

Pervasive Communities in the Internet of People

Eemil Lagerspetz*, Huber Flores*, Niko Mäkitalo*, Pan Hui*[†], Petteri Nurmi[‡], Sasu Tarkoma*, Andrea Passarella[§], Jörg Ott[¶], Peter Reichl^{||}, Marco Conti[§], Markus Fiedler^{**}, Jatinder Singh^{††}, Thorsten Strufe^{‡‡}, Tobias Hossfeld^x, Anders Lindgren^{xi}, Daniele Quercia^{xii}

*University of Helsinki, Finland, {firstname.lastname}@helsinki.fi

[†]The Hong Kong University of Science and Technology, Hong Kong, panhui@cse.ust.hk

[‡]Lancaster University, United Kingdom, p.nurmi@lancaster.ac.uk

[§]IIT-CNR, Italy, {firstname.lastname}@iit.cnr.it

[¶]Technical University of Munich, Germany, ott@in.tum.de

^{||}University of Vienna, Austria, peter.reichl@univie.ac.at

^{**}Blekinge Institute of Technology, Sweden, markus.fiedler@bth.se

^{††}University of Cambridge, United Kingdom, js573@cam.ac.uk

^{‡‡}Technische Universität Dresden, Germany, thorsten.strufe@tu-dresden.de

^xUniversity of Duisburg-Essen, Germany, tobias.hossfeld@uni-due.de

^{xi}RISE SICS, Sweden, anders.lindgren@ri.se

^{xii}Bell Labs Cambridge, United Kingdom, daniele.quercia@nokia-bell-labs.com

Abstract—The Internet has traditionally been a device-oriented architecture where devices with IP addresses are first-class citizens, able to serve and consume content or services, and their owners take part in the interaction only through those devices. The Internet of People (IoP) is a recent paradigm where devices become proxies of their users, and can act on their behalf. To realize IoP, new policies and rules for how devices can take actions are required. The role of context information grows as devices act autonomously based on the environment and existing social relationships between their owners. In addition, the social profiles of device owners determine e.g. how altruistic or resource-conserving they are in collaborative computing scenarios. In this paper we focus on *community formation* in IoP, a prerequisite for enabling collaborative scenarios, and discuss main challenges and propose potential solutions.

I. INTRODUCTION

Smart devices have become pervasive in our environment, providing access to computing capabilities that can enrich our daily lives and support our activities. Examples of relevant application domains include sports, health, transportation, and localization. At the same time, users are starting to own several smart devices, like, smartwatches, smart glasses, smartphone, and sensors and actuators are being embedded and deployed in the environments the user visits, e.g., smart homes, smart buildings, and autonomous cars. Since smart devices are constrained by energy and processing resources, they must rely on strategies for improving their resource use. One promising strategy is collaboration whereby devices share the cost of executing the applications to improve the performance of the applications that are being executed [1], [2].

Previous research has investigated the importance of context-awareness [3] for devices to engage into collaboration through the formation of communities. However, it is difficult to envision the adoption of those solutions in reality as many social aspects, e.g., lack of incentives, security, privacy,

and usability constrain the interactivity between devices. For instance, a user can be reluctant to allow other devices to consume its energy without compensation [4]. As another example, user's private information can be stolen when engaging into cooperation by malicious users. Therefore, creating collaborative communities of devices is challenging.

The Internet of People [5], [6] (IoP) is an emerging paradigm where devices become representatives of their owners and can act on their behalf. They become proxies or interfaces to people in IoP [7]. Devices may also be associated to a particular owner, such that for addressing a particular device, it is necessary to know the owner's identity. This opens up possibilities beyond simple traditional one-to-one sharing of data between users, such as communities being able to access a limited set of each others' data without explicit need to share to each member of the community. This may create incentives for users to share the resources of their devices, and provide services to their friends (cf. [8]). Sharing their computing or communication resources, they can maximize the overall performance or minimize total battery drain among a group of participating devices.

This IoP vision requires a new level of context awareness, as devices now act on behalf of their owners in a rich, multi-device, multi-user environment, as illustrated in Figure 1a. Devices also need to obtain information about the social context they are operating in, so they can share resources as their owners would. In addition, devices can follow social profiles defined by their owners. The profiles define properties such as the level of altruism towards friends' devices low on battery, the extent of Internet bandwidth or CPU power sharing to friends and acquaintances, etc (cf. Figure 1b). Such profiles for the IoP have been studied in previous work [7], [6].

