Temporal Coverage Analysis of Router-based Cloudlets Using Human Mobility Patterns

Christian Meurisch*, Julien Gedeon*, Artur Gogel*, The An Binh Nguyen*, Fabian Kaup*,
Florian Kohnhäuser*, Lars Baumgärtner†, Milan Schmittner*, Max Mühlhäuser*
*Technische Universität Darmstadt, Germany
Email: {meurisch@tk, gedeon@tk, gogel@tk, nguyen@kom, fkaup@ps,
kohnhaeuser@seceng.informatik, mschmittner@seemo, max@tk}.tu-darmstadt.de
†Philipps-Universität Marburg, Germany
Email: lbaumgaertner@mathematik.uni-marburg.de

***Abstract*—Responsive applications such as augmented reality require nearby computational offloading units with low latency. The concept of cloudlets is one promising approach that satisfies these requirements. However, due to their wireless range restrictions cloudlets have always been faced with deployment issues of achieving high spatial coverage. In this paper, we look at the coverage issue from an end-user's perspective instead of an established provider perspective: we first argue that temporal coverage of cloudlet accessibility is preferable to spatial coverage for an end-user's experience considering his daily mobility behavior. Next, we investigate what is necessary to achieve a high temporal coverage for an individual user and to what extent is temporal coverage realizable with concepts like router-based cloudlets. To show our hypothesis and understanding the temporal coverage aspect, we collected two comprehensive datasets, an access points dataset with estimated location information and a human mobility dataset consisting of mobility traces from 30 participants within a major city over 4 weeks. Our analysis results show that high temporal coverage can be achieved by a relatively small set of router-based cloudlets since students mainly stay at two places, their homes and university, which represent a large part of the temporal coverage. The remaining rate at which coverage increases heavily depends on the user's mobility patterns. Our findings can be used to place router-based cloudlets at the right locations and estimate the number needed to achieve a certain temporal coverage in urban environments.**

I. Introduction

Cloudlets are small-scale decentralized computation units located in the near of mobile users to perform resource-intensive tasks [1]. Due to the proximity, cloudlets overcome latency and network traffic issues, which benefit immersive applications such as augmented reality [2]. However, depending on the cloudlet deployment (e.g., hosted by an Internet service provider, a local business or private households), performances and range restrictions vary [3]. Especially, the low spatial coverage of the used short-range wireless technologies (e.g., WiFi) requires a dense deployment which blocks the breakthrough of cloudlets as ubiquitous computing resources.

As an example, to achieve a nearly complete spatial coverage of cloudlets in an urban area like Paris more than one million computation units need to be deployed assuming a circular WiFi range of $25m$ and a homogeneous distribution. However, this deployment approach is far from the reality due to the vast amount of required computation units to be installed and operated. Moreover, pursing a high spatial coverage does not inherently consider either popular areas with crowd of people or rural areas with a low density of potential end-users.

This small thought experiment demonstrates one of the practical deployment challenges of cloudlets to achieve a nearly complete spatial coverage. Many research works already studied the spatial coverage in different fields such as participatory sensor networks [4], mobile data offloading to the cloud [5], or cloudlet deployments [6]. Only a few works (e.g., [7]) rethink the deployment challenge of range-restricted technologies from a user's perspective to achieve a high end-user's satisfaction, which can be achieved by a high temporal coverage of cloudlet services. However, these works focus on a deployment scheme of access points (AP) for delayed mobile data offloading and mainly rely on simulations.

In the context of cloudlet deployment, temporal coverage is a metric about quality of service delivered by cloudlets and crucial to achieve a high individual user satisfaction. We believe that the concept of *router-based cloudlets* - upgrading wireless home routers as cloudlet [8] - are a promising lightweight approach to reach this temporal service coverage for responsive mobile applications such as augmented reality due to router's dense distribution in urban environments.

In this paper, we study temporal coverage on the concept of router-based cloudlets by considering real-world access points and human mobility data. For that, we collect two comprehensive datasets: (i) an access points dataset from a major city systematically gathered by over 20 volunteers through *wardriving* techniques [9], and (ii) a human mobility dataset consisting of mobility traces from 30 participants living in the same city over 4 weeks by using our tracking application [10]. Using both datasets, we first understand daily human mobility to answer the questions *what is required to achieve a high temporal coverage for an individual user* and *to what extent is this temporal coverage realizable with router-based cloudlets*. Our analysis results show that high temporal coverage can be achieved by a relatively small set of cloudlets at the right locations such as home and university (respectively work place). We also confirm that user's mobility patterns have a high impact on the rate at which temporal coverage increases.

