

<https://helda.helsinki.fi>

IoT Maps : Charting the Internet of Things

Shaw, Peter

ACM
2019

Shaw , P , Mikusz , M , Nurmi , P & Davies , N 2019 , IoT Maps : Charting the Internet of Things . in HOTMOBILE '19 : Proceedings of the 20th International Workshop on Mobile Computing Systems and Applications . ACM , New York, NY , pp. 105-110 , International Workshop on Mobile Computing Systems and Applications , Santa Cruz , California , United States , 27/02/2019 . <https://doi.org/10.1145/3301293.3302375>

<http://hdl.handle.net/10138/311560>
<https://doi.org/10.1145/3301293.3302375>

cc_by
publishedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

IoT Maps: Charting the Internet of Things

Peter Shaw

School of Computing and Communications
Lancaster University
Lancaster, U.K.
p.shaw@lancaster.ac.uk

Petteri Nurmi

University of Helsinki &
Lancaster University
Lancaster, U.K.
p.nurmi@lancaster.ac.uk

Mateusz Mikusz

School of Computing and Communications
Lancaster University
Lancaster, U.K.
m.mikusz@lancaster.ac.uk

Nigel Davies

School of Computing and Communications
Lancaster University
Lancaster, U.K.
n.a.davies@lancaster.ac.uk

ABSTRACT

Internet of Things (IoT) devices are becoming increasingly ubiquitous in our everyday environments. While the number of devices and the degree of connectivity is growing, it is striking that as a society we are increasingly unaware of the *locations and purposes* of such devices. Indeed, much of the IoT technology being deployed is invisible and does not communicate its presence or purpose to the inhabitants of the spaces within which it is deployed. In this paper, we explore the potential benefits and challenges of constructing *IoT maps* that record the location of IoT devices. To illustrate the need for such maps, we draw on our experiences from multiple deployments of IoT systems.

CCS CONCEPTS

• **Human-centered computing** → **Ubiquitous and mobile computing**.

KEYWORDS

mobile computing; internet of things; map schemas

ACM Reference Format:

Peter Shaw, Mateusz Mikusz, Petteri Nurmi, and Nigel Davies. 2019. IoT Maps: Charting the Internet of Things. In *The 20th International Workshop on Mobile Computing Systems and Applications (HotMobile '19)*, February 27–28, 2019, Santa Cruz, CA, USA. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3301293.3302375>

1 INTRODUCTION

The Internet of Things (IoT) is growing rapidly with an estimated 23 billion connected devices deployed worldwide in 2018 [28]. These devices range from expensive infrastructure components, such as actuators in smart cities, through to low-cost commodity devices such as radio frequency beacons (e.g. iBeacons). Deployment strategies for such IoT devices range from carefully controlled large-scale

rollouts with significant organisational support through to ad-hoc deployments by individuals. While the number of devices, and the degree of connectivity is growing, it is striking that as a society we are increasingly unaware of the *locations and purposes* of such devices. In keeping with Weiser’s vision of technology that fades into the background, much of the IoT technology being deployed is essentially designed to be invisible. This lack of awareness both limits the services that can be provided and raises concerns for users and system owners.

Fully harnessing the capabilities of IoT deployments while avoiding potential disadvantages, e.g. related to privacy and security concerns, requires knowledge about available devices, their locations, and capabilities. In other words, IoT devices should be *mapped*. While there have been previous attempts at cataloguing IoT devices (e.g. [14]) these have mostly focused on registering networked devices without providing detailed information about the locations and capabilities of devices. In this paper, we explore the potential benefits of constructing comprehensive maps of IoT devices, identifying key research challenges and describing partial solutions to these issues. To illustrate the need for such maps, we draw on our experiences of deploying multiple IoT systems. We make three contributions:

- (1) We highlight the importance of producing and maintaining maps of the IoT using illustrative examples drawn from real-world case-studies.
- (2) We present an example map schema designed to capture data on a broad class of IoT devices, and highlight some remaining challenges.
- (3) We discuss the challenges in populating and maintaining maps of the IoT.

