A Multi-Cloudlet Infrastructure for Future Smart Cities: An Empirical Study

Julien Gedeon, Jeff Krisztinkovics, Christian Meurisch, Michael Stein, Lin Wang, Max Mühlhäuser
Telecooperation Lab, TU Darmstadt
Darmstadt, Germany
{gedeon,krisztinkovics,meurisch,stein,wang,max}@tk.tu-darmstadt.de

ABSTRACT

The emerging paradigm of edge computing has proposed cloudlets to offload data and computations from mobile, resource-constrained devices. However, little attention has been paid to the question on where to deploy cloudlets in the context of smart city environments. In this vision paper, we propose to deploy cloudlets on a city-wide scale by leveraging three kinds of existing infrastructures: cellular base stations, routers and street lamps. We motivate the use of this infrastructure with real location data of nearly 50,000 access points from a major city. We provide an analysis on the potential coverage for the different cloudlet types. Besides spatial coverage, we also consider user traces from two mobile applications. Our results show that upgrading only a relatively small number of access points can lead to a city-scale cloudlet coverage. This is especially true for the coverage analysis of the mobility traces, where mobile users are within the communication range of a cloudlet-enabled access point most of the time.

CCS CONCEPTS

• **Networks → Wireless access points, base stations and infrastructure**; **Mobile networks**; *Public Internet*; • **Human-centered computing → Empirical studies in ubiquitous and mobile computing**;

KEYWORDS

Cloudlets, Edge Computing, Fog Computing, Smart Cities, Coverage, Mobile Applications

ACM Reference Format:
Julien Gedeon, Jeff Krisztinkovics, Christian Meurisch, Michael Stein, Lin Wang, Max Mühlhäuser. 2018. A Multi-Cloudlet Infrastructure for Future Smart Cities: An Empirical Study. In *EdgeSys'18: 1st International Workshop on Edge Systems, Analytics and Networking , June 10–15, 2018, Munich, Germany.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3213344.3213348

1 INTRODUCTION

Edge computing has gained tremendous attention and aims at bringing storage [12] and computation capabilities [7, 25] closer to mobile end users, resulting in lower latencies and a reduction of bandwidth utilization in the core network. In the next decades, urban spaces will be populated by a plethora of mobile devices, sensors and actuators. These devices are constrained in terms of battery lifetime and processing power and hence, they need to make use of surrogates to offload data and computations [8, 10]. While offloading has mostly been done through cloud computing infrastructures, this approach has some drawbacks, such as high-latencies, lack of mobility support and location awareness [33]. For that reason, we can observe a trend in moving computations and data storage away from the cloud into the network and closer to the end devices and users. In line with this trend, the concept of cloudlets has been proposed [23]. Research in this direction—often labeled fog computing [11, 33] or edge computing [5, 22]—has recently made progress in terms of defining architectural principles, proposing offloading mechanisms with different granularities and addressing network issues. The actual benefits of using cloudlets at the edge of the network have also been investigated [15]. However, the question of infrastructural support, i.e., where those cloudlets should be deployed to reach a maximum number of users or cover a large area, has not been addressed in detail. This question becomes even more interesting if we turn our attention to smart city applications [21, 26, 34] that aim to provide services to citizens. Examples include traffic management and optimization, emergency response or environmental monitoring. Those kinds of applications can be realized by deploying cloudlets close to mobile users in urban areas. This obviously raises the question where to place those cloudlets and what coverage can be achieved. Coverage as a metric for the quality of service that can be delivered by a network comes from the domain of wireless sensor networks but has not been analyzed in detail when it comes to the coverage of urban cloudlets.

In this paper, we propose to upgrade existing infrastructure to host cloudlets on a city-scale, being able to provide new services to citizens in the context of smart cities. Specifically, we consider hosting cloudlets on three types of access points: cellular base stations, routers and street lamps. This re-use of existing infrastructure has obvious cost advantages in comparison to a dedicated cloudlet infrastructure. To motivate our approach, we collect location data for the different types of access points and analyze the potential coverage if only a subset of those is upgraded to provide cloudlet capabilities. First, we analyze spatial coverage, i.e., which areas of the city are covered. Second, we further consider real-world traces from two mobile applications to show that urban cloudlets have the potential to cover the demands of mobile users.