In summary, the contributions of this paper are twofold:

- We collect two comprehensive datasets: (i) nearly 20,000 wireless access points data with estimated location information within a major city; and (ii) human mobility traces consisting of over 5 million location values from 30 participants over 4 weeks.
- To the best of our knowledge, this is the first work that investigates and understand the temporal coverage of router-based cloudlets from a user's perspective using real-world data, and shows up their vast potential for a large-scale deployment in urban environments.

The remainder of this paper is organized as follows. *First*, we give an introduction in router-based cloudlets and provide an overview of the related work. *Second*, we report the conducted studies and characterize the collected datasets. *Third*, we present our extensive analysis results. The paper closes with discussion and conclusion.

II. Background and Related Work

The emergence of mobile devices such as smartphones as daily companion and the increasing use of *mobile computing* require solutions to overcome problems such as resource scarcity, or mobility. *Mobile cloud computing* can address these problems by leveraging resourceful data centers that are distant (aka *the cloud*) or closely located (aka *edge servers*) [11]. One approach is the concept of *cloudlets* which are small-scale computation units installed in the public infrastructure (e.g., within coffee shops) [12]. However, since the introduction of cloudlets in 2009 [1], this concept is faced with coverage issues due to their inherently range restrictions [3].

In [8], the authors propose upgrading wireless home routers as cloudlets to enable a dense offloading infrastructure, which is maintained and operated by household owners. Various incentives mechanisms (e.g., [13]) can further increase the owner's willingness to provide this functionality to others. Equipped with a micro uninterrupted power supply [14], the router-based cloudlet concept can also be used in emergency situations, e.g., as communication bridge [15] in infrastructure-less wireless networking [16], for executing remote procedure calls in delay-tolerant networks [17], as broker for surrogate discovery [18], or for distributed in-network processing [19]. Since we believe that *router-based cloudlets* are a promising decentralized approach for a large-scale deployment, we use this concept for our analysis.

Finding the ideal placement of cloudlets to achieve a high temporal coverage instead of spatial coverage is extremely challenging. In literature, the deployment of wireless APs has been studied extensively (e.g., [7], [20]), especially in the context of mobile data offloading [5]. For instance, the authors of [7] investigate the wireless access point deployment from a user's point of view. However, this work focuses on offloading mobile data to the cloud and relies on simulations. The same is true for other existing works, e.g., the authors of [6] investigate a large-scale deployment of cloudlets in terms of achieving high spatial coverage, but do not consider real-world conditions and the temporal aspect.

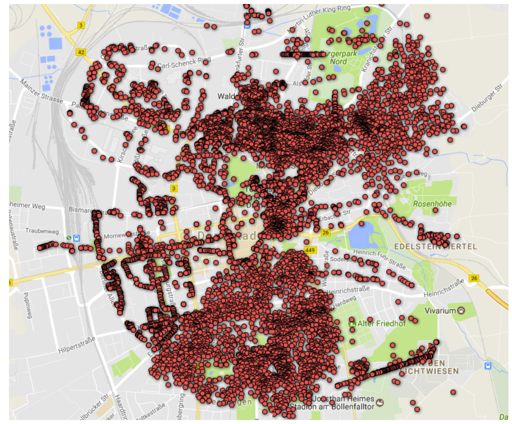

Fig. 1. Locations of wireless APs collected in the city of Darmstadt, Germany

Inspired by scientific works like [7], we investigate the deployment of router-based cloudlets from a user's perspective. For that, we first need to understand user's mobility [21], and user's daily routines [22]. Considering human mobility is essential for building smarter systems or deploying communication technologies in a useful way since the user's mobility has an impact on the technology performances (e.g., [23]), or can be used to further optimize them (e.g., [24]). For instance, the authors of [24] propose *BreadCrumbs*, a system that utilizes user's mobility, which is highly predictable (e.g., examined by [25], [26]), to forecast network connectivities. These predictive knowledge of connectivity changes can then be used by mobile devices to schedule their network usages more intelligently and conserve valuable resources.