In this paper, we focus on community formation within IoP, a necessary first step for many IoP applications and a




Fig. 1: (a): Interaction of devices based on user preferences, and (b): Social profiles for devices

fundamental enabler for harnessing the potential of IoP. We focus our attention on a scenario where devices owned by different entities interact, taking into account other activities taking place, the social profiles of involved users, and environmental factors such as traffic, weather, and network conditions. We highlight the challenges and potential solutions that lend themselves for the orchestration of pervasive devices into a community, analyzing the drawbacks and benefits of creating pervasive communities in the wild.

II. MOTIVATION: SMART THINGS FOR DAILY LIFE

Let us consider how multi-device setups taking into account social relationships can enrich the daily life of users. John is working in his office when his boss calls to ask him to give a presentation to some important customers later. After booking the conference room for the event, John's phone proactively proposes sending the meeting information to the attendees.

While John continues working, he can hear how the printer at the office automatically starts printing copies of the meeting information for everybody who has accepted the meeting invitation, just like he was expecting. John's mobile phone also communicates with the sociological profile of the other attendees and asks if they want to have coffee in the meeting. Almost every phone respond that their owners drink coffee, and the coffee machine at the office gets commanded to make sure that it is ready to make coffee at the time of the meeting. However, Jane who is one of the important customers, prefers chai tea instead, and thus John can now make a good impression by ordering her good chai tea.

While the clients are on their way from airport to the office by taxi, the car informs their phones that the speed is abnormally low and that they are basically in a traffic jam. Furthermore, the clients phones can now inform John's phone that they are not going to arrive on time. By saying out loud John's phone informs him that the guests are little late, and that John could spent the time by replying to some emails. At the same time, the coffee maker is proactively commanded to delay the time to make the coffee. When the guests finally are near the office building, John's phone proposes him to go to the lobby to welcome the guests, and also reminds John to get the chai tea from the cafeteria for Jane.

Although the meeting goes well, it started late and the clients are becoming hungry. As a good host John offers

to go to a restaurant together with the clients, and makes a reservation for them from a web page of his favorite restaurant. Based on the reservation confirmation, his phone realizes that John will not get home in time to see his favorite TV show, and hence his Smart TV system is commanded to record tonight's episode. At the same time, the sociological profile on John's phone also knows that he typically watches the TV show together with his wife, and thus asks her not to talk about and spoil the episode when he comes home after the dinner.

The device of one of John's friends responds to say that her owner will not be able to see the show either and perhaps they could watch the new episode together the next day. Both devices show an invitation to their owners to watch the episode together and, if both accept the invitation, John's device will send a message to his fridge to add popcorn to his shopping list for the next day.

III. RELATED WORK

Opportunistic strategies [9] for self-organization of devices into a decentralized mesh of resources has been investigated in many areas, such as wireless sensor networks [10], Mobile ad-hoc networks [11], peer-to-peer systems [12], [4], and transient infrastructure [13]. Given that the self-organization of devices is goal oriented, e.g., energy, performance, and context influence its outcome. There is a rich body of research on *context awareness* and its impact in practice [14], [3]. Context aware applications are those that automatically adapt their functionality based on the application's and user's actions, which are detected by the device.

Context has typically included awareness of the location, time, interactions and behavior patterns of its user. However, with the proliferation of devices partly brought on by IoT, nearby devices and the relationships between the application user and the owners of those devices, have become important and critical for conservation of resources via collaboration, e.g., energy. Several work has investigated the inclusion of social relationships within device's context to facilitate autonomous interaction and foster proactive intercommunication between devices [15], [16]. For example, movies can be recommended to groups of friends taking collective preferences into account. Similarly, sharing of information can be limited to a number of hops in the social network, or to contacts above a set level of familiarity. Social-aware traffic management utilizes social information to optimize Internet traffic management wrt. traffic load, energy consumption, or Quality of Experience [17].

However, sharing context information, granting resources access to other devices and exchanging data between devices, is a complex problem that cannot be generalized just by including social information into the context. Multiple devices from a particular user can have different constraints when engaging in device to device cooperation. Social relationships change over time, with flexibility beyond context awareness. In addition, the ownership of devices from users extend beyond their local context, e.g., cars, drones, homes, etc. These devices also have

autonomous behaviors, which should reflect the preferences of users when working in isolation without supervision.

