

Optimal Strategies for Floating Anchored Information with Partial Infrastructure Support

Gianluca A. Rizzo

HES-SO Valais, Switzerland

gianluca.rizzo@hevs.ch

Marco Ajmone Marsan

IMDEA Networks Institute, Spain & Politecnico di Torino, Italy

marco.ajmone@imdea.org

Torsten Braun

University of Bern, Switzerland

braun@inf.unibe.ch

Gaetano Manzo

University of Bern & HES-SO Valais, Switzerland

gaetano.manzo@hevs.ch

Abstract

Floating Content (FC) is a communication paradigm to locally share ephemeral content without direct support from infrastructure. It is based on constraining the opportunistic replication of content in a way that strikes a balance between minimizing resource usage and maximizing content availability among the intended recipients. However, existing approaches to management of FC schemes are unfit for realistic scenarios with non-uniform user distributions, resulting in heavy overdimensioning of resources allocated to FC. In this work, we propose a new version of FC, called Cellular Floating Content (CFC), which optimizes the use of bandwidth and memory by adapting the content replication and storage strategies to the spatial distribution of users, and to their mobility patterns. The main idea underlying our approach is to partition users into small “local communities”, and to optimally weight their contributions to the FC paradigm according to their specific mobility features, and to the resources required to achieve a target performance level. We characterize numerically the properties of the optimal strategies in a variety of mobility patterns and traffic conditions, showing

the accuracy of our approach, and the significant savings it enables in the amount of resources necessary to run FC, which in a realistic setup can be as high as 27% with respect to traditional FC dimensioning strategies.

1. Introduction

The growing relevance of context in communications, the increasing importance of information with local validity, and the need to serve an ever-rising traffic demand make opportunistic communications a key enabler of future wireless access networks. This is particularly true for cellular networks in dense urban areas [1], where opportunistic communications enable the offloading of those information transfers with mainly local relevance. In such scenarios, the static infrastructure manages and coordinates opportunistic information exchange, and it injects content into the opportunistic network whenever needed. Moreover, in order to implement reactive or proactive strategies for optimizing its operations, future wireless networks will need context data (such as floating car data and floating phone data) from users [2]. Such data are already being regularly collected by Mobile Network Operators (MNO) and service providers in order to implement services such as traffic management, and for urban sensing applications such as occupancy statistics of shops or points of interest [3].

Among the available opportunistic communication schemes for the local dissemination of information to end users through direct terminal-to-terminal connectivity, a special role is played by Floating Content (FC) [4]. FC was conceived to support infrastructure-less distributed content sharing over a certain geographic area called Anchor Zone (AZ). The objective of FC (and of similar schemes, such as Hovering Information [5], Locus [6], Linger [7], among others) is to spread content to the minimum amount of users within the AZ sufficient for the content to persist over time in the AZ, while achieving a target performance, typically in terms of *success ratio*, i.e. of the average fraction of nodes with content *entering* a given location (the *Zone of Interest* or ZOI). The purpose of constraining the opportunistic replication of a given content within the AZ is to minimize resource usage (bandwidth, memory), while keeping content available only where it is mostly relevant to users. However, in existing formu-

lations (e.g., see [4, 5, 6, 8, 9]) such optimization of resources usage is performed in a very crude, coarse manner. Specifically, they are based on estimating the contribution of each node only as a function of its distance from the ZOI. The resulting dimensioning strategies are vastly inefficient in realistic, heterogeneous settings, frequently involving in the FC scheme a large number of nodes with little or no contribution to the floating content scheme, thus often leading to the impossibility to satisfy the quality of service (QoS) requirements [8]. Moreover, such approaches, designed for settings where infrastructure is not available, are suboptimal in urban scenarios, in which infrastructure is ubiquitous. Given the growing complexity of these settings, strategies for network planning and optimization which rely heavily on the collection of a large amount of data at the user level (e.g. on user mobility and user perceived performance [10, 11]) are progressively taking over.

In this paper, we consider the issue of resource-optimal FC dimensioning. We propose Cellular Floating Content (CFC), an entirely new scheme for localized opportunistic information exchange, based on partitioning space into *cells* and on modulating content replication and storage strategies according to local requirements and resource availability. We propose a dimensioning approach for the derivation of optimal replication and storage strategies, which minimize the use of bandwidth and user memory according to a flexible cost function, while achieving a target performance. Our approach thus involves in the CFC scheme only those sets of users that, for a same resource cost, maximize their contribution to the performance of the FC scheme.

Specifically, the main contributions of this paper are the following:

- We propose a new scheme for localized opportunistic information exchange, denoted Cellular Floating Content (CFC), which allows modulating content replication and storage within each cell according to the contribution of the users within the cell to the overall performance of the scheme;
- We elaborate a mean field model of CFC, and we derive analytical expressions which relate the evolution over time of the main performance parameters of CFC to system parameters, both over finite time intervals, and in the stationary regime;
- We formulate the problem of determining the replication and storage strategies

which minimize the overall resource consumption of CFC, in a way which guarantees a specified target success ratio. We show that the problem can be solved efficiently.

- We characterize numerically the properties of the optimal strategies in a variety of settings. We validate our approach through numerical simulations on measurement-based traces, showing its accuracy under several realistic mobility patterns and traffic conditions. We show that our approach leads to a much more efficient resource utilization with respect to existing FC dimensioning strategies, based only on distance from the ZOI (savings up to 27% were observed in realistic settings).

The paper is structured as follows. In Section 1.1 we summarize the state of the art of FC dimensioning, and in Section 2 we introduce the system model. In Section 3 we present an analytical model of CFC performance, based on mean field theory. In Section 4, we formulate the CFC resource optimization problem, and in Section 5 we assess numerically the optimal strategies resulting from our approach. Finally, Section 6 concludes the paper.

1.1. Related Work

The Floating Content concepts has long been present in the literature, under different denominations. In [12] the authors present the *hovering information* paradigm, whose aim is to share content with spatio-temporal locality constraints using pure ad hoc communications, in a similar way as FC. The FC paradigm was first introduced by [13], which derived a *criticality condition*, i.e. a sufficient condition for a content to float indefinitely over time. In [14] authors validate the analytical results of [13], and they demonstrate the feasibility of FC (in terms of capacity of sustaining content within the anchor zone for a given period of time) under the presence of a modest number of nodes. [15] considers the performance of context-aware applications that use FC as a communication service. A simple analytical model is presented for computing success probability for two different kinds of context-aware applications, under a simple Random Direction mobility model. In [16] we presented the performance of context-aware

applications using FC under different synthetic mobility models. In [17], Hagihara et al. propose a delivery control method for FC called PFCS (Proportional control for Floating Content Sharing). The authors assume that there are multiple different content items present in the network and that replicating all of the content items using FC may lead to overloading the wireless network and/or the storage capacity of nodes. In order to avoid such situations, PFCS controls the number of copies of a particular content item (called *possession ratio*) within an anchor zone. Through a stability analysis, the authors derive a *stability condition* under which the possession ratio can be controlled to a target value. Different from the present paper, all these works assume synthetic mobility models, and a user distribution which is either uniform, or which can be approximated as being uniform.

Another set of works address those issues deriving from applying the FC paradigm to realistic setups. They propose approaches which take advantage, in various forms, of local variations of several system parameters. The paper [18] proposes a scheme in which, in scenarios with several different floating contents and overlapping AZ, users prioritize the contents to exchange based on such measures as anchor zone size, or total amount of users with content in each AZ. The authors of [19] propose a heuristic for controlling content replication within the anchor zone in a distributed manner, based on a circular shape for the anchor zone. For scenarios with nonuniform user distributions, the paper [20] derives a scheme in which the AZ radius is adjusted dynamically according to the user density in the area, while [21] proposes a method for choosing the optimal location for a circular AZ according to specific properties of user spatial distribution. The work in [22] applies the FC paradigm to the issue of building and delivering knowledge about traffic conditions, in order to enable drivers to optimally plan their route. In [23] authors propose another synthetic mobility model, the *city square mobility model*, in which nodes move within a rectangular area, and in which the number of nodes in the AZ is not constant over time. The paper derives simple conditions for content availability and persistence. However, all these works are based on rigidly circular or rectangular anchor zones, which in practical scenarios inevitably brings to include users which, despite being close to the ZOI, do not contribute to the scheme. Moreover, and all decisions on whether to replicate or drop a content are still based on

some notion of distance from the center or the border of the anchor zone. which often results into coarse, highly conservative FC dimensioning strategies.

Another set of works does away with regular anchor zone shapes. The work [24] considers a variant of FC, and it proposes an heuristic to control content availability in a predefined region in a distributed fashion, based on local estimations of availability by each user. Similarly to other schemes in which the anchor zone coincides with the ZOI, the main drawback of the proposed scheme is that in realistic user distribution it often brings to the impossibility of achieving the target performance. The authors in [25] and [26] evaluate experimentally the feasibility of the FC paradigm in an office and in a campus scenario, respectively. These works show that existing analytical approaches, based on non-spatial models are still inadequate to accurately predict the system performance in realistic scenarios, and that important issues need to be addressed in order to effectively apply FC schemes in practical scenarios.

The works in [27], [28] and [29] try addressing these issues, by elaborating mobility models and dimensioning strategies which take into account the specific features of mobility patterns in pedestrian and vehicular scenarios, as they emerge from measurements. However, these works still rely on a distance-based strategy for modulating the user behavior within a FC scheme, which delivers good results only when mobility patterns produce a uniform user distribution within the considered area. As a consequence, in the vast majority of practical scenarios these approaches are still overly conservative, requiring a substantial amount of overprovisioning to deliver a target performance level. Moreover, the work in [28] suggests that such distance-based strategies may often bring to the impossibility to satisfy the QoS constraints (e.g. in the form of minimum success probability).

The vast majority of available analytical approaches to FC performance adopt a *fluid approximation* similar to a mean field (MF) approximation (see [30, 31, 32] and references therein), which has been applied to several problems in opportunistic communications. However, these applications of the MF approach neglect several important issues. Most importantly, and unlike the present paper, existing works do not clarify the convergence properties of the model they propose, and hence the relationship between the proposed model and the actual behavior of the modeled system. To

the best of our knowledge, the only other approach to FC modeling which explicitly applies MF theory is in [33, 34]. These works adopt a stochastic susceptible-infected-recovered (SIR) epidemic model to model the content diffusion process. However they focus on content lifetime and on sufficient conditions for content persistence, while the present paper considers FC performance from the point of view of content availability within a predefined area (the ZOI), hence in a way which is related to the performance of those applications which rely on FC as a communication service.

