

Impact of Mobility on the Performance of Context-Aware Applications Using Floating Content

Shahzad Ali^{1,2}, Gianluca Rizzo³, Marco Ajmone Marsan^{1,4},
and Vincenzo Mancuso¹(✉)

¹ Institute IMDEA Networks, Madrid, Spain

² Universidad Carlos III de Madrid, Madrid, Spain

{shahzad.ali,marco.ajmone,vincenzo.mancuso}@imdea.org

³ HES-SO Valais, Sierre, Switzerland

gianluca.rizzo@hevs.ch

⁴ Politecnico di Torino, Torino, Italy

Abstract. The growth of mobile computing and the evolution of smart user devices are progressively driving applications towards “context-awareness”, i.e., towards behaviors that change according to variations in context. Such applications use information that is restricted in space and time, making their communication requirements very different from those of conventional applications, so that opportunistic schemes are better suited to this case than more conventional communications. In this work we consider an opportunistic communication scheme called “Floating Content” (FC), which was specifically designed for server-less distributed context-aware applications, and we refine our previous investigation of the viability of FC for context-aware mobile applications, by considering the impact of different mobility models on the performance of FC. In particular, we consider four different mobility models, and, by using extensive simulation experiments, we investigate the performance of three different categories of context-aware applications that use FC. We also compare the simulation results to the performance predictions of our previously proposed simple analytical model. Results show that good performance can be achieved in content distribution by using FC under a variety of mobility models. They also show that a simple analytical model can provide useful performance predictions even for complex and realistic mobility models, although some application-specific characteristics might call for specialized models to improve the accuracy of performance estimates.

1 Introduction

Context is defined as “any information that can be utilized to understand the situation of an entity” and the applications which adapt their behavior according to changes in context are called “context-aware” [10]. With the pervasiveness of smart devices in the environment, such applications are becoming increasingly

popular. One of the best examples of context and one of the most commonly used variables for context-aware applications is spatial and temporal locality. Locality plays an important role in a variety of applications. As an example, for a context-aware restaurant-finding application [7], information about a nearby restaurant may be of interest to an area close to the restaurant where the likelihood to find users interested in that piece of information is high, and also for a limited time, i.e., until the restaurant is open. Similar context-aware applications encompass an ever-expanding set of applications that make use of spatio-temporal locality, and wireless communications to deliver a variety of services. For many location-based context-aware applications, the scope of the generated content itself is local. This locally relevant content may be of little concern to the rest of the world, therefore moving this content from the user device to store it in a well-accessible centralized location and/or make this information available beyond its scope represents a clear waste of resources (connectivity, storage), and it may lead to the WORN (write-once, read never) problem. All these reasons make the communication requirements for context-aware applications significantly different from ordinary applications. Therefore, a careful design of the communication layer is necessary to serve such applications in the most efficient and scalable way. In this domain, opportunistic communications can play a special role. The benefit of using opportunistic communications is that they naturally incorporate context, as spatial proximity is not only associated to connectivity, but also, at the application layer, to correlation at several levels between communicating peers, between their needs, interests, etc. (the fact that they are in proximity of each other might be because they share interests and views: a same restaurant might mean same tastes for food, etc.). Indeed, connectivity to the infrastructure as a prerequisite is often limiting due to cost and capacity concerns, especially for mobile users for whom using such applications may be problematic due to high roaming charges, unavailability of data services, or simply no network coverage.

In this work we focus on a specific context-aware communication service, known as “Floating Content” (FC) [7], conceived to support server-less distributed context-aware applications. FC is a fully distributed version of ephemeral content sharing, purely based on opportunistic communications. It aims at ensuring the availability of data within a certain geographic area, and for a given duration in time (see Sect. 2 for more details).

The authors of [5] introduced the concept of FC and provided an analytical model for analyzing its feasibility. They derived a condition called “criticality condition”, which can guarantee the availability of information within a given region with high probability. In [7], the authors validated the analytical results presented in [5] with extensive simulations, and showed that the criticality condition behaves well, so that floating content is feasible even when a modest number of nodes is present in the network.