The remainder of this paper is organized as follows. In Section 2 we present our vision of a multi-cloudlet infrastructure for smart cities. We then analyze the potential coverage of those cloudlets in

EdgeSys'18, June 10–15, 2018, Munich, Germany

ACM ISBN 978-1-4503-5837-8/18/06…$15.00
https://doi.org/10.1145/3213344.3213348

Section 3. We conclude the paper and discuss future challenges in Section 4.

2 MULTI-CLOUDLET INFRASTRUCTURE

Since the deployment of new infrastructure is costly—especially on a city-scale—we propose to make use of existing infrastructures that are present in cities. More specifically, we suggest to upgrade cellular base stations, routers and street lamps to host cloudlets. Common to all these three types of access points is their ubiquity and their function as a 1-hop wireless gateway for mobile users. For cellular base stations and routers, we can further assume a powerful backhaul connection in terms of bandwidth and physical space close-by to colocate additional hardware. Figure 1(a) illustrates an example for such an urban cloudlet architecture with the different types of access points. Specifically, we will investigate this multi-cloudlet infrastructure in Darmstadt (Germany), a city with a population of around 150,000. For the later analysis, we focus on the inner region of the city as highlighted in Figure 1(b), because most of our collected access point data is located in that area. This area is $14.57km^2$ in size. The following subsections describe the types of cloudlet infrastructures in more detail.

(a) Example cloudlet infrastructure (b) Darmstadt inner city

Figure 1: Cloudlet types and considered area of Darmstadt

2.1 Cellular Base Stations

Each major city today features widespread cellular coverage, albeit at varying quality and speed. Nevertheless, cellular base stations represent a viable location for the deployment of extra resources that can be leveraged to place cloudlets. First, existing radio access networks have a high-bandwidth backlink on-site. In the case of cloudlets, this could be important if there is the need to retrieve large amounts of data from the cloud. Another characteristic of cellular base stations is their high reliability [9]. Leveraging resources co-located with the radio access network is commonly referred to as mobile edge computing [1, 31]. For network operators and cellular service providers, mobile edge computing is a future business opportunity, as they will be able to rent out computational resources located at the base stations. This will become even more viable with the adoption of 5G networks [27] and femtocells [6, 14]. By making use of dynamic network management through Software Defined Networking (SDN) [28], operators can quickly reconfigure their networks and make them adaptable to user demands and load levels. According to German laws, the *Bundesnetzagentur* (Federal Network Agency) as regulating body maintains a database of all licensed radio installations. This database can be viewed on a website[1]. As the site itself provides no export function, we manually crawled the website and exported all cellular base stations that are within the city limits of Darmstadt.

2.2 Routers

In urban areas, the density of WiFi routers is very high. This includes both privately owned devices as well as public access points offered by businesses like cafes or restaurants. The latter is a service increasingly valued by customers. Using WiFi routers as cloudlets has the obvious advantage of ubiquity of these devices. However, two major drawbacks have to be mentioned: First, as these devices are mostly privately owned, the owners need incentives to open up their devices in order to make them usable as cloudlets. However, recent efforts have shown that such a sharing economy has potential. Several initiatives already promote the sharing of ones WiFi to give others internet access (e.g., Freifunk[2] in Germany). We argue that going one step further—from providing free access to free computations and/or storage—is the next logical step. Existing works have investigated the use of commercial off-the-shelf (COTS) routers, but have either not analyzed the coverage [17, 18], focused only on spatial coverage without considering actual user mobility [13] or have not taken into account other types of access points [16]. Others assume that the routers can be freely placed in order to maximize offloading [2, 3]. Contrary to that, we assume we cannot freely place the access points and have to choose among existing ones. To obtain the locations, we relied on a wardriving approach and used an Android application that records available WiFi networks with metadata such as the SSID, BSSID and signal strength. Volunteers captured the data by either walking or driving through the city. From different measurements of the same access point (identified by its BSSID) and the respective signal strength, we can estimate the position of that access point via trilateration. We performed a manual filtering as to exclude manufacturers that do not produce routers. While we are aware that this dataset might include wrong data and inaccuracies wrt. exact positions, we argue it still gives a good impression of the overall number of available routers that can be leveraged for offloading. Furthermore, by collecting the access point data from the street and not indoors, we mimic the usage contexts of mobile users in a city.