Overall, we aim to stimulate new work by the mobile computing community to begin to create general purpose IoT maps.

2 USE CASES FOR IOT MAPS

2.1 Overview of Case Studies

We motivate the need for IoT maps by highlighting three case studies drawn from our previous research.

Using Maps for Personalisation of Smart Spaces. Personalisation in the context of the IoT and smart spaces evokes privacy concerns. For the digital signage personalisation system ‘Tacita’, a



This work is licensed under a Creative Commons Attribution International 4.0 License.

HotMobile '19, February 27–28, 2019, Santa Cruz, CA, USA
© 2019 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-6273-3/19/02.
<https://doi.org/10.1145/3301293.3302375>

deployment which has motivated our work, privacy was a primary focus at an architectural level [5]. Typical display personalisation systems require viewers to install applications that transmit their preferences to personalisable displays [16] – enabling displays to observe and track viewers and therefore imposing a significant privacy risk. In contrast, Tacita relies on users downloading maps of nearby displays. The detection of user proximity to displays is determined on the user’s device, and requests to personalise content on nearby displays are issued accordingly.

Supporting Privacy Awareness. Future pervasive computing environments are likely to include large numbers of sensors such as cameras and microphones embedded in the physical environment and capable of capturing personal data. Such data can be used for a wide range of applications ranging from augmented cognition through entertainment to personalised advertising. In previous work, we found indications that user attitudes to sensing and data collection in smart environments depends highly on the intended purpose [27] – security was deemed the most acceptable while applications that appeared to only benefit the collector (e.g. for personalised ads) were widely disliked. Interestingly, participants were strongly opposed to any form of covert data collection and highlighted the importance of the ability to determine the nature of data collection taking place in any given space. Maps can be used as a technique that allows users (and applications) to identify instrumented spaces.

While using maps requires manual effort and does not protect users against deliberate covert surveillance it does provide a mechanism for owners of physical spaces to inform occupants of the data capture. We created a prototype application in which maps could be downloaded to a smartphone, using regions to represent areas of surveillance with an associated list of devices (e.g. cameras). Additional data about regions can also be included such as: size of regions, legal agreements and their time validity periods, devices in use, and the owner of the data collected. Upon approaching a region, users were presented with a notification listing devices in use and offering an option to accept the data collection (suppressing future notifications for the same region) or decline (providing an opportunity for future work on surveillance consent).

Instantiating Cloudlets for Privacy Preservation. Spaces are increasingly often equipped with IoT sensing and processing capabilities, potentially capturing sensitive data about individuals present in the space. Privacy mediation is an approach allowing individuals to control and mediate the captured data using Cloudlets before it is released beyond the immediate physical area [26]. For example, in the context of smart meeting rooms that consist of microphones and video cameras capturing events taking place in the space, users can configure privacy mediators to prevent or limit the amount of data collected. In order to ensure that privacy mediation takes place at the point at which the users enter the space, the corresponding Cloudlets need to be ready – imposing the requirement for detecting users before entering the space. This can be achieved through the specification of ‘trigger zones’, a separately monitored region purely serving the purpose of instantiating appropriate processes prior to the user entering smart space. In maps, active regions of data capture, locations of individual sensors, and trigger zones can be defined separately in the form of geo-fences or

proximity-based location descriptions, in addition to configuration parameters to support privacy mediation.

2.2 Example Implementation

We implemented all of the use cases as part of a common demonstrator system consisting of three core components: a map generator, a map repository, and mobile clients. The map generator is a simple tool that creates a map of IoT devices using a common schema based on a set of source lists that hold information on the devices available, their locations (in the form of proximity-based locations using iBeacon identifiers and their absolute locations in latitude/longitude) and capabilities. The maps created are stored within the map repository that serves as a centralised storage and distribution space. Within the map repository, we support the storage of multiple versions of a single map supporting different scenarios and deployment stages. Mobile clients can access the map repository to retrieve specific maps and use the information within the map to instantiate functionalities. For example, we developed a mobile phone application that utilises ‘trigger zones’ to prepare geo-fences and track the user’s location in order to instantiate appropriate cloudlets for privacy preservation, or in order to request personalised content on displays nearby.