Recently, [35] and [36] proposed new approaches which do away with a mean-field based analysis and (similarly to the present work) with regular geometries for the AZ, in order to further optimize the amount of user resources involved in the FC scheme. More specifically, they propose a Deep Learning strategy, which employs a Convolutional Neural Network (CNN) to capture the relations between patterns of users mobility, of content diffusion and replication, and FC performance in terms of resource utilization and of content availability within a given area. The performance of such learning based approaches is however tightly linked to the characteristics of the training dataset, in terms of amount of data and of different input combinations present in it. In the specific case, such approaches inevitably perform poorly (leading to heavy suboptimal configurations, if not possibly altogether to the unfeasibility of the dimensioning problem) in those setups and configurations – in terms of node mobility, or of service constraints, such as minimum target Quality of Service (QoS), or ZOI shape, size and position – which are least represented in the training dataset. For instance, this might represent an issue in scenarios with changing service requirements or service demand spatio/temporal patterns, and/or with unusual patterns of user mobility (e.g. due to emergency conditions or to such events as road accidents or blockages). This last feature is a significant drawback, as communications in contexts of emergency or during spontaneous gatherings have always been among the key applications of the FC paradigm, and the issue of how to mitigate such a problem is currently still open. In this paper, we take the first steps at tackling this issue: We focus on FC in VANETs (vehicular area networks) and we analyze how FC performance and resource efficiency may benefit from the availability of a centralized entity, which interacts with moving vehicles, collects relevant mobility data, and build estimates of those mobility parame-

Table 1: Main notation used in this paper (i is the index of the i -th cell)

Notation	Description	Units
r	Transmission range	m
τ_0	Connection setup time	s
C_i	Mean contact rate	s^{-1}
f_i	PDF of contact time in cell i	
F_i	CDF of contact time in cell i	
λ_i	Mean user density	m^{-2}
S	Cell area	m^2
$\sigma_{i,j}$	Mean arrival rate from cell i to j	s^{-1}
d_i	Content infectivity	
γ_i	Recovery rate	s^{-1}
L	Content size	<i>bytes</i>
D	Replication strategy vector	
Γ	Caching strategy vector	
n_i	Number of nodes with content	
N_i	Number of nodes	
a_i	Availability at stationary state	
B_i	Infection rate at the stationary state	s^{-1}
R_i	Death rate at the stationary state	s^{-1}
\mathcal{C}	Zone of Interest	
P_0	Target success ratio	
k	Cost function coefficient	

ters for optimal FC dimensioning. Most importantly, and differently than [35] and [36], our approach derives optimal solutions, whose performance is thus insensitive to how common are the spatio-temporal patterns of mobility or the FC service requirements of a given scenario.

2. System Model

We consider a set of wireless nodes with transmission range r , moving on the plane according to an arbitrary mobility model. We assume two nodes come in contact when they are in range of each other, i.e. when their distance is not larger than r . We define a *contact event* an event consisting of two nodes getting in range of each other. Let τ_0 denote the *connection setup time*, i.e., the time interval between the event of two nodes coming in contact, and the beginning of the exchange of application-level information; τ_0 models the delay with which a link is established. We assume each node knows its position in space, e.g. via GPS or triangulation-based methods.

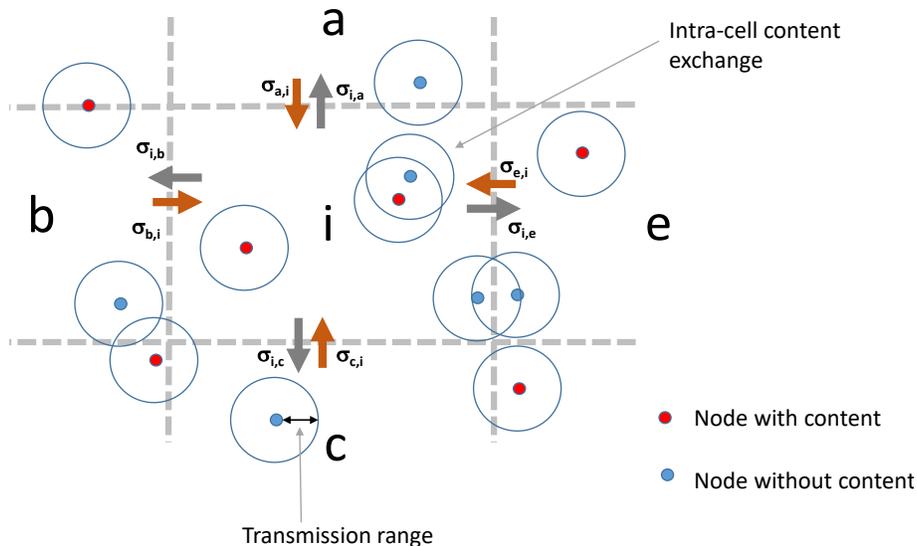


Figure 1: Schematic representation of the considered setup. Each cell is characterized by a content infectivity and a recovery rate, denoted for cell i , as d_i , and γ_i , respectively.

2.1. Cellular Floating Content operation

In what follows, we describe the operation of the Cellular Floating Content communication paradigm.

We consider the case in which a piece of content (e.g. a picture, a video, or a traffic warning message) is to be delivered and made available to nodes within a specific region (which we call the *Zone of Interest* or ZOI), for a predefined time interval, which we call the *floating interval*. Typical floating interval durations range from 30 minutes up to a few hours, during which we assume the traffic patterns do not vary significantly¹.

Given the ZOI, we consider a region of the plane, denoted as *floating region* (FR), which includes the ZOI. We assume the floating region to be partitioned into a set of I cells, which can be of arbitrary shape and size. For instance, each cell may correspond to the coverage range of a cellular base station (BS), or several contiguous BSs, or

¹If longer time intervals must be considered, over which traffic variations cannot be neglected, it is possible to divide the longer time interval into several subintervals over which traffic variations can be considered negligible.

it could be defined geographically, as composed by a set of contiguous city blocks. In what follows, without loss of generality, we assume the FR to be rectangular, and the cells to be rectangular and equally sized, with area S . In such conditions, in each cell i the mean capacity between a user and its neighbor can be expressed as $E(r)/\sqrt{\lambda_i}$ where λ_i is the mean user density in the i -th cell, and $E(r)$ is function of transmission radius [37, 38].

For every pair of contiguous cells i, j , let $\sigma_{i,j}$ denote the mean amount of nodes per unit length which move from i to j , averaged over the floating interval, so that the total rate at which nodes move from i to j is given by $\sigma_{i,j}\sqrt{S}$. We assume $\sigma_{j,i} = 0$ when i and j are not adjacent. For every cell i , for the considered floating interval, with f_i we denote the PDF of contact time, and with C_i the mean rate of contact events per unit area within the cell. C_i is a function of the specific mobility pattern of nodes within the cell during the floating time interval. Hence, typically C_i is a function of the node speed distribution and node density within the cell. Note that in the computation of C_i , every contact between a node in cell i and a node in another cell is counted in both cells with weight 0.5.

It is worth observing that with the present features of smartphones, and with the coming of age of the Smart City paradigm, the quantities necessary for the parametrization of our model are already available in a number of cases, and those that are not yet available are likely to be so in the near future. Indeed, many applications are currently using smartphones as traffic probes, since they move with pedestrians and vehicles, and tracing their movements enables the crowdsourced collection of accurate traffic data. Even more data will come with the hordes of sensors deployed in smart cities for the optimization of movements of people and vehicles [39].

A CFC scheme is identified, for every cell i , by the *content infectivity* d_i , and the *recovery rate* γ_i , as follows. We assume that at the beginning of the floating interval, a piece of content of size L is generated by one or more nodes in the FR (the *seeders*). Then, for the whole duration of the floating interval, every time a node with content in cell i comes in contact with a node without it, the node with content transmits it with probability d_i .

When a node with content enters cell i , or when it acquires content in that cell, it starts

a timer. After x seconds, with x exponentially distributed with mean $1/\gamma_i$, if the node has not moved out of that cell, it removes the content from its memory. The recovery rate $\gamma_i \geq 0$ is, therefore, the mean rate at which a node *with content* in cell i drops the content. Finally, we assume that content is never replicated outside the FR, and that all nodes with content drop it once they are out of the FR. If the content disappears from the FR before the end of the floating interval, we assume it is seeded again. Such *content reseeding* may be performed by nodes in the FR or by the communication infrastructure (e.g. by roadside units, or by the cellular network). In what follows, without loss of generality, we assume that reseeding is made by seeding the content to a node in the same cell in which the last content was lost.

As a result, for a given CFC scheme, the vectors $\mathbf{D}=\{d_i\}$ and $\mathbf{\Gamma}=\{\gamma_i\}$ describe the replication and caching strategies, respectively, which enable the content to persist probabilistically in the FR during the floating interval, despite the mobility of nodes.

In this paper, we consider the ideal case in which nodes do not replicate content when it is not needed (i.e., when both nodes in contact already possess the content – this is doable by proper advertising of available content). Moreover, we assume that content exchanges are always unicast (one-to-one), though results can be easily extended to the case in which broadcast mechanisms are in place.

As we have already mentioned, the goal of CFC is to make the floating content available in a subregion of the FR, the *Zone of Interest* (or ZOI), which is, therefore, the region of space which is taken as reference for CFC performance. The main performance parameter of CFC is the *success ratio*, i.e. the fraction of users which enter the ZOI with content, averaged over the whole content lifetime. The target value of the success ratio, as well as the choice of size, shape, and location of the ZOI, depend entirely on application-level performance requirements. For instance, for an advertisement application, the application level performance target could be to deliver a given advertisement (e.g., info on a sale) to at least a given percentage of the people who enter a shopping center. In such case, the borders of the ZOI should include all entrances to the shopping center, and that percentage would constitute the target success ratio for

such application ².

3. A Mean Field Model of CFC

The CFC scheme we have described can be seen as a system of interacting objects, in which each object, at any time instant, can take only two states (i.e., it can possess the content or not), and in which interactions bring to changes in the state of the objects via content replication. Classical analytical approaches to performance modeling of such systems (e.g. via Markov Chain analysis) are unpractical already for a relatively small number of objects, due to an exponential explosion in the size of the state space. In what follows, we adopt an approximation technique which allows deriving some meaningful insights in the performance of a CFC system. It is based on an approximate model of the interactions between nodes, which goes under the name of *mean field interaction model* or *fluid limit* (see e.g., [40] and references therein), and originated in statistical physics.

If I is the total number of cells, for a given cell area S at any time instant t we associate to the system the state vector $(n_1(t, S), \dots, n_i(t, S), \dots, n_I(t, S))$, such that $n_i(t, S)$ is the number of nodes in cell i with content at time t . Such a system can be modelled as a *multidimensional birth-death process*, that we assume evolves according to Markovian dynamics, and it is therefore a Continuous Time Markov Chain (CTMC). For each cell, births correspond to nodes without content getting it from nodes with content, and to nodes with content entering the given cell, while deaths are due to nodes with content moving out of the cell or dropping the content. In addition, as part of such approximate modeling approach, we assume that the system satisfies the following *homogeneous conditions*:

Definition 1 (Homogeneous conditions). We say that the considered system satisfies the homogeneous conditions when at any time instant, users with content are uniformly

²Note however that, depending on the specific application, the success ratio can be defined differently (e.g. as the ratio of people *getting out* of the ZOI with content) [9]. Our approach can be adapted to such variants of the definition of success ratio.

distributed within each cell, and the probability of a user to have the content is independent from the probability of any other user to have the same content.