2.3 Street Lamps

Municipalities might not want to leave the provisioning of a cloudlet infrastructure to private citizen, businesses and traditional service providers alone. Instead, we expect them to upgrade parts of their infrastructure in order to provide advanced services to their citizens in the context of smart city applications. Besides high-speed wireless, this should also include processing power for those applications. A lamp post is a viable location to place a cloudlet for two reasons: First, there is a large number deployed in every city, sometimes with the distance between two lamp posts being only a couple of meters. Therefore, especially in densely populated areas, they can very well complement cell towers and routers. Second,

[1] http://emf3.bundesnetzagentur.de/karte/

[2] https://freifunk.net/

from the perspective of users moving on a city street, the communication range is less obstructed compared to WiFi access points that are typically placed in buildings. Third, existing cabling makes it easy to connect them to high-speed backbone networks. Some companies already market smart street lamps that extend the mere lighting functionality with different sensors and wireless access points. An example is *SM!GHT*[3] from the German company *EnBW*. Extending this definition of smartness, we propose not only to upgrade street lamps with sensory capabilities and wireless gateways, but to colocate processing power at lamp posts. We obtained the positions of all street lightings in Darmstadt from *e-netz Südhessen*[4], a company in charge of maintaining electrical infrastructures. The dataset includes different types of lighting. Besides the obvious lamp posts there are also city lights taut over streets by supporting cables. We only include street lighting in a fixed lamp post in our analysis as this ensures the space required for the deployment of a cloudlet in a safe enclosure.

For each type of cloudlet infrastructure, Table 1 summarizes the total number of access points we collected and how many of these are located in the inner city.

Table 1: Number of access points collected

	Cellular Base Stations	Routers	Street Lamps
Total	205	34,699	14,331
Inner City	66	31,974	5,608

3 COVERAGE ANALYSIS

In this section, we analyze the coverage that can be achieved by using a certain percentage of all available access points. Coverage is a term that originates from the domain of wireless sensor networks and describes how well an area of interest can be monitored [19, 29]. It is therefore a metric for the quality of service that can be delivered by the network. Mapped to our problem, coverage indicates how well mobile users can be served by nearby cloudlets.

3.1 Spatial Coverage

First, we look at the spatial coverage of our city cloudlets when we only select a subset of those. For a given percentage of access points, we randomly select the corresponding number. Assuming a unit-disk model for the communication range, we calculate the area formed by the union of all communication ranges. Given this definition, the coverage is given by the ratio between the area of that union and the total size of the city area, as defined by our inner city boundary (see Figure 1(b)). We not only choose access points that are located within the inner city boundary, but also those whose communication ranges intersect with this boundary. Since the access points are chosen randomly, we run the experiment 5 times for each step size of 10 percent. We plot the results for the analysis on spatial coverage for the different types of access points with different communication ranges in Figure 2. The plots also contain error bars for the different simulation runs, although they are very small for the routers and street lamps compared to cellular base stations. This is because we have less cell towers, but they have a greater communication range. Thus, when randomly selecting access points, overlaps are more likely and the gain in coverage might vary more significantly. From the results we can observe that assuming a medium-conservative communication range, we can achieve a relatively high degree of coverage compared to the percentage of chosen access points. For instance, if we assume a sensing range of 40 meters for routers, already 30 percent of all access points cover about 65 percent of the total area. For street lamps the numbers are slightly lower, since we have less of them compared to routers. Because we only have very few cellular base stations compared to other types of access points, the coverage ratio is lower for small percentages (e.g., 10 or 20 percent) but quickly surpasses the coverage ratio of routers and street lamps for greater percentages because of the high communication range of cell towers. Already from this analysis, we can see that deploying cloudlets is feasible for a large-scale coverage within a city and that only a fraction of access points are required.

