



Routine-based network deployment for data offloading in metropolitan areas

Eduardo Mucelli Rezende Oliveira, Aline Carneiro Viana

► To cite this version:

Eduardo Mucelli Rezende Oliveira, Aline Carneiro Viana. Routine-based network deployment for data offloading in metropolitan areas. WCNC - Wireless Communications and Networking Conference - 2014, Apr 2014, Istanbul, Turkey. hal-00862916v1

HAL Id: hal-00862916

<https://inria.hal.science/hal-00862916v1>

Submitted on 17 Sep 2013 (v1), last revised 17 Jan 2014 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Routine-based network deployment for data offloading in metropolitan areas

Eduardo Mucelli Rezende Oliveira
Ecole Polytechnique and INRIA

Aline Carneiro Viana
INRIA

ABSTRACT

This paper tackles the WiFi hotspot deployment problem in a metropolitan area by leveraging mobile users' context, i.e., their trajectories and scenario interaction. The careful deployment of hotspots in such areas allow to maximize WiFi offloading, a viable solution to the recent boost up of mobile data consumption. Our proposed strategy considers the restrictions imposed by transportation modes to people trajectories and the space-time interaction between people and urban locations, key points for an efficient network planning. Using a real-life metropolitan trace, we show our strategy guarantees high coverage time with a small set of deployed hotspots.

1. INTRODUCTION

WiFi offloading seems to be a viable solution to the recent boost up of mobile data consumption that is straining 3G cellular networks in metropolitan areas. [1, 7, 11]. The idea consists in shifting the traffic from cellular networks to WiFi inexpensive ones. Carefully deploying WiFi hotspots can both be cheaper than upgrade the current cellular network structure and can concede substantial improvement in the network capacity [7]. Nevertheless, one question remains: *how WiFi hotspots should be deployed?* The following factors make the answer to this question a challenge task.

The expansion of metropolitan areas increased the possibility of moving around [9]. This fact together with the increase of smartphone use results in highly dynamic links, which may significantly affect the performance of the network [5]. Moreover, people may use different transportation modes, which significantly impacts their trajectories: e.g., a person riding a bike or walking can decide the path to follow contrarily to someone inside a bus. Finally, it is also important to take into account the space-time interaction between people and urban locations, a key point for an efficient network planning. Such considerations can reveal fundamental insights in terms of network usability. Popular sites for instance, are the source of the most of the traffic on the network [6].

In order to consider these issues, *this work tackles the WiFi hotspot deployment problem in a metropolitan area*

by leveraging mobile users' context, i.e., their trajectories and scenario interaction. Our objective is to define what are the best places to receive WiFi hotspots in order to maximize the coverage on an urban scenario. This is a convenient solution for both cellular operators and users: The former can see the traffic being shifted to inexpensive networks while the latter can take advantage of higher data rates and less monetary costs than using cellular networks. We claim that unplanned deployment of hotspots may lead to both under-utilized and over-utilized network areas.

To the best of our knowledge, [13] proposes the only approach to deploy hotspots that considers user mobility characteristics. Our proposal differs from this one in the following main points. First, we consider an urban scenario, which presents significantly higher complexity than the campus scenario considered in [13]: i.e., higher densities, many popular areas, diverse types of mobility (imposed by a variety of transportation modes), bigger area, etc. Second, our objective is not to provide continuous coverage: In an urban scenario, this is prohibitively expensive since it would require the deployment of hotspots over the whole area, where most of them may be under-utilized. Finally, our approach is not restricted to the consideration of only one popular area: In an urban scenario, diverse popular areas may exist and their features may vary according to space-time issues.

Some other related works provide solutions for data offloading but not related to hotspot deployment e.g., delegating the data offloading for people's devices [1, 4, 11].

To accomplish our objective, we study the mobility context of people in a metropolitan area of a major city and identify a set of locations to well deploy WiFi hotspots (cf. Section 2). Our strategy (cf. Section 3) is methodologically structured as follows. First, we create a time dependent graph to represent the interaction between people mobility and locations suitable to receive a hotspot. Then, we measure how much coverage a location can contribute to. For this, we use a metric herein proposed to rank which locations are suitable to support more: Better positioned hotspots are likely to provide better coverage, and therefore, are likely able to offload more data (cf. Section 4).