The focus of both [5] and [7] was to evaluate the general feasibility of FC. However, an open issue of those works is that they do not shed light on the performance of an application that uses FC as a communication service. In contrast, in our previous work [1], the focus was on evaluating the performance of

context-aware applications that use FC as a communication service. We defined *success probability*, i.e., the probability that a user obtains the content in which it is interested when it passes through the FC area, as the primary performance parameter. In [1], we assumed that content floats, i.e., that the criticality condition is satisfied, then considered the Random Direction Mobility Model (RDMM) [3] and developed a simple approximate analytical model for computing the success probability, with key parameters the node density, the node transmission range, and the radius of the area within which the content floats. However, RDMM is a very simple mobility model and does not capture the complexity of realistic movement patterns. Hence, it is important to evaluate the FC performance under different and more realistic mobility settings.

In this work our goal is to investigate the impact of different mobility models on the performance of context-aware applications using FC. In particular, in addition to RDMM, we also simulate RPGM, to account for group mobility [4], MGMM, which provides a simplistic model for vehicular mobility [2], and a synthetic mobility trace based on real vehicular traffic statistics collected in the frame of the TAPASCOLOGNE project [8] in the city of Cologne, Germany.

With our experiments, we want to verify by simulation how well FC behaves in realistic mobility settings, and how closely the values of success probability predicted by our simple analytical model match those obtained with complex mobility models. Our results show that FC is a very useful communication paradigm for a variety of context-aware applications, capable of providing good performance in terms of success probability for different mobility models. The success probability values predicted by our simple analytical model are quite close to the values obtained from simulations for very realistic mobility models. However, we also show that specific features of particular mobility models might require more accurate FC models to better predict the system performance.

The rest of the paper is organized as follows. In Sect. 2 we present the floating content service and also introduce the considered families of context-aware applications. Section 3 outlines the analytical model that was proposed in [1]. In Sect. 4 we give a brief introduction of the considered mobility models. In Sect. 5 we present a performance analysis of context-aware applications using floating content. In Sect. 6 we present our conclusions.

2 Floating Content Service

In this section we describe our system model, and the basics of operation of the Floating Content communication service, which we refer to as FC. We consider a scenario with nodes moving on a plane, and communicating with each other in ad hoc mode. Figure 1 summarizes the operation of FC. We assume that at a given point in time, a user (via a context-aware application) issues a message that is of interest only for those users that are in a given region in space, called *Anchor Zone* (AZ), and for a given duration of time. This content is spread using opportunistic communications: Whenever a node possessing this content comes within the transmission range of some other node not having that content,

the content is replicated. When a node possessing the content moves out of the spatio-temporal limits for that message, we assume that it deletes the content. In this way, the content may be available on a set of nodes which moves and varies over time within the AZ, even after the node that generated the content has left the AZ, i.e., the content ‘floats’ within the AZ.

The basic idea behind FC is to store a given content in a spatial region without any fixed infrastructure, making it available through opportunistic communications to all users traversing the region. For this reason the performance metric we consider for the FC service is the probability that a user entering the AZ receives the floating content *timely*. We call this parameter the *success probability* of the FC service. The exact definition of this parameter depends on the way we define the time by which the content should be replicated to the new user who enters the AZ. The determination of this time is application specific, and is made with reference to a subregion of the AZ called the range of interest (ROI). We consider three cases, corresponding to three different categories of context-aware applications, and to three different definitions of success probability:

Baseline application: In this case, ROI and AZ are coincident, and the success probability is the probability that a new user entering in the AZ gets the content before leaving the AZ.

Application category 1: For this category, the message must be delivered to the new user by the time it exits the ROI. Typically, in these applications the message is expected to trigger some specific actions once the user is outside of the ROI. One example of such application can be advertising, when the fact of traversing a given area makes a user very likely to be interested in a specific offer/discount. For such applications, success probability is defined as the percentage of times a node gets content before exiting the ROI.

Application category 2: For this category of applications, the content must be delivered to users before they enter the ROI. Examples of such applications can be accident or traffic jam warnings, when a user should be notified in time to take informed decisions about alternative paths. Here success probability is the probability of getting the content before entering the ROI.

For all applications, the success probability is influenced by node density, by size, shape and relative position of the AZ and of the ROI, and node transmission range. In what follows we consider only circular anchor zones and ranges of interest, and we assume they are concentric. While ROI is strictly related to the application level definition of performance, the AZ radius can be tuned in order to get the desired success probability for all of the proposed applications.

3 An Analytical Model for Success Probability

In this section we briefly recall the main available result for the computation of the success probability for a generic application relying on FC. The derivation of this result in [1] assumes that nodes are distributed according to a planar Poisson point process with intensity λ users/Km², and that users move according to the Random Direction Mobility Model (RDMM) [3].