[3]https://smight.com/en/

[4]https://www.e-netz-suedhessen.de/

Table 2: Evaluation scenarios

Scenario#	Cellular Base Stations	Routers	Street Lamps
SC1	75%	10%	25%
SC2	75%	25%	10%
SC3	50%	25%	25%
SC4	50%	50%	5%
SC5	25%	25%	10%
SC6	25%	10%	50%

Next, we define different scenarios and mix the percentage of the different access points as listed in Table 2. These scenarios reflect different deployment models that could occur, depending on the underlying business models and incentives. For instance, by giving more incentives to small business and private citizens to open up their routers, we may get a higher percentage of those. Similarly, the efforts from a city administration to upgrade street lamps might vary. Network operators on the other hand might choose to upgrade their cell towers based on the average user density or demand at certain points in the city. For the communication ranges, we again assume a unit disk model and randomize the ranges for each type of access point for each evaluation run within a certain range. For cellular base stations, we randomly select a range of 300 to 800 meters. Regarding routers, some measurements indicate average ranges in practice of around 50 to 60 meters [4, 20]. Since routers are most often placed inside buildings and obstructed, we select a conservative range between 20 to 70 meters. Since cloudlets on street lamps will probably also rely on WiFi for communication but are more unobstructed, we increase the range to 30-80 meters. We use a uniform continuous distribution for all communication ranges and again run each experiment 5 times. The results are shown in Figure 3. For each scenario, the bars represent the different types of access points. In Figure 3(a) we consider 1-coverage, i.e., we assume an area to be covered if it is within the range of at least one access point. For 1-coverage, one can therefore think of

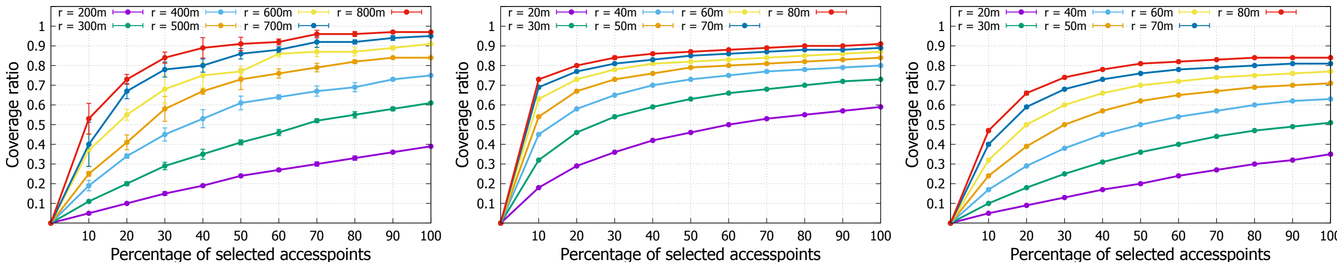

(a) Cellular Base Stations (b) Routers (c) Street Lamps

Figure 2: Spatial coverage for different access point types

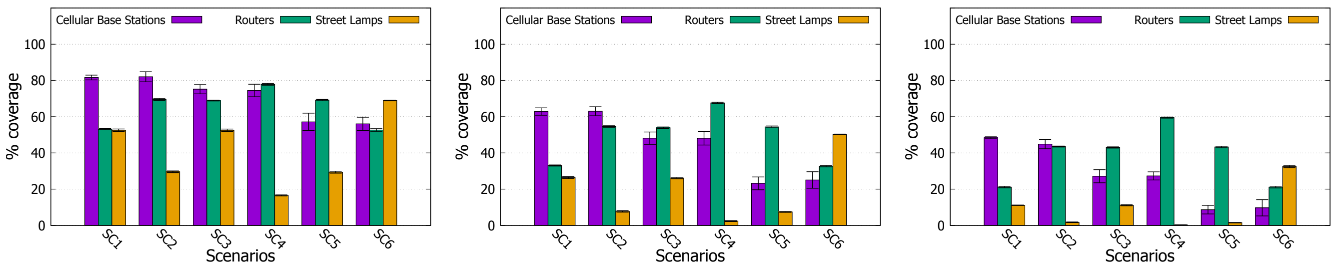

(a) $k = 1$ (b) $k = 2$ (c) $k = 3$

Figure 3: Scenario-based evaluation of spatial coverage for different values of k