Through extensive experiments on a real-life trace, we evaluate the performance of our routine-based network deployment in terms of network coverage, by varying the number of hotspots deployed. We also compare our solution with an unplanned deployment. The results reveal that with, on average, 16% of deployed spots, our strategy provides about 75% of coverage time, contrarily to 73% of deployed spots in the unplanned deployment case (cf. Section 5). Finally, Section 6 concludes this work.

2. RATIONALE

Before presenting our strategy for hotspot placement, we provide in this section the insights considered in our solution and discuss our system model.

2.1 General view

People are routinary semi-rational entities, they have regular circles of actions guided by their decisions but unexpected situations may interfere on their directions [3]. A person may change their itinerary due to a traffic jam, problems on the public transportation, etc. When choosing an itinerary, people tend to use the shortest-path to reach their destination, also known as "*desire line*". The desire line is the shortest line between origin and destination, and expresses the way a person would like to go, if such a way were available [12]. Furthermore, the people's itinerary is characterized by its *confinement*, i.e., despite of choosing the shortest-paths, people will roam close by their main physical address [2].

Inspired by these observations, our hotspot deployment proposal considers people context described by how they move. For this, we consider a real dataset describing movements in a large metropolitan area of almost two hundred people and several transportation modes.

In general terms, we want to investigate the possibility of providing a better networking service for people based on their *context*: i.e., their trajectories and scenario interactions. Specifically, in the work herein proposed, we are going to instantiate such service as a data offloading architecture. Therefore, our main question is: *how to provide a context-aware data offloading architecture based on the people's daily trajectories?*

2.2 System model

Our system model represents a fairly real urban scenario composed of 182 people and their routes. Besides, we use data describing more than two hundred thousand real locations in a metropolitan area.

Urban scenario: In this work, we are using GeoLife dataset on its latest version [14]. GeoLife is considered to be unique in the literature. This is due to the fact that it provides a rich view of people mobility using 11 different transportation modes in an urban area for a long period of time. It provides geolocalized and timestamped

points from 182 people during a 4 year span, from 2007 to 2011, mostly in Beijing. For each person, the dataset provides a set of geolocalized points ascendingly sorted by timestamp, i.e., a GPS trajectory.

Trajectories: A trajectory represents how people moves around and is described as a set of points representing GPS coordinates periodically collected. In GeoLife dataset, people may move around building their trajectories using at most ten different transportation modes such as taxi, bike, run, bus, walk, train, subway, car, boat, and motorcycle.

People and points of interest: People move, build their trajectories, and carry mobile devices capable of WiFi communication and able to receive GPS information. While walking by, people may "interact" with points of interest, e.g., bar, bus station, supermarket, etc. Those points of interest describe more than mere locations in the map but they reflect a social aspect: e.g., students are frequently going to meet their colleagues in a coffeehouse close to the university they attend to. Indeed, this represents a routinary behavior that involves not only people but also their interaction with points of interest in a city. In our scenario, points of interest are geolocalized physical venues, e.g., bank, cafe, school, stadium, train station, university, etc.

Coverage: While moving around, a person passes by many points of interest and sometimes, may stop by. We consider that the interaction between a person and a point of interest lasts as long as the former is inside the interaction range of the latter. The interaction range is defined as the WiFi range on a urban scenario. Taking into account the interferences caused by buildings, vehicles, or any other obstacle, we consider a interaction range of 50 meters in our experiments: we consider that if a person is, at most, 50 meters from a point of interest they are able to wirelessly communicate and consequently, the latter covers the former. It is important to enhance that initially a point of interest is not compulsory considered as a hotspot but as potential place to receive a WiFi hotspot structure.

3. PROPOSAL

This section presents how mobility is mapped into a time dependent weighted graph and introduces our hotspot placement strategy.