3 MAP SCHEMAS

Currently no standardised solution for IoT maps exists. In this section we present an example of a partial solution that we adopted to support the privacy mediator use case.

3.1 Example of IoT Mapping

We designed a map schema to support privacy mediation using cloudlets (fig. 1). The schema consists of three high-level entities: *unique identifier*, *meta*, and *domains*. The meta objects consist of the global description of the map including its temporal validity (in the form of start and expiration dates), the publication date of the map and its version. The domains object consists of a list of active areas of data capture that are owned by the same entity (e.g. a commercial organisation). Domains provide the name of the organisation, the API end-point for privacy mediation requests at the organisational level, and a list of data capture zones in which sensors and devices can capture potentially privacy invasive data. Each data capture zone consists of lists of regions, trigger zones and capabilities. Regions specify the geographical locations or boundaries of the space in which data capture is taking place. Such locations can be described, for example, by specifying circular regions (latitude, longitude and radius) or proximity-based locations (based on Bluetooth Low Energy beacons). Trigger zones represent descriptions of monitored areas which, when entered by the user, trigger an immediate notification to the backend system. Similar to regions, trigger zones can be described using geo-fences or proximity-based location descriptions. Capability objects define the available mediation services of the data capture zone, such as audio and video processors in the context of smart meeting rooms. Each capability is described by a universally unique identifier object, a human-readable name, and a list of supported cloudlet applications – in this case representing privacy mediators. Each cloudlet application consists of an object that defines the application programming interfaces used to

notify the service of the user's presence in the data capture area (callback_url), a URL to the user-facing configuration interface, and a set of additional metadata objects including the name of the cloudlet application, an icon and a URL to the homepage.

3.2 Further Considerations

We note that whilst we aimed that our proposed map schema can be reused in different domains, a generalised schema that goes beyond the support of privacy mediation leads to a number of challenges that need to be considered.

2D vs 3D. Maps have traditionally represented two dimensional top down perspectives of the world with some extensions in recent years to support the three-dimensional modelling of structures in prominent cities. With the increased popularity of indoor mapping and tracking, we will likely see a growing demand to model more complex structures. In our proposed map schema, we currently follow a two-dimensional approach, ignoring the altitude of active regions and trigger zones. Supporting three-dimensional models of the world introduces challenges regarding representation, visualisation and interaction. We note that some of these challenges, however, have been addressed with systems such as Building Information Modelling (BIM) that are designed to provide digital representations of places and structures for planning and construction purposes [7].

Access Control. Our map schema assumes that users have full access to the data stored within the map. However, in some cases certain sections of a map may be considered confidential and may require access restrictions. Supporting such access control will likely add a new layer of complexity to the map schema definition in order to supply detailed information on access permissions for individual users or groups. We note that initial work has been carried out to incorporate security features (e.g. access control lists) in the context of BIM [18].

Moving Objects. Traditionally, mapping has been used to represent stationary devices and structures such as buildings and sensors embedded in the environment. However, modern IoT devices are often mobile and move frequently within and across spaces (e.g. wearable IoT sensors). The challenges for representing moving objects within a map lie particularly in the processes to maintain and report frequent location updates (e.g. at the level of seconds, minutes or hours) – and are highly depended on the location technique that has been chosen to describe the location of moving objects.

Indoor vs. Outdoor. While outdoor mapping and location tracking is well established, indoor location tracking is an active area of research and imposes a number of challenges regarding reliability and accuracy of location tracking techniques [4]. Researchers have already worked on utilising existing Wi-Fi infrastructure to support indoor location tracking and on improving the accuracy of such tracking techniques [15, 19]. However, in order to support commercial-grade applications, the wide deployment and adoption of accurate indoor location tracking techniques and appropriate maps of IoT devices that are situated indoors will become an important challenge in future.