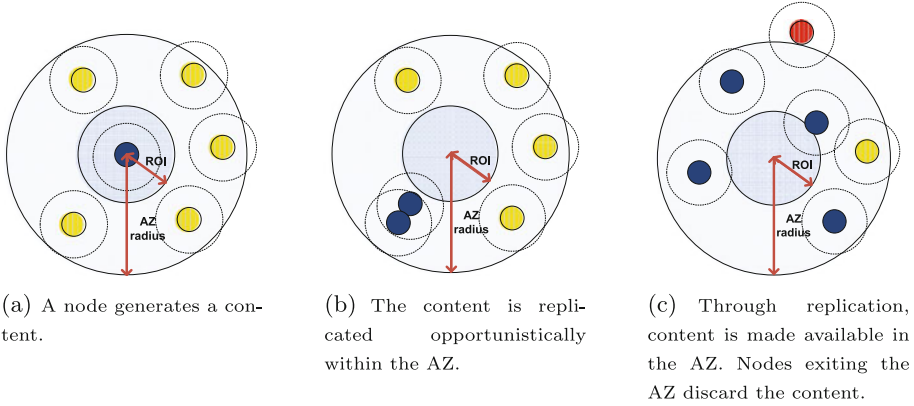


Fig. 1. Operation of floating content service.

Result 1. Consider an AZ with radius R , and nodes with transmission range r and speed v . Let Q denote the probability that two nodes successfully transfer the content while they are in contact. Then an approximated formula for the probability P_s that a node entering the AZ at time $t = 0$ gets the content by time $\tau \leq 2R/v$ is given by

$$P_s(\tau) = \int_0^{2R} \frac{\ell^2}{\pi R^2 \sqrt{4R^2 - \ell^2}} \sum_{k=1}^{\infty} \left[1 - \left(1 - \frac{Q\bar{n}}{(\bar{m} + \bar{n})} \right)^k \right] \frac{\rho^k e^{-\rho}}{k!} d\ell, \quad (1)$$

where $\rho = 2r\lambda \cdot \min(\ell, v\tau)$, $\bar{m} = \min\left(\frac{v}{Q\nu R}, \lambda\pi R^2\right)$, $\bar{n} = \lambda\pi R^2 - \bar{m}$, and $\nu = \frac{2rv^2}{(\pi R^2)}$. \bar{n} and \bar{m} are respectively the average number of nodes with and without content within the anchor zone.

The expression, which we use to compute an estimate of success probability for the three classes of applications presented in this paper, has two main parts. The first one is the probability of meeting k nodes along a path of length ℓ within the AZ, and is computed as the product of the pdf of the path lengths and the conditional pdf of the number of contacts, for a given path length. The second part is the probability that at least one out of k encounters brings to a successful transfer of content.

4 Mobility Scenarios

One of the most important aspects impacting the performance of the FC service is the way in which users move in space. In this paper we investigate this aspect, considering three different mobility models and a set of realistic vehicular traffic traces, and assessing the relationship between their characteristics and the performance of FC through extensive simulation experiments. The first

mobility model we consider is the above mentioned RDMM, one of the most commonly used, and the one underlying the derivation of Result 1. In RDMM, nodes independently travel along a straight line, with an angle of movement uniformly distributed between 0 and 2π . This mobility model is simple and easily tractable analytically because the spatial node distribution remains uniform at all time instants [3]. The second model is the Manhattan Grid Mobility Model (MGMM), used to describe the mobility of vehicles in an urban area [2]. It uses a grid road topology for modeling the movements of vehicles. At each road junction, each vehicle may turn left, turn right or continue straight according to some given probability, which can be tuned to obtain different mobility behaviors. We chose it in order to analyze the impact of a grid topology, typical of a city, on the performance of FC when used by applications residing on vehicles. The third model is the Reference Point Group Mobility model (RPGM), a group mobility model [4]. We have chosen it to evaluate the impact on the performance of FC of clustering and of correlation in user mobility patterns. In RPGM, nodes move in the form of a group and each group has a geographical scope. Nodes belonging to a group are uniformly distributed within its geographical scope. Each group has a *logical center* and all the nodes belonging to that group follow the *logical center*. This logical center moves according to a group motion vector \vec{V}_g . For individual movement of nodes, each node is assigned a reference point which follows the group motion vector. After time τ , a new reference point is calculated by adding a random motion vector \vec{RM} to the group motion vector \vec{V}_g . The length of \vec{RM} is uniformly distributed within a certain radius centered at the reference point, and the direction is uniformly distributed between 0 and 2π . Adding a random motion vector enables a random motion behavior for each individual node. Different mobility scenarios can be modeled with RPGM. One example is groups of tourists visiting some famous attractions in a city. Another example is mobility in a disaster recovery area where different medical assistant teams, rescue teams, firemen teams are randomly moving in the area for the help and rescue operation.