3.1 Graph creation

Let S be a set of n points of interest $S = \{s_1, s_2 \dots s_n\}$ uniquely identified by their geographical position: $\forall s_{(x,y)} \in S$, where (x, y) is a pair of coordinates of a point of interest s . Since we are using real data traces containing GPS-based positioning, x and y are respectively, latitude and longitude. Fig. 1(b) illustrates a set of spots of interest presented as stars on map. Initially, we assign all

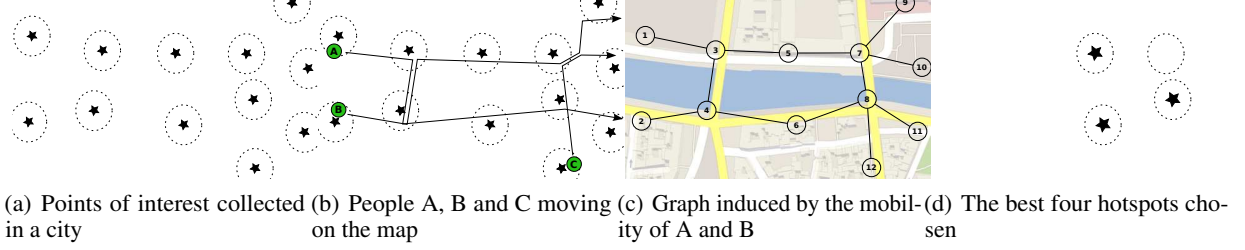


Figure 1: A general view of our proposed methodology.

points of interest with a hypothetical range illustrated as dashed circles¹.

Our graph $G(V, E)$ represents the interaction between people and the urban scenario. The graph should provide a simple and easy-to-see representation of people’s mobility. While walking, people create trajectories and their interactions with the urban scenario may point out important places to start planning the deployment of a network infra-structure. Fig. 1(b) illustrates the three trajectories created by three hypothetical people (A, B, and C) and their respective interactions with points of interest. Let P_i be a set of all geographic coordinates on the trajectory of a person i , i.e., $P_i = \{p_1, p_2 \dots p_o\}$. A vertex $v \in V$ is created on the same coordinates of the point of interest if the latter covers at least one person passing by. Note that a point of interest that does not respect this condition is not considered as vertex in the graph. A vertex coverage of a person i by point of interest s with a range r is expressed as $\forall p \in P_i, \exists p | (p_x - s_x)^2 + (p_y - s_y)^2 \leq r^2$, i.e., i is inside s ’s circular range.

An edge $e \in E$ is created between two vertices if their corresponding points of interest sequentially cover a person during its trajectory. Fig. 1(c) shows the graph when three people are covered by points of interest.

3.2 Hotspot placement strategy

A routine centrality metric herein proposed is calculated on the graph, based on two intuitions: *Desire lines* and *confinement* of people’s mobility.

Desire lines and Stress Centrality: In GeoLife dataset, people may move around building their trajectories using at most ten different transportation modes such as taxi, bike, run, bus, walk, train, subway, car, boat, and motorcycle. Therefore, to capture this urban behavior a trajectory is divided into *legs*. A leg is a subset of points from a trajectory with an unique transportation mode. The concept of “desire lines” states that people tend to choose the shortest-paths to arrive on their destinations. In order to verify that, we have compared the length of each traveled leg against the length of the corresponding shortest path considering the same initial and final points of the original

leg. Dividing the length of the original leg by the length of the shortest path allows us measuring how longer the path made by a person is from the shortest path. We have used Google Directions API² to compute the shortest path. The API receives the coordinates of both initial and final points and a *travel mode*, i.e., transportation mode. Then, it returns the shortest path considering the restrictions imposed by the existing routes and obstacles in the city for a specific transportation mode. Note that, we only have considered transportation modes where people have decision control of their paths. This excludes for example, buses, boats or trains.

Legs traveled by walking, running, and biking had their lengths divided by their respective shortest paths computed while using the API in *walking* mode. Google Directions API indeed has a *bicycling* travel mode, but at the moment, it does not contain routes in Beijing. Legs traveled by taxi, car, and motorcycle had their lengths divided by results of the API in the *driving* mode. Fig. 2(a) shows the CDF of the ratio between the original legs length and the Google shortest path, by transportation mode. It is possible to see that regardless of the transportation mode 69% of the legs measure, at most, half longer than the shortest path. This result shows that the users in our scenario tend to choose the shortest paths when traveling.

Considering such feature, vertices are ranked with the *stress centrality metric* [10]. Vertices with high stress are those that lay on most of people’s shortest routes and may become well positioned hotspots. Vertices with high stress values are likely to be places for hotspot deployment. Indeed, centrality metrics are widely used in the literature, but not yet explored in the problem we are considering.