Taking human mobility into account, we study and understand the temporal coverage of router-based cloudlets from a new user's perspective for further increasing user satisfaction. With our study, we aim to improve the placements of router-based cloudlets and support their deployments by considering human's demands. In summary, to the best of our knowledge this is the first work that investigates the temporal coverage aspect of router-based cloudlets from an end-user's point of view using real-world access points and human mobility data.

III. Collected Datasets

In this section, we briefly report our conducted studies to collect two comprehensive real-world datasets, an *access points dataset* and a *human mobility dataset* within a German major city, namely Darmstadt, with a population of around 150,000 inhabitants. We further clean and characterize both datasets to serve as the basis for our experimental analysis.

A. Access Points Dataset

To get a consistent up-to-date dataset of visible real-world access points (AP) from the city of Darmstadt (Germany), over 20 volunteers systematically walked through the city and collected the signals from wireless access points (aka *wardriving* [9]). We estimate the actual location of an access point by using trilateration methods on multiple data samples for this AP [15]. To further increase the data quality, we apply

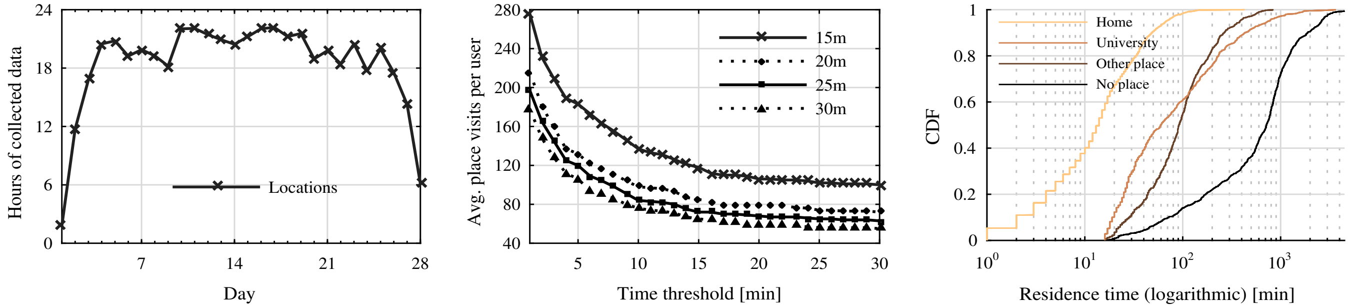

(a) Avg. temporal coverage of collected locations (b) Avg. number of significant places per user (c) Time spent in significant locations

Fig. 2. Analysis of location values collected from all 30 participants

some basic filters on the resulting data to eliminate duplicated MAC addresses, or remove temporary mobile access points by looking up the manufacturers from the *Organizationally Unique Identifier* part of a MAC address.

The resulting dataset consists of 19,311 AP locations in an area of approximately $63\,km^2$ covering core districts of the city of Darmstadt (cf. Fig. 1), i.e., an average density of about 3 routers per 10,000 square meters (100m by 100m area). As can be seen in Figure 1, the distribution of wireless access points is really dense in uptown, and sparse in parks or industrial areas, which reflects the common router deployment pattern in urban environments [15]. Assuming all 60,000 households[1] and shops located in the mentioned area are equipped with wireless home routers, we were able to collect one out of three existing access points. Since the access points are collected while walking through the city and not inside (tower) buildings or private locations, our dataset still reflects the same usage context as a mobile user, which makes the dataset adequate for our analysis.

B. Human Mobility Dataset

To get a human mobility dataset, we conducted a user study with 30 participants living in Darmstadt (Germany) over four weeks. All participants are students of Technische Universität Darmstadt and pursue their Master of computer science. The average age of the participants is 25.4 years. In terms of gender, 24 are male (80%) and 6 are female (20%). We use a light version of our Android data collection app (cf. [10]) to automatically track required location data in the background.

The dataset consists of over 5 million location values $(173,109 \pm 100,369$ per user) with a median accuracy of 30 meters. Figure 2a shows the average temporal coverage of our dataset. On average, we have data accounting for approximately 77.1% of the time since the phones have been deployed, i.e., 18.5 ± 4.8 hours of location values per day. Since the temporal coverage of the first two and last two days of the user study is really sparse, we remove these days from the following investigations. More precisely, we use data from 24 valid days out of 28 study days, which results in a daily collection coverage of 20.1 ± 1.6 hours (83.8%).