The homogeneous conditions assumption allows writing simple expressions for the birth and death rates of the process. Moreover, we assume that cell dimensions and the transmission range of each node are such that intra-cell contacts largely dominate over inter-cell contacts, in such a way that the effects of the latter can be neglected. In Section 5 we assess the accuracy of our approach in application scenarios where these assumptions do not always hold.

Let $N_i(t, S)$ denote the number of nodes in cell i at time t for a cell area S , and let $\bar{N}_i(S)$ denote its mean. As we have assumed the system to be stationary for what concerns node dynamics, and all cells to be of equal size, $\bar{N}_i(S)$ does not depend on time. Moreover, given the homogeneous assumption, we can write $\bar{N}_i(S) = \lambda_i S$. In order to apply the mean field approximation, instead of the number of nodes with content for each cell i at time t , we consider the ratio between this quantity and $\bar{N}_i(S)$, which we denote as $A_i(t, S)$. $\bar{N}_i(S)$ is therefore the parameter used for normalizing the state occupancy in every cell³. The set of all the $A_i(t, S)$, $\forall i \in I$, defines a normalized state space for our system, also known as *occupancy measure* [41].

A key parameter of a CTMC model is the mean dynamics (also called *drift*), which describes the average local variation of the CTMC model with respect to time [41, 42]. The following lemma derives its expression.

Lemma 1. *When the system satisfies the homogeneous conditions, for every cell $i \in 1, \dots, I$ the drift of the CTMC is given by the sum of two components. The first, which*

³Note that choosing $\bar{N}_i(S)$ as a normalization factor instead of $N_i(t, S)$ is necessary in order to apply the results available for mean field convergence of CTMC models. However, it produces a state occupancy which at some time instants might take values above 1, as the actual number of nodes fluctuates around its mean. However, for growing S the variance of the mean number of nodes in cell i tends to zero, hence the state occupancy at the limit for $S \rightarrow \infty$ tends to a quantity which is not larger than 1.

we denote as birth rate (or infection rate), is

$$2 \frac{C_i}{\lambda_i} d_i A_i(t, S) \frac{\bar{N}_i(S)}{N_i(t, S)} \left(1 - A_i(t, S) \frac{\bar{N}_i(S)}{N_i(t, S)} \right) F_i \left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)} \right) + \sum_j \frac{\sigma_{j,i}}{\lambda_i \sqrt{S}} A_j(t, S) \frac{\bar{N}_i(S)}{N_j(t, S)} \quad (1)$$

where F_i is the complementary CDF of contact duration τ in cell i . The second component of the drift, denoted as death rate, is given by

$$A_i(t, S) \frac{\bar{N}_i(S)}{N_i(t, S)} \left(\sum_j \frac{\sigma_{i,j}}{\lambda_i \sqrt{S}} + \gamma_i \right) \quad (2)$$

For the proof, please refer to Appendix A.

In order to study the mean field convergence properties of our system, to each value of S we associate a CTMC model. Indeed to each value of S it corresponds a mean number of nodes in each cell, $\lambda_i S$, which is therefore the *size* of our model. We consider then a sequence of increasing values of cell area S , to which it corresponds a sequence of CTMC models, each with the features described so far. The following result shows that for increasing model size, in the limit the variance of the system becomes zero, so that the approximation made by considering the average behavior of the system is good. Therefore, the average behavior can be modeled considering the unique solution of a system of Ordinary Differential Equations (ODEs) defined by using the limit mean dynamics (i.e. the drift given by Lemma 1 at the limit $S \rightarrow \infty$) of the family of CTMC models.

The expression of the drift allows deriving the differential equations governing the evolution in time, over a finite horizon, of the mean field limit $a_i(t)$ of the occupancy measure for each cell i at time t . Let $\mathbf{A}(0, S) = A_0(0, S), \dots, A_i(0, S), \dots, A_I(0, S)$ denote the array of the initial values for the occupancy measures of the considered system.

Theorem 1. *In the given system, for any initial condition $\mathbf{A}(0, S)$, for $S \rightarrow \infty$, over*

any finite horizon, for any cell i the occupancy measure $A_i(t, S)$ converges almost surely to the solution $a_i(t)$ (the mean field limit) of the following ODE:

$$\frac{da_i(t)}{dt} = 2\frac{C_i}{\lambda_i}d_ia_i(t)(1 - a_i(t))F_i\left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)}\right) + \sum_j \frac{\sigma_{j,i}}{\lambda_i\sqrt{S}}a_j(t) + a_i(t)\left(\sum_j \frac{\sigma_{i,j}}{\lambda_i\sqrt{S}} + \gamma_i\right) \quad (3)$$

For the proof, please refer to Appendix B. Theorem 1 states that the probability of observing a difference between any point of the trajectory of the given system and the solution of the ODE goes to zero as S (and hence the mean total number of nodes in each cell) grows. That is, in the limit, the error made by considering for each cell i a deterministic system characterized by $a_i(t)$ instead of the actual system goes to zero. Moreover, at any time t and for every i, S , $a_i(t)$ is the expected value of $A_i(t, S)$.

Theorem 1 characterizes the evolution of the system over finite time intervals, during which the effects of the initial conditions (in terms of initial distribution of nodes with content among cells) on the occupancy measure for each cell are non-negligible. However, in a large amount of practical settings and applications, the content is set to float for time intervals which are long enough for the initial transient not to affect significantly the system behavior. The following result characterizes the asymptotic behavior of the system, i.e. the configuration of the system after "enough time" has passed since the initial seeding of the content. To this end, it establishes a relation between the stationary state of the system and the steady-state solutions of the ODEs from Theorem 1.

Theorem 2. *For the considered system, the steady-state solutions of Eq. (3) are an approximation of the state distribution for $t \rightarrow \infty$. Moreover, if $(a_1, \dots, a_i, \dots, a_I)$ is a steady-state solution of Eq. (3), then $\forall i$ the corresponding infection rate at the stationary state is*

$$B_i = 2\frac{C_i}{\lambda_i}d_ia_i(1 - a_i)F_{\tau,i}\left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)}\right) \quad (4)$$

and the death rate at the stationary state is

$$R_i = \sum_j \frac{\sigma_{j,i}}{\lambda_i \sqrt{S}} a_j - a_i \left(\sum_j \frac{\sigma_{i,j}}{\lambda_i \sqrt{S}} - \gamma_i \right) \quad (5)$$

Proof. The considered system is a CTMC, which can also be seen as a multidimensional birth-death process. As birth-death processes belong to the class of reversible stochastic processes [43], from Theorem 1.2 in [44] it follows that their stationary behavior is completely determined by the steady-state solutions of Eq. (3). \square

The steady-state solutions of Eq. (3) are the solutions of a system of I equations of second degree. Hence, in general, there are 2^I of them. Let us rewrite the system of equations using the quadratic formula. $\forall i$, we have

$$a_i = \frac{-\beta_i \pm \sqrt{\beta_i^2 + 4\eta_i\theta_i}}{2\eta_i} \quad (6)$$

with

$$\eta_i = 2SC_i d_i F_i \left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)} \right) \quad (7)$$

$$\beta_i = \eta_i - \sum_j \sigma_{i,j} \sqrt{S} - \gamma_i \lambda_i S \quad (8)$$

$$\theta_i = \sum_j \sigma_{j,i} \sqrt{S}. \quad (9)$$

Since a_i cannot take negative values, from Eq. (6) we have that there exists a single feasible value for each a_i . Hence, the system admits only a single feasible solution in which at least some of the variables are nonzero. In addition, the case in which $\forall i, a_i = 0$, is another feasible solution of the system (the *empty system* solution). By studying the gradient of the system of equations, one can see that the empty system solution is not stable, while the other steady-state solution is an attractor for all trajectories which do not start from an empty system. Hence, for all initial conditions in which the system is not empty, the system evolves to the same steady-state solution, given by Eq. (6).

In the rest of the paper, we focus on the stationary state of the system. This amounts

to assuming the floating interval to last much longer than the time during which the dynamics of initial content diffusion are prevalent in the system. In the stationary state, a parameter which plays a key role in determining the performance of CFC is the content *availability*, which is the expectation of the fraction of nodes in cell i with content, given as the steady-state solution of (3). Typically, a high value of availability is correlated with low likelihood of content disappearance, and with high probability of transferring content to nodes traversing the cell.

4. Derivation of the optimal CFC configuration

As we observed, CFC aims at limiting content replication and storage to the minimum amount sufficient to achieve a given target success ratio. In this section, we formulate the problem of finding, for a given content, the replication and caching strategies (that is, the elements of matrices \mathbf{D} and $\mathbf{\Gamma}$) which achieve a target mean success ratio while minimizing a given cost function, related to the amount of memory and bandwidth resources consumed by the floating content scheme.

The following result derives an upper bound on the total amount of bits transferred at the same time within a cell, derived from the finite capacity of the wireless channel in the considered system.

Lemma 2. *In the given scenario, in the stationary state in every cell i the mean total amount of bits transferred at the same time is given by $B_i^\infty d_i \frac{E(r)}{\sqrt{\lambda_i}}$, where*

$$B_i^\infty = 2SC_i a_i (1 - a_i) \tau_t \quad (10)$$

where τ_t is the mean duration of a content exchange (successful or not), given by

$$\tau_t = \int_{\tau_0}^{\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)}} (\tau - \tau_0) f_i(\tau) d\tau + \frac{L\sqrt{\lambda_i}}{E(r)} \left[1 - F_i \left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)} \right) \right] \quad (11)$$

Moreover, it holds $B_i^\infty d_i \leq S\lambda_i$.

For the proof, please refer to Appendix B.1. Without loss of generality, let us assume that the ZOI coincides with a set of cells (not necessarily adjacent), which we

denote with \mathcal{C} , and let $\bar{\mathcal{C}}$ its complement. According to its definition, the success ratio is computed as the average availability over all nodes entering the ZOI. The success ratio P can be expressed as

$$P = \frac{\sum_{i \in \mathcal{C}, j \in \bar{\mathcal{C}}} a_j \sigma_{j,i}}{\sum_{i \in \mathcal{C}, j \in \bar{\mathcal{C}}} \sigma_{j,i}} \quad (12)$$

Using Eq. (10), from Theorem 2 for any cell i the balance equations which give the mean availability in the stationary state a_i take the form

$$F_{\tau,i} \left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)} \right) \frac{B_i^\infty d_i}{\tau_t} + \sum_j \sigma_{j,i} \sqrt{S} a_j - a_i \left(\sum_j \sigma_{i,j} \sqrt{S} - \gamma_i \lambda_i S \right) = 0 \quad (13)$$

In what follows, we model resource consumption in CFC through a cost function which is the sum of two contributions:

$$k \sum_i a_i \lambda_i S L + \sum_i B_i^\infty d_i \quad (14)$$

The first is the cost of storing the content on a set of nodes, or *storage cost*. It is proportional to the mean number of nodes involved, and to the size of the content being floated. The second component is the *replication cost*, and it is proportional to the total number of content transmissions taking place in the FR, assuming one-to-one content exchanges. Note that such cost function is a linear combination of the two cost components (memory and communications), with coefficient $k \geq 0$ as the relative weight between the two cost components. The value of k can be chosen as a function of the resource constraints in the considered setting. For instance, when the nodes participating in the CFC scheme are memory constrained (as it can be the case in an IoT setup), k can be chosen to make the overall contribution to the cost of the two components of the same order of magnitude. Conversely, in vehicular/smartphone applications, storage is typically not an issue. Hence, one can choose k in a way that the minimization of bandwidth is privileged over memory optimization.