For the fourth considered scenario, we use synthetic mobility traces from the city of Cologne. The Cologne dataset is one of the largest freely available realistic traces capturing both macroscopic and microscopic features [8]. It is realistic from a microscopic point of view because it captures the realistic movement of individual drivers in presence of other vehicles, traffic signals, road junctions, etc. From a macroscopic point of view, it mimics the evolution of large traffic flows across a metropolitan area over time.

5 Performance Evaluation

For all simulation experiments, we use the OMNeT++ based framework called INET [9]. The end-user transmission range is always taken to be 50m. We evaluate the performance under the four different mobility scenarios previously described. When considering RDMM, MGMM and RPGM, each node moves

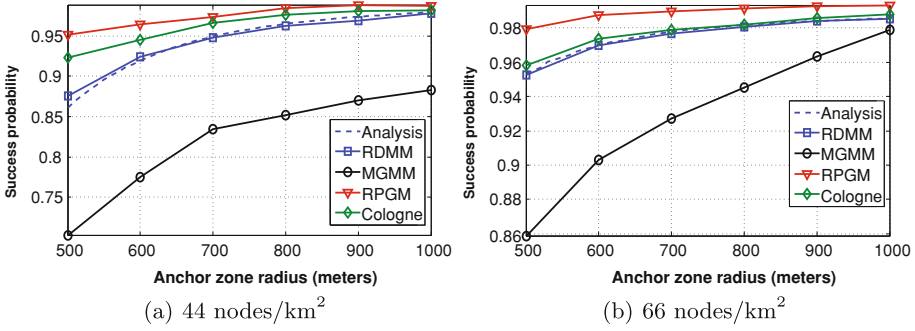


Fig. 2. Success probability for baseline application.

with a constant speed of 5 m/s, while in the Cologne dataset vehicles move at variable speed. For the case of MGMM, a block size of 200 m × 150 m is used. We simulate various values for the AZ radius, while we keep the ROI constant and equal to a circle or radius 200 m at the center of the AZ. As the Cologne dataset is very large, covering a region of 400 Km², for our simulations we considered an area of 9 Km² at the center of the city, and a two-hour time interval (from 6 AM to 8 AM) [6]. For the MGMM, the probability of turning right, turning left and going straight are, respectively, 0.25, 0.25 and 0.5, mimicking typical behavior of cars in city centers.

Figure 2 shows values of success probability as a function of the AZ radius for all the considered mobility models for the baseline application, for two values of nodes density. It can be seen that an increase in either the AZ radius or the node density improves success probability for all scenarios. The reason behind this behavior is that a larger AZ radius increases the average time a node spends inside the AZ, while a higher node density increases the contact rate, both resulting in more chances of meeting a node having content, and thus in higher success probability. For a given AZ radius and node density, RPGM yields the highest success probability, showing that clustering has a beneficial impact on the propagation of the content within the AZ and on its availability. The success probability predicted by our analytical model is very close to the ones by simulations for RDMM, for which the model was developed. We note that MGMM yields a lower success probability than RDMM in all cases. There are two main reasons behind this. First of all, if we look at the path length distribution within the AZ for MGMM and RDMM (see Fig. 3), we see that, unlike in RDMM, in MGMM shorter path lengths have a high probability as compared to relatively longer ones. For the considered block size, a high percentage of nodes traverse shorter paths inside AZ, which reduces the probability of meeting a node with content. The second reason is that, assuming that block size is much larger with respect to the node transmission range, and nodes move with a constant speed, a node can meet another node only if the other node is moving in the opposite direction (if both of them are on same road) or at the road junctions (where a

node can meet other nodes traveling in other directions). This reduces the contact rate, resulting in decreased chances of meeting a node with content. MGMM and Cologne mobility traces are somewhat similar, in what they are both based on a grid of streets in a urban area. However, unlike MGMM, in a realistic setting like the one in the Cologne dataset, vehicles stop at intersections due to traffic signals, and also move according to car following model, which represents a realistic driver behavior [8]. Moreover, nodes move with variable speed, unlike in MGMM, where speed is constant. This also results in increased contact rate, and larger probability of meeting a node with content, resulting in increased success probability in the case of Cologne mobility. Moreover, MGMM keeps nodes uniformly distributed on all the area, while we have verified that mobility patterns in the Cologne dataset exhibit some correlation between vehicles mobility patterns, and some degree of clustering (traffic jams, traffic lights, etc.), which, as it happens for RPMM, improve the performance of FC.

