaningful insight into the performance of such system, in what follows we adopt a technique based on the *mean field interaction model* or *fluid limit* [53], hence on an approximate model of the interactions between nodes. The first approximation step of such approach is based on assuming that the following *homogeneous conditions* hold:

Definition 1 (**Homogeneous conditions**). We say that a floating system satisfies the homogeneous conditions if:

- at $t = 0$ the mean number of nodes per unit surface possessing a given content is the same for all contents;
- at any time instant, nodes possessing a given content are uniformly distributed within the RZ for that content;

5. The Markovian characteristics of the system dynamics descend from the Poisson distribution of nodes, the random choice of the exchanged content, and the independence assumptions that will be introduced later in the paper.

- the probability of a node to have a given content is independent from the probability of any other node to have the same content.

The homogeneous conditions assumption (equivalent to, e.g., the "well stirred" assumption in chemistry [54], and to the assumption of stochastic equivalence of nodes within a same class in [55]) allows deriving simpler expressions for the evolution of the main performance parameters of the system, at the cost of neglecting spatial inhomogeneities. Note however that our approach can be extended to account for spatial variations as well as for possible nonuniform seeding strategies (e.g., through the notion of node class, as in [55]), though at the cost of an increase in complexity of the analytic expressions.

In order to model the dynamics of content diffusion and availability in our system, we focus on the temporal evolution of four classes of node populations. The first and second class are composed by those static (resp. moving) nodes possessing a given content at a given time instant. The other two classes are composed by those static (resp. moving) nodes which are *busy*, i.e., exchanging content, at a given time. Note that the busy state is not associated with a given content, but with the fact that the node is involved in an exchange of content at a given time instant, and that each node can be part of both classes of populations at the same time.

Let $\bar{N}(R)$ be the mean number of nodes in a RZ. As we have assumed nodes to be uniformly distributed at any point in time, and all RZ to be of equal size, $\bar{N}(R)$ does not depend on time, and it is the same for every RZ. In order to apply the mean field approximation, instead of the number of (static or dynamic) nodes with content for each RZ at time t, we consider the ratio between these quantities and $\bar{N}(R)$, which is therefore the parameter used for normalizing the state occupancy in every RZ. Specifically, we consider the following parameters:

- For every radius R and content k, the fraction of static (resp. dynamic) nodes which possess k (henceforth denoted as the *availability* of content k) at time t, denoted with $a_s^k(t, R)$ (resp. $a_d^k(t, R)$).
- The fraction of static (resp. dynamic) nodes which are *busy*, i.e. exchanging content, at a given time t for a RZ radius R, denoted with $b_s(t, R)$ (resp. $b_d(t, R)$).

For ease of notation, in what follows we drop the indication of the dependency of variables on time t and RZ radius R, unless required by the context.

Parameters b_s and b_d are key in order to model the decrease of the rate at which contents get replicated successfully when the mean time required to transfer the exchangeable contents is comparable to or larger than the mean contact time. Indeed, in those conditions (which correspond to the conditions in which the system is approaching the maximum amount of information it can sustain) a significant amount of content exchanges do not take place or is delayed because nodes are still engaged in a content transfer at the time in which they come in contact with other nodes. Moreover, note that the busy state is not associated with a specific content, but with the fact that the node is involved in an exchange of content at a given time instant.

As a consequence of the homogeneous conditions and of the random scheduling of content transfers, for both static and dynamic nodes at any time instant the stochastic process of the fraction of nodes possessing a given content within the RZ for that content (i.e., of the *availability* of the given content) has the

same distribution for all contents, in both distributed and localized floating systems.

In what follows, for a given choice of RZ radius R, we indicate with a_i, $i \in \{s, d\}$, the mean availability at time t, averaged over all contents, for static and dynamic nodes respectively. The following result derives the PDF of the number of contents possessed by a node.

Lemma 1. *With the given assumptions, in a floating system (distributed or localized), the number of contents possessed by a static (resp. moving) node at time t, denoted as m_i, is distributed as a binomial $Bin(n, p)$, with parameters $n = \lfloor \gamma \pi R^2 \rfloor$ and $p = a_i$, and truncated in $\lfloor \frac{M}{L} \rfloor$.*

For the proof, please refer to Appendix A. The fact that the distribution is binomial should not be a surprise, given the independence assumption included in the homogeneous conditions, and given the Markovian assumptions necessary for the development of a mean field model. The truncation is necessary to account for the limit in the memory of end user devices, hence for the maximum number of contents that can be simultaneously stored at a node.

With $\bar{m}_i$ we denote the mean number of contents possessed by nodes.

On the occurrence of a contact event, a relevant parameter is the amount of *exchangeable* contents, i.e., of contents which are possessed only by one of the two nodes, and for which there is enough free memory at the receiving node. This parameter is key in determining the likelihood of a content to be exchanged on a contact event.

Let $i, j \in \{s, d\}$, and let x_{ij} denote the random variable modeling the number of exchangeable contents from a node i (static or moving) to a node j (static or moving) at time t and for a RZ radius R.

Lemma 2. *In a floating system (distributed or localized), the PDF of x_{ij} is given by the expectation, with respect to m_i and m_j, of a binomial PDF $Bin(n, p)$, with parameters $n = \lfloor m_i \rfloor$ and $p = 1 - a_j$, and truncated in $\lfloor \frac{M}{L} - m_j \rfloor$.*

For the proof, please refer to Section B in the Appendix.

Let $E[x_{ij}|M = \infty]$ be the mean number of exchangeable contents, for the case in which host memory is infinite. A key parameter for modeling content diffusion within a RZ is the probability of successful transfer of a *single* content during a contact of duration τ. The following result gives an analytical expression for it, for the considered system.

Lemma 3. *The probability of successful transfer of a single content from a node i (static or moving) to a node j (static or moving) during a contact is well approximated by*

$$S_{ij} = \frac{E[x_{ij}]}{E[x_{ij}|M=\infty]} \int_{\tau_0}^{+\infty} \min\left(1, \left\lfloor \frac{(\tau - \tau_0)}{\frac{L}{C_0}} \right\rfloor \frac{1}{E[x_{ij}+x_{ji}]}\right) f_i(\tau) d\tau \quad (2)$$

The lemma tells that the success of a transfer when there is infinite host memory is just the probability that the duration of a contact initiated in the RZ be larger than what needed to exchange the data (this is expressed by the term in the integral). Whereas, in case the memory is bounded to $M < \infty$, it is enough to rescale the success probability according to the average number of exchangeable contents.

In order to correctly model the dynamics of information exchange over time, another key parameter is the mean time spent exchanging contents during a contact, which varies according to whether only one or both of the nodes in contact are moving.

Lemma 4. *The mean time spent exchanging content during a contact between a static node and a dynamic node (denoted with T_s), and between two dynamic nodes (denoted with T_d), are well approximated by*

$$T_i = \int_0^{+\infty} \min\left(\tau, \frac{E[x_{ij} + x_{ji}]L}{C_0} + \tau_0\right) f_i(\tau) d\tau \quad (3)$$

The lemma expresses the fact that the average amount of time spent exchanging data is obtained by averaging the contact time, whose distribution conditional to a meeting in the RZ is f_i, under the constraint that data transfer duration is upper-bounded by the volume of data to exchange, as expressed in (3).

For the proofs of Lemma 3 and 4, please refer to Section C in the Appendix.

We observe that, in the spirit of the mean field approach, the derivation of these expressions is also based on a deterministic approximation, which neglects the stochastic nature of channel throughput and of the amount of content to transfer. In Section 6 we verify the impact of this as well as other approximations on the accuracy of the model.

The following result models the asymptotic dynamics of our system over finite time intervals, for large RZ areas and hence for a large amount of nodes involved in the process of diffusion of each content.

Theorem 1. *In a floating system (distributed or localized), for any initial condition $(a_s(0, R), a_d(0, R), b_s(0, R), b_d(0, R))$ with $b_s(0, R) = b_d(0, R) = 0$, for large R, the quantities (a_s, a_d, b_s, b_d) converge almost surely over any finite horizon to the solution $(\bar{a}_s, \bar{a}_d, \bar{b}_s, \bar{b}_d)$ (the mean field limit) of the following ordinary differential equations (ODEs):*

$$\begin{cases} \dfrac{d\bar{a}_s}{dt} = \dfrac{\bar{b}_s}{\bar{T}_s} \bar{a}_d (1 - \bar{a}_s) \bar{S}_{ds} \\[2ex] \dfrac{d\bar{a}_d}{dt} = (1 - \bar{a}_d)\left[\aleph \dfrac{\bar{b}_s}{\bar{T}_s} \bar{a}_s \bar{S}_{sd} + \dfrac{\bar{b}_d - \aleph \bar{b}_s}{\bar{T}_d} \bar{a}_d \bar{S}_{dd}\right] - \dfrac{2\alpha}{\psi DR} \bar{a}_d \\[2ex] \dfrac{d\bar{b}_s}{dt} = \dfrac{g_s}{(1-\psi)D}(1 - \bar{b}_s)(1 - \bar{b}_d) - \dfrac{\bar{b}_s}{\bar{T}_s} \\[2ex] \dfrac{d\bar{b}_d}{dt} = \dfrac{g_d}{\psi D} 2(1 - \bar{b}_d)^2 + \aleph \dfrac{d\bar{b}_s}{dt} - \dfrac{\bar{b}_d - \aleph \bar{b}_s}{\bar{T}_d} - \dfrac{4\alpha}{\psi DR} \bar{b}_d \end{cases}$$

$$(4)$$

with $\bar{b}_d \geq \aleph \bar{b}_s$, and with initial condition $(\bar{a}_s(0), \bar{a}_d(0), \bar{b}_s(0), \bar{b}_d(0)) = (a_s(0, R), a_d(0, R), b_s(0, R), b_d(0, R))$. With $\bar{S}_i, \bar{T}_i, i \in \{s, d\}$ we indicate the expressions in Lemma 3, respectively, with $\bar{a}_s$ and $\bar{a}_d$ instead of a_s and a_d.