Therefore, an optimization problem can be formulated, where the objective function to be minimized is Eq. (14), and where the optimization variables are, for each cell i , d_i and γ_i . In order to study the convexity of the problem, we consider the equivalent

problem formulation obtained by introducing the variables $\psi_i \geq 0$ and $\mu_i \geq 0$, as follows:

$$\begin{aligned} B_i^\infty d_i &= B_i^\infty - \mu_i \\ a_i \gamma_i \lambda_i S &= \psi_i \end{aligned} \quad (15)$$

Substituting into the balance equations (13), we get

$$\frac{F_{\tau,i} \left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)} \right)}{\tau_t} (B_i^\infty - \mu_i) + \sum_j \sigma_{j,i} \sqrt{S} a_j - a_i \sum_j \sigma_{i,j} \sqrt{S} - \psi_i = 0 \quad (16)$$

With these new variables, and letting $\mathbf{A} = \{a_i\}$, $\mathbf{M} = \{\mu_i\}$, and $\mathbf{\Psi} = \{\psi_i\}$, the formulation of the *CFC resource optimization problem* takes the following form:

Problem 1 (CFC resource optimization problem).

$$\underset{\mathbf{M}, \mathbf{\Psi}}{\text{minimize}} \quad k \sum_i a_i \lambda_i S L + \sum_i (B_i^\infty - \mu_i) \quad (17)$$

Subject to:

$$\frac{\sum_{j \in \bar{c}, i \in \mathcal{C}} a_j \sigma_{j,i}}{\sum_{j \in \bar{c}, i \in \mathcal{C}} \sigma_{j,i}} \geq P_0; \quad (18)$$

$\forall i$

Eq. (16)

$$B_i^\infty - \mu_i \geq 0; \quad (19)$$

$$B_i^\infty - \mu_i \leq S \lambda_i; \quad (20)$$

$$B_i^\infty = 2SC_i a_i (1 - a_i) \tau_t;$$

$$0 \leq a_i \leq 1;$$

$$\mathbf{A} \neq \mathbf{0}_v \quad (21)$$

$$\mu_i \geq 0, \quad \psi_i \geq 0 \quad (22)$$

Constraint (18) imposes that the success probability be larger than or equal to the target value P_0 deriving from the QoS requirements of the specific application supported by CFC. Condition (19) imposes that the component of infection rate due to

opportunistic replication within the cell be nonnegative. Condition (20) is derived from Lemma 2, and it imposes an upper bound to the maximum amount of concurrent content transfers, due to the finite capacity of the wireless channel. Finally (21) excludes the empty state $\mathbf{0}_v$ from the feasible solutions of the system of equations (13).

Such optimization problem formulation is convex, and hence it can be solved efficiently. Indeed, Eq. (13) can be cast as

$$B_i^\infty - \mu_i = \frac{\tau_t}{F_{\tau,i} \left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)} \right)} \left(- \sum_j \sigma_{j,i} \sqrt{S} a_j + a_i \sum_j \sigma_{i,j} \sqrt{S} + \psi_i \right) \quad (23)$$

By substituting this expression for $B_i^\infty - \mu_i$ in (17), (19) and (20) we get a convex problem.

Let $(\mathbf{M}^*, \mathbf{\Psi}^*)$ be the solution of the problem, and let a_i^* be the values of availability in cell i at the optimum. The corresponding values of parameters γ_i and d_i at the optimum (henceforth denoted as γ_i^* and d_i^*), which identify the optimal strategy of content replication and storage, can be derived as follows:

$$\gamma_i^* = \frac{\psi_i^*}{a_i^* \lambda_i S} \quad (24)$$

$$d_i^* = 1 - \frac{\mu_i^*}{2a_i^*(1 - a_i^*)F_{\tau,i} \left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)} \right) C_i S} \quad (25)$$

As the fraction in Eq. (25) is always positive, $d_i^* \leq 1$ is always verified, while inequality (19) guarantees that $d_i^* \geq 0$. Therefore, parameters $(\mathbf{D}^*, \mathbf{\Gamma}^*)$ describe the optimal storage and replication strategies deriving from solving Problem 1.

4.1. Comments on the proposed approach

The approach presented so far optimizes the replication and storage parameters over a single floating interval. As we have seen, all the parameters of our model are derived from forecasts and/or from historical trends and patterns of mobility, which in urban scenarios have indeed a very regular and predictable behavior over time. Such parameters are then averaged over the floating interval, and taken as inputs to our model.

The resulting optimal storage and replication strategies are then applied over the whole interval.

However, in applications in which the content is expected to float for several hours, it is very likely that mobility patterns experience strong variations in the mean density of vehicles and users over space, or in the rate at which nodes move between contiguous cells. In those cases, our approach can be applied by dividing the overall floating interval into smaller subintervals, each of a duration which should be short enough to have only minor variations in mobility patterns, but not too short in order to avoid parametrizing the system with noisy estimates. In each subinterval, the optimal storage and replication strategies would be computed as shown in this section. Moreover, at the beginning of all subintervals except the first, the content which has been floating in a subinterval can be used to initiate content diffusion in the following subinterval, possibly with additional seeding from infrastructure if, where, and when required.

It is worth observing here that among the factors which affect the result, a key role is played by the choice of the floating region. Indeed, the FR defines the region of the plane involved in a CFC scheme, and therefore it identifies the set of resources which are available for implementing it. A first basic requirement on the FR is that it should be large enough to guarantee the feasibility of Problem 1. In general, larger regions may bring to more efficient CFC schemes, as increasing FR size (while keeping constant cell size shape and location) corresponds to increasing the optimization parameter space of Problem 1. In practical settings, the choice of a FR might be influenced also by considerations of computation complexity and of communication overhead due to the collection of data on user mobility.

Moreover, note that our approach can be easily expanded to account for different resource costs in different cells, and possibly taking into account the specific composition of the population of nodes (i.e. vehicles, smartphones, UAVs) participating in the content diffusion and replication process. In scenarios where the floating interval is split in several subintervals, our model can also take into account changes in these costs, which could reflect variations in the availability of resources over time, due to, e.g., other competing services based on opportunistic replications, or higher interference levels, etc.

Table 2: Default values for system parameters in the baseline scenario.

Parameter	Value
Target success ratio P_0	0.7
Node speed [m/s]	16.6
Simulation time [s]	7200
Mean node density [nodes/m ²]	$2.77 \cdot 10^{-6}$

Finally, as it is common when dimensioning a system under uncertainty, when the mobility of at least part of the agents in a scenario is highly uncertain (such as in scenarios with a large number of UAVs, with highly irregular patterns of mobility) a conservative approach can be taken. For instance, unpredictability may be accounted for by assuming in our approach a lower density of unpredictable users with respect to the actual one. The use of other, possibly less conservative techniques, depends however on the availability of some form of characterization of the uncertainty on node trajectory.

5. Numerical assessment

We evaluate our approach for the derivation of optimal strategies for replication and storage, obtaining numerical results in a variety of settings and mobility patterns. Unless otherwise specified, we assume $k=1$ (i.e. we give equal weight to storage and replication costs), and $\forall i, F_{\tau,i} \left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)} \right) = 1$. That is, we assume that the mean contact time to be much larger than the mean time required to complete a content transfer.

5.1. Baseline

We first evaluate the accuracy of our model by means of simulations, on a scenario with a synthetic mobility model. We consider a square FR of side 3000 m. In order to obtain an uniform random distribution of nodes in the scenario, we assumed that nodes move according to the Random Direction mobility model [45]. Cells are square, with a side of 600 m, and thus form a 5×5 grid. The ZOI is cell (3, 3), i.e., the cell at the center of the grid. We chose a target success ratio of 0.7, suitable for non-critical, best effort message delivery services such as traffic warnings or advertisements. Note however that every scenario (and in particular, node density, node speed, and transmission

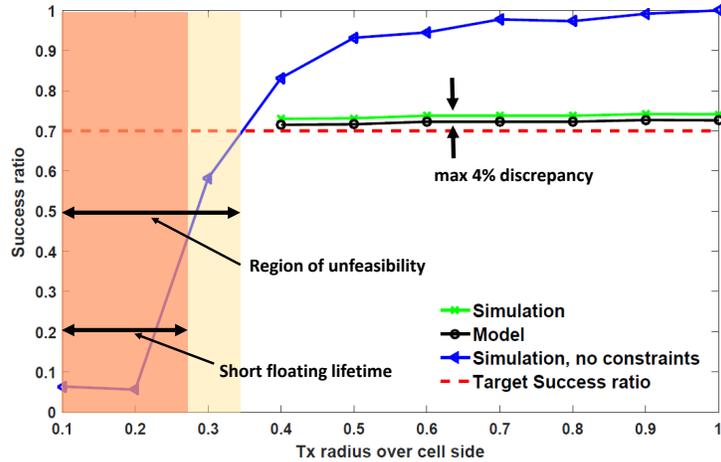


Figure 2: Success ratio vs transmission range, for $P_0 = 0.7$ and cell side of 600 m. All simulation points are within a 95% confidence interval of 5.1%.