Confinement and Closeness Centrality: People mobility is generally periodic and confined. Even if people are not using the shortest routes, they are at least not going far from their home location. To check how that premise occurs on our scenario, we have measured how confined the trajectories are. People with confined mobility tend to repetitively visit the same areas. For this purpose, Beijing map was divided into cells of 50 square meters, in a grid shape. Then, we have calculated the number of unique cells (*NUC*) and total cells (*NTC*) that each trajectory

¹Unit disk graph model.

visited. To quantitatively express the confinement of a trajectory, a metric called *Repetitiveness* was conceived and is represented as the following: $\frac{NTC-NUC}{NTC} * 100$

Repetitiveness represents how repetitive is the route described by a trajectory. Fig. 2(b) shows the results for repetitiveness for all trajectories grouped by period of the day. It is possible to see that regardless of the period of the day, trajectories present a considerable amount of repetitivity: Almost 90% of the performed trajectories present at most 50% of the visited cells repeated.

On our proposal, we consider *Closeness Centrality metric* [8] as a way to measure how likely a place will be to receive repeated visits. Closeness is calculated based on the geodesic distance between all pairs of vertices on the network. It assigns higher values for vertices closer to the rest of the network, that is, on a city those are probably hospitals, and markets, i.e., places planned to be close to most of the people's trajectories.

Centralized computation: The centrality metrics are calculated on a offline centralized manner, i.e., the solution herein proposed may be applied on data previously collected by the interested entity, e.g., a Telecom operator. This premise is reasonable since network deployment planning generally uses historical data as input. The result is a set of points of interest ranked by how good they are in providing people coverage. From that point, the operator needs to decide how many hotspots they are willing to deploy and the trade-off between the monetary cost and the predicted amount of WiFi coverage.

Spots ranking: To couple with all the above intuitions, we introduce *Routine centrality* (ρ), a centrality metric that aims to rank spots based on their potential to offload data: $\rho(v) = C_s(v) + C_c(v)$ where $v \in V$, C_s is the normalized Stress Centrality of v and C_c is the normalized Closeness Centrality of v . All normalized values are in the interval $[0, 1]$. Fig. 1(d) shows an example of the best four hotspots (stars) chosen from ρ .

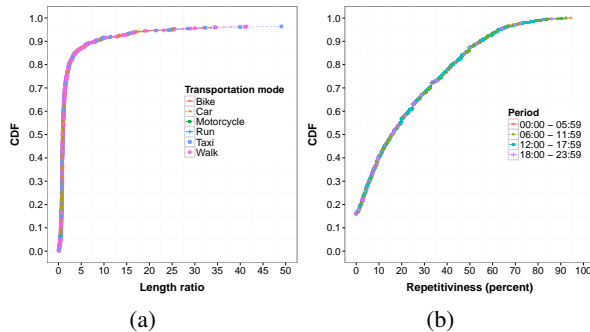


Figure 2: (a) CDF of the ratio between the length of traveled legs and the shortest paths calculated using Google Directions API. (b) Repetitiveness of trajectories by period of the day.

4. EXPERIMENTAL SETUP

This section describes technical details from two components of our experimental setup: trajectories and points of interest.

General Information: All components are based on geolocalized information, i.e., latitude and longitude coordinates within a 1177 km² central area in Beijing. Besides, all evaluations were made based on a subset of Geolife's dataset. Since our work bases its premises on routinary behavior present on the mobility of people, a subset of Geolife data is already enough to show how we can explore routine to provide a better user experience. Therefore, the data of the two most active months are selected. This subset spans from 1st November to 31th December of 2008 and contains 39 users and 2203 trajectories. Moreover, to better understand specific behaviors inherent from different periods of the day, every day is divided into four periods of 6 hours, every day is divided into four periods of 6 hours, from 00:00 to 05:59, from 06:00 to 11:59, from 12:00 to 17:59, and from 18:00 to 23:59.

Trajectories: As previously mentioned, a trajectory is a set of geolocalized points describing a route traveled by a person using at least one transportation mode. In GeoLife dataset the average sample rate between each trajectory point falls between 5s and 10s. Furthermore, every point of a trajectory may contain extra information: transportation mode label that describes which transportation mode was used on that point of the trajectory.