For the proof, please refer to Section D in the Appendix. Note that $\forall t$, the constraint $\bar{b}_d \geq \aleph \bar{b}_s$ must be satisfied. I.e., the number of busy nodes which move cannot be inferior to the number of busy nodes which are static, as for every busy static node there is (at least) a busy moving node with which the static node is exchanging content.

This result states that the probability of observing a difference between any point of the trajectory of the given system and the

solution of the ODEs goes to zero as R (and hence the mean total number of nodes in each RZ) grows. That is, in the limit, the error made by considering a deterministic system characterized by $\bar{a}_i$, and $\bar{b}_i$, $i \in \{s, d\}$, instead of the actual system, goes to zero. Moreover, at any time t, for any R, variables $\bar{a}_i$, and $\bar{b}_i$ are the expected values of a_i, and b_i, respectively. Note that we consider $b_i(0, R) = 0$ as we assume that at $t = 0$ no node has initiated any content transfer yet.

Finally, note that our focus is on settings for which no strongly connected component exists in the network graph of static nodes (i.e., static nodes are out of reach of one another). This makes sense in those setups in which static nodes are deployed in a deterministic manner, e.g., to implement content seeding, to guarantee a rapid content diffusion and thus a minimum level of content availability. Nonetheless our approach, and the ODEs in Theorem 1 in particular, can be easily extended to account for direct exchange of content among static nodes. We observe however that the only effect of these direct exchanges is a more rapid diffusion of content among static nodes, without affecting the system performance once the diffusion transient is exhausted.

4.2 Quasi-stationary regime

In this work, we are mainly interested in the performance of the system after *enough time* has passed from the initial seeding of the content, i.e., at times in which the dynamics of initial content diffusion are exhausted. However, in a finite system there is always a non-zero chance of having a content item disappear from its RZ due to random fluctuations in the population of nodes which possess that content. The possibility of such an event is present in any finite floating system, which therefore for $t \to \infty$ inevitably tends towards the empty state. In our system in particular, this event may happen when there are no static nodes, or when a given content is not possessed by any static node at $t = 0$. As the CTMC of the system has an absorbing state (the empty state), its only equilibrium state is the empty system.

Therefore, in what follows we focus on the performance of the system in its *quasi-stationary* regime, i.e., at a time from content seeding which is *large enough* for the initial dynamics of seeding to be exhausted, and at the same time *small enough* for content not to have been absorbed yet. Indeed, this is the performance regime which is most relevant, as in any practical setting the amount of time during which a content should be stored and made available is not infinite (e.g., due to day/night patterns in vehicles and pedestrians mobility) but it is often long enough for the initial transient to have only a marginal impact on the overall performance.

The computation of an estimate of the time to content extinction based on the original CTMC is unfeasible due to the explosion in the number of the states which should be considered. In this section, under some mild assumptions, we derive a condition for the content to float for an indefinitely long amount of time (i.e., for the quasi-stationary regime to exist), as well as expressions for the main performance parameters at times from content seeding in which the initial transient effects no longer influence the system behavior. In the numerical section we will assess the accuracy of our model, and we will investigate the conditions under which the decay of the system towards the empty state has a significant impact on system performance.

Given the existence of the absorbing state for the original system (which we denote as $\mathcal{S}$), in order to apply the mean field approach and compute the performance parameters for the time before absorption we consider a slightly different system (denoted with $\mathcal{S}'$). This new system is obtained from the original one by assuming that for each content, when the system is empty, content is re-seeded. Specifically, we assume that, when a content is absorbed in $\mathcal{S}$, in $\mathcal{S}'$, at rate ϵ (constant and independent on any parameter of the system) a management function selects one node (or a few nodes) at random in the RZ, and injects the content in that node(s). Note that, in the time period between content seeding and content absorption, $\mathcal{S}$ and $\mathcal{S}'$ are indistinguishable. This is confirmed by the following result, relative to the mean field limit over finite time intervals:

Theorem 2. *For any $\epsilon \geq 0$ the mean field limit of systems $\mathcal{S}$ and $\mathcal{S}'$ are the same.*

Proof. This result is due to the fact that the rate ϵ, being constant with RZ radius R, vanishes with increasing RZ radius at a rate R^{-2}, and is hence negligible in the mean field approximation, in which all neglected terms are $\mathcal{O}(R^{-2})$. $\qquad\square$

Theorem 2 implies that the content seeding process has a vanishing impact on system performance in the mean field regime, as modeled in the previous section. As a result, the mean field approximation results of Theorem 1 hold also for $\mathcal{S}'$. This result allows exploiting Theorem 1 in order to derive a mean field approximation for the stationary state, which is only defined for $\mathcal{S}'$. Then, thanks to the fact that the two systems are indistinguishable at times between content generation and absorption, in a regime in which content absorption is relatively infrequent (i.e., in a regime where the quasi-stationary state is defined) the mean field approximation for the stationary state of $\mathcal{S}'$ is also an approximation of the quasi-stationary regime of $\mathcal{S}$. In Section 6 we assess the accuracy of our approach.

The following result establishes a relation between the stationary state of $\mathcal{S}'$ and the steady-state solutions of the problem in (19).

Theorem 3. *For any $\epsilon > 0$, for large R the steady-state solutions of Equation (19) are an approximation of the state distribution of $\mathcal{S}'$ for $t \to \infty$.*

Proof. The CTMC associated to $\mathcal{S}'$ has no absorbing states, and its state diagram presents no cycles. Hence it belongs to the class of reversible stochastic processes [56]. Therefore, from Theorem 1.2 in [57] it follows that the stationary behavior of the CTMC associated to $\mathcal{S}'$ is completely determined by the solutions of the ODEs (19). $\qquad\square$

The values of the variables $\bar{a}_i$, and $\bar{b}_i$, $i \in \{s, d\}$ at the steady state are therefore the solutions of the following system of equations derived from the ODEs of Theorem 1:

$$\frac{\bar{b}_s}{\bar{\bar{T}}_s} \bar{a}_d (1 - \bar{a}_s) \bar{S}_{ds} = 0 \tag{5}$$

$$(1 - \bar{a}_d) \left[\aleph \frac{\bar{b}_s}{\bar{\bar{T}}_s} \bar{a}_s \bar{S}_{sd} + \frac{\bar{b}_d - \aleph \bar{b}_s}{\bar{T}_d} \bar{a}_d \bar{S}_{dd} \right] - \frac{2\alpha}{\psi D R} \bar{a}_d = 0 \tag{6}$$

$$\frac{g_s}{(1 - \psi)D}(1 - \bar{b}_s)(1 - \bar{b}_d) - \frac{\bar{b}_s}{\bar{\bar{T}}_s} = 0 \tag{7}$$

$$\frac{g_d}{\psi D} 2(1 - \bar{b}_d)^2 - \frac{\bar{b}_d - \aleph \bar{b}_s}{\bar{T}_d} - \frac{4\alpha}{\psi D R} \bar{b}_d = 0 \tag{8}$$

with the constraints

$$\bar{b}_d - \aleph \bar{b}_s \geq 0 \tag{9}$$

$$0 \leq \bar{a}_i \leq 1, \qquad 0 \leq \bar{b}_i \leq 1, \quad i \in s, d \tag{10}$$

$$\tag{11}$$

By solving the system of equations directly, it can be verified that the solutions which satisfy constraints (9) and (10) are only two. Specifically, if the system starts from the empty state (i.e., one in which a_s and a_d are both zero), the system persists in the empty state. Consider instead the case in which the system starts from a non-empty state (i.e., a state in which at least one among $\bar{a}_s(0)$ and $\bar{a}_d(0)$ are nonzero). Let

$$y = \aleph \frac{\bar{b}_s}{\bar{T}_s} \bar{S}_{sd}$$

$$z = \frac{\bar{b}_d - \aleph \bar{b}_s}{\bar{T}_d} \bar{S}_{dd}$$

$$x = y + \frac{2\alpha}{\psi DR} - z$$

then the system converges to the unique feasible solution given by

$$\bar{a}_d = \frac{-x + \sqrt{x^2 + 4yz}}{2z} \tag{12}$$

only if the *criticality condition* $x \geq 0$ is satisfied, with $\bar{b}_d$, $\bar{b}_s$ given by Equation (7) and (8). Otherwise, the system converges to $\bar{a}_d = 0$. Thus, as can be seen from the study of the gradient of the system of equations, if the criticality condition is satisfied, any trajectory starting from a nonempty state converges to the unique nonempty feasible steady-state solution.