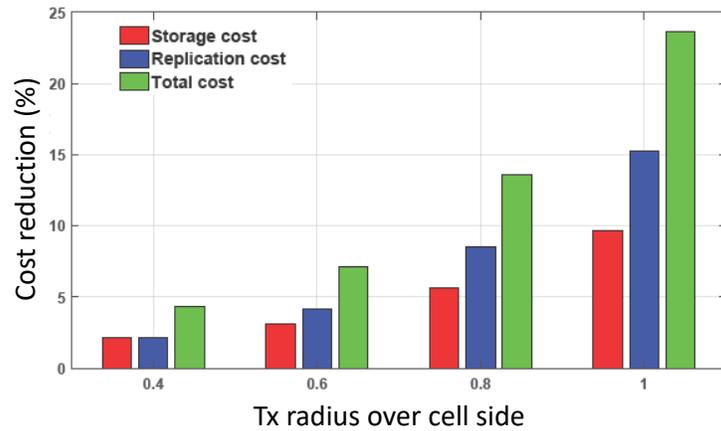


Figure 3: Reduction in cost of the CFC scheme with respect to unconstrained replication and storage, for $P_0 = 0.7$ and a cell side of 600 m.

range) poses an upper limit to the maximum success ratio it can sustain. We assume 25 nodes are present in the scenario, moving at a speed of 60km/h. By modulating transmission range, this choice of mean node density and node speed allows emulating both settings where sparse node distributions and low contact rates make content diffusion more difficult, and scenarios with high contact rates. Table 2 shows the values of the system parameters for the considered setup. For each run, we chose a simulation

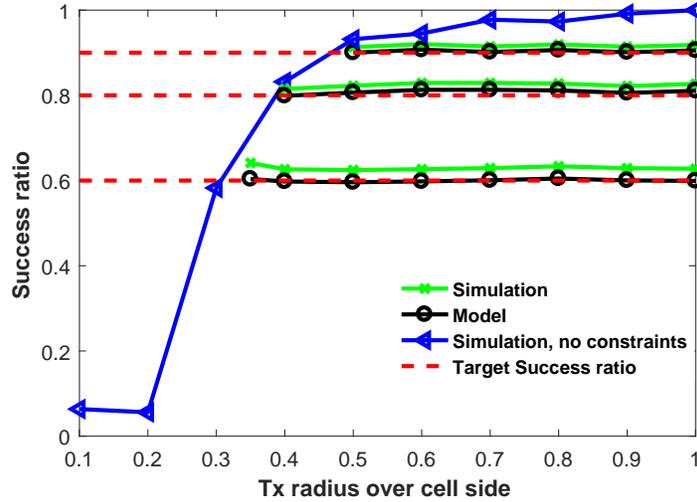


Figure 4: Success ratio vs transmission range, for different target success ratios. Cell side: 600 m. Simulation points are within a 95% confidence interval of 5.1%.

time of two hours, which corresponds to a typical message validity for such vehicular applications as road accident or traffic warnings. In each simulation run, at $t=0$, all nodes possess the content, and are placed at random in the FR. After discarding the initial transient, we measured the success ratio averaged across that part of the remaining simulation time during which the content keeps floating, and then we further averaged across simulation runs.

In the first set of experiments, we solved the optimization problem for a transmission range varying between 60 m (comparable with the range of Bluetooth Low Energy) and 600 m (i.e. up to the length of the cell side). For each value of transmission range, we compared the success ratio deriving from the solution of Problem 1, with the one resulting from simulating the optimal strategy $(\mathbf{D}^*, \mathbf{\Gamma}^*)$. As shown in Fig. 2 and 4, by tuning $(\mathbf{D}, \mathbf{\Gamma})$ within the FR, our approach allows controlling the success ratio with a high level of accuracy over different values of transmission range and target values of success ratio. As the two figures show, the discrepancy never exceeds 4%, and it decreases with increasing target success ratio.

Significantly, these results shows that such effects as node clustering and inter-cell con-

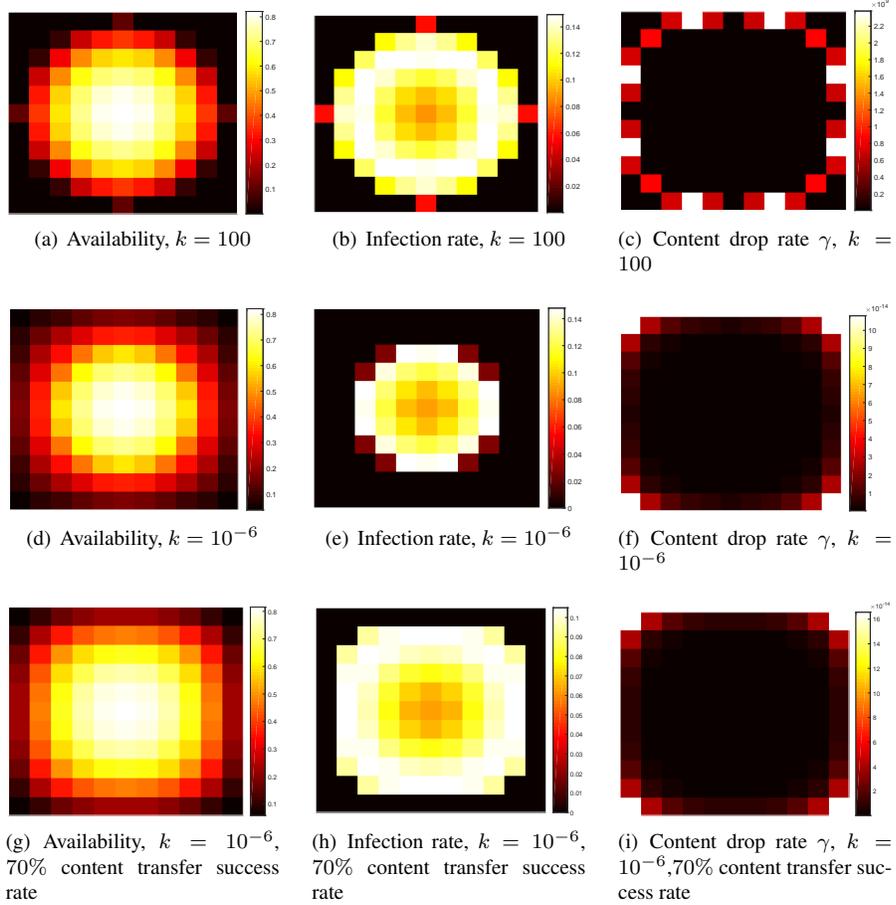


Figure 5: Optimal strategies in a homogeneous setup, for different values of the relative weight k between storage cost and replication cost. Cells form a square grid of 11×11 cells. ZOI is cell (6,6). $\forall i, C_i = 0.3$ contacts/s; cell area 10 m^2 ; $\forall i, j, \sigma_{i,j} = 1$ node/s. Target success ratio: 0.8.

tent exchanges, whose impact on content diffusion grows with increasing transmission range, and which have not been included in our model, have a limited impact on the accuracy of our model, as long as the cell side is larger than the transmission radius. In general, the choice of the cell size is a compromise between accuracy in the way in which the local mobility features are taken into account in the problem, and computational cost. The transmission range however puts a lower limit to the spatial resolution of our approach. Indeed, cells with sides smaller than the transmission range imply a

preponderant amount of inter-cell content exchanges which are not accounted for in our model. Moreover, the smaller the cell, the lower the mean amount of users per cell, and the larger the discrepancy between our model, based on a mean field approximation, and the actual behavior of the system. In both Fig. 2 and 4, the blue curve shows the success ratio from simulations achievable when content replication and memorization are not constrained within the FR (i.e. when $\forall i, d_i = 1, \gamma_i = 0$). This represents the performance achievable when the floating content scheme employs all of the available resources within the floating region in the considered setup, hence it constitutes, for a given setup, an upper bound on the performance achievable by CFC.

Therefore, with reference to Fig. 2, for a given target P_0 (0.7 in the figure), the values of transmission range for which the blue curve takes values inferior to P_0 correspond to settings in which the target success ratio is not achievable in the system. When transmission range is further decreased, contact rate decreases, lowering both the maximum achievable success ratio, and the average content lifetime. In Fig. 2, the orange region corresponds to those values of transmission range for which content lifetime lasts, on average, less than two hours.

In order to estimate the potential resource savings achievable with our approach, albeit in a homogeneous setting, for each value of transmission radius of Fig. 2, in Fig. 3 we plot the decrease in cost of the optimal solution with respect to the case of unconstrained content replication and memorization. As expected, with increasing transmission range, the contact rate increases, and so does the difference in resource utilization between the constrained (i.e., CFC) and the unconstrained scheme, reaching values over 20%.

In order to evaluate the spatial configuration of the optimal replication and storage schemes, in Fig. 5 we depict the spatial distribution of availability, as well as of drop rate γ and of infection rate, resulting from the optimal strategies. These distributions are displayed for the case in which the resource cost of the CFC scheme is dominated by storage cost ($k=100$), and for the case in which it is dominated by content replication cost ($k=10^{-6}$).

As expected, in a homogeneous setup in which the ZOI coincides with the center cell, the optimal schemes induce spatial distributions which tend to the classical circular

shape. Indeed, in a homogeneous scenario where all directions of movements are equally likely, the contribution of each resource to the CFC scheme at a given point in time is only a function of its distance from the ZOI, producing a geometry which tends towards a radial symmetry.

Such property of our optimal schemes allows our approach to satisfy the regression test to the main result available in the state of the art. It has been proposed in [9], and it is based on a random direction mobility model, on a circular, concentric ZOI, and on a circular AZ, whose radius is tuned in order to achieve a given success ratio.

For the same target success ratio, we compared the performance of the two schemes, in terms of total cost, adopting for the scheme of [9] the same cost function as our approach. For our scheme, we have taken the central cell as the ZOI, and we have varied the ZOI radius of the scheme from [9] in order to have a circular ZOI with a same area as the one in our scheme. Our tests have shown that, for a variety of target success ratios and of values of coefficient k , by decreasing the cell area (and hence by refining the spatial quantization of our scheme) the cost of our optimal schemes tend to the cost of the scheme from [9], which is radially symmetric and not affected by spatial quantization effects. This shows that our approach can be regarded as a generalization of [9] to more realistic, heterogeneous scenarios and hence to generic (and possibly not synthetic) mobility patterns.

Results in Fig. 5 also show how our approach produces spatial strategies which adapt to resource cost. When storage is expensive ($k = 100$), content is dropped just out of the area where the infection rate is high. Conversely, when replication costs prevail, the area where replication takes place is smaller and more concentrated around the ZOI, while content (which can be stored and hence fluctuates almost "for free") is practically never dropped within the FR.

Finally, we have considered the case in which connection setup time and content transfer time are such that not all content transfers have enough time to complete successfully during the contact time. Fig. 5 (g) to (i) shows how the optimal spatial strategies change when only 70% of content transfers are successful. As expected, this brings to an average increase of the infection rate, and of the overall number of replications. Moreover, keeping constant the target success ratio, the decrease in effectiveness of

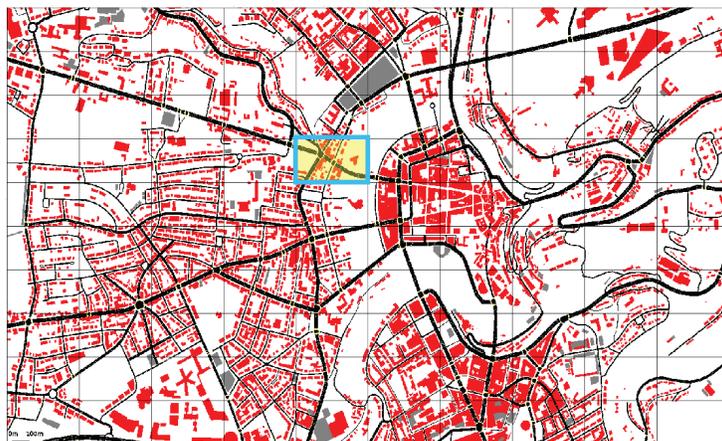


Figure 6: The floating region within the Luxembourg City (4 km by 2.5 km), and the 10×10 cell grid considered. The ZOI is cell (4,5) (in thick borders in the figure).

content transfers increases the amount of time in which, on average, a node must be exposed to contacts with other nodes with content. As shown in Fig. 5, this translates into an increase of the size of the areas where infection rate and availability are high.