[1]https://www.darmstadt.de/standort/statistik-und-stadtforschung/datenreport-2016/bevoelkerung

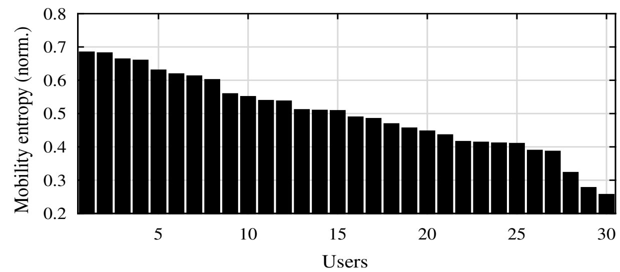

Fig. 3. Normalized mobility entropy of participants

IV. Experimental Analysis

In this section, we first analyze location traces of users to *understand human mobility*. Considering these new insights in human behavior, we study and *understand temporal coverage*, as well as investigate how end-users can achieve a high temporal coverage using router-based cloudlets.

A. Understanding Human Mobility

Based on the collected location values, we first extract significant places and place visits [27] of each individual user using a spatiotemporal clustering algorithm proposed by [28]. For that, we need to choose suitable distance th_{dist} and time thresholds th_{time} as clustering parameters. Thus, we investigate how the average number of significant places per user changes as a function of the distance and the time thresholds. In Figure 2b, we can see that the average number of significant place visits decreases when the time threshold increases. The same is true for the distance threshold: if the distance threshold increases, the average number of place visits decreases. As suitable choices, we choose $th_{dist} = 25\,m$ since the curves make not much of a difference when the distance threshold further increases, and $th_{time} = 15\,min$, where the curve converges to an asymptote of a fixed number of significant place visits. Using these clustering parameters, we extract 631 places (21.0 ± 8.2 per user) and 2,353 place visits (78.4 ± 19.2 per user). However, most places (e.g., restaurant) are only visited once within the study, which make them unsuitable for our analysis. Thus, we cluster these places and identify four significant place categories over all users,

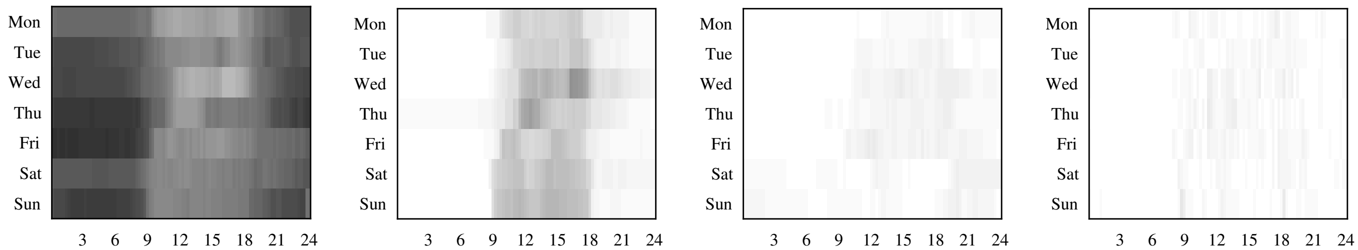

(a) Home (b) University (c) Other place (d) No place / Transition

Fig. 4. Temporal distribution of participant's daily mobility patterns revealed from our collected dataset (the intensity reflects the number of participants over the time of day, where *black* is a high population and *white* a low population)

namely *home*, *university*, *other places*, and *no places* (i.e., a transition between two places or the user is on the way).

Figure 2c shows the cumulative distribution function of the resulting residence times for the particular places (i.e., how long a user stays there). We see that a student is mainly at *home* (M = 860.5min, SD = 850.2min, MD = 718min), which is obvious since students primarily sleep and learn at home. The median residence time at the *university* including canteen or library is about 1.5h (M = 128.6min, SD = 121.5min, MD = 92min), which represents the normal duration of a lecture or exercise group. While the median residence time at *other places* (M = 181.5min, SD = 339.1min, MD = 60min) such restaurant/bar, gym, or friend's home is lower than the residence time at the university, students still stay there much longer on an average. As expected, residence times at *no places* or transitions between places are very short (M = 21.1min, SD = 25.5min, MD = 14min) since most students live in the city where the university is also located. Comparing to other scientific works (e.g., [22]), our participants show similar behavior patterns.