The significance of the mean availability of content over moving nodes in steady state derives from the fact that, in the mean field limit, it coincides with the mean availability of those nodes that enter the ZOI, and is therefore key in determining the success probability.

5 THE STORAGE CAPACITY OF FLOATING CONTENT

As we already noted, the FC paradigm can be seen as a way to implement, through opportunistic replications, a distributed information storage service, enabling *probabilistic* content persistence and retrieval in a limited area, typically with no direct support from infrastructure (except possibly for the initial seeding of content). In this section we characterize the *storage capacity* of a FC storage system, i.e., the maximum expected amount of information which can be stored probabilistically. We first derive a powerful model of FC storage based on information theory, which allows the computation of the storage capacity of a FC system. Then, driven by the results of the information theoretic model, we obtain an extremely simple explicit expression for the FC storage capacity in saturation conditions.

5.1 An information theoretic storage model of FC

In order to derive a model for the FC storage capacity in a way which is analogous to classical information storage systems, we start by considering a single floating element belonging to a floating system (be it localized or distributed). For this system, we define the read and write operations as follows.

The **write operation** consists in the initial seeding of the content within its RZ, with the goal of enabling content to persist probabilistically even after the transient of content diffusion has passed. From Section 4.2 we know that a seeding strategy which attributes each content to at least one node within the RZ enables the system to converge to a nonempty steady state, though in finite systems more conservative seeding strategies are often necessary in order to decrease the likelihood of content disappearance during the initial transient of content diffusion.

In order to define the **read operation**, we recall that the main goal of FC is to deliver the content to a given population of users, in a probabilistic fashion, by proactively populating the local memory of the target hosts in a distributed, collaborative manner, based on opportunistic content replications between the target hosts and all the other users in the RZ. In this context, every content delivery is a read operation, as it makes available to the node the stored content. For instance, the ZOI might correspond to the location of a movie theater, and the content item to a movie trailer which we assume will start being requested when users enter the theater. Therefore, to each node entering the ZOI corresponds a content request, and hence a read operation, which is considered as failed if the node does not possess the given content. In what follows, we consider the case in which the ZOI radius is much smaller than that of the RZ, and the RZ radius large enough for content availability to be well approximated as uniform across the whole RZ. In this case P_{succ} is well approximated by the mean content availability within the RZ.

In information theory, a storage system is modeled as a communication channel [58], [59]. In this view, information is transmitted over the channel through a write operation, and it is received through a read operation. According to the operational definition given by Shannon [59], the channel capacity is the largest amount of bits *per channel use* at which information can be sent on the considered channel with arbitrarily low error probability. For the specific case in which FC is used as a storage technology, a channel use is the operation of setting a content to float in the floating area, i.e., a write operation for a content of size $L < M$ (remember that M is the node memory size). Let us denote as *stored information* of a FE the mean amount of information which can be recovered (i.e., read). Then, similarly to communication channels, if we consider the maximum of this quantity over all channel uses (and therefore content sizes), we have the following definition of storage capacity of a floating element:

Definition 2 (Storage capacity). The *storage capacity* of a floating element with radius R is the maximum of the stored information, over all content sizes $L \leq M$.

In what follows, we consider the floating element and its floating system to be in a stationary state. With $\bar{a}_d(R, \gamma, L)$ we denote the mean field limit of the availability of moving nodes for the stored content, for a RZ radius R. With $P_{succ}(R, \gamma, L)$ we denote the success probability at the mean field limit, which we assume is a monotonically increasing function of $\bar{a}_d$.

Theorem 4. *In a floating system in quasi-stationary regime, the mean storage capacity of a floating element is*

$$C_{FE}(R, \gamma) = \max_{L \leq M} L P_{succ}(R, \gamma, L) \tag{13}$$

Proof. At the mean field limit, every node in the RZ has the same probability $\bar{a}_d(R, \gamma, L)$ to possess the content, independently

from its position. In the communication channel model of storage systems, a FE can be modeled as a *packet erasure channel* [60], with packet size equal to the size of the content, L, and packet erasure probability $1 - P_{succ}$. In such a model, every channel use is a write operation, consisting in setting a content of size L bits to float in the floating region, and erasures in the channel derive from the fact that the read operation is not deterministic. The amount of bits which can be recovered, on average, on a single channel use is hence $LP_{succ}(R, \gamma L)$. The maximum of this quantity over all channel uses, and hence over all content sizes $L \leq M$, is the storage capacity of the system. $\qquad \square$

In a floating system composed by more than one floating element, the overall amount of information stored is also a function of how the RZs overlap, hence of both FE intensity γ and RZ radius R. The following result gives the capacity per unit area of a floating system for a given FE intensity, and a given RZ area.

Corollary 1 (Area capacity of a floating system). In a floating system (localized or distributed) in the stationary regime, the mean storage capacity per unit area is given by

$$C_{FS}(\gamma, R) = \gamma \max_{L \leq M} LP_{succ}(R, \gamma, L) \qquad (14)$$

Proof. The system can be modeled as a set of γ parallel, independent packet erasure channels per unit area, one per distinct content. The independence between the channels holds at the mean field limit and it derives from the "propagation of chaos" result, by which at the mean field limit for each user the probability to have a content is independent from the probability of having another content. For each content, the expression of the capacity of the associated packet erasure channel is given by Theorem 4. $\qquad \square$

We observe that the mean amount of information possessed by a node is given by $\max(M, \gamma \pi r^2 L \bar{a}_d(R, \gamma, L))$. Hence, when the success probability is defined as the probability to transit in a given location of the RZ with a copy of the content, *the mean amount of content possessed by a node is equal to the product of the area capacity and of the area within the transmission range of the node*, and upper bounded by the node storage space M. It is therefore equal to the minimum between the amount of information which can be read within the node transmission range, and the available storage at the node.

In a floating system, the RZ radius modulates the average amount of users, hence of system resources, dedicated to a given FE, while the FE intensity tells how many floating contents on average are sustained by the system per unit area. In what follows we are interested in how to modulate these parameters in order to maximize the amount of information per unit area stored in a floating system. We have therefore the following optimization problem:

Problem 1 (Maximum area capacity of a floating system).

$$\max_{R, \gamma, L} \quad \gamma LP_{succ}(R, \gamma, L)$$

subject to :

Equations (2), (3)

Equations (5), (6), (7), (8)

Constraints (9), (10)

$$R \geq 0 \qquad (15)$$

$$1 \leq L \leq M \qquad (16)$$

$$0 \leq \gamma \leq \frac{DM}{L} \qquad (17)$$

Constraints (5) to (10) allow deriving the expression of the steady-state variables $\bar{a}_i$, and $\bar{b}_i$. Constraint (17) derives from the fact that an upper bound to the intensity γ, hence to the average number of different contents floating in a given area, is given by the average number of nodes present in that area, multiplied by the number of contents which each node can store. As a consequence, an upper bound to the amount of stored information per unit area is DM, and it corresponds to the case in which, for each content, a single copy exists in the system, so that the system has no redundancy. As for the content size L, we have the following result:

Proposition 1. If (γ^*, R^*, L^*) is a solution of Problem 1, then $L^* = 1$.

This result derives from the fact that, holding constant everything else, the lower the content size, the lower the amount of contact time wasted in content transfers which do not complete due to the end of contact time. Moreover, the lower the content size, the higher the amount of information which each node can store in its finite memory, as the memory size is not always an exact multiple of content size. As a consequence, Problem 1 becomes a maximization problem over R and γ only, with content size equal to its minimum value $L = 1$. [6]

Finally, it can be easily shown that, for any choice of the other system parameters, there exists always a RZ radius beyond which availability decreases monotonically with increasing R. Summing up, despite Problem 1 is nonconvex and nonlinear, it is function of two variables over finite intervals, hence can be solved efficiently by brute force approaches.

5.2 A simplified model for capacity under saturation

As we will see in numerical results, floating systems can be in one of two regimes with respect to the amount of injected information. In what can be called the *linear regime*, stored information increases linearly with injected content, and availability is very close to one. In the *saturation regime* instead, the available bandwidth (or equivalently, the contact time available for content exchanges) becomes the system bottleneck, since it is fully used to transfer content. That is, on a contact between two nodes, the likelihood that the two nodes terminate the exchange before the end of the contact is very low, and possibly the contact duration is not sufficient to transfer all contents. In a system with infinite host memory, bandwidth (paired with contact duration) thus becomes

---

6. Note that in real systems, in which per-content communication overhead is non negligible, the optimal content size is not the minimal one ($L = 1$). For those systems however, the capacity computed in the ideal case represents an upper bound to what achievable when overheads are accounted for.

the storage capacity limiting factor. Our mean field model of FC capacity shows that the transition between the two regimes is quite sharp, so that FC performance can be well approximated by considering each of them separately.

In order to derive a first-order approximate expression of the capacity in saturation conditions, we assume contact duration between static and moving nodes (resp. between moving nodes) to be constant and equal to the mean amount of time taken by a node to cross a distance equal to $2r$ (where r is the node transmission radius), given by $2r/v_i$ where v_i is the mean node speed for the contact between static and moving nodes, and the mean relative node speed for contacts between moving nodes. We further assume that (since the system is in saturation) increasing the amount of injected contents does not change the performance of the system.