Single Common Map vs. Bespoke Maps. Our example schema was mostly bespoke to our systems with attempts at generalisation

```
{
  "id": "0db4c16a-f225-4a71-8e05-fbb7d4619c99",
  "meta": {
    "description": "Cloudlets",
    "start_date": "01/10/2018",
    "expiration_date": "02/10/2019",
    "publication_date": "01/10/2018",
    "map_version": "1.2",
    "agreement": "n/a"
  },
  "domains": [
    {
      "name": "Company Office",
      "server": "https://example.com/privacy_mediator_request",
      "data_capture_zones": [
        {
          "use_capturezone_as_triggerzone": false,
          "regions": {
            "circular_regions": [
              {
                "lat": "41.367149",
                "long": "-37.580631",
                "radius": "30m"
              }
            ]
          },
          "trigger_zones": {
            "circular_regions": [
              {
                "lat": "41.367149",
                "long": "-37.580631",
                "radius": "50m"
              }
            ]
          },
          "proximity_beacons": [
            {
              "beacon_major": "10",
              "beacon_minor": "5",
              "beacon_type": "iBeacon",
              "beacon_uuid": "41fbe746-8e66-46a2-95bd-a1e1fb2b0783"
            }
          ]
        }
      ],
      "capabilities": {
        "uuid": {
          "cloudlet_id": "company-meeting-room-5",
          "cloudlet_name": "Meeting Room 5"
        }
      },
      "cloudlet_apps": [
        {
          "name": "Privacy Mediator",
          "callback_url": "https://example.com/tacita_callback",
          "description": "Mediates sensors on behalf of the user.",
          "icon_url": "https://example.com/privacy_mediator_logo.png",
          "homepage": "https://example.com",
          "config_url": "https://example.com/config"
        }
      ]
    }
  ]
}
```

Figure 1: Example JSON map used for proximity triggered cloudlet provisioning.

(e.g. supporting generic ‘capabilities’ that can represent privacy mediation and personalisable display applications at the same time). It is unclear if a ‘one map fits all’ approach is desirable, or even possible in the context of mapping ‘the IoT’. However, designing bespoke map schemas would likely lead to a large number of heterogeneous maps that become unusable beyond their original context – or incur a heavy cost to be integrated into other contexts.

We note that the set of challenges presented is not exhaustive and that a number of other challenges are likely to emerge as more research is conducted regarding the design and implementation of generalisable mapping schemas – potentially leading to new IoT mapping standards.

4 POPULATING MAPS

The population and maintenance of maps is a further challenge in our proposed approach. Maps can be populated from three key sources: (i) authorities, (ii) ordinary users, and (iii) data provided by infrastructures. Based on these available sources, we provide a set of example population techniques that can be applied in order to create and maintain maps of the IoT.

Authoritative. The obvious solution is to employ an authority (e.g. system administrator or owners of spaces) that collects and supplies information about available IoT devices. The main advantage is that the map is likely to have a high accuracy, and capabilities of the devices can be easily identified. The main drawback is that the collection of required information can be laborious, particularly if large IoT deployments have to be mapped from scratch. Additionally, IoT devices present in the same space may be owned by distinct authorities, leading to only partially complete IoT maps of spaces. Nevertheless, authoritative maps are likely to serve as starting points for IoT maps, but should not be interpreted as absolute ground truths.

An example application for the authoritative approach is the ‘Using Maps for the Personalisation of Smart Spaces’ use case in which the locations, capabilities and interfaces of personalisable displays (and other devices) are populated by a trusted entity that controls the deployment, and made accessible in the form of a centrally hosted map.

Crowdsourcing. In the crowdsourcing approach, a number of regular users provide the necessary data to create maps, akin to OpenStreetMap. This reduces the burden of a dedicated party responsible for populating a map. However, information quality is likely to decrease as users may not be aware of the device capabilities and exact locations. Similarly, coverage may suffer as users fail to identify all relevant devices. The usefulness of crowdsourced information can be potentially increased using a two-phased approach where the information provided by users is verified in a second phase conducted by domain experts or through automated analysis (e.g. by comparing multiple reports of the same device, or based on computer vision techniques whereby images taken by users are matched against a device database to populate the relevant parts of the map schema).