this figure as a combination of the values shown in Figure 2. As with our previous results, we see that a relatively small number of access points already provides good coverage. However, this significantly differs for the different types of access points. For instance, using only a small number of street lamps (as in scenario SC4) is clearly not feasible in practice. The notion of 1-coverage can be generalized to k-coverage. To give us an idea of how these coverage metrics differ in our scenario, we also plot the results for $k = 2$ and $k = 3$ in Figures 3(b) and 3(c), respectively. The motivation to consider values of $k > 1$ is to explore the possibility of choosing between multiple cloudlets to optimize user experience in terms of connection bandwidth or resources. For instance, in areas within the city where many users are present, one cloudlet might not be sufficient to satisfy all offloading demands of users. The biggest drop in coverage for increased values of k occur with the street lamps, since they are spaced out evenly and therefore, overlaps in the communication ranges seldom occur. Since the communication ranges of cellular base stations and routers often overlap, we see that for $k = 2$ we still get good coverage. In future work, instead of randomly selecting access points, one might want to optimize the selection for a desired value of k.

3.2 Coverage of Mobile Users

While the results of the previous section give an indication of how well an area is covered, one might argue that this does not necessarily represent the cloudlet coverage a mobile user experiences, since actual user locations are not uniformly distributed throughout the area. Furthermore, users are mobile and change their location frequently. Therefore, in this section, we analyze the actual coverage mobile users can expect. For this purpose, we use two datasets from mobile applications, namely Kraken.me and Ingress.

Kraken.me: Kraken.me [24] is a tracking suite aimed at collecting users' activity in order to enable personal assistance. The framework runs on different kinds of devices and large-scale user study was conducted using Android phones. Over the course of several weeks, about 20 million data points were collected. Although the framework collects data from different soft and hard sensors, we only use the timestamped location data of the users. The data is fine-grained as the position of users is reported every 30 to 60 seconds on average.

Ingress: Next, we use data crawled from the mobile augmented reality game *Ingress.*[5] Ingress is the predecessor of the popular game *Pokémon Go*. In the game, users have to physically visit so-called portals that are located at certain points of interest in a city. Hence, the game records the position of users at the portals. Through their website, one can see the current state of the game and the activity of players. We obtained the game data by building a crawler using Python and *Selenium*, a tool to automate browsers. This allowed us to automatically request changes in the game state, e.g. players visiting a portal, every second. Compared to the Kraken data, there is more difference in time between two user locations, since the location is only recorded at a portal and not in between.

[5] https://www.ingress.com/

Table 3: Mobile application traces

	Kraken.me		Ingress	
	points	users	points	users
Total	19,615,130	225	1,886,546	2,435
Inner City	437,417	205	520,409	1,401

As with the access point data, Table 3 summarizes the number of data points and distinct users for both datasets. For our evaluation we again only consider data points that are within our inner city boundary. Here, we define coverage as the ratio between the sum of data points that are within the communication range of at least one cloudlet and the sum of all data points per dataset. We analyze the results of this type of coverage using the same scenario-based approach as previously mentioned with the same communication ranges per access point type and $k = 1$. Again, we perform each analysis five times. The results with the corresponding mean values and error bars are depicted in Figure 4(a) for Kraken.me and in Figure 4(b) for the Ingress dataset. For both datasets, we can observe that the coverage is higher than in the spatial coverage analysis. This is because spatial coverage also takes into account areas where people are less likely to be (e.g. in-between factory buildings, at parking spaces, more sparsely populated residential areas). In contrast to that, the mobility data from Kraken.me and Ingress represent the actual locations of people. This allows to more accurately estimate the coverage, and therefore the expected quality of service from a user's point of view. While quite distinct from each other, both applications represent realistic future applications that could leverage cloudlets deployed in the city. For Kraken.me, this includes the offloading of (personal) sensor data for real-time analysis and guidance. As the trend in gaming moves towards AR and VR, we can easily imagine that games like Ingress with have such features in the future. In both cases, timely processing of data is crucial for the user experience and cloudlets can provide a way to achieve this. For the Ingress data, the coverage values are slightly higher, because different users aggregate at the same locations (i.e. where the game portals are located). Therefore, in many scenarios, we get coverage values of over 80 percent for cellular base stations and routers. Even though the Kraken.me data is more fine-grained (in the sense of capturing more data points along a user's path), we see that our cloudlet infrastructure can support offloading demands. As an example, scenarios SC1, SC3 and SC6 consistently lead to 60 percent coverage or more, regardless of the access point type. As with the previous analysis, because of their grid-style layout and less overlap, we need more street lamps to achieve the same coverage. We however believe that street lamps can still be beneficial to complement the other types of access points, for instance when other cloudlets are overloaded or as relays for the migration of data across cloudlets.