Given these assumptions, capacity in saturation conditions can be approximated as the maximum amount of information that can be transferred upon nodes contact, given by

$$C_{FS}^{sat} = C_0 2r \left(\frac{(1-\psi)}{v_s} + \frac{\psi}{v_d} \right) \tag{18}$$

The results obtained with this approximation will be discussed in the next section, together with those obtained with the information theoretic model. Numerical results will show that the simplified model provides quite close an approximation of the floating system capacity in saturation conditions predicted by the detailed model, and estimated by simulations.

It is important to stress that the derivation of this extremely simple and useful result was possible only thanks to the insight into the floating system behavior provided by the accurate mean field model developed in this paper.

6 NUMERICAL ASSESSMENT

In this section, we evaluate numerically the accuracy of our models by means of simulations, and we characterize the storage capacity of a floating system as a function of the main system parameters and of node mobility.

We assume nodes move according to the Random Direction Mobility Model, with reflection at the boundary of the simulation area. When two nodes are in contact, we assume the channel data rate is constant over time and equal to 10 Mb/s. Unless otherwise specified, nodes have a transmission radius of 30 m, and they have unlimited memory. For localized floating systems, the simulation area is a square of side 500 m, the RZ has a default radius of 100 m and it is located at the center of the area.

At the beginning of each simulation run, nodes are distributed uniformly at random. In order to minimize the probability of content loss during the content diffusion transient, all nodes within the RZ for a given content possess it at start time. When node memory is finite, the set of contents possessed by each node at start time is a random subset (different for each node, and of total size equal to the node memory) of all those contents in whose RZ the node is located. Simulated time has been divided into equally sized slots, and their duration has been chosen so as to minimize the effects of quantization in time on accuracy of simulations, and in particular on the errors in detecting when two nodes are in range or when a node is within a given RZ. Content size has been set in such a way to have the start and end of time slots coincide with the start and end of content transfers. Specifically, it has been set to 5 Mb in setups with node speed of 0.2 m/s, and to 100 kb for those with a node speed of 10 m/s. Simulations have been run
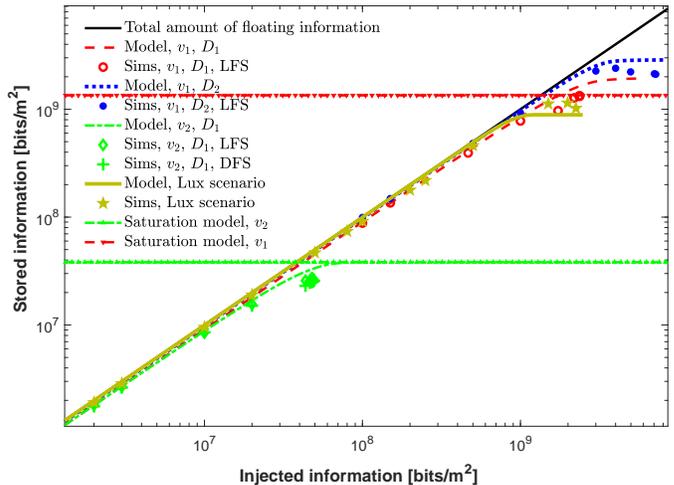
Fig. 3: Information stored per unit area in a distributed (DFS) and localized (LFS) floating system versus injected information per unit area. $v_1 = 0.2 \ m/s$, $v_2 = 10 \ m/s$, $D_1 = 800 \ users/km^2$, $D_2 = 4000 \ users/km^2$. Simulation are with a 95% confidence interval of at most 0.35%.
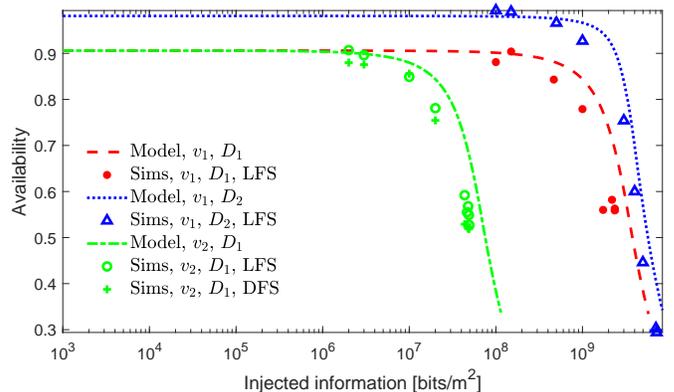
Fig. 4: Mean availability in a distributed (DFS) and localized (LFS) floating system versus injected information per unit area. $v_1 = 0.2 \ m/s$, $v_2 = 10 \ m/s$, $D_1 = 800 \ users/km^2$, $D_2 = 4000 \ users/km^2$. Simulation are with a 95% confidence interval of at most 0.35%.

for a duration of 10000 time slots, which proved to be sufficient to observe the system out of any transient, and data affected by transient effects have been discarded.

6.1 Baseline

In a first set of simulations, we considered scenarios with only dynamic nodes, and we measured the amount of information stored per unit area in both distributed and localized floating systems, as a function of the *injected information per unit area*, i.e., of the product of the content size times the mean number of contents per unit area which persist in the system for the whole duration of the simulation (which coincides with the intensity γ, when the resulting density of floating contents can be sustained by the system according to conditions (5) to (10)). As Figs. 3 and 4 suggest, two regimes can be observed. For low amounts of injected information and with infinite host memory, resource contention among different floating contents is weak, as the mean contact time is larger than the mean amount of time required
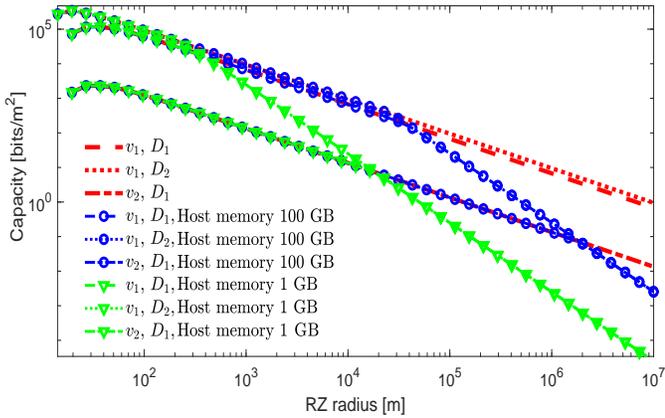
Fig. 5: Maximum storage capacity of a floating system over FE intensity, as a function of RZ radius R, for different configurations of host memory, and with content size $L = 1$. $v_1 = 0.2\ m/s$, $v_2 = 10\ m/s$, $D_1 = 800\ users/km^2$, $D_2 = 4000\ users/km^2$.

to exchange contents. Therefore, each FE performs almost as if in isolation. Indeed, mean availability remains constant with increasing injected information, while stored information grows proportionally to it. For larger values of injected information, the effects of contention (mainly, on contact time) start to appear, the mean content availability decreases, and the stored information saturates. As expected, decreasing the node speed (and therefore increasing the mean contact time) and increasing the node density (and therefore the rate of contacts between nodes, and the opportunities for content replication) have the effect of increasing the saturation value of the stored information, which coincides with the capacity of the system.

As Figs. 3 and 4 show, the estimates of mean availability and of the amount of stored information derived with our mean field based approach are accurate across different values of node density, of injected information, and of mean node speed. When the amount of injected information gets close to the maximum amount which can be sustained by the system, our mean field model yields slightly optimistic results, due to the difference between finite systems and their mean field limit. Indeed, as the plots show, increasing the mean number of nodes in the RZ improves accuracy.

Specifically, a reason for such discrepancy is that while contact events are uniformly distributed within the RZ, in finite systems, nodes with content are slightly more numerous around the center of the RZ [10]. As a consequence, there are overall less opportunities for content replications within a RZ (e.g., at the border of the RZ). This is also the reason for the slightly more pessimistic simulation results of distributed floating systems with respect to localized ones. Indeed, in systems with distributed RZs, and particularly for low densities of FEs, much of the overlapping involves mainly the border regions of each RZ. As shown in the figures, such difference tends to decrease as the FE intensity grows.

Another reason for the observed discrepancy is the effect of random fluctuations in a finite population of nodes. Indeed, when the system gets close to those conditions in which the steady state solutions of the ODEs in Theorem 1 do not satisfy the constraints in Equation (9) and Equation (10) (that is, the conditions in which contents cannot persist in their RZs), the effects of even small random perturbations in a finite population of nodes with

content gets amplified, because of resource contention and the consequent loss of efficiency of the content replication process. This brings to an average decrease in content availability, and to the difficulty in achieving those values of maximum injected information forecasted by the model.

In Fig. 3 we also report the predictions of storage capacity in saturation conditions obtained with the simplified approximate expression (18). We can see that predictions are extremely accurate, obviously only in the range of validity of the approximation, i.e., under saturation. However, before saturation, stored information coincides with injected information, so that it is possible to approximate the curves obtained from the information theoretic model with two lines, one where stored and injected information coincide, and one with constant stored information, at the level of the saturation capacity.