Crowdsourcing may be the most appropriate population technique for the ‘Supporting Privacy Awareness’ use case. Users could collectively report IoT sensors that are visible in the environment

and, for example, collect sensitive information about individuals (e.g. cameras and Bluetooth Low Energy beacons used to support indoor location tracking). The crowdsourced map then serves as a foundation to make other individuals aware of the potential data collection taking place in the space.

Infrastructure Sensing. The population of a map can also be delegated to the infrastructure itself. For example, analysis of electric signals has been used to identify home appliances [10] and network traffic signatures have been used to identify IoT devices in a specific administrative network [22] removing the need for human effort. However, the underlying sensing techniques may not generalise sufficiently across diverse environments and provide only limited device information.

The ‘Instantiating Cloudlets for Privacy Preservation’ use case is an example in which cloudlet components (e.g. cameras and microphones that can be found in smart meeting rooms) can automatically report their capabilities (e.g. the support of video and audio recordings) and their location to a centralised map generator – reducing the need for manual updates and enabling the support of larger-scaled deployments.

Opportunistic Crowdsensing. Opportunistic crowdsensing can be seen as a halfway point between user- and system-generated maps. The idea is to use sensors available on mobile devices to identify IoT devices as users navigate across spaces. For example, magnetometers are capable to identify signatures of specific displays [23] and could be used to detect other IoT sensing devices [8]. However, this approach would also identify non-IoT devices, such as elevators or ticketing machines [9], i.e. leading to noisy results. Additionally, the signatures of devices would depend on the orientation of the magnetometer relative to the device [3]. Discovered devices can then be associated with locations either by the user manually supplying the location (suffering from the need to label locations) or automatically using appropriate localisation techniques (potentially resulting in biases in estimated locations [17]).

Similarly to crowdsourcing, the use of opportunistic crowdsensing can be an appropriate technique to gather information on IoT devices for privacy awareness purposes (as described as part of the ‘Crowdsourcing’ approach).

We note that in our demonstrators we adopted the authoritative approach as system and available client components are owned and managed by a single entity. Of course, the choice of the map population technique is highly dependent on use cases and requirements. In some cases, for example, the use of multiple techniques can provide a number of benefits including the ability to verify and validate existing maps.

5 DISCUSSION

In addition to populating IoT maps there are a number of additional challenges that can be identified:

IoT Maps for Connected Devices. Existing research on mapping IoT devices has predominantly focused on identifying connected devices such as networked IoT temperature and air quality sensors. While useful, these systems only provide a partial solution as they do not consider the full spectrum of sensing capabilities, actuators,

and interfaces connected IoT devices incorporate. The use cases presented are examples of applications where full-fledged maps are required.

Maintenance. IoT maps are only useful if they accurately capture the devices that are currently available in the environment. This requires active maintenance, particularly within larger administrative regions. Infrastructure sensing could potentially be used to identify “anomalies” that serve as starting points to investigate changes in the environment. Note that this covers both the appearance of new devices, and disappearance of existing ones.

Global or local scale. Existing initiatives to catalogue IoT devices tend to focus on providing a *global* index of IoT devices. While having such a map would certainly be desirable, for many applications it is sufficient to have a *local* view that captures all available devices and their capabilities within a specific administrative region (e.g. a building or a room). Attempting to construct global-scale maps introduces unnecessary overhead – both regarding access and maintenance of such maps. Local solutions that describe IoT sensors and capabilities of instrumented environments in which users are currently present can be made more dynamic (making it easier for administrators and developers to provide updates) and therefore of higher quality to the user.