5.2. Luxembourg Scenario

In order to perform a more realistic assessment of the performance of our approach, we considered a second scenario, shown in Fig. 6. It consists in a rectangular floating area of 4 km by 2.5 km, which includes the city center of Luxembourg, as well as part of its surroundings. The street grid and the measurement-based mobility traces for this scenario for a 24-hour period were derived from [46]. We considered a transmission range of 360 m, which is typical of technologies such as IEEE 802.11p in realistic urban environments [47]. We chose rectangular shaped cells, with sides 400 m by 250 m, and hence of the same order of magnitude as the transmission range, in order to minimize the impact of inter-cell contacts on the accuracy of our approach.

We considered as ZOI cell (4,5) placed around a crossroad between two large avenues in the city center. The rationale of this choice has been to reproduce a scenario in which, on the occurrence of an accident on that crossroad, cars get a warning before reaching the location of the accident. Network simulations have been performed using the VEINS framework [48], based on OMNET++ [49] for network simulation, and on

Table 3: Success Ratio in the Luxembourg City scenario.

Target value (P_0)	Time interval	Value at $(\mathbf{D}^*, \mathbf{\Gamma}^*)$ (model)	Simulation at $(\mathbf{D}^*, \mathbf{\Gamma}^*)$	Simulation, no constr.
0.7	4 - 6 AM	0.704	0.716	0.954
	8 - 10 AM	0.703	0.721	0.981
0.9	4 - 6 AM	0.912	0.896	0.954
	8 - 10 AM	0.918	0.925	0.981

SUMO [50] for road traffic simulation.

We considered two different time intervals. One, from 8 AM to 10 AM, corresponding to a period of peak traffic, due to people commuting to work, and a second from 4 AM to 6 AM. These two time periods correspond to a best and worst case scenario, respectively. Indeed, the performance of FC is highly sensitive to node density, as the denser the nodes, the easier the content spreads and persists in the given area. In particular, the early morning hours, being typically a period of very low car density in the whole city, represents a worst case scenario for CFC performance, as it makes it hard to get the critical mass of vehicles for guaranteeing a mean content floating lifetime of a few hours. For both time intervals, we have chosen a duration of two hours, which is compatible with typical periods of validity of traffic or accident warnings. As Fig. 7 (a) and (d) shows, the two time intervals are characterized by a different spatial distribution of contact rate. Moreover, as expected, the peak traffic scenario has also, on average, a higher contact rate with respect to the low-traffic scenario.

For each interval, we have computed the optimal strategies $(\mathbf{D}^*, \mathbf{\Gamma}^*)$ by solving Problem 1 with $k = 1$, parametrized with data from the measurement-driven mobility traces and averaged, for each cell, for the whole duration of the time interval.

In Table 3 we have compared the success ratio $P(\mathbf{D}^*, \mathbf{\Gamma}^*)$ computed by solving Problem 1, with the one derived from simulating the optimal CFC schemes $(\mathbf{D}^*, \mathbf{\Gamma}^*)$. In each simulation, we have assumed all nodes within the FR possess the content at the beginning of the given time interval. This can be achieved with an initial seeding of the content through a broadcast transmission over the existing infrastructure. In each configuration, the number of simulation runs has been such as to achieve a 95% confidence interval of $\pm 5\%$. These results show that even in a highly heterogeneous scenario our approach is able to control CFC performance with a very high degree of accuracy.

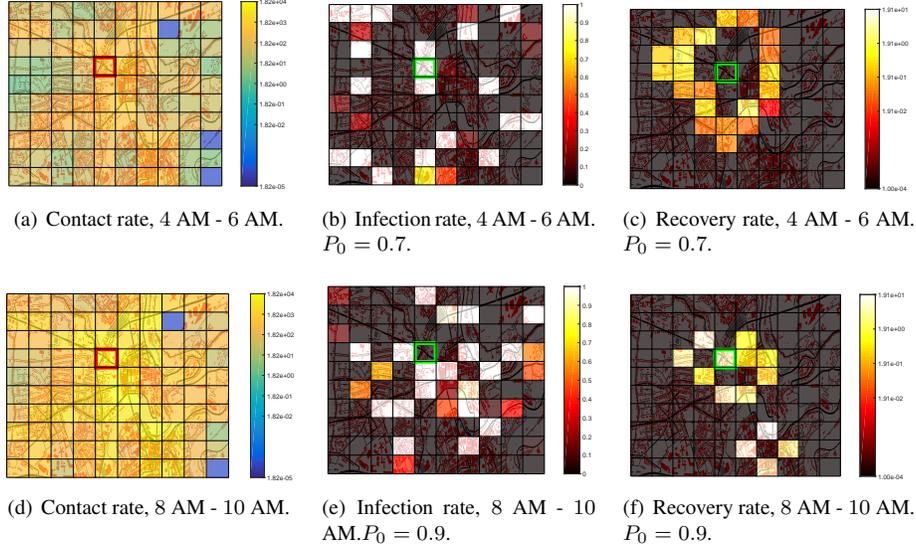


Figure 7: Optimal replication and storage strategies, in the Luxembourg scenario, for two considered time intervals, and for different target values of success ratio. The ZOI is cell (4, 5).

The last column in Table 3 shows the performance of the CFC scheme when content replication and storage are unconstrained everywhere inside the considered FR. As it can be seen, even in settings with low traffic density, the success ratio with unconstrained CFC is very high. This suggests that in the considered settings, and for the given choice of FR shape and size, there are margins for substantially reducing resource utilization, while still achieving the target performance for CFC.

This is illustrated by Fig. 7, which shows the structure of the replication and storage strategies in the time interval (4 AM-6 AM) for $P_0=0.7$, and in the period (8 AM-10 AM) for $P_0=0.9$. The plots show how in both setups our optimal strategies substantially reduce the communication traffic with respect to the unconstrained case, allowing message exchange in few cells only, and dropping the content around the ZOI. For $P_0=0.9$, the optimal scheme enables content replication over those cells with high contact rate around the ZOI (Fig. 7 (e)), corresponding to the main roads which send traffic towards the ZOI, while avoiding those geographical areas where vehicles are not present, such as parks, or rivers. As for the recovery rate, the optimal strategies in Fig. 7 seem to concentrate content dropping on those locations around the ZOI where

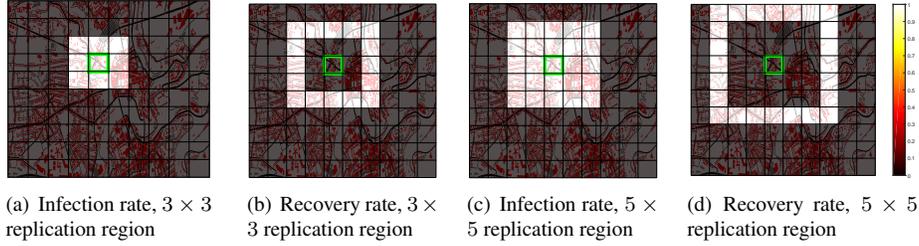


Figure 8: Replication and storage strategies, for the distance-based approach, in the Luxembourg scenario. $P_0 = 0.7$; Time period: 8 AM - 10 AM.

there is a prevalence of traffic flows exiting the ZOI. Indeed, once out of the ZOI, vehicles are less likely to play a relevant role in achieving the target performance for a CFC scheme.

Overall, these results show the capacity of our approach to selectively modulate in space the use of resources of the CFC scheme, adapting to the specific features of mobility of a given scenario, and to the resulting spatial distribution of nodes. Nonetheless, these results indicate that an important property of the optimal strategies is their being radically different from those resulting from intuitive approaches based simply on distance from ZOI, such as those proposed so far [9]. Indeed, the optimal spatial distributions of content replication and dropping in Fig. 7 exhibit several counterintuitive features, such as zones of high content replication (or with high rate of content dropping) which are scattered in the FR.

In order to better assess this feature of the optimal strategies, we have compared their performance with those of a simple, intuitive strategy. It consists in allowing unconstrained content replication and storage in all cells within a given distance (counted in number of cells) from the ZOI, and on dropping the content in those cells which lay immediately out of such replication region. We have run a set of simulations in the time interval 8 AM-10 AM, for the same ZOI as in Fig. 6, for $P_0 = 0.7$, and for different choices of the size of the region of unconstrained replication (Fig. 8). For a replication area of size 3×3 (Fig. 8 (a) and (b)), the mean content lifetime resulting from simulations is of the order of a few minutes only (less than 10), during which success ratio remains low, and very far from the target value. Making the replication area larger (5×5 , Fig. 8 (c) and (d)), allows the content to float for the whole observation window

of two hours, while achieving a success ratio of 0.87, and therefore larger than the target value. However, this second strategy implies a cost which is 37% higher than the one incurred with our approach (which therefore saves 27% of resource costs) for the same target success rate, same ZOI, same time interval, and same cell size. This shows the effectiveness of our approach in minimizing resource costs of a CFC scheme, while guaranteeing the target performance.

5.2.1. Content seeding

On a new set of simulations, we have investigated the impact of various strategies for content seeding on the time required by the system to reach the configuration resulting from the solution of Problem 1. The strategies we considered differ in the nodes which possess the content at the beginning of a CFC scheme. As expected, from Fig. 9 we see that the strategies which bring to a very short transient are those in which at least one node has the content in each cell. Indeed, these strategies do not require the content to diffuse across long distances in order to reach those regions in which it should persist. Despite this observation, seeding the content only to all nodes in the ZOI brings to a performance which is slightly worse than the previous strategies. Finally, we have considered the case in which initial seeding is restricted to those cells which, according to the optimal configuration resulting from the solution of Problem 1, are likely to host several nodes with content once the transient is extinguished. Specifically, it assumes all nodes in cells with (optimal) availability larger than 0.5 have the content at the beginning of the given interval. From Fig. 9 we see that this strategy brings to a slightly longer transient (around 30 min), during which success ratio is lower than the target value.

In conclusion, the choice of the seeding strategy ultimately depends on how critical (for the specific application relying on CFC) is the performance of the CFC scheme in the first 30 minutes from content seeding. Overall, these results show that our approach, despite bringing to spatial strategies which in heterogeneous settings are often quite fragmented in space requires a limited (in time, and in space) support from infrastructure to content replication.