One of the key parameters affecting the mean area capacity of a floating system is RZ radius, which is related to the total amount of node memory and content exchanges dedicated to storing probabilistically a single content. Fig. 5 shows the maximum area capacity of a floating system (Problem 1) as a function of RZ radius. As the figure shows, the solution of Problem 1 is achieved for values of RZ radius only slightly larger than the minimum values below which contents do not persist in the RZ. For lower values of RZ radius, content availability, hence maximum storage capacity, decrease rapidly, as smaller RZs imply less opportunities for contents to replicate. For values of RZ radius larger than the optimum instead, the benefits of a larger RZ (in terms of a larger amount of time spent in the RZ by nodes, bringing more opportunities for content replication, hence higher availability) are offset by the fact that a much larger amount of nodes is involved in content replication, so that the marginal utility of adding more users to each RZ is negative. By further increasing the RZ radius, the maximum stored information reaches a regime where host memory limits start to affect system performance, further decreasing the marginal utility of adding more users to each RZ.

As we have seen, for a given content size, node density and RZ radius, when the density of contents seeded in a floating system is larger than the maximum amount which can be sustained by the system in steady state (i.e., if the system operates in conditions in which there is no feasible solution to Problem 1), according to our model the system converges to the empty state. In practical terms, as Fig. 6 shows, this implies that a process of content disappearance from the system is initiated, and it continues until the contents which remain in the system can be sustained in the mean field regime. Similarly, when host memory is finite, if the system is seeded with a higher number of contents than those which can be stored in host memory, the total number of floating contents decreases rapidly during the initial transient of content diffusion, until it coincides with the maximum number of contents which can be stored in host memory.

6.2 Impact of static nodes

In the FC scheme, a key role is played by node mobility. On the one side, mobility induces volatility of content, since nodes that exit a RZ are allowed to discard the content associated to that RZ (and we assume they do so in the present paper). One consequence of such volatility is an increase of the amount of resources (in terms of user memory, and of rate of content replications) required to effectively store a given amount of information. On the other
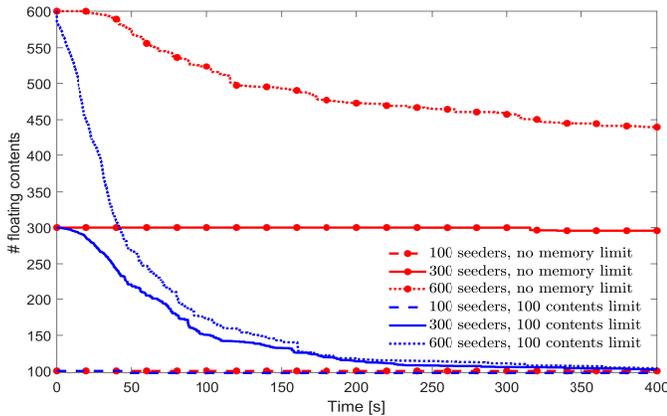
Fig. 6: Average over 20 simulation runs of the number of floating contents over time, in a localized floating system, for different initial number of seeders. Node speed is equal to 10 m/s, and node density is 800 $users/km^2$.

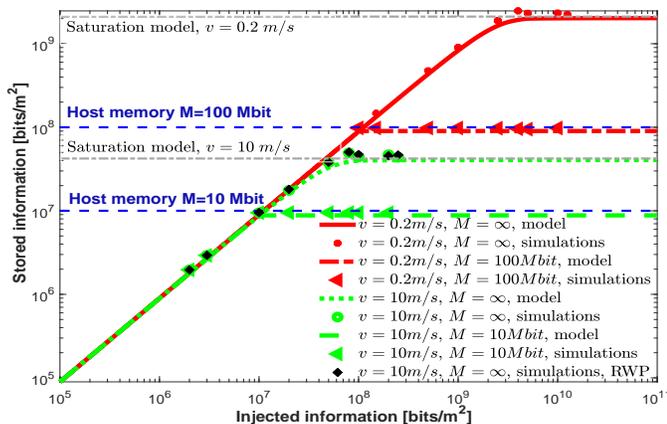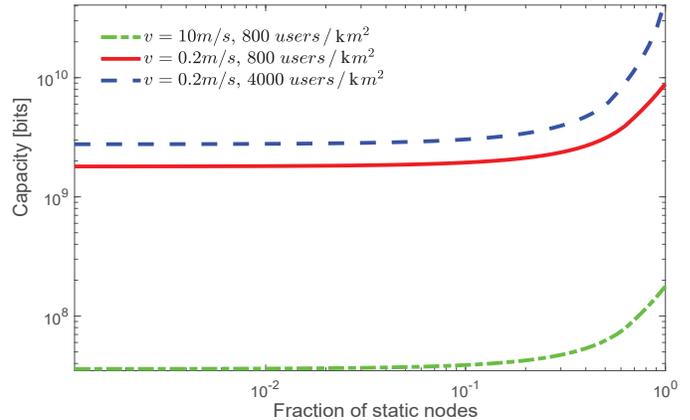Fig. 8: Information stored per unit area in a distributed (DFS) and localized (LFS) floating system versus injected information per unit area, with 15% of static nodes. Simulation are with a 95% confidence interval of at most 0.35%.
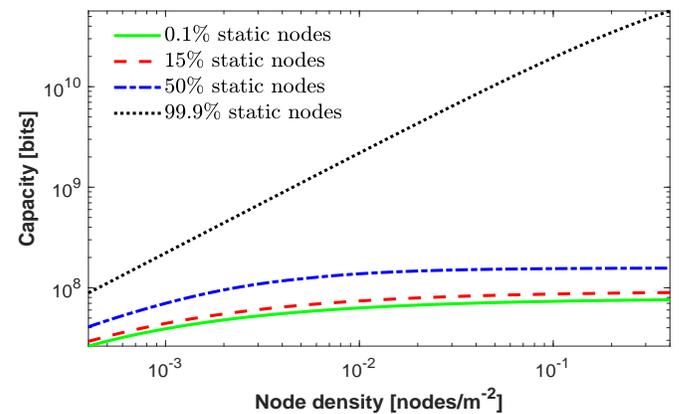

Fig. 7: Information stored per unit area in a distributed (DFS) and localized (LFS) floating system versus injected information, with 15% of static nodes, for the Random Direction as well as the Random Waypoint (RWP) mobility model. Simulations are with a 95% confidence interval of at most 0.35%. Node density is 800 $users/km^2$.

Fig. 9: Capacity of a distributed (DFS) and localized (LFS) floating system versus node density, for different percentages of static nodes. $v = 10m/s$, tx radius $30m$, no memory limit, $R = 100m$.

side, mobility and the induced pattern of contacts allow content to replicate and thus persist in the RZ, and to be delivered to the target users. The question thus arises of what is the impact of that portion of nodes which are not mobile on the performance of a FC probabilistic storage system.

To address this issue, we performed a set of experiments in which a varying fraction of nodes is static. Fig. 7 shows the amount of stored information in a floating system, when a portion of the nodes is static. Also in this new set of experiments, our analytical approach proves very accurate across a variety of settings and for different choices of system parameters. By comparison with Fig. 3, it emerges clearly that the presence of static nodes allows at least a fraction of contents to persist deterministically in the RZ. As a consequence, in a system with infinite host memory and at the mean field limit, there is no upper bound to the amount of injected contents which can be kept floating. However, similarly to the case with no static nodes, by increasing the injected information (e.g., by increasing the amount of contents injected while keeping constant content size), the amount of stored information reaches a saturation value.

The point of saturation corresponds to a value of injected

information at which the mean amount of time spent exchanging content during a contact is equal to the mean duration of a contact. Indeed, as the plots show, to a higher node speed corresponds a lower value of stored information at saturation (i.e., of capacity). As these results suggest, when host memory is limited, the capacity of the floating system is the minimum between the saturation value without memory limits, and the host memory.

Also in this case that includes static nodes, we report the predictions of storage capacity in saturation conditions obtained with the simplified approximate expression (18), which are again extremely accurate.

The impact of node speed and therefore of the mean contact time is also clearly visible from Fig. 8, which shows how capacity varies as a function of the percentage of static nodes in the system. As the plots show, when the fraction of static nodes is small, their effect on system capacity is negligible. When the fraction of static nodes is significant, however, the mean duration of a contact increases, and therefore, as observed already, so does the value of saturation of the amount of stored information.