Figure 4 shows the temporal distribution of participant's daily movement patterns as function of weekdays and daytimes. We can observe that students are mainly at *home*, especially in the period between 6pm and 9am (*evening* and *night epoch*) (cf.Fig. 4a). There is one exception on Friday and Saturday evening and night, where some students go out, and stay at *other places*. During the *day epoch*, which is between 9am and 6pm, students also leave their homes to go to university (cf. Fig. 4b). Note that some students stay at university on weekends, e.g., to learn in the library. Since most students do not have organized daily routines, the stays at other places are not well-defined and concentrate to the day epoch (cf. Fig. 4c). This is also true for *transitions* between places, which are irregular distributed through the day and evening period (cf. Fig. 4d).

To further quantify the mobility of participants, we apply an information entropy metric on our mobility dataset. In information theory, the amount of randomness in a signal corresponds to its entropy as defined by Shannon [22]:

$$H(x) = -\sum_{i=1}^{n} p(i)\ log_2\ p(i) \tag{1}$$

Using Equation 1, we calculate individual weekly entropies

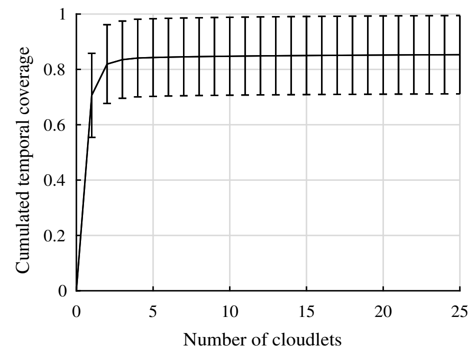

Fig. 5. Temporal coverage over upgraded router-based cloudlets

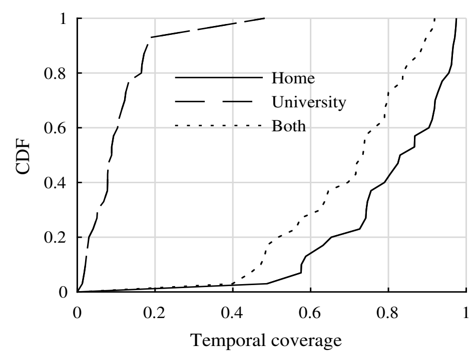

Fig. 6. Distribution of specific temporal coverages at both most visited places (home and university) in student's daily life

for users' mobility patterns by drawing 144 samples for each day. In our case, the maximal theoretical entropy of mobility is about 2.32. We normalize the entropy values between 0 and 1 to make it comparable to other scientific works (cf. Fig. 3). We see that most of our participants live entropic lives which tend to be more variable [22]. Moreover, we do not find a significant gender and entropy dependency in our dataset.

B. Understanding Temporal Coverage Using Router-based Cloudlets

We define *temporal coverage* T^i for a single user i as a temporal union $\bigcup^T$ of all time ranges covered by all mobility traces M^i, where the user has access to an upgraded router as cloudlet. Thus, $T^i = \bigcup_j^T T_{m_j}$, where T_{m_j} is the time range covered by mobility trace m_j. In general, coverage is

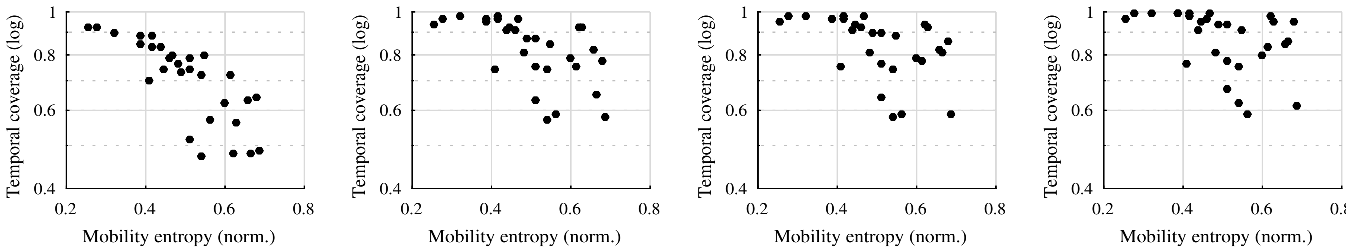

(a) n=1 (b) n=2 (c) n=3 (d) n=max

Fig. 7. Temporal coverage as a function of normalized mobility entropy and n upgraded router-based cloudlets