Continuing the comparison between the results for MGMM and RDMM, an interesting observation is that, in case of application categories 1 and 2, the path length of users entering the ROI cannot be shorter than the difference between the AZ radius and the ROI radius. Therefore, we can expect that the success probabilities for application categories 1 and 2 are not impaired by the path length distribution shown for MGMM. Indeed, as shown in Figs. 4 and 5, the success probability of application categories 1 and 2 under MGMM is closer to that of RDMM, as compared to the baseline case. Under the topological settings used for the experiment reported in the figure, the minimum path length for application categories 1 and 2 is 300 m. This means that for application categories 1 and 2, the success probability is computed for paths inside the AZ with length greater than 300 m. This leads to considering longer paths as compared to the baseline case, and as a result the success probability increases and approaches the one given by the simple analytical model.

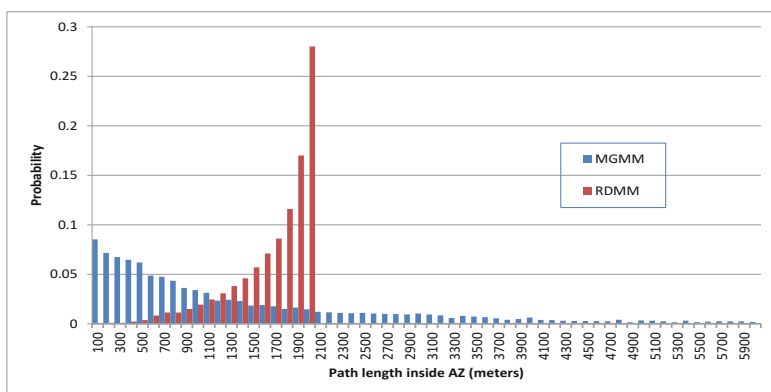


Fig. 3. Distribution of path lengths inside AZ for MGMM (1000 m AZ radius).

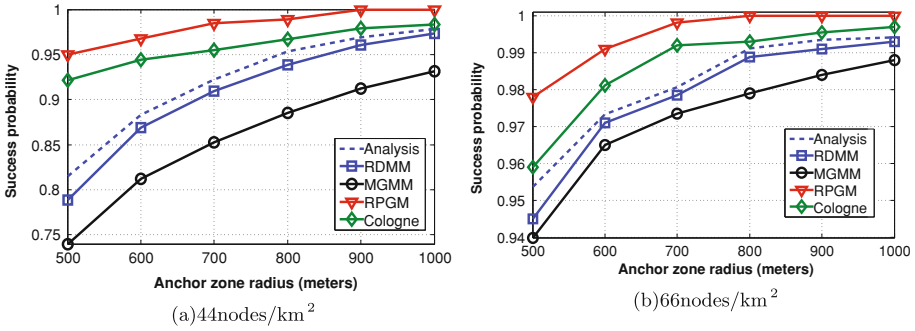


Fig. 4. Success probability for application category 1.

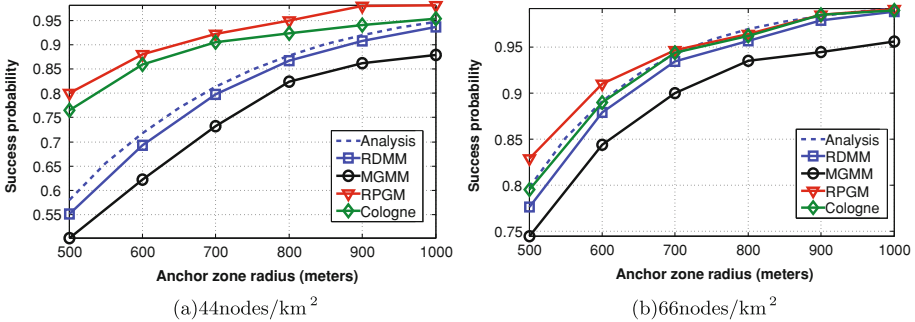


Fig. 5. Success probability for application category 2.

