Evaluation of a Location Coverage Model for Mobile Edge Computing

Michele Girolami
Information Science and Technologies Institute
National Research Council, ISTI-CNR
Pisa, Italy
michele.girolami@isti.cnr.it

Teodorico Pacini
Dep. of Computer Science
University of Pisa
Pisa, Italy
t.pacini2@studenti.unipi.it

Stefano Chessa
Dep. of Computer Science and ISTI-CNR
University of Pisa and CNR
Pisa, Italy
stefano.chessa@unipi.it

***Abstract*—The Mobile Edge Computing paradigm shifts the computation back to places where it is required. A traditional MEC architecture comprises a number of Edge Data Centers (EDC) in charge of seamlessly providing services to users with wireless network technologies. In this scenario, it becomes crucial to deploy the EDCs in strategic locations, such as highly visited places. In this paper we focus on the deployment phase of an EDC. In particular, we propose a probabilistic model designed to measure the location converge, namely the probability that a candidate location for an EDC is visited by users. Our model is based on the analysis of user's trajectories and on the probability of detouring towards the target locations for the EDS. The information returned by our model offers the possibility of implementing mobility-aware deployment strategies in urban environments. We test the model with two real-world mobility data sets, evaluating its applicability of realistic settings.**

***Index Terms*—Mobile Edge Computing, offloading, mobility**

I. Introduction

The Mobile Edge Computing (MEC) paradigm introduces the possibility of moving the computation close to the place where it is required [1], [2]. The MEC enhances a traditional Cloud-based architecture, with the addition of a number of Edge Data Center (EDC) deployed according to a locality principle. The idea of moving the computation to the edge of the network is suitable for different kinds of services, ranging from content offloading to on-demand computation [3]. In particular, we argue that content offloading represents a promising approach to reduce the required bandwidth of traditional broadband Internet connections. The distinguishing feature of the MEC paradigm is the existence of a further layer, composed by a network of EDC deployed in specific locations and connected to the Cloud. EDC generally interacts with devices e.g. smartphone, smart watches and other portable units, with wireless interfaces. In this context, the choice of the target locations for the EDCs represents a strategic choice, as such units are supposed to provide services to users roaming at *short* distances from the location of the EDC [4], [5].

In this work we focus on a heuristic driving the deployment of a set of EDCs according to knowledge extracted from user's mobility. More specifically, we introduce the notion of location coverage which measures the probability that a user visits a location or that a user detours towards such location. Detouring implies that the original trajectory is altered, passing close to the target location. Therefore, given a set of possible EDC locations (also referred to as target locations), we propose a probabilistic model that estimates how much such locations will be visited by users. In turn, such assessment can be used to decide the optimal EDC location. More specifically, our model analyzes the past user's trajectories and, for every target location, it computes the cumulative probability of visiting it by adopting an exponential distribution modeling the detour probability. We test our model with two real-world mobility data sets, namely Gowalla [6] and epfl/mobility (Cabspotting) [7]. In the first case, the data set is obtained from user's check-ins marked as GPS location of the visited places, the data set covers many world-wide countries and we cropped the analysis to the most crowded area, namely Austin, Texas, USA. In the second case, we consider a data set of taxi rides, in this case the collected trajectories report the paths followed by taxis in San Francisco Bay Area, CA, USA. Also in this case, we cropped to most active area, namely San Francisco city. The considered data sets allow us to evaluate the model at different but realistic conditions. We first analyze some mobility features of the experimental data sets, and then we execute our model by studying the effect of the distribution's scale to the target locations. Our experiments show that the model allows to clearly identify visited and non-visited locations with both of the data sets, hence providing a easy-to-use heuristic for the deployment of EDC units in urban areas.

In Section II we survey the state-of-the art concerning deployment strategies for MEC paradigm. Section III introduces our location coverage model and Section IV describes the experimental data sets. Section V describes how we executed our model and the obtained results.