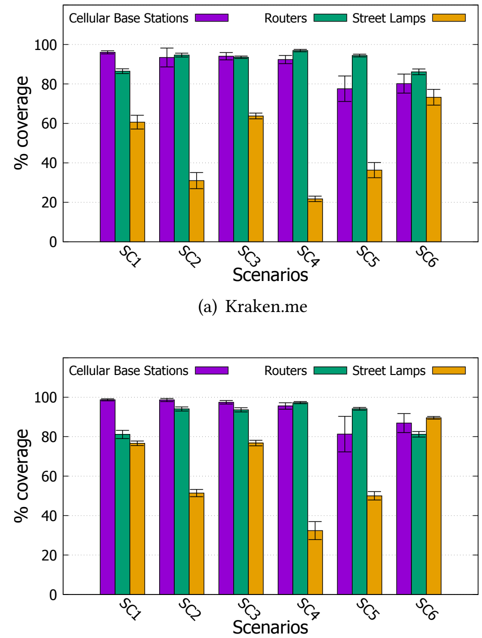

(a) Kraken.me

(b) Ingress

Figure 4: Mobile user coverage

4 OUTLOOK

In this paper, we have outlined our vision of a multi-cloudlet architecture that makes use of existing urban infrastructure to provide new services in the context of smart cities. We motivated this approach by analyzing the potential coverage of urban cloudlets hosted on different kinds of existing infrastructures. Through this analysis, we have shown that it is feasible to re-use cellular base stations, routers and street lamps as supporting infrastructure for a city-wide deployment of cloudlets. In detail, we first analyzed the spatial coverage for the individual types of access points and then defined six different scenarios that represent the incentives for the different stakeholders to upgrade their infrastructures. For both the spatial coverage and user traces from two mobile applications, we were able to show that a relatively small percentage of all access points is sufficient to provide high coverage. Our preliminary work in this paper opens up several future research questions that we plan to investigate in the future.

Incentive mechanisms and business models: In the description of our infrastructure we have already mentioned the different stakeholders relevant to establish a city-wide cloudlet infrastructure, such as private citizens, network operators, internet service providers, small businesses and hardware manufacturers. For some of these stakeholders, city cloudlets are a business opportunity, while for others, non-monetary incentives might suffice to be willing to offer cloudlets at their access points. Especially since cloudlets often work in a cooperative way and many users compete for resources, new business and pricing models are required.

Coverage definitions and communication model: In this paper, we considered spatial coverage and the coverage of individual data points of mobility traces. However, other definitions can be thought of, e.g., taking into account the entire path of a user and the total time one has connection to a cloudlet. Also, since urban areas have various obstacles for the communication range of a wireless medium, more realistic models than assuming a unit-disk range might be required for even more accurate results.

Optimal cloudlet selection: Given a certain number of available cloudlet locations as modeled by our percentage-based scenarios, we have to make a decision where to actually deploy them, i.e. which access points to upgrade in order to provide cloudlet capabilities. For this decision, the heterogeneity of the different types of access points in terms of communication ranges, costs and available resources should be taken into account. As an example, cloudlets colocated at cellular base stations can easily accommodate server-grade hardware and thus provide a huge amount of resources. However, they might be more expensive than using cloudlets hosted by private citizens or the city. One interesting question is whether the latter can compensate for their lack of resources and individual coverage by their price and their greater number. In future work, we will explore this tradeoff and design efficient algorithms for the decision on where to place cloudlets.

User-to-cloudlet assignment: As we have seen in our k-coverage analysis, a user might have connection to more than one cloudlet at a given moment. For instance, as cellular coverage is widespread, users in the city might always have cellular connection, but at the same time could connect to a cloudlet hosted on a router or a street lamp. When making the decision which cloudlet to connect to, one should take into account both the user requirement and their mobility as well as the current load of the cloudlets. This line of research has recently attracted some attention [30, 32].