Another key factor affecting the capacity of a floating system is node density, which determines the amount of resources involved in the process of information storage. As shown in Fig. 9 and Fig. 10, by increasing node density, capacity increases (as ex-
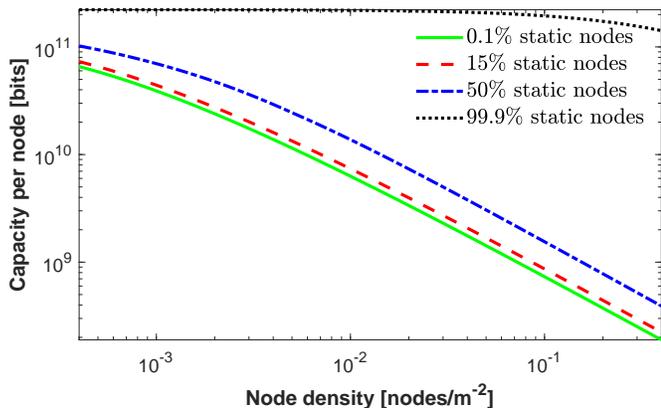
Fig. 10: Capacity per node of a distributed (DFS) and localized (LFS) floating system versus node density, for different percentages of static nodes. $v = 10m/s$, tx radius $30m$, no memory limit, $R = 100m$.
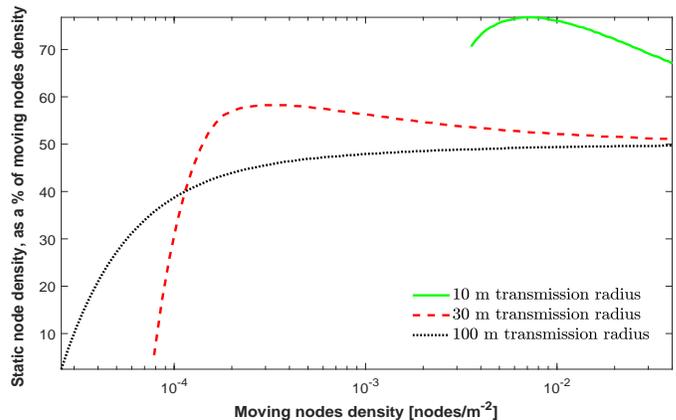

Fig. 11: Density of static nodes required, in a floating system without content exchange among moving nodes, to achieve the same capacity as in the case in which content is exchanged among moving nodes and no static nodes are present. The density of static nodes is expressed as a percentage of the density of moving nodes. $v = 1\ m/s$, no memory limit, $R = 100\ m$.

Fig. 12: The area of Luxembourg City considered in the simulations. The circle denotes the RZ, with a radius of 800 m.

pected) until it reaches a saturation value which is function of the fraction of static nodes in the system. The initial increase is due to the fact that by increasing node density, the rate of contact between nodes increases (quadratically, for the considered mobility model). However, for scenarios with a majority of moving nodes, an upper bound to the amount of exchanges which can take place at the same time is given by the fact that, if a node is busy exchanging contents with another one, it is not available for other exchanges until the current exchange ends. At high node densities, this decreases the likelihood for a node that just came in contact with another node to successfully engage in a content exchange.

In these settings, the reason for which the scenario with mainly static nodes performs better does not reside only on the fact that mean contact duration is larger. When dynamic nodes are only a small percentage of the total nodes in the scenario, as the rate of contacts between static nodes is zero, the overall contact rate is drastically reduced, and the saturation effects described above start to appear for much larger values of node density. Finally, note that our model does not account for the bandwidth saturation which might take place at high densities. Indeed, for large enough node densities (or equivalently, large enough transmission rates) nodes form clusters, thus boosting the likelihood of content persistence and successful replication, further improving capacity beyond the bounds given by our model. However, as already stated, as in those conditions store-carry-and-forward is not any more the main mode of communication, more efficient paradigms of communications and content diffusion than FC are available.

One of the main drawbacks of FC as a probabilistic information storage system is its reliance on the user's willingness to make resources available for the FC scheme. A possible alternative implementation of such a system may consist in delegating to a set of static nodes (*totems*) the task of content dissemination to moving nodes. This approach would avoid dealing with incentives for cooperation, at the cost of deploying dedicated infrastructure for spatial information storage and diffusion. For instance, *road-side units* (RSUs) could play the role of totems in a vehicular communication scenario. In Fig. 11 we plot the density of totems required to achieve the same capacity as in the case in which content is exchanged among moving nodes without any infrastructure. Reported curves start at the density of moving nodes below which the content does not float. As the figure shows, a very large amount of totems would be required to achieve the same capacity

achievable with opportunistic content replications among dynamic users. In practice, the number of totems should be comparable with the number of dynamic users. E.g., one totem should be deployed every two or three users, in the cases considered in Fig. 11. The cost of such infrastructure would clearly be unsustainable for any operator in any dense scenario like the previously mentioned vehicular communication scenario.

6.3 Impact of non-uniform mobility models

The results we discussed so far were derived using the random direction mobility model, which enjoys the property of generating a uniform distribution of users, as requested by our modeling assumptions. However, real mobility patterns are more complex, and they do not share the above property, thus violating one of the assumptions used for modeling.

In order to validate the robustness of the predictions generated by our model, we now discuss results obtained with two other mobility models. The first one is the well-known random waypoint (RWP) mobility model, that is known to generate clustering of

users in the center of the simulated area. The second one is a more realistic mobility model, generated using the SUMO (Simulation of Urban MObility) tool [61], considering measurement-based mobility traces of the the city of Luxembourg [62]. We feed SUMO with vehicular traffic data to produce realistic traces of vehicle movements over the road grid of the city. Specifically, we have considered a square area of side 4 km near the centre of Luxembourg City (see Fig. 12), and the vehicular traces relative to the time interval going from midnight to 3 AM, so as to observe a low node density condition. We have considered a RZ radius of 800 m, a transmission radius of 240 m (such as the one achievable by DSRC [63]), and a data rate of 10 Mb/s. An average of 22.02 cars were present in the RZ during the given time interval, exiting the RZ at a mean rate of 0.105 cars per second, with a mean contact time of 37.5 s and generating an average of 0.959 contacts per second within the RZ.

Results in the case of the RWP mobility model are reported in Fig. 7, assuming the presence of 15% of static nodes. Our modeling approach can be seen to be quite accurate in spite of the nonuniform user distribution generated by mobility.

Results in the case of the SUMO mobility model over the streets of a portion of Luxembourg city are reported in Fig. 3, for moving nodes only. Also in this case, results can be seen to be quite accurate in comparison to model predictions generated with the corresponding parameters.

We can conclude from these experiments that our model is quite robust to variations of the node mobility pattern, and in particular, that it performs well on realistic mobility patterns which generate non-uniform user distributions.

7 CONCLUSION

In this paper, we have developed a first analytical model to study the amount of information which can be stored in opportunistic distributed edge storage systems such as Floating Content. Our approach enables a first order characterization of the relation between storage capacity and the main design parameters of such systems, which is crucial for resource-efficient and QoS-aware dimensioning.

The insight provided by our model allowed us to obtain an explicit approximate expression of the FC storage capacity, which can be a very simple and useful tool for analysis and design of FC systems.

REFERENCES

[1] M. Conti and M. Kumar, "Opportunities in Opportunistic Computing," *IEEE Computer*, vol. 43, no. 1, pp. 42–50, 1 2010.

[2] G. Gorbil, "No way out: Emergency evacuation with no internet access," in *IEEE PerCom, PerNEM Workshop*, St. Louis, MO, Mar. 2015.

[3] N. Pérez Palma, V. Mancuso, and M. Ajmone Marsan, "Infrastructureless Pervasive Information Sharing with COTS Devices and Software," in *IEEE WoWMoM 2018*. Chania, Greece: IEEE, 6 2018, pp. 1–9.

[4] I. C. S. L. M. S. Committee *et al.*, "Wireless lan medium access control (mac) and physical layer (phy) specifications," *ANSI/IEEE Std. 802.11-1999*, 1999.

[5] S. Bluetooth, "Specification of the bluetooth system-covered core package version: 4.0," *Bluetooth SIG*, 2010.

[6] J. Schlienz and A. Roessler, "Device to device communication in lte whitepaper," *ROHDE & SCHWARZ: Munich, Germany*, 2015.

[7] R. Molina-Masegosa and J. Gozalvez, "LTE-V for Sidelink 5G V2X Vehicular Communications: A New 5G Technology for Short-Range Vehicle-to-Everything Communications," *VTMag*, vol. 12, no. 4, pp. 30–39, 12 2017.

[8] A. Bazzi, G. Cecchini, M. Menarini, B. M. Masini, and A. Zanella, "Survey and perspectives of vehicular Wi-Fi versus sidelink cellular-V2X in the 5G era," *Future Internet*, vol. 11, no. 6, p. 122, 2019.

[9] Z. Zhang and R. Krishnan, "An overview of opportunistic routing in mobile ad hoc networks," in *IEEE MILCOM 2013*, San Diego, CA, Nov. 2013.

[10] E. Hyytia, J. Virtamo, P. Lassila, J. Kangasharju, and J. Ott, "When does content float? Characterizing availability of anchored information in opportunistic content sharing," *IEEE INFOCOM*, Apr. 2011.

[11] G. Manzo, M. A. Marsan, and G. Rizzo, "Performance modeling of vehicular floating content in urban settings," in *ITC 29*, vol. 1, Sept. 2017, pp. 99–107.

[12] L. Chancay-García, E. Hernández-Orallo, P. Manzoni, C. T. Calafate, and J. Cano, "Evaluating and enhancing information dissemination in urban areas of interest using opportunistic networks," *IEEE Access*, vol. 6, pp. 32 514–32 531, 2018.

[13] N. Thompson, R. Crepaldi, and R. Kravets, "Locus: A location-based data overlay for disruption-tolerant networks," in *Proceedings of the 5th ACM workshop on Challenged networks*. ACM, 2010, pp. 47–54.