II. Related Work

The MEC paradigm has been described and officially standardized by the European Telecommunications Standards Institute (ETSI) and Industry Specification Group (ISG). A comprehensive review of the MEC paradigm from a research perspective can be found in [8], [9]. For the purpose of this work, we review solutions addressing the problem of deploying EDC in a efficient way. The ETSI standard refers such problem as NFVI-PoP Network Functions Virtualisation Infrastructure Point of Presence (NFVI-PoP).

The work reported in [10] refers to the placement of cloudlet in wireless metropolitan areas. Authors adopt a strategy aiming at minimizing the response time of clients accessing a service provided by a cloudlet in terms of latency. The paper described in [11] proposes a model to define the optimal number of EDC to meet a specific requirement in a multi-layered MEC architecture. The model is based on finding the optimal place by considering both the user population and the distances from the Base Station with a simulative-based approach. A similar approach is implemented in [12] in which authors propose two deployment strategies, namely DDA and Mobility-Aware Deployment Algorithm (MDA) applied to a cellular network. In the Distributed Deployment Algorithm (DDA) solution the deployment is only based on the distance between the EDC and the BS leading to possible high underused or overload EDC. In the second case, authors also consider the mobility to better balance the deployment of EDC. Differently from the previous works, our approach consists of analyzing the user's mobility and to measure how much a set of target location are covered. The proposed model quantifies the coverage probability of such locations and, in turn, such information can be used to drive the deployment of EDC in highly visited locations.

III. The Location Coverage Model applied to a MEC Scenario

In this work, we refer to a MEC distributed architecture characterized by a back-end server, a set of EDC and a set of users. The MEC architecture is structured to efficiently provide services to users, i.e. by moving the computation *close* to the location where the service is actually required. The MEC paradigm can be applied to two specific use cases to which we refer to. On the one hand, we refer to the possibility of offloading the computation to a nearby EDCs. Users have limited resources, such as computational capability and battery life time, therefore high demanding tasks can be offloaded to a local EDC. This is the case of content adaptation through which a device adapts a multimedia content according to its hardware features (e.g. downscaling a video stream). In this case, the device delegates to the EDC the task of adapting the video to a lower screen resolution, instead of fetching a standard content and adapting it locally. On the other hand, we refer to the possibility of providing low-latency services. More specifically, some services need to be provisioned according to stringent Quality of Services (QoS) requirements. This is the case of augmented-reality-based contents and on-line gaming designed to perform responsively. In these two use-cases, we assume that users can interact with an EDC which, in turn, offloads the required resources from the Cloud or perform local computation on behalf of a user.

The previously aforementioned use-cases require investigating where to deploy the EDCs units in a region of interest, e.g. an urban area. Indeed, the MEC architecture is grounded on the locality principle, therefore it becomes mandatory to adopt an efficient approach to evaluate the target location of EDC. This work focuses on this last aspect. In particular, we propose a heuristic which measures how much a set of target locations are visited by users. Given a mobility data set and a set of target locations, we measure the probability that at least a user will visit each of such target locations. Our probabilistic model not only considers the past user's trajectories, but also the possibility of detouring from the original destination towards the target location. The notion of *detour* allows us to evaluate a deviation of the final destination towards a target location where an EDC can be deployed.

More formally, our model considers the set of target locations L and a set of user's trajectories. The set L defines those places where it is admitted to deploy an EDC, while trajectories represent an ordered sequence of way-points of a specific user. Our model measures for each location $l \in L$, if at least a user *visits* or *detours* towards l. In particular, we consider the possibility for a user of deviating from its original trajectory $A \longrightarrow B$, passing though l, $A \xrightarrow{l} B$. Of course, the closer the user to l, the higher the probability of detouring. We report in Fig. 1 three user's trajectories, the red circle shows the target location l for an EDC. We compute the distances x and x' corresponding to the minimum distance between trajectory #2 and l, and between trajectory #1 and l, respectively. Since $x \leq x'$, user following trajectory #2 more likely will accept a detour towards l, with respect to user following trajectory #1 as closer to the target location l.