Proximity vs. Absolute Location. We recognise that locations (e.g. to describe regions, trigger zones and the actual location of sensor deployments) can be captured through a range of techniques – using relative or absolute descriptions. Depending on the location of sensors and actuators (e.g. indoor vs. outdoor), certain location techniques may be not sufficient or appropriate. We specifically proposed a map schema that is flexible to accommodate different location tracking technologies depending on the use case of the IoT device or service that is defined. Additionally, the flexibility allows us to model both stationary and portable IoT devices without the limitations of a specific localisation technology.

Not just points. While knowing the exact location is useful for some sensors and other objects, it is not always the only piece of geographical data that is useful. For example, the locations of the effective range of devices within an environment may also be of importance such as the field of vision of a video camera, or the maximum capture distance of microphones. We note that including additional regions and metadata is not a novel concept but often overlooked in current IoT maps.

Evaluation. The quality of mapping approaches are commonly assessed using measures such as freshness (or timeliness), coverage (or recall), and accuracy. However, these measures provide only a partial solution for IoT maps due to each device having multiple pieces of information that need to be captured. Indeed, coverage can refer to the fraction of devices or the total set of capabilities of devices. Similarly, freshness depends on whether the unit of assessment is a device or specific functionality. This suggests that new metrics for evaluating the quality of IoT maps will be required.

6 RELATED WORK

Many platforms have been proposed for solving the challenge of describing and cataloguing IoT devices, sensors, and services [6, 13, 20]. While having this data is clearly important, with no commonly

adopted standard many of these systems appear not to be used beyond the initial period of research. ‘Thingful.net’ is an example of a service that aimed to become the “Google of the IoT” [2], i.e. providing a search engine for IoT devices – yet at the time of writing objects listed appear to have not been updated in over a year. Many platforms tend to feature simplistic approaches to describe locations (longitude, latitude, and sometimes altitude) suitable for overlaying data layers on top of base maps. However, using different frames of reference (e.g. metric measurements in relation to an origin point) is not supported by these platforms. This limitation highlights that although the issue of mapping the IoT is not a new problem, additional work is required to converge on a standard and to ensure correct scoping of attempts to cataloguing the IoT.

User contributed data for maps is also an area that has seen much interest. An early example of this is ‘War driving’, the act of locating and recording geo-locations of Wi-Fi networks from a moving vehicle [1, 17]. More recently, OpenStreetMaps [11] is an example of a successful attempt to crowd source a base map and additional data layers of objects and places. Some weaknesses are still present when trying to employ this strategy. For example, it is unrealistic for users to report very frequent location changes of devices (e.g. every minute). In addition, a critical mass of users is required – the WikiBeacon service [25] appears to lack sufficient engagement to be considered a viable source of an up-to-date map of beacons around the world.

If user contributed data is not sufficient to build a map of the IoT strategies for automating the process will be required. While mapping of the outside world has recently reached a level of detail sufficient to support a wide range of applications, the mapping of indoor spaces are still lacking the higher level of accuracy these smaller scale areas require. Automated solutions infer indoor maps from the use of wearables that report movement traces [24] or in combination with accelerometer and magnetometers [29]. Another approach is to have a robot perform the mapping [21]. If indoor mapping and location tracking reaches sufficient accuracy for common commercial devices, we are likely to see a significant increase in applications and usage.

The value of maps of IoT devices in pervasive environments was perhaps best illustrated in the Active Bat project [12] that created models of smart spaces and tracked the movement of people and objects within these spaces. The Active Bat system even provided an API to enable applications to be developed that utilise the ability of spatial triggers.

7 CONCLUSIONS

In this paper we have highlighted reasons why mapping the IoT is an important area of research with potential benefits for multiple stakeholders. Despite its importance and significant prior work in the field of IoT, there are no standardised or widely adopted solutions to achieve such mapping. Developing appropriate solutions will require a coordinated effort from the research community – the challenges are diverse and we would anticipate a federated solution in order to accommodate the full range of mapping scenarios. It is also vitally important that any mapping technologies or repositories are open and not owned by any single entity.

Our concrete proposal for the way forward is that the community begins to develop a comprehensive set of requirements derived from a broad collection of use cases. These requirements can drive the selection of an appropriate set of open mapping technologies that can, initially, be tested in the context of a single use case or technology, e.g. mapping Bluetooth Low Energy beacons.