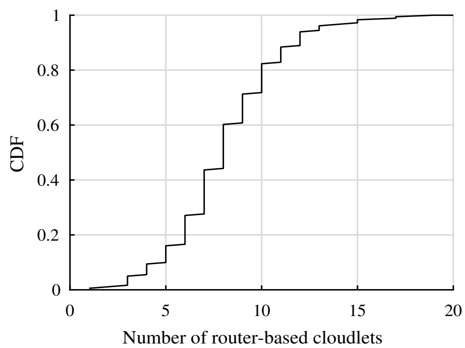

Fig. 8. Required router-based cloudlets for achieving $\geq 99.9\%$ temporal coverage at *walking transitions* $\geq$ *10min*

often defined as *k-coverage*, i.e., a time frame is said to be covered if the user has access to at least *k* cloudlets [4]. For our study, we assume *1-coverage* (k=1), i.e., a single router-based cloudlet covering the time is sufficient. Given this definition, we are able to study and understand the temporal coverage of router-based cloudlets in our real-world data set.

Figure 5 shows the cumulated temporal coverage averaged over all users as function of accessible router-based cloudlets. We can see that only two cloudlets located at the right places (home and university) are sufficient to achieve a high temporal coverage over 80% (MD = 81.9%, SD = 14.2%). Each additional router-based cloudlet only increases the coverage minimally. On average, we only reach 85.3% of a complete temporal coverage after considering 25 cloudlets.

As expected from the mobility data, placing two stationary cloudlets at the user's home (e.g., upgrading his home router) and at the university (e.g., upgrading his laptop computer as router substitute) are crucial. For that, we further investigate the temporal coverage, which can be provided by these two cloudlets. Upgrading user's home router provides over 80 % temporal coverage in three out of four cases (cf. Fig. 6). This finding is highly relevant for (service) providers, which offer responsive applications that rely on offloading resource-intensive tasks to nearby cloudlets. More importantly, upgrading the user's own router is straightforward, as well as does not require advanced privacy and authentication mechanisms. On the contrary, using a public router upgraded as cloudlet (e.g., located at the university) requires a more sophisticated security concept, which is an important aspect for our future research. Placing a router-based cloudlet at the university can achieve up to 20% temporal coverage gain for most students. A temporal coverage up to 97.4% for an individual user can be achieved by placing cloudlets at both crucial places.

Next, we understand the daily mobility impact on the temporal coverage of router-based cloudlets (cf. Fig. 7). As expect, it is more difficult to achieve temporal coverage for users with a higher entropic lifestyle (cf. Fig. 7a). In contrast, upgrading user's home router is sufficient to achieve more than 90% temporal coverage over an entire day for a user with low mobility entropy ($H_m = 0.25$).

Considering a higher number of router-based cloudlets, we observe a slight gain of the temporal coverage (cf. Fig. 7b-7d). This results from the fact that much more cloudlets are needed if the user is on the way; due to the user's mobility and the cloudlet's range restriction the connect time is shorter, and thus the temporal coverage of a single cloudlet is smaller compared to the stationary case. For instance, Figure 8 shows the cumulative density function for achieving more than 99% temporal coverage at a walk of 10min or longer; for 80% of all transitions it is required to access less than 10 cloudlets.

V. Discussion and Limitations

We conducted an automatic self-tracking user study with 30 participants over four weeks to collected human mobility data. On average, we have data accounting for more than 80% of the time since the phones have been deployed. The missing data can be explained due to data corruption or powered-off devices. Especially, the error of powered-off devices is hard to avoid since users often turn off their phones at night or the battery dies. Therefore, we only consider time ranges in our study where enough location values of the user are available.

We saw that most of these participants live an entropic lifestyle, which is typical of students. Selecting another target group, for example, office workers with better organized daily routines and lower mobility entropies compared to students, we would expect similar results; cloudlets located at their homes and work places are crucial for achieving high temporal coverages, which we will investigate in future research.

Moreover, the conducted user studies and our findings are limited to urban environments, where we find a dense distribution of wireless home routers. In future we plan to investigate the temporal coverage aspect of cloudlets in rural areas as well, and show up ways to achieve an adequate temporal coverage there, e.g., by upgrading cellular antenna towers or street lamps as cloudlets with low latencies.