We define $L = \{l_i \cdots l_H\}$ the target locations, and users are denoted with set $C = \{c_i \cdots c_K\}$ Users move along a set of GPS-based trajectories, e.g. the j-th trajectory of c_i is represented as $t_{i,j} \in T_i$. We can compute the minimum distance between trajectory $t_{i,j}$ and the location $l_h \in L$, denoted as $\bar{x}_{i,j,h}$. We define the random variable X_i^h modeling the event that user c_i accepts detouring towards l_h up to distance $\bar{t}$ from l_h. Events modeled by X_i^h are continuous in $\mathbb{R}^+$, with $f_{X_i^h}$ the probability density function. We also denote with $\phi^{i,j,h}$ the random variable modeling the event that c_i detours from $t_{i,j}$ to location l_h. It is worth to notice that, if c_i detours up to distance $\bar{t}$ from l_h, with a given probability, then c_i will also detour at distance $t \leq \bar{t}$ from l_h, as the user is closer to the target location. Therefore, the probability $P(\phi^{i,j,h} = 1)$ of detouring at distance $\bar{t}$ from l_h is given by: $P(\phi^{i,j,h}) = \int_{\bar{x}_{i,j,h}}^{\infty} f_{X_i^h}(t)dt$.

We adopt an exponential distribution to reproduce the idea of increasing the detour probability with the reduction of the distance from the target location: $f_{X_i^h}(t) = \lambda * \exp(-\lambda * t)$.

Given a user and all its trajectories, we can define the probability that c_i visits location l_h from any of the followed trajectories. Such probability is given by $P(\eta^{i,h}) = 1 - \prod_{\forall t_{i,j} \in T_i}(1 - P(\phi^{i,j,h}))$.

Finally, we define the location coverage α_h as the probability that l_h is visited from any user following any trajectory. To this purpose, the random variable R^h models the events that at least one user detours towards l_h and its opposite event. Recalling ϕ and η, the probability $P(R^h)$ is therefore given by:

$$P(R^h) = 1 - \prod_{\forall c_i \in C} (\prod_{\forall t_{i,j} \in T_i} (1 - \exp(-\lambda * t_{i,j}^h))) \quad (1)$$

obtained by the opposite probability that none of the users in C detour towards l_h. We evaluate in this work the adopted lo-

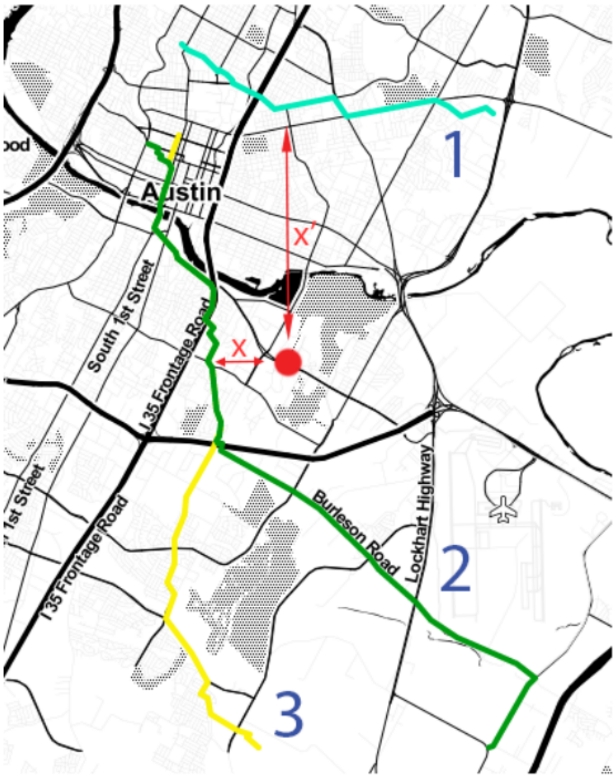

Fig. 1. Example of trajectory's detour (Map data copyrighted OpenStreetMap contributors and available from https://www.openstreetmap.org).

cation coverage model by analyzing trajectories extracted from two representative data sets: Gowalla [6] and epfl/mobility (Cabspotting) [7]. Section IV introduces the experimental data sets with a mobility features analysis, while Section V shows our model in action.