Points of interest: The proposed graph is built based on a set of points of interest in Beijing. Points of interest are places generally present in most of major cities, e.g., bars, universities, supermarkets, etc. In order to be able to points of interest in Beijing, databases of places (e.g., Google Places³) were used. Such databases are growing and are the most accurate source of public information about points of interest. To avoid to be biased by the characteristics of one unique database, we have collected data from multiple sources namely Google Places, Nokia Maps, and Foursquare, counting more than 202 thousand real unique points of interest with their respective IDs, latitudes and longitudes. While Google Places and Nokia Maps databases provide information about points of interest collected from city hall, owner of venues, etc, Foursquare provides only information from places where its users checked in, generally places related with leisure and social relationships. For each set of places collected from a source, repeated ones were removed by keeping an unique occurrence of each ID. The set of collected points of interest contains 77919 from Google Places, 119346 from Nokia Maps, and 5059 from Foursquare totalizing 202324 places distributed in 98 categories, e.g., market,

³<https://developers.google.com/maps/documentation/directions>

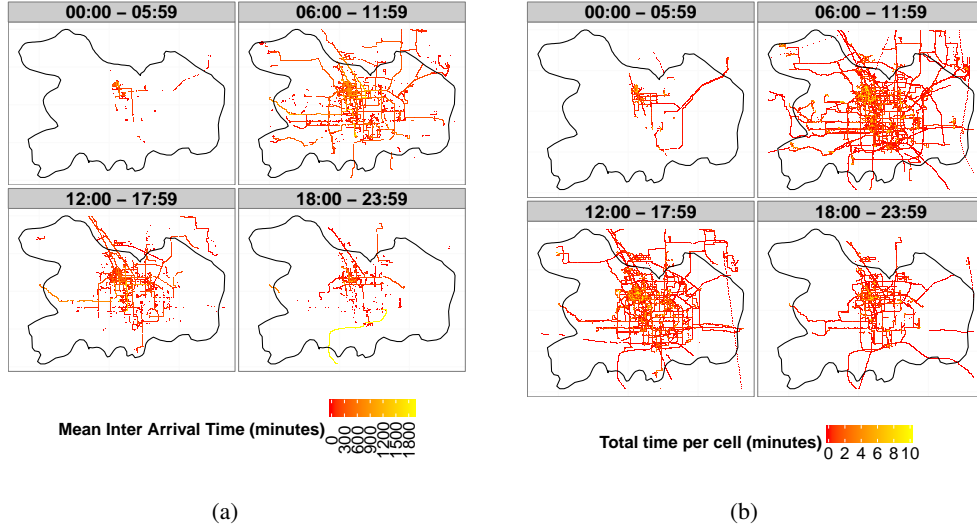


Figure 3: (Better visualised in colors) (a) Heatmap of the inter arrival time on the cells. (b) Heatmap of the time spent inside cells. For better visualization, this picture shows only visits lasting at most 10 minutes, i.e., 98% the cells are depicted.

library, school, etc.

5. PERFORMANCE EVALUATION

To better understand the results of our proposal, hereafter, we have first studied users’ mobility characteristics and then assessed our main results.

5.1 Characterization

To understand how people explore the regions, central Beijing was divided into square-shaped cells of $50 m^2$ summing up 181 thousand cells. First, we have analyzed the inter arrival time (IAT). The IAT of a cell A represents the time between visits of people on A and may show how frequently A is visited. Therefore, this metric also may point out how popular a cell is considering that small values of IAT indicates that a cell is frequently visited and it is likely to be a cell inside a popular area. On the other hand, high values of IAT indicates that a cell is not frequently visited and probably does not belongs to a popular area.

Fig. 3(a) (better visualized in colors) shows a heatmap of cells’ mean IATs in central Beijing, and by period of the day. It is important to enhance that cells showed in this figure are a subset of all visited cells: A cell must have been visited at minimum twice, i.e., at least one value of IAT was computed for it to appear in the heatmap (125 thousand cells). Regardless of the period of the day, two conclusions can be made from this figure. First, most of the cells have small IAT values: 81% of the cells present

mean IAT less than 1 hour. Therefore, it shows a dynamic city in which lots of people are frequently passing by the cells. Second, a north central area consistently shows a high concentration of activity, which is due to the fact that Microsoft Research Asia headquarter is located there. GeoLife experiment was conducted mostly with Microsoft members that were constantly walking nearby the working place.