VI. Conclusions and Future Work

In this paper, we investigated the coverage issue of cloudlets from an end-user's point of view. We argue that temporal coverage is more important for the user's experience than spatial coverage considering his daily mobility behavior. We further collected two comprehensive datasets to study and understand human mobility as well as temporal coverage of router-based cloudlets. The results showed that students mainly stay at two places, their homes and university, which represent a large part of the temporal coverage. The remaining rate at which coverage increases heavily depends on the user's mobility pattern. Our findings can be used to place router-based cloudlets at the right locations and estimate the number needed to achieve a certain coverage in urban environments.

In future work, we plan to investigate whether smart street lamps can be upgraded as cloudlet and use for a large-scale deployment by analyzing the coverage aspect. Utilizing stationary street lamps - one of the densest powered infrastructure in urban environments - we would further extend the router-based cloudlet infrastructure. Moreover, we will take greater account of the mobility of end-users to deliver computational results through the proposed computing infrastructure.

Acknowledgment

This work has been co-funded by the LOEWE initiative (Hessen, Germany) within the NICER project, by the German Research Foundation (DFG) as part of project B03 within the Collaborative Research Center (CRC) 1053-MAKI, and by the German Federal Ministry of Education and Research (BMBF) Software Campus project "PersonalAssistant" [01IS12054].

References

[1] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, "The Case for VM-based Cloudlets in Mobile Computing," *IEEE Pervasive Computing*, vol. 8, no. 4, 2009.

[2] T. Verbelen, P. Simoens, F. De Turck, and B. Dhoedt, "Leveraging Cloudlets for Immersive Collaborative Applications," *IEEE Pervasive Computing*, vol. 12, no. 4, pp. 30–38, 2013.

[3] U. Shaukat, E. Ahmed, Z. Anwar, and F. Xia, "Cloudlet Deployment in Local Wireless Networks: Motivation, Architectures, Applications, and Open Challenges," *Journal of Network and Computer Applications*, vol. 62, pp. 18–40, 2016.

[4] J. Gedeon and I. Schweizer, "Understanding Spatial and Temporal Coverage in Participatory Sensor Networks," in *40th Local Computer Networks Conference (LCN'15): Workshop*. IEEE, 2015, pp. 699–707.

[5] K. Lee, J. Lee, Y. Yi, I. Rhee, and S. Chong, "Mobile Data Offloading: How Much Can WiFi Deliver?" in *6th Intl. Conf. on Emerging Networking Experiments and Technologies (CoNEXT'10)*. ACM, 2010, pp. 26–38.

[6] L. Tawalbeh, Y. Jararweh, F. Dosari *et al.*, "Large Scale Cloudlets Deployment for Efficient Mobile Cloud Computing," *Journal of Networks*, vol. 10, no. 1, pp. 70–77, 2015.

[7] E. Bulut and B. K. Szymanski, "Rethinking Offloading WiFi Access Point Deployment from User Perspective," in *12th Intl. Conf. on Wireless and Mobile Computing, Networking and Communications (WiMob'16)*. IEEE, 2016, pp. 1–6.

[8] C. Meurisch, A. Seeliger, B. Schmidt, I. Schweizer, F. Kaup, and M. Mühlhäuser, "Upgrading Wireless Home Routers for Enabling Large-scale Deployment of Cloudlets," in *7th Intl. Conf. on Mobile Computing, Applications, and Services (MobiCASE'15)*. Springer, 2015, pp. 12–29.

[9] M. Kim, J. J. Fielding, and D. Kotz, "Risks of Using AP Locations Discovered Through War Driving," in *4th Intl. Conf. on Pervasive Computing (Pervasive'06)*. Springer, 2006, pp. 67–82.

[10] C. Meurisch, B. Schmidt, M. Scholz, I. Schweizer, and M. Mühlhäuser, "Labels - Quantified Self App for Human Activity Sensing," in *17th Intl. Conf. on Ubiquitous Computing (UbiComp'15): Adjunct*. ACM, 2015, pp. 1413–1422.

[11] N. Fernando, S. W. Loke, and W. Rahayu, "Mobile Cloud Computing: A Survey," *Future Generation Computer Systems*, vol. 29, no. 1, pp. 84–106, 2013.

[12] S. Clinch, J. Harkes, A. Friday, N. Davies, and M. Satyanarayanan, "How Close is Close Enough? Understanding the Role of Cloudlets in Supporting Display Appropriation by Mobile Users," in *10th Intl. Conf. on Pervasive Computing and Communications (PerCom'12)*. IEEE, 2012, pp. 122–127.