IV. Analysis of the Mobility Data Sets

We now describe the mobility data sets that we adopted to test the model described in Section III. More specifically, Section IV-A describes how we prepare and filter the data sets, while Section IV-B analyzes the mobility features.

A. Description of the data sets

The Gowalla project [7] has been launched as an online social network in which the registered users were allowed to share their geographic location in real-time. The geographic position is obtained with *check-ins* through which friends were notified about new friends' positions. The Gowalla app also allowed to share digital contents. The Gowalla project last from 2007 to 2012 and the collected data are worldwide, mostly of them located in Europe and US. For the purpose of this work, we are interested in studying the location coverage on a specific region. Therefore, we cropped the data set to the most visited city, namely Austin, Texas (US), as reported in Fig. 2.

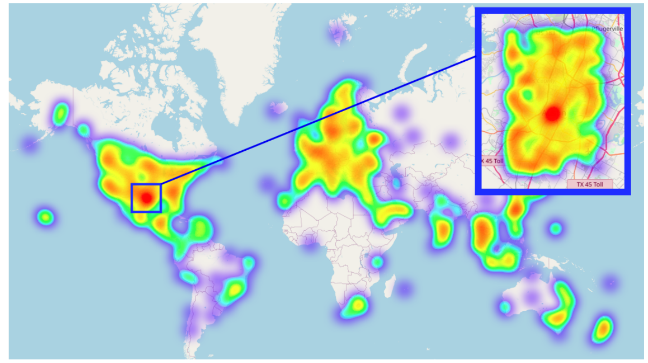

Fig. 2. Geographic representation of the selected region for the Gowalla data set (Map data copyrighted OpenStreetMap contributors and available from https://www.openstreetmap.org).

The data provided with the Gowalla data set are in the form: `[uid, latitude, longitude, timestamp, id location]`, where `id location` identifies the location where user `uid` checked-in. We filtered the data so that to cover 12 months, from November 2009 to October 2010, as a result we obtain 295.532 check-ins. Since the number of points is a relevant aspect for our location coverage model, we decided to extended the number of check-ins. To this purpose, we adopt a simple generative method based on the following steps:

- detecting the user's stop places, namely locations where users rest for a while. To this purpose, we adopt the methodology described in [13] and implemented with the scikit-mobility library [14];
- computing the medoids of the stop places. The previous step returns a set of stop places, we cluster them and we extract from each cluster the corresponding medoid. The medoid provides us a coordinate pair of the representative stop place for each cluster;
- generating a set of new trajectories with origins and destinations selected from the medoids previously identified.

The previous method extend the number of trajectories by only considering those places where people generally stop. As a result, the extended data set now contains 2.187.698 points with 7711 users and 129.113 distinguished trajectories.

The Cabspotting data set is obtained from May to June 2008 and it reports 11.219.955 GPS trajectories followed by the Yellow Cab in San Francisco, CA, USA. Similarly to Gowalla, the information provided by the data set are: `[uid, latitude, longitude, timestamp, occupancy]`. The `uid` identifies the taxi, while the `occupancy` column describes if the taxi is occupied by a passenger or not. Also for the Cabspotting data set, we applied some filtering mechanisms to restrict the geographic area. In particular, we restrict to San Francisco area as it shows the highest number of taxi rides. Fig. 3 shows the Cabspotting heatmap and the inset reports the cropped area.

Fig. 3. Geographic representation of the selected region for the Cabspotting data set (Map data copyrighted OpenStreetMap contributors and available from https://www.openstreetmap.org).