Data management across cloudlets: Since users are highly mobile, they might be connected to a given cloudlet for only a very short amount of time. Therefore, it might be necessary to migrate data across cloudlets in order for users to be able to retrieve it close to their new position. This both includes static data as well as results of computations that were not finished before the user lost connection to a cloudlet. Besides quick migration mechanisms, accurate ways to predict the user's next location are required. The latter could enable proactive migrations as opposed to reactive migration as mostly done today.

ACKNOWLEDGEMENT

This work has been cofunded by the German Research Foundation (DFG) and the National Nature Science Foundation of China (NSFC) joint project under Grant No. 392046569 (DFG) and No. 61761136014 (NSFC) and by DFG as part of subproject A1 of the Collaborative Research Center (CRC) 1053 - MAKI. The authors would like to thank *e-netz Südhessen GmbH & Co. KG* for providing the location data of lamp posts in Darmstadt and Patrick Felka for crawling the Ingress game data.

REFERENCES

[1] Abbas, N., Zhang, Y., Taherkordi, A., and Skeie, T. Mobile edge computing: A survey. *IEEE Internet of Things Journal 5*, 1 (2018), 450–465.
[2] Bulut, E., and Szymanski, B. K. Wifi access point deployment for efficient mobile data offloading. *Mobile Computing and Communications Review 17*, 1 (2013), 71–78.
[3] Bulut, E., and Szymanski, B. K. Rethinking offloading wifi access point deployment from user perspective. In *WiMob 2016* (2016), pp. 1–6.
[4] Burger, V., Seufert, M., Kaup, F., Wichtlhuber, M., Hausheer, D., and Tran-Gia, P. Impact of wifi offloading on video streaming QoE in urban environments. In *ICCW* (2015), pp. 1717–1722.
[5] Chandra, A., Weissman, J., and Heintz, B. Decentralized edge clouds. *IEEE Internet Computing 17*, 5 (2013), 70–73.
[6] Chandrasekhar, V., Andrews, J. G., and Gatherer, A. Femtocell networks: a survey. *IEEE Communications Magazine 46*, 9 (2008), 59–67.
[7] Chang, H., Hari, A., Mukherjee, S., and Lakshman, T. V. Bringing the cloud to the edge. In *INFOCOM Workshops* (2014), pp. 346–351.
[8] Cuervo, E., Balasubramanian, A., Cho, D.-k., Wolman, A., Saroiu, S., Chandra, R., and Bahl, P. Maui: Making smartphones last longer with code offload. In *MobiSys* (2010), pp. 49–62.
[9] Elmokashfi, A., Zhou, D., and Baltrünas, D. Adding the next nine: An investigation of mobile broadband networks availability. In *MobiCom* (2017), pp. 88–100.
[10] Fernando, N., Loke, S. W., and Rahayu, W. Mobile cloud computing: A survey. *Future Generation Computer Systems 29*, 1 (2013), 84–106.
[11] Gedeon, J., Heuschkel, J., Wang, L., and Mühlhäuser, M. Fog computing: Current research and future challenges. In *1. GI/ITG KuVS Fachgespräche Fog Computing* (2018), pp. 1–4.
[12] Gedeon, J., Himmelmann, N., Felka, P., Herrlich, F., Stein, M., and Mühlhäuser, M. vStore: A context-aware framework for mobile micro-storage at the edge. In *MobiCASE* (2018), pp. 1–18.
[13] Gedeon, J., Meurisch, C., Bhat, D., Stein, M., Wang, L., and Mühlhäuser, M. Router-based brokering for surrogate discovery in edge computing. In *ICDCS Workshops* (2017), pp. 145–150.
[14] Hasan, S. F., Siddique, N. H., and Chakraborty, S. Femtocell versus WiFi - A survey and comparison of architecture and performance. In *International Conference on Wireless Communication, Vehicular Technology, Information Theory and Aerospace Electronic Systems Technology* (2009), pp. 916–920.