[13] W. Dong, S. Rallapalli, R. Jana, L. Qiu, K. Ramakrishnan, L. Razoumov, Y. Zhang, and T. W. Cho, "iDEAL: Incentivized Dynamic Cellular Offloading via Auctions," in *32nd Intl. Conf. on Computer Communications (INFOCOM'13)*. IEEE, 2013, pp. 755–763.

[14] C. Meurisch, A. Yakkundimath, B. Schmidt, and M. Mühlhäuser, "Upgrading Wireless Home Routers as Emergency Cloudlet: A Runtime Measurement," in *7th Intl. Conf. on Mobile Computing, Applications, and Services (MobiCASE'15): Poster*. Springer, 2015, pp. 338–340.

[15] K. Panitzek, I. Schweizer, A. Schulz, T. Bönning, G. Seipel, and M. Mühlhäuser, "Can We Use Your Router, Please?: Benefits and Implications of an Emergency Switch for Wireless Routers," *Intl. Journal of Information Systems for Crisis Response and Management (IJISCRAM)*, vol. 4, no. 4, pp. 59–70, 2012.

[16] C. Meurisch, T. A. B. Nguyen, S. Wullkotte, S. Niemczyk, Kohnhäuser, and M. Mühlhäuser, "NICER911: Ad-hoc Communication and Emergency Services Using Networking Smartphones and Wireless Home Routers," in *18th Intl. Symp. on Mobile Ad Hoc Networking and Computing (MobiHoc'17): Poster*. ACM, 2017.

[17] A. Sterz, L. Baumgärtner, R. Mogk, M. Mezini, and B. Freisleben, "DTN-RPC: Remote Procedure Calls for Disruption-Tolerant Networking," in *IFIP Networking 2017 Conference and Workshops (Networking'17)*, 2017, pp. 1–8.

[18] J. Gedeon, C. Meurisch, D. Bhat, M. Stein, L. Wang, and M. Mühlhäuser, "Router-based Brokering for Surrogate Discovery in Edge Computing," in *Intl. Workshop on Hot Topics in Planet-Scale Mobile Computing and Online Social Networking (HotPOST'17)*. IEEE, 2017, pp. 1–6.

[19] T. A. B. Nguyen, C. Meurisch, S. Niemczyk, D. Böhnstedt, K. Geihs, M. Mühlhäuser, and R. Steinmetz, "Adaptive Task-oriented Message Template for In-Network Processing," in *Intl. Conf. on Networked Systems (NetSys'17)*. IEEE, 2017.

[20] R. Church and C. R. Velle, "The Maximal Covering Location Problem," *Papers in Regional Science*, vol. 32, no. 1, pp. 101–118, 1974.

[21] M. C. Gonzalez, C. A. Hidalgo, and A.-L. Barabasi, "Understanding Individual Human Mobility Patterns," *Nature*, vol. 453, no. 7196, pp. 779–782, 2008.

[22] N. Eagle and A. Pentland, "Reality Mining: Sensing Complex Social Systems," *Personal and Ubiquitous Computing*, vol. 10, no. 4, pp. 255–268, 2006.

[23] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of Human Mobility on Opportunistic Forwarding Algorithms," *IEEE Transactions on Mobile Computing*, vol. 6, no. 6, 2007.

[24] A. J. Nicholson and B. D. Noble, "Breadcrumbs: Forecasting Mobile Connectivity," in *14th ACM Intl. Conf. on Mobile Computing and Networking (MobiCom'08)*. ACM, 2008, pp. 46–57.

[25] 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.

[26] M. Lin, W.-J. Hsu, and Z. Q. Lee, "Predictability of Individuals' Mobility with High-resolution Positioning Data," in *14th Intl. Conf. on Ubiquitous Computing (UbiComp'12)*. ACM, 2012, pp. 381–390.

[27] R. Montoliu, J. Blom, and D. Gatica-Perez, "Discovering Places of Interest in Everyday Life from Smartphone Data," *Multimedia Tools and Applications*, vol. 62, no. 1, pp. 179–207, 2013.

[28] J. H. Kang, W. Welbourne, B. Stewart, and G. Borriello, "Extracting Places from Traces of Locations," *Mobile Computing and Communications Review, ACM*, vol. 9, no. 3, pp. 58–68, 2005.