As a result, the extended data set now contains 2.650.083 points with 525 users and 270.136 distinguished trajectories.

B. Mobility Features Analysis

We now analyze some mobility features of the considered data set. More specifically, we characterize both the users' profiles and the visited locations. For what concerns the users' profiles, we measure the typical traveled distance by computing the radius of gyration r_g [15], [16]. The radius is defined as:

$$r_g = \sqrt{\frac{1}{n_u}\sum_{i=1}^{n_u}\delta(r_i(u) - r_{cm}(u))^2} \tag{2}$$

where $r_i(u)$ represents the n_u locations of user u and r_{cm} is center of mass of trajectories of user u, and δ is a distance measure, e.g. the geodetic distance. The center of mass of a given user can be defined as the location where the majority of the points are located, more formally $r_{cm}(u) = \frac{1}{n_u}\sum_{i=1}^{n_u} r_i(u)$, as reported in [15]. The smaller r_g the shorter the distance traveled, while the higher r_g the higher the traveled distance. We report in Fig. 4 the distribution of the radius of gyration. As expected, the distributions differ as the data sets reproduce very different conditions. For what concerns Gowalla, we compute the $50th$ percentile corresponding to $3.5km$, with a maximum r_g of $14.46km$ and a standard deviation of $2.7km$. Differently, for what concerns Cabspotting we observe that the $50th$ percentile of r_g is $2.64km$, but with a maximum r_g of $3.3km$ and a standard deviation of $0.24km$. Users of the Cabspotting data set exhibit a characterizing radius, as shown with the peaked distribution of r_g reported in Figure 5a and 5b.

Concerning the visited locations, we analyze how much they can be predicted. To this purpose, we compute the real entropy

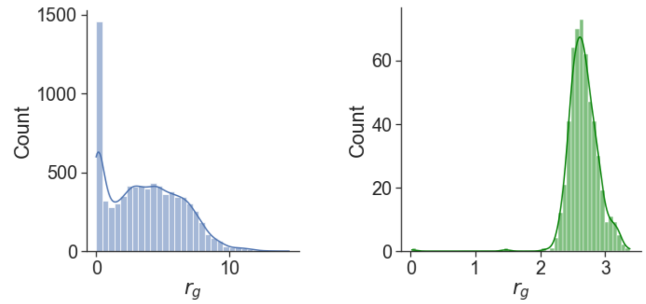

(a) Gowalla data set. (b) Cabspotting data set.

Fig. 4. (a), (b) Distribution of r_g of Gowalla and Cabspotting.

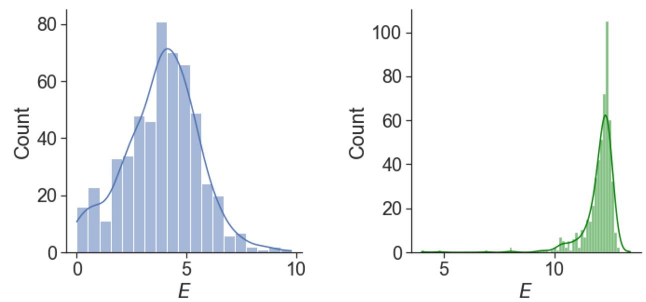

(a) Gowalla data set. (b) Cabspotting data set.

Fig. 5. (a), (b) Distribution of real entropy E of Gowalla and Cabspotting.

E [17] which considers the frequency and the order of the visits for a location, capturing the full spatio-temporal order of the mobility pattern of user. The real entropy $E(u)$ is defined as:

$$E(u) = -\sum_{T'_u} P(T'_u) log[P(T'_u)] \tag{3}$$

where $P(T'_u)$ is the probability of finding an ordered subsequence T'_u ordered along the time in T_u. As a representative example, given a certain value of $E(u) = x$, this implies that the next visited location of user u can be chosen between 2^x distinguished locations. We report in Fig. 5 the distribution of the real entropy for the two data sets. Also in this case, the considered data set shows a very different result. In the case of Gowalla's users, the $50th$ percentile of $E = 3.9$, while for the Cabspotting data set we measure a $50th$ percentile of $E = 12.21$, hence in this case the next predictable location can be found in a set of of $2^{12.3} \approx 5042$ different locations. This remarkable difference is caused by the fact that users of Cabspotting can select any location to stop with a taxi.