Drawing on our experience of deploying multiple IoT systems we hope to continue exploring the challenges in future works and encourage ‘mapping the IoT’ to be considered for its value and potential in future systems and services.

ACKNOWLEDGMENTS

This work is part funded by the UK EPSRC under grants EP/N028228/1 (PACTMAN) and EP/N023234/1 (PETRAS IoT Research Hub – Cybersecurity of the Internet of Things).

REFERENCES

- [1] Hal Berghel. 2004. Wireless infidelity I: War driving. *Commun. ACM* 47, 9 (2004), 21–26.
- [2] Jonathan Brandon. 2018. Thingful aims to be the Google of the Internet of Things. <http://telecoms.com/206211/thingful-aims-to-be-the-google-of-the-internet-of-things/>. (2018). Accessed: 2018-10-17.
- [3] Jaewoo Chung, Matt Donahoe, Chris Schmandt, Ig-Jae Kim, Pedram Razavai, and Micaela Wiseman. 2011. Indoor location sensing using geo-magnetism. In *Proceedings of the 9th international conference on Mobile systems, applications, and services*. ACM, 141–154.
- [4] Davide Dardari, Pau Closas, and Petar M Djuric. 2015. Indoor Tracking: Theory, Methods, and Technologies. *IEEE Trans. Vehicular Technology* 64, 4 (2015), 1263–1278.
- [5] Nigel Davies, Marc Langheinrich, Sarah Clinch, Ivan Elhart, Adrian Friday, Thomas Kubitzka, and Bholanathsingh Surajbali. 2014. Personalisation and privacy in future pervasive display networks. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2357–2366.
- [6] Suparna De, Tarek Elsaleh, Payam Barnaghi, and Stefan Meissner. 2012. An internet of things platform for real-world and digital objects. *Scalable Computing: Practice and Experience* 13, 1 (2012), 45–58.
- [7] Chuck Eastman, Paul Teicholz, Rafael Sacks, and Kathleen Liston. 2011. *BIM handbook: A guide to building information modeling for owners, managers, designers, engineers and contractors*. John Wiley & Sons.
- [8] Shane B. Eisenman, Emiliano Miluzzo, Nicholas D. Lane, Ronald A. Peterson, Gahng-Seop Ahn, and Andrew T. Campbell. 2010. BikeNet: A Mobile Sensing System for Cyclist Experience Mapping. *ACM Trans. Sen. Netw.* 6, 1, Article 6 (Jan. 2010), 39 pages.
- [9] Moustafa Elhamshary, Moustafa Youssef, Akira Uchiyama, Hirozumi Yamaguchi, and Teruo Higashino. 2016. TransitLabel: A crowd-sensing system for automatic labeling of transit stations semantics. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services*. ACM, 193–206.
- [10] Sidhant Gupta, Matthew S. Reynolds, and Shwetak N. Patel. 2010. ElectriSense: Single-point Sensing Using EMI for Electrical Event Detection and Classification in the Home. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing (UbiComp '10)*. ACM, New York, NY, USA, 139–148.
- [11] Mordechai Haklay and Patrick Weber. 2008. Openstreetmap: User-generated street maps. *Ieee Pervas Comput* 7, 4 (2008), 12–18.
- [12] Andy Harter, Andy Hopper, Pete Steggle, Andy Ward, and Paul Webster. 1999. The Anatomy of a Context-aware Application. In *Proceedings of the 5th Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '99)*. ACM, New York, NY, USA, 59–68.
- [13] Sehyeon Heo, Sungpil Woo, Janggwan Im, and Daeyoung Kim. 2015. IoT-MAP: IoT mashup application platform for the flexible IoT ecosystem. In *Internet of Things (IOT), 2015 5th International Conference on the*. IEEE, 163–170.
- [14] IoT Ecosystem Demonstrator Interoperability Working Group and Rodger Lea. 2013. *HyperCat: an IoT interoperability specification*. IoT ecosystem demonstrator interoperability working group.