[15] Hu, W., Gao, Y., Ha, K., Wang, J., Amos, B., Chen, Z., Pillai, P., and Satyanarayanan, M. Quantifying the impact of edge computing on mobile applications. In *SIGOPS Asia-Pacific Workshop on Systems* (2016), pp. 5:1–5:8.
[16] Meurisch, C., Gedeon, J., Gogel, A., Nguyen, T. A. B., Kaup, F., Kohnhäuser, F., Baumgärtner, L., Schmittner, M., and Mühlhäuser, M. Temporal coverage analysis of router-based cloudlets using human mobility patterns. In *GLOBECOM* (2017), pp. 1–6.
[17] Meurisch, C., Nguyen, T. A. B., Gedeon, J., Konhauser, F., Schmittner, M., Niemczyk, S., Wullkotte, S., and Mühlhäuser, M. Upgrading wireless home routers as emergency cloudlet and secure DTN communication bridge. In *ICCCN* (2017), pp. 1–2.
[18] Meurisch, C., Seeliger, A., Schmidt, B., Schweizer, I., Kaup, F., and Mühlhäuser, M. Upgrading wireless home routers for enabling large-scale deployment of cloudlets. In *MobiCASE* (2015), pp. 12–29.
[19] Mohamed, S. M., Hamza, H. S., and Saroit, I. A. Coverage in mobile wireless sensor networks (M-WSN): A survey. *Computer Communications 110* (2017), 133–150.
[20] Mota, V. F. S., Macedo, D. F., Ghamri-Doudane, Y., and Nogueira, J. M. S. On the feasibility of wifi offloading in urban areas: The paris case study. In *IFIP Wireless Days (WD)* (2013), pp. 1–6.
[21] Perera, C., Qin, Y., Estrella, J. C., Reiff-Marganiec, S., and Vasilakos, A. V. Fog computing for sustainable smart cities: A survey. *ACM Comput. Surv. 50*, 3 (2017), 32:1–32:43.
[22] Satyanarayanan, M. The emergence of edge computing. *IEEE Computer 50*, 1 (2017), 30–39.
[23] Satyanarayanan, M., Bahl, P., Cáceres, R., and Davies, N. The case for VM-based cloudlets in mobile computing. *IEEE Pervasive Computing 8*, 4 (2009), 14–23.
[24] Schweizer, I., and Schmidt, B. Kraken.me: Multi-device user tracking suite. In *UbiComp Adjunct* (2014), pp. 853–862.
[25] Shi, W., Cao, J., Zhang, Q., Li, Y., and Xu, L. Edge computing: Vision and challenges. *IEEE Internet of Things Journal 3*, 5 (2016), 637–646.
[26] Taleb, T., Dutta, S., Ksentini, A., Iqbal, M., and Flinck, H. Mobile edge computing potential in making cities smarter. *IEEE Communications Magazine 55*, 3 (2017), 38–43.
[27] Taleb, T., Samdanis, K., Mada, B., Flinck, H., Dutta, S., and Sabella, D. On multi-access edge computing: A survey of the emerging 5G network edge cloud architecture and orchestration. *IEEE Communications Surveys and Tutorials 19*, 3 (2017), 1657–1681.
[28] Truong, N. B., Lee, G. M., and Ghamri-Doudane, Y. Software defined networking-based vehicular Adhoc Network with Fog Computing. In *IFIP/IEEE International Symposium on Integrated Network Management* (2015), pp. 1202–1207.
[29] Wang, B. Coverage problems in sensor networks: A survey. *ACM Computing Surveys 43*, 4 (2011), 32:1–32:53.
[30] Wang, L., Jiao, L., He, T., Li, J., and Mühlhäuser, M. Service entity placement for social virtual reality applications in edge computing. In *INFOCOM* (2018), pp. 1–9.
[31] Wang, L., Jiao, L., Kliazovich, D., and Bouvry, P. Reconciling task assignment and scheduling in mobile edge clouds. In *ICNP* (2016), pp. 1–6.
[32] Wang, L., Jiao, L., Li, J., and Mühlhäuser, M. Online resource allocation for arbitrary user mobility in distributed edge clouds. In *ICDCS* (2017), pp. 1281–1290.
[33] Yi, S., Hao, Z., Qin, Z., and Li, Q. Fog Computing: Platform and Applications. In *HotWeb* (2015), pp. 73–78.
[34] Zanella, A., Vangelista, L., Bui, N., Castellani, A., and Zorzi, M. Internet of Things for Smart Cities. *IEEE Internet of Things Journal 1*, 1 (2014), 22–32.