V. Experimental Results

Our experimental campaign measures the location coverage for a set of target locations L. Target locations identify those places where it is allowed to deploy an EDC, e.g. locations provisioned with power line, Internet connection and other requirements for the hardware maintenance.

TABLE I
EXPERIMENTAL SETTINGS.

	Gowalla	Cabspotting
duration	18 months	1 month
#points	2.187.698	2.650.083
#users	7711	525
#trajectories	129.113	270.136
50th percentile r_g	3.5Km	2.6km
50th percentile E	3.9	12.21

A. Experimental Settings

Given the user's mobility and the set L, we execute the model described with 1 so that to quantitatively measure how much items $l \in L$ are actually visited by users of the two data sets. The output of our model is a probability map for every $l \in L$ which provides an indication to select the highly visited locations for the EDC deployment. Concerning the mobility data set, we analyze the trajectories of Gowalla and Cabspotting, with 7711 and 525 users, respectively. Concerning the choice of the actual target locations L, they strictly depend on local policies, e.g. admitted areas for the EDCs. Since the analyzed data sets do not provide any information about the target locations, we apply the following criteria to build the set L:

- target locations lay within the selected area of the data sets (see Fig. 2 and 3;
- target locations represent popular places of the urban area;
- target locations are freely accessible.

For what concerns Gowalla, we extract from OpenStreepMap the bus stops of Austin city (921 tags of type `bus stop`), as reported in Figure 6a. Differently, the target locations selected for Cabspotting correspond to historical places of San Francisco area, namely places of worship (794 tags of type `place of warship`), as reported in Fig. 6b. Figure 6 shows a heatmap of the GPS points for the two data sets and of the target locations, so that to qualitatively appreciate how items in L overlap with the GPS's points. It is worth to notice that the selection criteria of the target locations does not affect the validity of the proposed solution, as our primary goal is showing how the model defined in Section III can be used to drive the selection of the deployment locations of EDCs.

Concerning the settings of the model in 1, we analyze several values of the distribution's scale $1/\lambda$. The scale affects the dispersion of the distribution, modifying the attitude of users in accepting a detour or rejecting it. High values of the distribution scale decrease the slope of the curve, and in this case users tend to accept a detour also a high distances from the target location. Conversely, low values of the scale parameter increase the slope, reproducing the situation in which users accept a detour only at short distances from the target location. We test our model with values of $\lambda \in [10, 500]$ so that to capture two different user's behaviors.

(a) Target locations of Gowalla. (b) Target locations of Cabspotting.

Fig. 6. (a), (b) Target locations of the experimental data sets.

B. Numerical Results of the Location Coverage Model

We now report the numerical results of the location coverage model. Fig.s 7a and 7b report the results for Gowalla and Cabspotting data sets, respectively. We show the probability distribution obtained with the distribution's scale set to $1/\lambda = 10, 500$ and the resulting heatmap. The heatmaps show the probability obtained with model for the target locations. Concerning the Gowalla data, setting the scale to small values, reproduces the situation in which users accept a detour only if they are close to the target locations. As a result, the model clearly distinguishes between locations not visited (0-bin in Figure 7a) form locations that will be visited in high probability (1-bin). In this last case, we observe the effect of the scale parameter is to reproduce a situation in which users always accept a detour. The information provided by our model help deciding the location of a set of Mobile Edges from a list of possible target locations. More specifically, given k EDC, a reasonable choice is deploying them on the top-k visited location according to our model as they will represent places in which users probability will pass through. Similar considerations also apply for the Cabspotting data set reported in Fig. 7b. Differently from Gowalla, the heatmap are more clearly distinguishable.