- [15] Kasthuri Jayarajah, Rajesh Krishna Balan, Meera Radhakrishnan, Archan Misra, and Youngki Lee. 2016. LiveLabs: Building In-Situ Mobile Sensing & Behavioural Experimentation TestBeds. In *Proceedings of the 14th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '16)*. ACM, New York, NY, USA, 1–15. <https://doi.org/10.1145/2906388.2906400>
- [16] Rui José, Nuno Otero, Shahram Izadi, and Richard Harper. 2008. Instant places: Using bluetooth for situated interaction in public displays. *IEEE Pervasive Computing* 7, 4 (2008).
- [17] Minkyong Kim, Jeffrey J Fielding, and David Kotz. 2006. Risks of using AP locations discovered through war driving. In *International Conference on Pervasive Computing*. Springer, 67–82.
- [18] P. T. Kirstein and A. Ruiz-Zafra. 2018. Use of templates and the handle for large-scale provision of security and IoT in the built environment. (March 2018), 10 pages.
- [19] Manikanta Kotaru, Kiran Joshi, Dinesh Bharadia, and Sachin Katti. 2015. SpotFi: Decimeter Level Localization Using WiFi. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication (SIGCOMM '15)*. ACM, New York, NY, USA, 269–282. <https://doi.org/10.1145/2785956.2787487>
- [20] Fei Li, Michael Vögler, Markus Claeßens, and Schahram Dustdar. 2013. Efficient and scalable IoT service delivery on cloud. In *Cloud Computing (CLOUD), 2013 IEEE Sixth International Conference on*. IEEE, 740–747.
- [21] Ren C Luo and Chun C Lai. 2012. Enriched indoor map construction based on multisensor fusion approach for intelligent service robot. *IEEE Transactions on Industrial Electronics* 59, 8 (2012), 3135–3145.
- [22] Markus Miettinen, Samuel Marchal, Ibbad Hafeez, N Asokan, Ahmad-Reza Sadeghi, and Sasu Tarkoma. 2017. IoT Sentinel: Automated device-type identification for security enforcement in IoT. In *Distributed Computing Systems (ICDCS), 2017 IEEE 37th International Conference on*. IEEE, 2177–2184.
- [23] Rui Ning, Cong Wang, Chunsheng Xin, Jiang Li, and Hongyi Wu. 2018. DeepMag: Sniffing Mobile Apps in Magnetic Field through Deep Convolutional Neural Networks. In *2018 IEEE International Conference on Pervasive Computing and Communications, PerCom 2018, Athens, Greece, March 19-23, 2018*. 1–10.
- [24] Damian Philipp, Patrick Baier, Christoph Dibak, Frank Durr, Kurt Rothermel, Susanne Becker, Michael Peter, and Dieter Fritsch. 2014. Mapgenie: Grammar-enhanced indoor map construction from crowd-sourced data. In *2014 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 139–147.
- [25] Radius Networks. 2018. WikiBeacon by Radius Networks. <http://www.wikibeacon.org/>. (2018). Accessed: 2018-10-18.
- [26] Mahadev Satyanarayanan, Victor Bahl, Ramón Caceres, and Nigel Davies. 2009. The case for vm-based cloudlets in mobile computing. *IEEE pervasive Computing* (2009).
- [27] Peter Shaw, Mateusz Mikusz, Nigel Davies, and Sarah Clinch. 2017. Using Smartwatches for Privacy Awareness in Pervasive Environments. *HotMobile 2017* (2017).
- [28] statista. 2018. Internet of Things - number of connected devices worldwide 2015-2025). <https://www.statista.com/statistics/471264/iot-number-of-connected-devices-worldwide/>. (2018). Accessed: 2018-10-18.
- [29] Yiguang Xuan, Raja Sengupta, and Yaser Fallah. 2010. Crowd sourcing indoor maps with mobile sensors. In *International Conference on Mobile and Ubiquitous Systems: Computing, Networking, and Services*. Springer, 125–136.