[14] A. A. V. Castro, G. D. M. Serugendo, and D. Konstantas, "Hovering information–self-organizing information that finds its own storage," in *Autonomic Communication*. Springer, 2009, pp. 111–145.

[15] T. Nikolovski and R. W. Pazzi, "A Lightweight and Efficient Approach (LEA) for Hovering Information Protocols," in *DIVANet*. New York, NY, USA: ACM, 2017, pp. 31–38.

[16] A. Asadi, Q. Wang, and V. Mancuso, "A survey on device-to-device communication in cellular networks," *Communications Surveys Tutorials, IEEE*, vol. 16, no. 4, pp. 1801–1819, 2014.

[17] X. Zhang and Q. Zhu, "Distributed mobile devices caching over edge computing wireless networks," in *2017 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, May 2017, pp. 127–132.

[18] T. Xu, J. Jiao, X. Chen, and Y. Chen, "Social-aware d2d caching content deployment strategy over edge computing wireless networks," in *2018 27th International Conference on Computer Communication and Networks (ICCCN)*, July 2018, pp. 1–6.

[19] X. Zhang and Q. Zhu, "Collaborative hierarchical caching over 5g edge computing mobile wireless networks," in *2018 IEEE International Conference on Communications (ICC)*, May 2018, pp. 1–6.

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

[21] G. Rizzo, N. Pérez, M. Ajmone Marsan, and V. Mancuso, "A walk down memory lane: On storage capacity in opportunistic content sharing systems," in *IEEE WoWMoM*, Aug 2020.

[22] R. Ghebleh, "A comparative classification of information dissemination approaches in vehicular ad hoc networks from distinctive viewpoints: A survey," *Computer Networks*, vol. 131, pp. 15–37, 2018.

[23] E. Leme, N. Ivaki, N. Laranjeiro, and R. Moraes, "Analyzing gossip protocols for reliable manet applications," in *2017 IEEE International Conference on Edge Computing (EDGE)*, 2017, pp. 98–105.

[24] A. Datta, S. Quarteroni, and K. Aberer, "Autonomous gossiping: A self-organizing epidemic algorithm for selective information dissemination in wireless mobile ad-hoc networks," in *International Conference on Semantics for the Networked World*. Springer, 2004, pp. 126–143.

[25] C. Barberis and G. Malnati, "Epidemic information diffusion in realistic vehicular network mobility scenarios," 11 2009, pp. 1–8.

[26] S. Zhao, L. Chang, T. Zhao, and W. Yan, "Lcs-manet: A mobile storage architecture with location centric storage algorithm in manets," in *2013 International Conference on Computing, Networking and Communications (ICNC)*, 2013, pp. 973–977.

[27] H. Hsu and K. Chen, "Optimal caching time for epidemic content dissemination in mobile social networks," in *2016 IEEE International Conference on Communications (ICC)*, 2016, pp. 1–6.

[28] S. Glass, I. Mahgoub, and M. Rathod, "Leveraging manet-based cooperative cache discovery techniques in vanets: A survey and analysis," *IEEE Communications Surveys Tutorials*, vol. 19, no. 4, pp. 2640–2661, 2017.

[29] H. Khelifi, S. Luo, B. Nour, H. Moungla, Y. Faheem, R. Hussain, and A. Ksentini, "Named data networking in vehicular ad hoc networks: State-of-the-art and challenges," *IEEE Communications Surveys Tutorials*, vol. 22, no. 1, pp. 320–351, 2020.

[30] K. Machado, A. Boukerche, E. Cerqueira, and A. A. F. Loureiro, "A socially-aware in-network caching framework for the next generation of wireless networks," *IEEE Communications Magazine*, vol. 55, no. 12, pp. 38–43, 2017.

[31] E. Hernández-Orallo, C. Borrego, P. Manzoni, J. M. Marquez-Barja, J. C. Cano, and C. T. Calafate, "Optimising data diffusion while reducing local

of opportunistic wireless architectures and mobile broadband services.

Marco Ajmone Marsan (marco.ajmone@polito.it) is a full professor at the Politecnico di Torino in Italy, and a part-time research professor at IMDEA Networks Institute in Leganes, Spain. He obtained degrees in EE from the Politecnico di Torino and the University of California, Los Angeles (UCLA). He is a Fellow of the IEEE, and a member of the Academia Europaea and of the Academy of Sciences of Torino. He received a honorary degree in Telecommunication Networks from the Budapest University of Technology and Economics. He was the Vice-Rector for Research, Innovation and Technology Transfer at the Politecnico di Torino, and the Director of the Istituto di Elettronica e Ingegneria dell'Informazione e delle Telecomunicazioni of the Italian National Research Council. He was the Italian delegate in the ICT and IDEAS committees of FP7.

APPENDIX A
PROOF OF LEMMA 1

Proof. In distributed floating systems, the number of contents which a node can possess and replicate at a given point in time is equal to the minimum between the number of RZs in which the node is at that time, whose mean is $\gamma\pi R^2$, and the number of contents which can be stored in its memory. In localized floating systems, the number of RZs in which a node is located is $\gamma\pi R^2$. Since the probability of possessing a given content is well approximated by $a_s(t, R)$ (resp. $a_d(t, R)$), and since by the homogeneous assumption the probability for a node to possess a content is independent from other nodes, and the same for all nodes, the number of contents possessed by a node follows a binomial distribution, truncated at $\lfloor \frac{M}{L} \rfloor$. $\qquad\square$

APPENDIX B
PROOF OF LEMMA 2

Proof. Let us denote with m_i, $i \in \{s, d\}$ and m_j, $j \in \{s, d\}$ the amount of contents possessed by each of the two nodes in contact at time t for RZ radius R, and let x_{ij} be the amount of exchangeable contents at node i (i.e. the amount of contents possessed by node i but not by node j, and for which there is enough storage space at node j). The probability that node i has x_{ij} exchangeable contents is equal to the probability that x_{ij} out of the m_i contents are not possessed by the other node, and that the remaining contents are possessed by the other node. Moreover, x_{ij} is upper bounded by the available storage space at node j, equal to $\frac{M}{L} - m_j$. Note that m_i and m_j are random variables, whose distribution is given by Lemma 1. Therefore, the distribution of x_{ij} conditioned to m_i, m_j is a binomial $Bin(n, p)$, with parameters $n = \lfloor m_i \rfloor$ and $p = 1 - a_j(t, R)$, and truncated in $\lfloor \frac{M}{L} - m_j \rfloor$. The PDF of x_{ij} is therefore the expectation with respect to m_i and m_j of $P(x_{ij}|m_i, m_j)$. $\qquad\square$

APPENDIX C
PROOF OF LEMMA 3

Proof. Let us consider first the case of a contact between a static and a moving node. The amount of contacts between two nodes depend on the amount of exchangeable contents (i.e., of contents possessed by only one of the two nodes) that each of the two nodes has, as well as on the ratio between the contact time available for the exchange, and by the amount of free storage available at each node for storing the received contents. The probability that a content will be considered for exchange is the probability that the node is among those for which there is enough storage space at the other node, which is well approximated by the ratio between the mean amount of exchangeable contents for which the other node has storage space, and the mean amount of exchangeable contents when the other node has infinite storage space, i.e. by the ratio $\frac{E[x_{ij}]}{E[x_{ij}, M=\infty]}$.

For a contact of duration τ, the amount of time available for transferring contents is given by $\tau - \tau_0$. The mean time taken for transferring a content is given by $\frac{L}{C_0}$. Therefore, on average the maximum amount of contents which can be exchanged during a contact of duration τ is $\lfloor \frac{(\tau-\tau_0)}{\frac{L}{C_0}} \rfloor$. The mean amount of contents which have to be exchanged during a contact taking place at time t is given by $E[x_{ij} + x_{ji}]$. When this quantity is larger than zero, if $E[x_{ij} + x_{ji}] \geq \lfloor \frac{(\tau-\tau_0)}{\frac{L}{C_0}} \rfloor$, then there is enough time for

transferring all contents which can be exchanged. Otherwise, on average, the likelihood for a single content to be exchanged is equal to the ratio between $E[x_{ij} + x_{ji}] \geq \lfloor \frac{(\tau-\tau_0)}{\frac{L}{C_0}} \rfloor$ and $E[x_{ij} + x_{ji}]$. Averaging over contact duration, we get Equation (2). The derivation of T_i follows along the same line. $\qquad\square$

APPENDIX D
DERIVATION OF THEOREM 1

In order to apply the mean field approximation, a key step is the derivation of the expression of the mean dynamics (also called *drift*), which describes the average local variation of the CTMC with respect to time [53], [64].