VI. CONCLUSIONS

The problem focused in this paper addresses the deployment strategy of EDCs in a MEC architecture. In particular, we study the problem of measuring how much the candidate locations for a set of EDCs are actually visited by users. The model measures for each of such locations the visiting probability obtained by also considering a trajectory's detour towards such locations. We test our model with two real world experimental data sets, namely Gowalla and Cabspotting. They offer a very different perspective of user's mobility useful to evaluate the model with different scenarios.

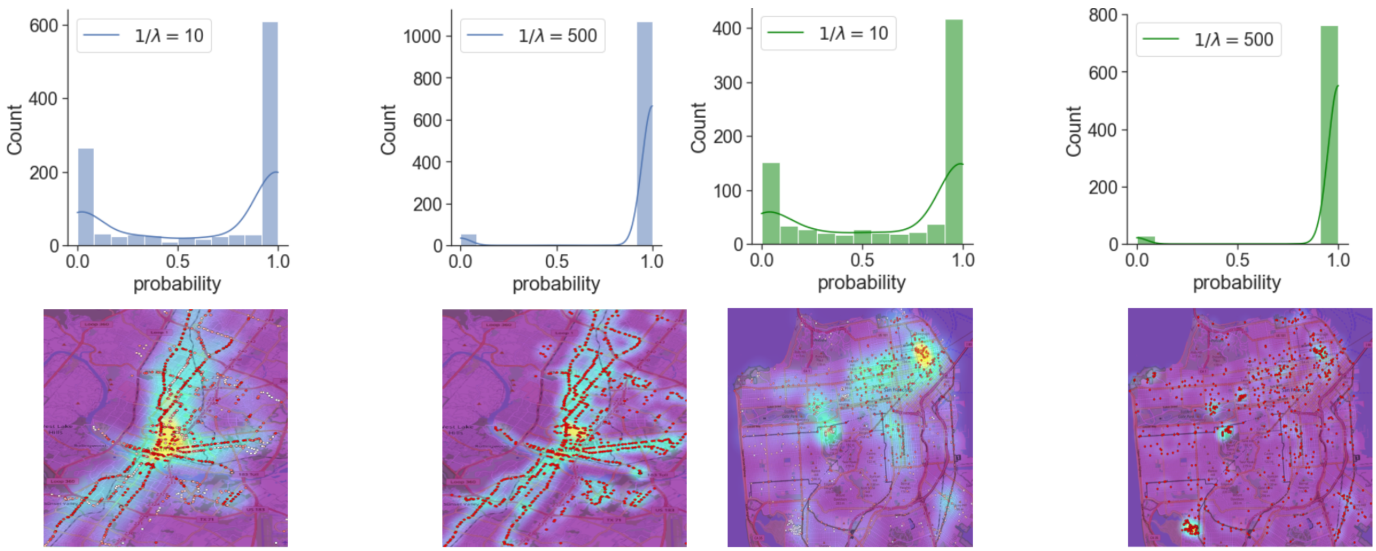

(a) Location coverage model applied to Gowalla data set.

(b) Location coverage model applied to Cabspotting data set.

Fig. 7. (a), (b) Results of the location coverage model applied to the experimental data sets.

The proposed model represents one of the building blocks for a complete deployment strategy for EDCs. We consider two further lines of investigation. On the one hand, we are interested in integrating our model with a predictive tool able to anticipate crowded locations giving the past user's trajectories, e.g. by adopting Recurrent Neural Network approaches. On the other hand, we are interested in investigating the possibility of migrating the computation according to movements of the crowd. This approach allows to further extend the concept of locality to the notion of *proximity* so that to emphasize the possibility of requesting a service to devices in close proximity.

References

[1] P. Mach and Z. Becvar, "Mobile edge computing: A survey on architecture and computation offloading," *IEEE Communications Surveys Tutorials*, vol. 19, no. 3, pp. 1628–1656, 2017.