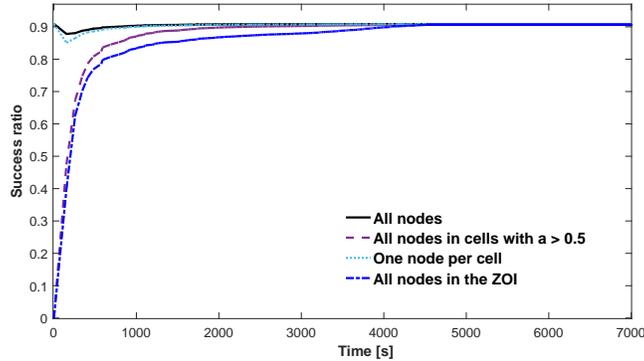


Figure 9: Success ratio vs time from initial content seeding, as a function of the seeding strategy, in the Luxembourg scenario. Time interval: 8 AM-10 AM. $P_0 = 0.9$.

6. Conclusions

In this paper, we have presented a new approach to floating content, based on a mean field model of the main performance measures of the floating content scheme, and on a flexible definition of the cost of resources employed by the FC scheme. This new approach allowed us to optimize the amount of resources required by the scheme, adapting it to the complexity and heterogeneity of real (vehicular, and not only) scenarios, and to a variety of setting with diverse resource constraints. The proposed formulation enables FC to support a wide range of applications, with diverse performance requirements, in a QoS-aware and resource-effective way. On a realistic setup, we showed that our approach achieves resource savings of 27% with respect to existing FC dimensioning strategies, based only on distance from the ZOI.

This work leaves open several avenues for further investigation. Among the possible extensions of this work, it would be important to account for scenarios with several different contents floating in (at least partially) overlapping floating regions. A challenge there would be represented by ensuring scalability of the approach for optimal management of the FC scheme over the number of different contents. Another interesting extension of the present work concerns heterogeneous scenarios where nodes are a combinations of vehicles, pedestrians, and UAVs. In those cases, the goal should be to account not only for the different resource costs to support the various node classes, but also for their specific mobility patterns. In particular, the approach presented in

this paper can be easily extended to scenarios with drones, in which the floating region (and the elements of the partition) are three-dimensional. For those scenarios in which drone trajectories can be configured, it can be interesting to address the issue of the joint optimization of resources over the drone trajectories and of the FC scheme, possibly accounting also for the energy budget of drones. Another important direction of investigation is the extension of the approach presented in this paper to settings with very short content lifetime, in which the initial dynamics of content diffusion play an important role in defining application-level performance.

7. Acknowledgments

This research was supported/partially supported by the Swiss National Science Foundation (SNSF, project CONTACT, no. 164205), by RCSO ISNET MOVNET, and Chist-Era ABIDI.

References

- [1] Q. C. Li, H. Niu, A. T. Papathanassiou, G. Wu, 5G network capacity: Key elements and technologies, *IEEE Vehicular Technology Magazine* 9 (1) (2014) 71–78.
- [2] N. Zhang, P. Yang, J. Ren, D. Chen, L. Yu, X. Shen, Synergy of Big Data and 5G wireless networks: Opportunities, approaches, and challenges, *IEEE Wireless Communications* 25 (1) (2018) 12–18.
- [3] T. Jeske, Floating car data from smartphones: What Google and Waze know about you and how hackers can control traffic, *BlackHat Europe* (2013) 1–12.
- [4] E. Hyytiä, J. Virtamo, P. Lassila, J. Kangasharju, J. Ott, When does content float? Characterizing availability of anchored information in opportunistic content sharing, *INFOCOM*, Shanghai, China.
- [5] A. A. V. Castro, G. D. M. Serugendo, D. Konstantas, Hovering information–self-organizing information that finds its own storage, in: *Autonomic Communication*, Springer, 2009, pp. 111–145.

- [6] N. Thompson, R. Crepaldi, R. Kravets, Locus: A location-based data overlay for disruption-tolerant networks, in: CHANTS'10, ACM, 2010, pp. 47–54.
- [7] D. Borsetti, M. Fiore, C. Casetti, C.-F. Chiasserini, Cooperative support for localized services in vanets, in: ACM MSWiM, 2009.
- [8] G. Manzo, R. Soua, A. Di Maio, T. Engel, M. R. Palattella, G. Rizzo, Coordination mechanisms for floating content in realistic vehicular scenarios, in: IEEE INFOCOM - MobiWorld'17, 2017.
- [9] S. Ali, G. Rizzo, B. Rengarajan, M. A. Marsan, A simple approximate analysis of floating content for context-aware applications, in: MOBIHOC, ACM, 2013, pp. 271–276.
- [10] K. Zheng, Z. Yang, K. Zhang, P. Chatzimisios, K. Yang, W. Xiang, Big data-driven optimization for mobile networks toward 5G, *IEEE Network* 30 (1) (2016) 44–51.
- [11] C. Zhang, P. Patras, H. Haddadi, Deep learning in mobile and wireless networking: A survey, *IEEE Communications Surveys & Tutorials*.
- [12] A. A. V. Castro, G. Di Marzo Serugendo, D. Konstantas, Hovering Information - Self-Organising Information that Finds Its Own Storage, in: SUTC '08, 2008, pp. 193–200.
- [13] E. Hyttiä, J. Virtamo, P. Lassila, J. Kangasharju, J. Ott, When does content float? characterizing availability of anchored information in opportunistic content sharing, in: INFOCOM, Shanghai, China, 2011, pp. 3123–3131.
- [14] J. Ott, E. Hyytia, P. Lassila, T. Vaegs, J. Kangasharju, Floating content: Information sharing in urban areas, in: PerCom 2011, 2011, pp. 136 –146.
- [15] S. Ali, G. Rizzo, B. Rengarajan, M. Ajmone Marsan, A simple approximate analysis of floating content for context-aware applications, in: MobiHoc '13, 2013, pp. 271–276.

- [16] S. Ali, G. Rizzo, M. Ajmone Marsan, V. Mancuso, Impact of mobility on the performance of context-aware applications using floating content, in: ICCASA 2013, 2013.
- [17] R. Hagihara, Y. Yamasaki, H. Ohsaki, On delivery control for floating contents sharing with epidemic broadcasting, in: CCNC, 2017, pp. 353–356.
- [18] J. Ott, E. Hyttiä, P. Lassila, T. Vaegs, J. Kangasharju, Floating content: Information sharing in urban areas, in: IEEE PerCom, 2011, pp. 136–146.
- [19] R. Hagihara, Y. Yamasaki, H. Ohsaki, On delivery control for floating contents sharing with epidemic broadcasting, in: IEEE CCNC, 2017, pp. 353–356.
- [20] R. Yamamoto, A. Kashima, T. Yamazaki, Y. Tanaka, Adaptive contents dissemination method for floating contents, in: 2019 IEEE 90th Vehicular Technology Conference (VTC2019-Fall), 2019, pp. 1–5.
- [21] F. Massalino, A. L. L. Aquino, Identification of anchor zones for floating content in vanets based on centrality measures, in: Proceedings of the 33rd Annual ACM Symposium on Applied Computing, SAC 18, Association for Computing Machinery, New York, NY, USA, 2018, p. 21172124.
- [22] A. M. de Souza, N. L. da Fonseca, L. Villas, A fully-distributed advanced traffic management system based on opportunistic content sharing, in: 2017 IEEE International Conference on Communications (ICC), IEEE, 2017, pp. 1–6.
- [23] M. S. Desta, E. Hyttiä, J. Ott, J. Kangasharju, Characterizing content sharing properties for mobile users in open city squares, in: 2013 10th Annual Conference on Wireless On-demand Network Systems and Services (WONS), IEEE, 2013, pp. 147–154.
- [24] E. Koukoumidis, L.-S. Peh, M. Martonosi, Regres: Adaptively maintaining a target density of regional services in opportunistic vehicular networks, in: IEEE PerCom, IEEE, 2011, pp. 120–127.

- [25] S. Ali, G. Rizzo, V. Mancuso, V. Cozzolino, M. A. Marsan, Experimenting with floating content in an office setting, *IEEE Communications Magazine* 52 (6) (2014) 49–54.
- [26] S. Ali, G. Rizzo, V. Mancuso, M. A. Marsan, Persistence and availability of floating content in a campus environment, in: 2015 IEEE conference on computer communications (INFOCOM), IEEE, 2015, pp. 2326–2334.
- [27] G. A. Rizzo, V. Mancuso, S. Ali, M. A. Marsan, Stop and forward: Opportunistic local information sharing under walking mobility, *Ad Hoc Networks* 78 (2018) 54–72.
- [28] A. D. Maio, R. Soua, M. R. Palattella, T. Engel, G. A. Rizzo, A centralized approach for setting floating content parameters in vanets, in: CCNC, 2017, pp. 712–715.
- [29] G. Manzo, M. A. Marsan, G. Rizzo, Performance modeling of vehicular floating content in urban settings, in: 2017 29th International Teletraffic Congress (ITC 29), Vol. 1, 2017, pp. 99–107.
- [30] M. Benaïm, J. W. Weibull, Deterministic approximation of stochastic evolution in games, *Econometrica* 71 (3) (2003) 873–903.
- [31] M. Benaïm, J.-Y. Le Boudec, A class of mean field interaction models for computer and communication systems, *Performance evaluation* 65 (11-12) (2008) 823–838.
- [32] A. Chaintreau, J.-Y. Le Boudec, N. Ristanovic, The age of gossip: spatial mean field regime, in: ACM SIGMETRICS Performance Evaluation Review, Vol. 37, ACM, 2009, pp. 109–120.
- [33] L. Pajevic, G. Karlsson, Modeling opportunistic communication with churn, *Computer Communications* 96 (2016) 123–135.
- [34] L. Pajevic, V. Fodor, G. Karlsson, Ensuring persistent content in opportunistic networks via stochastic stability analysis, *ACM TOMPECS* 3 (4) (2018) 16.

- [35] G. Manzo, J. S. Otalora, M. A. Marsan, G. Rizzo, A deep learning strategy for vehicular floating content management, *ACM SIGMETRICS Performance Evaluation review* 46 (3) (2019) 159–162.
- [36] G. Manzo, S. Otálora, T. Braun, M. A. Marsan, G. Rizzo, H. Nguyen, Deepfloat: Resource-efficient dynamic management of vehicular floating content, in: *2019 31st International Teletraffic Congress (ITC 31)*, IEEE, 2019, pp. 46–54.
- [37] M. Grossglauser, D. Tse, Mobility increases the capacity of ad-hoc wireless networks, in: *INFOCOM 2001*, Vol. 3, IEEE, 2001, pp. 1360–1369.
- [38] M. Garetto, P. Giaccone, E. Leonardi, On the capacity of ad hoc wireless networks under general node mobility, in: *INFOCOM*, IEEE, 2007, pp. 357–365.
- [39] C. Lim, K.-J. Kim, P. P. Maglio, Smart cities with big data: Reference models, challenges, and considerations, *Cities* 82 (2018) 86 – 99.
- [40] J.-Y. Le Boudec, Mean field methods for computer and communication systems: A tutorial., EPFL Technical report, 2012.
- [41] A. Kolesnichenko, V. Senni, A. Pourranjabar, A. Remke, Applying mean-field approximation to continuous time markov chains, in: *ROCKS*, 2012.
- [42] L. Bortolussi, J. Hillston, D. Latella, M. Massink, Continuous approximation of collective system behaviour: A tutorial, *Performance Evaluation* 70 (5) (2013) 317–349.
- [43] J. Virtamo, *Queuing Theory - Markov processes*, Course 38.3143.
URL <https://www.netlab.tkk.fi/opetus/s383143/kalvot/>
- [44] J.-Y. L. Boudec, The stationary behaviour of fluid limits of reversible processes is concentrated on stationary points, arXiv preprint: 1009.5021.
- [45] B. Christian, H. Hannes, P. C. Xavier, Stochastic Properties of the Random Way-point Mobility Model, *Wireless Networks* 10 (5) (2004) 555–567.