Lemma 5. *When the system satisfies the homogeneous conditions, the* drift *of the CTMC is given by (notice that, for ease of notation, we drop the indication of the dependency on time t and RZ radius R from all the variables):*

$$
\begin{cases}
\dfrac{da_s}{dt} = \dfrac{b_s}{T_s} a_d(1 - a_s)S_{ds} \\[2mm]
\dfrac{da_d}{dt} = (1 - a_d)\left[\aleph \dfrac{b_s}{T_s} a_s S_{sd} + \dfrac{b_d - \aleph b_s}{T_d} a_d S_{dd} \right] - \dfrac{2\alpha}{\psi DR} a_d \\[2mm]
\dfrac{db_s}{dt} = \dfrac{g_s}{(1-\psi)D}(1 - b_s)(1 - b_d) - \dfrac{b_s}{T_s} \\[2mm]
\dfrac{db_d}{dt} = \dfrac{g_d}{\psi D}2(1 - b_d)^2 + \aleph\dfrac{db_s}{dt} - \dfrac{b_d - \aleph b_s}{T_d} - \dfrac{4\alpha}{\psi DR}b_d
\end{cases}
\tag{19}
$$

with $0 \leq a_i, b_i \leq 1$, $\aleph = \frac{1-\psi}{\psi}$, and $b_d \geq \aleph b_s$.

Proof. (Lemma 5) Let N_s denote the mean number of static nodes in a RZ of radius R possessing a given content at time t, averaged across all j. We have therefore $N_s = a_s \bar{N}(R)(1 - \psi) = a_s(1 - \psi)\pi R^2 D$, and similarly, for dynamic nodes, $N_d = a_d \bar{N}(R)\psi = a_d \psi \pi R^2 D$. Let us compute the rate at which N_s varies over time. As nodes are static, this quantity can only increase over time. The increase is due solely to static nodes which exit from the busy state due to completion of content transfers. The mean rate at which nodes exit the busy state is given by the ratio between the mean number of static busy nodes at time t in the RZ, given by $\bar{N}(R)(1-\psi)b_s$, and the mean time taken by an exchange between a static and a dynamic node, T_s. Moreover, let us consider one of these terminating exchanges. The probability that the j-th content was transferred to the static node during such exchange is equal to the probability that the dynamic node had the content and that the static node did not have it, given by $a_d(1 - a_s)$, multiplied by the probability that the j-th content was transferred during the contact time, given by S_{ds}.

Summing up, we have

$$
\frac{dN_s}{dt} = \frac{\bar{N}(R)(1-\psi)b_s}{T_s} a_d(1 - a_s)S_{ds}
\tag{20}
$$

Normalizing this expression by $\bar{N}(R)(1 - \psi)$ we get the first differential equation in Equation (19).

The rate of change of N_d is given by the sum of three components. The first is given by those nodes which complete a content transfer in a contact with a static node. It is derived in a similar way as in the previous point, and it is given by:

$$
\frac{\bar{N}(R)(1 - \psi)b_s}{T_s} a_s(1 - a_d)S_{sd}
$$

This exploits the fact that the mean number of busy static nodes coincides with the mean number of dynamic nodes busy in contact with static nodes. The second component is the increase in N_d due to contacts among dynamic nodes. The number of busy dynamic nodes involved in such contacts is given by the difference between the total number of busy dynamic nodes, and the total number of busy static nodes, $\bar{N}(R)\psi b_d - \bar{N}(R)(1-\psi)b_s$.

When a couple of such nodes ends being busy, only one of the two has acquired the given content. Hence the rate of these events is given by $\frac{\bar{N}(R)\psi b_d - \bar{N}(R)(1-\psi)b_s}{2T_d}$. The probability that the given content has been exchanged during a contact is given by the probability that only one of the two dynamic nodes in contact has the content, $2a_d(1-a_d)$ multiplied by the probability that the content is successfully exchanged between the two nodes, given by S_{dd}. Hence, the second contribution takes the form

$$\bar{N}(R)\frac{\psi b_d - (1-\psi)b_s}{T_d}2a_d(1-a_d)S_{dd}$$

The third contribution is given by those dynamic nodes which move out of the RZ for the given content. The rate of this type of events is $4\alpha\pi R b_s$. Note that this is computed by doubling the rate at which busy dynamic nodes exit the RZ. Indeed, if a busy node exits the RZ, two busy nodes are not busy anymore, even if the other node remains in the RZ. Putting all together, we have

$$\frac{dN_d}{dt} = \frac{\bar{N}(R)(1-\psi)b_s}{T_s}a_s(1-a_d)S_{sd}+$$

$$+\bar{N}(R)\frac{\psi b_d - \bar{N}(R)(1-\psi)b_s}{T_d}2a_d(1-a_d)S_{dd} + 4\alpha\pi R b_s \tag{21}$$

Normalizing by the mean total number of dynamic nodes in a RZ, we obtain the second differential equation in Equation (19).

Let us now consider the rate at which the mean number of busy static nodes in a RZ at time t changes over time. Their increase is due to contacts between a static and dynamic node (an event which happens with a rate $g_s\pi R^2$ contacts per second), both of which must be non-busy (with a probability $(1-b_s)(1-b_d)$). Their decrease is due to static nodes which complete the process of content exchange with another node. Such an event takes place with a mean rate $\frac{\pi R^2 D(1-\psi)b_s}{T_s}$. Putting all together, and normalizing by $\pi R^2 D(1-\psi)$, we get the third differential equation in Equation (19).

Finally, we consider the rate at which the mean number of busy dynamic nodes in a RZ at time t changes over time. The first contribution is given by those dynamic nodes which come in contact with static nodes. The contribution due to this population of nodes is equal to the rate at which the mean number of busy static nodes in a RZ at time t changes over time, because each of these contacts involves a static and a dynamic node. The second contribution is given by content exchanges among dynamic nodes. With a similar reasoning as above, its expression is given by $g_d\pi R^2 2(1-b_d)^2$. The decrease in the number of busy nodes is due to two effects. The first is the end of content exchanges between two nodes in contact, due to either completion of all the pending transfers, or to the fact that the two nodes are not in contact anymore, given by $\pi R^2 D\frac{\psi b_d - (1-\psi)b_s}{T_d}$. Finally, the number of busy dynamic nodes in the RZ decreases when these nodes exit the RZ. The rate of this type of events is $4\alpha\pi R b_d$, i.e twice the rate at which busy nodes exit the RZ. This is due to the fact that a busy node exiting the RZ stop exchanging contents (and therefore being busy), hence also the other node involved in the exchange is not busy anymore, even if it remains in the RZ. Normalizing by the mean number of dynamic nodes in a RZ, $\pi R^2 D\psi$, we get the fourth differential equation in Equation (19). □

We can prove now the main result.

Proof. (Theorem 1) First, we show that, for any initial conditions $\mathbf{I}(0,R) = (a_i(0,R), b_i(0,R))$, there exists an array $\mathbf{I}_0$ such that $\lim_{R\to\infty}\mathbf{I}(0,R) = \mathbf{I}_0$ (*convergence of initial conditions* [53]). Let us choose $\mathbf{I}_0 = \mathbf{I}(0,R)$. Then for each content k, if $N_i^k(0,R)$ ($N_k^T(0,R)$) is the number of nodes with the k-th content in the RZ of content k (respectively, the total number of nodes in the RZ of content k) at time $t = 0$ in the RZ, choosing $N_i^k(0,R) = \lfloor a_i(0,R)N_k^T(0,R)\rfloor$, and setting to zero the number of busy nodes at $t = 0$ allows satisfying the convergence condition. Given the assumption of stationarity of the mobility patterns, and of uniform node distribution, the mean total number of nodes in a RZ for a content k is equal for each content (given that all RZ have the same shape and size) and we denote it with $N(R)$. In order to apply the mean field approximation approach, we start by assuming $N_k^T(t,R) = N(R)$, for any content k an any time $t \geq 0$. As a consequence of the homogeneous condition, $N(R)$ grows proportionally to R. With these properties, the considered system can be modeled as a Population Continuous Time Markov Chain (PCTMC) [53]. Specifically, to each value of R we can associate a PCTMC model with a total number of nodes $N(R)$. As for the *size* of the model (i.e., as for the parameter used for normalizing the state occupancy), we choose the parameter $N(R)$ itself. Let us consider now a sequence of increasing values of R, to which we can associate a sequence of PCTMC models, each with the features described so far. By the nature of the system, one can easily verify that for any state transition, the state change vector (i.e., the difference between the state occupancy before and after the state transition) is independent of R and hence of the size of the model.

From Lemma 5 it is easy to see that the drift of the generic PCTMC is continuous. Let $\mathbf{I}(t,R) = (a_i(t,R), b_i(t,R))$ and $\mathbf{I}(t) = (\bar{a}(t), \bar{b}(t)) = \lim_{R\to\infty}\mathbf{I}(t,R)$. As our sequence of PCTMC models satisfies these properties, by Theorem 1 in [53] we have that for any finite time horizon $T \leq \infty$, $\mathbb{P}\{\lim_{R\to\infty}(\sup_{0\leq t\leq T}\|\mathbf{I}(t,R) - \mathbf{I}(t)\|) = 0\} = 1$. That is, the sequence of population models associated to R converges *almost surely* to the dynamics of the ODEs in Theorem 1. Finally, in the case in which the number of nodes in each RZ is not constant, one can follow the same approach and derive an additional differential equation for the mean number of nodes in each RZ. Indeed, as we assumed the mobility is stationary, such differential equation would be a balance equation, giving a mean number of nodes in each RZ which does not vary over time, and which is not affected by the evolution over time of the other two variables of the system. Given such decoupling, the mean field approach can be applied separately to the mean number of nodes in each RZ, and to the two variables we have considered so far, obtaining again the ODEs in (19). □