[2] Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, "A survey on mobile edge computing: The communication perspective," *IEEE Communications Surveys Tutorials*, vol. 19, no. 4, pp. 2322–2358, 2017.

[3] J. Almutairi and M. Aldossary, "A novel approach for iot tasks offloading in edge-cloud environments," *Journal of Cloud Computing*, vol. 10, no. 1, 2021.

[4] W. Lim, J. Ng, Z. Xiong, J. Jin, Y. Zhang, D. Niyato, C. Leung, and C. Miao, "Decentralized edge intelligence: A dynamic resource allocation framework for hierarchical federated learning," *IEEE Transactions on Parallel and Distributed Systems*, vol. 33, no. 3, pp. 536–550, 2022.

[5] C. Li, Y. Zhang, Q. Sun, and Y. Luo, "Collaborative caching strategy based on optimization of latency and energy consumption in mec," *Knowledge-Based Systems*, vol. 233, 2021.

[6] E. Cho, S. A. Myers, and J. Leskovec, "Friendship and mobility: User movement in location-based social networks," in *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '11. New York, NY, USA: Association for Computing Machinery, 2011, p. 1082–1090. [Online]. Available: https://doi.org/10.1145/2020408.2020579

[7] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser, "CRAWDAD dataset epfl/mobility (v. 2009-02-24)," Downloaded from https://crawdad.org/epfl/mobility/20090224/cab, Feb. 2009, traceset: cab.

[8] C. Mouradian, D. Naboulsi, S. Yangui, R. H. Glitho, M. J. Morrow, and P. A. Polakos, "A comprehensive survey on fog computing: State-of-the-art and research challenges," *IEEE Communications Surveys Tutorials*, vol. 20, no. 1, pp. 416–464, 2018.

[9] N. Abbas, Y. Zhang, A. Taherkordi, and T. Skeie, "Mobile edge computing: A survey," *IEEE Internet of Things Journal*, vol. 5, no. 1, pp. 450–465, 2018.

[10] M. Jia, J. Cao, and W. Liang, "Optimal cloudlet placement and user to cloudlet allocation in wireless metropolitan area networks," *IEEE Transactions on Cloud Computing*, vol. 5, no. 4, pp. 725–737, 2017.

[11] J. Martín-Pérez, L. Cominardi, C. J. Bernardos, A. de la Oliva, and A. Azcorra, "Modeling mobile edge computing deployments for low latency multimedia services," *IEEE Transactions on Broadcasting*, vol. 65, no. 2, pp. 464–474, Jun 2019.

[12] P. Vitello, A. Capponi, C. Fiandrino, G. Cantelmo, and D. Kliazovich, "The impact of human mobility on edge data center deployment in urban environments," in *Proc. of IEEE GLOBECOM*, 2019, pp. 1–6.

[13] R. Hariharan and K. Toyama, "Project lachesis: Parsing and modeling location histories," in *Geographic Information Science*, M. J. Egenhofer, C. Freksa, and H. J. Miller, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 106–124.

[14] L. Pappalardo, F. Simini, G. Barlacchi, and R. Pellungrini, "scikit-mobility: a python library for the analysis, generation and risk assessment of mobility data," 2019.

[15] M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi, "Understanding individual human mobility patterns," *Nature*, vol. 453, no. 7196, pp. 779–782, June 2008.

[16] L. Pappalardo, S. Rinzivillo, Z. Qu, D. Pedreschi, and F. Giannotti, "Understanding the patterns of car travel," *The European Physical Journal Special Topics*, vol. 215, no. 1, pp. 61–73, Jan 2013. [Online]. Available: https://doi.org/10.1140/epjst/e2013-01715-5

[17] C. Song, Z. Qu, N. Blumm, and A.-L. Barabási, "Limits of predictability in human mobility," *Science*, vol. 327, no. 5968, pp. 1018–1021, 2010.