As expected, the late night period which comprises from 00:00 to 05:59 shows few activity. Periods from 06:00 to 11:59 and 12:00 to 17:59 show, as expected, the highest levels of activity. That is due to the fact that both comprise the most active hours of the day, i.e., people going from home to work and vice-versa, lunchtime, etc. On both periods, cells along roads that connect the central city with the periphery present high IAT values. It is due to the fact that those roads are used to bring people to the city and then, few hours later, take them back when going to home. This behavior is enhanced on the period from 18:00 to 23:59, a specific road on the north central area has almost 25 hours of IAT. That is due to a specific person that went to the city center one day and came back on the day after using the same road.

Fig. 3(b) (better visualized in colors) shows a heatmap of the total time that people spent (TTS) on the cells by period of the day. Regardless of the period of the day, most of cells present low values of TTS, mainly cells along roads. Indeed, 95.8% of the cells present TTS value less or equal 3 minutes. As for IAT results, Fig. 3(b) confirms that the mobility in Beijing is very dynamic, i.e., people are mostly passing by than staying still for long periods of time.

³<https://developers.google.com/maps/documentation/places>

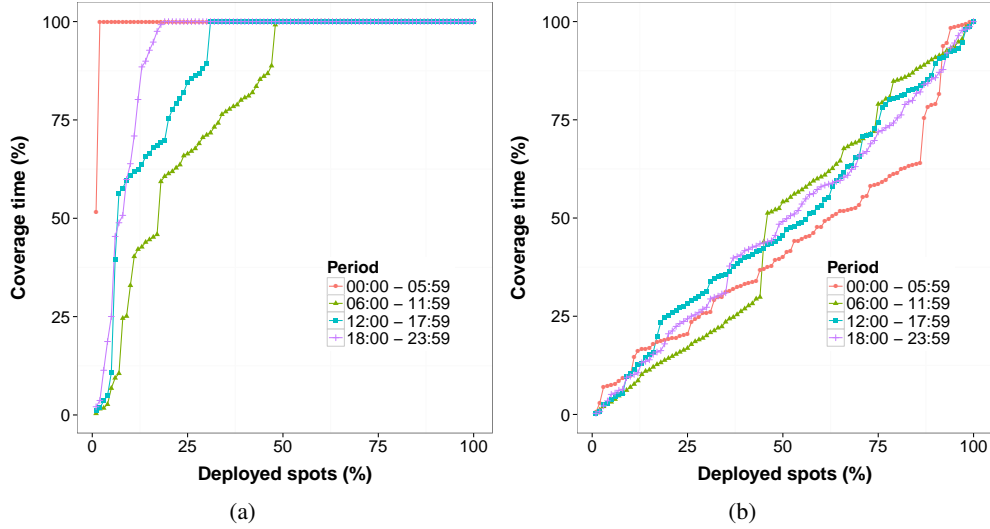


Figure 4: Percentage of time of users’ trajectories covered upon deposition of hotspots (a) ranked by Routine Centrality and (b) selected uniformly at random.

Although this dynamic behavior, people’s mobility in Beijing presents high repetitiveness. Therefore, people will likely have a chance to be covered for longer periods if considering the aggregation of their several repeated short-period coverage.

5.2 Routine Centrality

We consider coverage time as the amount of time people were inside the range of a spot. To understand how covered people are while moving, we have measured the total amount of coverage time the hotspots provided. To assess hotspots effectiveness on providing coverage, we have computed the percentage of coverage time. For this, after hotspots are ranked by Routine Centrality, we deploy them accumulatively, e.g., best 1%, best 2% and verified what was the percentage of coverage time they provided: By dividing the total time people have been covered in their trajectories by the time they were uncovered.

Fig. 4(a) shows the percentage of coverage time as a function of the percentage of deployed hotspots according to the rank given by the Routine Centrality. Furthermore, this result is presented on four curves representing the periods of the day. Indeed, every period of the day has its own characteristics, e.g., late night presents less activity and people are concentrated in areas surrounding night clubs, bars, etc., while working hours are likely to have people scattered throughout the city. Periods from 00:00 to 05:59 and from 18:00 to 23:59 have less people and using low percentage of hotspots, 2% and 11% respectively, it was possible to cover about 75% of the time. Periods from 06:00 to 11:59 and from 12:00 to 17:59 present more activity and consequently require more efforts to guarantee high coverage rates. They require 34% and

20% of deployed hotspots to achieve 75% of coverage.