- [46] L. Codeca, R. Frank, T. Engel, Luxembourg SUMO Traffic (LuST) Scenario, in: IEEE VNC, 2015, pp. 1–8.
- [47] A. Böhm, K. Lidström, M. Jonsson, T. Larsson, Evaluating CALM M5-based vehicle-to-vehicle communication in various road settings through field trials, in: IEEE LCN, 2010, pp. 613–620.
- [48] C. Sommer, R. German, F. Dressler, Bidirectionally coupled network and road traffic simulation for improved IVC analysis, IEEE Transactions on Mobile Computing 10 (1) (2011) 3–15.
- [49] A. Varga, The OMNeT++ Discrete Event Simulation System, ESM 2001.
- [50] D. Krajzewicz, G. Hertkorn, C. Rössel, P. Wagner, Sumo (simulation of urban mobility)-an open-source traffic simulation, in: MESM, 2002, pp. 183–187.
- [51] P. Ballarini, M. Beccuti, E. Bibbona, A. Horvath, R. Sirovich, J. Sproston, Analysis of timed properties using the jump-diffusion approximation, in: P. Reinecke, A. Di Marco (Eds.), Computer Performance Engineering, Springer International Publishing, Cham, 2017, pp. 69–84.
- [52] N. Gast, B. Van Houdt, A refined mean field approximation, Proceedings of the ACM on Measurement and Analysis of Computing Systems 1 (2) (2017) 33.

Appendix A. Proof of Lemma 1

To compute the drift, we first derive the rate at which nodes with content in cell i increase (the birth rate) and decrease (the death rate). In order to compute the birth rate, we observe that the rate at which contact events take place within the i -th cell is $C_i S$. As we assume homogeneous conditions, at any time t the probability for a random node within the cell to have the content is $A_i(t, S) \frac{\bar{N}_i(S)}{\bar{N}_i(t, S)}$. At any time t , after a contact between two nodes the probability that a transfer of content can take place is the product of the probability that only one of the two nodes has the content ($2A_i(t, S) \frac{\bar{N}_i(S)}{\bar{N}_i(t, S)} (1 - A_i(t, S) \frac{\bar{N}_i(S)}{\bar{N}_i(t, S)})$), and the probability d_i that the exchange takes

place, and it exploits the fact that these events are independent. The probability that there is enough time for content exchange while the two nodes are in contact can be approximated by the probability that contact duration is larger than the connection establishment time τ_0 , plus the average transfer time $\frac{L\sqrt{\lambda_i}}{E(r)}$. Hence, if we assume that the dynamic of the population of nodes with content is much slower than the mean time to complete a content transfer, the rate at which the mean number of nodes with content in cell i increases due to content replication is

$$2C_i S d_i A_i(t, S) \frac{\bar{N}_i(S)}{N_i(t, S)} \left(1 - A_i(t, S) \frac{\bar{N}_i(S)}{N_i(t, S)} \right) F_i \left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)} \right)$$

The second contribution to the infection rate is given by the sum, over all adjacent cells, of the rate at which nodes with content arrive in cell i . Hence, for every couple of cells i, j , by the definition of $\sigma_{j,i}$, this second contribution is $\sum_j \sigma_{j,i} \sqrt{S} A_j(t, S) \frac{\bar{N}_i(S)}{N_j(t, S)}$.

These expressions describe the rate with which the number of nodes with content increase over time. Normalizing by the model size $\lambda_i S$, we obtain the expression in Eq. (1).

The death rate is the sum of two components. The first one is the rate at which nodes with content within cell i leave the cell, given by

$$A_i(t, S) \frac{\bar{N}_i(S)}{N_i(t, S)} \sum_j \sigma_{i,j} \sqrt{S}$$

The second is the rate at which content is dropped within the considered cell, and it is equal to $\gamma_i < \lambda_i S A_i(t, S) \frac{\bar{N}_i(S)}{N_i(t, S)}$. Again, normalizing by the model size $\lambda_i S$, we obtain the expressions in Eq. (13).

Appendix B. Proof of Theorem 1

First, we show that, with the given assumptions, the *convergence of initial conditions* condition [41]) is verified. This amounts to showing that it always exists an array $\mathbf{a}(0)$ such that

$$\lim_{S \rightarrow \infty} \mathbf{A}(0, S) \frac{\bar{N}_i(S)}{N_i(t, S)} = \mathbf{a}(0)$$

Indeed, for any $\mathbf{a}(0)$, choosing the number of nodes with content at $t = 0$ in such a way that $\forall i, n_i(0, S) = \lfloor a_i(0)N_i(0, S) \rfloor$ allows satisfying this condition.

We now consider the expression of the drift from Eq. (1), and note that the terms which account for the exchange of nodes among cells vanish for $S \rightarrow \infty$. Let $T_0(t, S)$ denote the drift without such terms, given by

$$T_0(t, S) = 2 \frac{C_i}{\lambda_i} d_i A_i(t, S) \frac{\bar{N}_i(S)}{N_i(t, S)} \left(1 - A_i(t, S) \frac{\bar{N}_i(S)}{N_i(t, S)} \right) F_i \left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)} \right) + A_i(t, S) \frac{\bar{N}_i(S)}{N_i(t, S)} \gamma_i \quad (\text{B.1})$$

If we associate the drift $T_0(t, S)$ to the CTMC model, the sequence of CTMC models parametrized by S is a *density dependent* sequence of models [51, 42]. Indeed, from the homogeneous conditions assumption follows that the size of the model (i.e. the mean number of nodes in a cell) grows linearly with the parameter S (the cell area). Moreover, the drift $T_0(t, S)$ depends on S only in terms of normalization. Indeed, the term $\frac{\bar{N}_i(S)}{N_i(t, S)}$ has a mean value equal to 1, and it tends towards 1 for $S \rightarrow \infty$. Hence, the sequence of CTMC models with drift $T_0(t, S)$ is density dependent according to Definition 4 in [41].

$\forall i$, the drift $T_0(t, S)$ is a Lipschitz continuous vector field, as each component has bounded first derivatives. Then by the Mean Field Approximation theorem for density dependent CTMC [41, 42] for any finite time horizon $T \leq \infty$, it holds

$$\mathbb{P} \left\{ \lim_{S \rightarrow \infty} \left(\sup_{0 \leq t \leq T} \|\mathbf{A}(t, S) - \mathbf{a}(t)\| \right) = 0 \right\} = 1 \quad (\text{B.2})$$

That is, the sequence of CTMC models associated to S and with drift $T_0(t, S)$ converges *almost surely* to the dynamics of the solutions of the following ordinary differential equations:

$$\frac{da_i(t)}{dt} = 2 \frac{C_i}{\lambda_i} d_i a_i(t) (1 - a_i(t)) F_i \left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)} \right) - a_i(t) \gamma_i \quad (\text{B.3})$$

The steady-state solutions of Eq. (B.3) are the solutions of a system of I equations of second degree. Hence, in general, there are 2^I of them. Using the quadratic formula,

$\forall i$, we have

$$a_i = \frac{\gamma_i \lambda_i S}{\eta_i} - 1 \quad (\text{B.4})$$

with

$$\eta_i = 2SC_i d_i F_i \left(\tau_0 + \frac{L\sqrt{\lambda_i}}{E(r)} \right)$$

Hence, when $\forall i, \frac{\gamma_i \lambda_i S}{\eta_i} > 1$ there exists a single feasible value for each a_i , and the system admits only a single feasible solution in which at least some of the variables are nonzero. The case in which $\forall i, a_i = 0$, is another feasible solution of the system (the *empty system* solution). From the gradient of the system of equations, one can see that the empty system solution is not stable, while the other steady-state solution is an attractor for all trajectories which do not start from an empty system. Therefore, the ordinary differential equations of Eq. (B.3) with a nonempty initial value have a unique global attractor. As a consequence, by Theorem 3.1 in [52], for a finite system the drift converges to

$$T_0(t, S) + \sum_j \frac{\sigma_{j,i}}{\lambda_i \sqrt{S}} a_j(t) - a_i(t) \left(\sum_j \frac{\sigma_{i,j}}{\lambda_i \sqrt{S}} \right) + \mathcal{O} \left(\frac{1}{S} \right) \quad (\text{B.5})$$

Note that Theorem 3.1 would allow us to further refine this expression. However, here we neglect the terms in $\mathcal{O} \left(\frac{1}{S} \right)$, as the terms in Eq. (B.5) which are proportional to $\frac{1}{\sqrt{S}}$ (and which derive from Lemma 1) are those which model the interactions among neighboring cells. Hence, in a finite system the occupancy measure converges to the solutions of the ODE in Eq. (3).

Appendix B.1. Proof of Lemma 2

Proof. The total amount of content exchanges taking place at the same time in a cell is upper bounded by the finite capacity of the wireless channel. The duration of a content transfer can be approximated by the minimum between the mean amount of time necessary to transmit the whole content, $\frac{L\sqrt{\lambda_i}}{E(r)}$, and the amount of contact time $\tau - \tau_0$ which is available for content transfer. As $f(\tau)$ is the PDF of τ , the mean amount of time spent transmitting the content (not counting the time for connection

setup), denoted with τ_t , is given by Eq. (11).

In every cell i , the quantity $2SC_i d_i a_i (1 - a_i)$ in Eq. (3) represents the rate of infecting contacts, i.e. the rate at which new content exchanges are initiated. The mean number of total ongoing content exchanges at any time instant in a cell can be derived from such rate using Little's result as $2SC_i d_i a_i (1 - a_i) \tau_t$. That is, as the product of the rate of infecting contacts in cell i , seen as the average arrival rate, and the mean amount of time spent transmitting the content.

As $E(r)/\sqrt{\lambda_i}$ is the mean capacity per user, the average aggregate throughput in a cell is hence $2SC_i d_i a_i (1 - a_i) \tau_t \frac{E(r)}{\sqrt{\lambda_i}}$. Such aggregate throughput is upper bounded by the mean aggregate capacity in an ad-hoc network where, as in our case, only one-hop transfers take place (given by $E(r)\sqrt{\lambda_i}$ [37]), times the cell area S , which gives $B_i^\infty d_i \leq S\lambda_i$. □