Specifically, Fig. 4 shows curves for success probability versus the AZ radius, for application category 1 under node densities of 44 and 66 nodes/km² respectively. For the Cologne traces, the plots have been obtained by individuating two time intervals, of 1000 s each, during which the average node density in the considered area is equal to the values of node density previously mentioned. As expected, increasing the AZ radius results in higher success probability for application category 1, under all mobility models. The reason is that increasing the AZ radius results in longer average amounts of time a node spends in AZ, resulting in increased chances of meeting a node with content. For all the considered mobility models, our analytical model predictions of success probability become more accurate for higher node densities. This is due to the assumptions underlying its derivation, which hold for a large number of nodes in the AZ.

Figure 5 shows curves for the success probability versus the AZ radius, for application category 2 under node densities of 44 and 66 nodes/km² respectively. The behavior is similar to Fig. 4. For all the considered mobility models, larger AZ radiuses translate in increased success probability for application category 2. If we consider an accident warning application, where the objective is to notify

the nodes entering an area close to the accident location, so that a driver can make an informed decision, we can observe from Fig. 5 that FC is capable of providing a reasonably high success probability.

From our evaluation, we can conclude that FC is a very useful communication paradigm that can be used for a variety of context-aware applications. If parameters are carefully tuned/configured, it is capable of providing a reasonable success probability for a variety of applications and of user mobility patterns. The success probability values predicted by our simple analytical model are quite close to the values obtained from simulations for the case of RDMM. A better representation of the path lengths within the AZ is necessary to obtain comparable accuracy for the other considered mobility models, especially MGMM.

6 Conclusions

In this paper, we focused on the impact of end-user mobility models on the performance of context-aware applications using floating content as a communication paradigm. We considered three different categories of context-aware applications, and four different user mobility models. We found that FC can provide very effective performance to a variety of context-aware applications under quite diverse mobility patterns. Comparing simulation results to the performance predictions of a simple analytical model that was developed for RDMM, we found a very good agreement in the case of RDMM, as already observed in [1]. Other mobility models call for some model re-working to achieve similarly accurate estimates. For all the considered mobility models, high success probabilities can be achieved by tuning the anchor zone radii for a variety of context-aware applications, which justifies the viability of FC as an enabler for context-aware applications.

References

1. Ali, S., Rizzo, G., Rengarajan, B., Marsan, M.A.: A simple approximate analysis of floating content for context-aware applications. In: Proceedings of the Fourteenth ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc '13, pp. 271–276. ACM, New York. <http://eprints.networks.imdea.org/535/> (2013)
2. Bai, F., Sadagopan, N., Helmy, A.: IMPORTANT: a framework to systematically analyze the Impact of Mobility on Performance of Routing Protocols for Adhoc Networks. In: Proceedings of IEEE INFOCOM 2003, Twenty-Second Annual Joint Conference of the IEEE Computer and Communications (2003)
3. Bettstetter, C.: Mobility modeling in wireless networks: categorization, smooth movement, and border effects. SIGMOBILE Mob. Comput. Commun. Rev. **5**(3), 55–66 (2001)
4. Hong, X., Gerla, M., Pei, G., Chiang, C.-C.: A group mobility model for ad hoc wireless networks. In: Proceedings of the 2nd ACM International Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems, MSWiM'99, pp. 53–60. ACM, New York (1999)

5. Hyytiä, E., Virtamo, J., Lassila, P., Kangasharju, J., Ott, J.: When does content float? characterizing availability of anchored information in opportunistic content sharing. In: INFOCOM, Shanghai, China, April 2011 (2011)
6. Cologne mobility traces. <http://kolntrace.project.citi-lab.fr/>
7. Ott, J., Hyytiä, E., Lassila, P., Vaegs, T., Kangasharju, J.: Floating content: information sharing in urban areas. In: PerCom 2011, March 2011 (2011)
8. Uppoor, S., Trullols-Cruces, O., Fiore, M., Barcelo-Ordinas, J.M.: Generation and analysis of a large-scale urban vehicular mobility dataset. *IEEE Trans. Mob. Comput.* **99**(PrePrints), 1 (2013)
9. Varga, A.: The OMNeT++ discrete event simulation system. In: ESM 2001, June 2001 (2001)
10. Zimmermann, A., Lorenz, A., Oppermann, R.: An operational definition of context. In: Kokinov, B., Richardson, D.C., Roth-Berghofer, T.R., Vieu, L. (eds.) CONTEXT 2007. LNCS (LNAI), vol. 4635, pp. 558–571. Springer, Heidelberg (2007)