In order to measure how effective our routine-based network deployment approach, we have performed comparison with a simpler strategy that deploy hotspots select at random. Fig. 4(b) shows the percentage of coverage time as a function of percentage of deployed hotspots when selected uniformly at random. On contrary to our proposal, the random deployment provides poor coverage time for people in the city: For all periods, it was needed about 73% of hotspots to provide coverage of about 75%. Finally, our results show that a routine-based spot deployment can provide with lower costs a high percentage of coverage time compared to a random deployment. In addition, the coverage is dependent on the period of the day which make us think in a dynamic strategy. Indeed, in order to save energy certain hotspots could be turned off during some periods of the day in which they would provide low or no coverage. This makes possible a adaptative solution in which each period of the day would have only the best hotspots turned on.

6. CONCLUSIONS AND NEXT STEPS

In this work, we have presented to the best of our knowledge, the first analyses of a metropolitan-wide hotspot deployment. We have proposed a graph model to represent the relationship between people and the city infra structure. Furthermore, based on common behaviors presented on people’s real routines, we have proposed a centrality-based metric to rank points of interest in a city that are capable of providing network coverage. Our results show that with a small quantity of spots it is possible to provide high percentage of coverage time.

As future work, we intend to further investigate the mobility characteristics and aggregate a realistic traffic generation model. Therefore, we are going to be able to evaluate traffic offloading and verify the relation between coverage and the amount of offloaded data.

7. REFERENCES

- [1] BARBERA, M. V., STEFA, J., VIANA, A. C., DE AMORIM, M. D., AND BOC, M. VIP delegation: Enabling VIPs to offload data in wireless social mobile networks. In *Distributed Computing in Sensor Systems and Workshops (DCOSS), 2011 International Conference on* (June 2011), IEEE, pp. 1–8.
- [2] BROCKMANN, D., HUFNAGEL, L., AND GEISEL, T. The scaling laws of human travel. *Nature* 439, 7075 (Jan. 2006), 462–465.
- [3] HAMERMESH, D. S. Routine. Working Paper 9440, National Bureau of Economic Research, January 2003.
- [4] HAN, B., HUI, P., KUMAR, V., MARATHE, M., SHAO, J., AND SRINIVASAN, A. Mobile data offloading through opportunistic communications and social participation. *Mobile Computing, IEEE Transactions on* 11, 5 (May 2012), 821–834.
- [5] JUDD, G., WANG, X., AND STEENKISTE, P. Efficient channel-aware rate adaptation in dynamic environments. In *Proceedings of the 6th international conference on Mobile systems, applications, and services* (New York, NY, USA, 2008), MobiSys '08, ACM, pp. 118–131.
- [6] KIM, M., KOTZ, D., AND KIM, S. Extracting a mobility model from real user traces. In *INFOCOM 2006. 25th IEEE International Conference on Computer Communications. Proceedings* (April 2006), pp. 1–13.
- [7] LEE, K., RHEE, I., LEE, J., YI, Y., AND CHONG, S. Mobile data offloading: how much can wifi deliver? *SIGCOMM Comp. Comm. Rev.* 41, 4 (Aug. 2010).
- [8] SABIDUSSI, G. The centrality index of a graph. *Psychometrika* 31, 4 (1966), 581–603.
- [9] SETO, K. C., FRAGKIAS, M., GNERALP, B., AND REILLY, M. K. A meta-analysis of global urban land expansion. *PLoS ONE* 6, 8 (August 2011), e23777.
- [10] SHIMBEL, A. Structural parameters of communication networks. *The bulletin of mathematical biophysics* 15, 4 (1953), 501–507.
- [11] THILAKARATHNA, K., PETANDER, H., AND SENEVIRATNE, A. Performance of content replication in mobitribe: A distributed architecture for mobile ugc sharing. In *Local Computer Networks (LCN), 2011 IEEE 36th Conference on* (October 2011), pp. 554–562.
- [12] THROGMORTON, J. A., AND ECKSTEIN, B. Desire lines: The chicago area transportation study and the paradox of self in post-war america, Nov. 2000.
- [13] WANG, T., XING, G., LI, M., AND JIA, W. Efficient wifi deployment algorithms based on realistic mobility characteristics. In *IEEE MASS* (Nov. 2010), pp. 422–431.
- [14] ZHENG, Y., XIE, X., AND MA, W.-Y. Geolife: A collaborative social networking service among user, location and trajectory. *IEEE Data Eng. Bull.* 33, 2 (2010), 32–39.