Unlike other work, we overcome the problems of collaboration of devices in pervasive communities by exploring the IoP paradigm. The key insight is that any device carried by users or deployed in the wild is identified by a social profile, which defines different levels of interaction with others depending on the preferences of the user. For instance, the behavior of appliances in a smart home could be configured by the devices of relatives, while friends or acquaintances just have access to a limited set of functionality while at the smart home.

IV. CHALLENGES FOR COMMUNITY FORMATION FROM THE INTERNET OF PEOPLE

IoP builds on a vision where devices act to a higher degree of autonomy than what they are currently able to, sharing resources with devices belonging to users that are friends or relatives of the owner. Resources to share can include computational power, sensor data, or even connectivity, instead of merely sharing content between related devices. In the IoP Manifesto [6], a set of four *core principles* are presented as a basis for the interactions between devices: (P1) *Be Social* requires that the devices should consider and reflect the social relationships of their owners to the interactions. (P2) *Be Personalized* points that the preferences of not only one user, but the preferences of all the participating users should reflect to the collective interactions. (P3) *Be Proactive* means that the devices should be enabled to initiate the interactions between users and with other devices in an automated manner. (P4) *Be Predictable* principle highlights the fact that devices contain a lot of personal and sensitive information, so it is essential to protect the user, and hence be careful what kind of interactions are allowed with whom, and in which context. Common to all these principles is the need for devices to interact together. In this paper we target specifically *community formation*, a prerequisite for devices to interact together, and challenges posed to it by these four core principles of IoP.

A. Control

Devices can interact with each other rapidly, even many times per second. It is therefore crucial that control of these interactions is transparent and that mechanisms for changing access control preferences are accessible to their owners in a clear way (P4). It should be possible to not only allow or deny interactions with individual devices, but rather authorize the owner, social group, and context basis. Instead of allowing all nearby devices to interact with the Internet connection and a home music player during a party, IoP will authorize family members (and solely their devices) to access certain resources, like a security camera (P1 and P2).

B. Transparency

When devices are allowed to act on behalf of their ‘owners’ in an autonomous way (P3), we need to establish which devices initiated actions and which of their owners’ preferences they followed when doing so. This keeps owners aware of the

activities their devices, and therefore the owners themselves, are responsible for, while giving owners some confidence that their devices are operating appropriately on their behalf (P4). Further, recording the policies applied and actions undertaken will assist investigation in situations of failure. One challenge in this space is where a device is shared or services a group of people, whereby a myriad of individual preferences may need to be managed (P1).

C. Context Awareness

With a rich number of sensors and personal devices available in the environment, the importance of context awareness increases when sharing resources (P1). Multi-device interactions pose challenges for context awareness. For instance, when sharing communication interfaces to a nearby device to accelerate apps, it might share these resources further. Therefore, we need to discover also these other nearby devices and agree on shared resources in a context-aware fashion (P4). The way resources are shared with other devices in the same temporal or spatial context may be further restricted or increased by the social context of those devices as outlined in the next section.

D. Social Context

The natural graph of interpersonal relationships presents both opportunities and challenges (P1, P2, and P3 vs. P4). Accessing and intelligently using the graph beyond single-hop relationships may be questionable with respect to privacy, and also difficult on resource restricted devices. Access to large scale data processing facilities, like fog and cloud computing, however, may alleviate this situation.

Social relationships are also an opportunity to automate resource sharing when the local context presents us with unknown devices (P1). Device owners can specify groups such as friends and friends-of-friends with different levels of access to their devices at the home or workplace; various groups such as coworkers can have access to devices only at the workplace, but not at the home, unless also members of the friends group. There are plenty of possibilities from combining context awareness with social relationships.

V. FORMING PERVASIVE COMMUNITIES IN IOP

Smart devices are constrained by energy and performance issues when they operate in isolation. Yet, several studies have shown that smart devices are frequently co-located in proximity to at least one other device throughout the day, suggesting that devices can potentially collaborate to reduce the effort of resource intensive tasks, e.g., sensing [1], offloading [4], networking [13], storage, etc. Pervasive communities are formed then when devices in-situ shared their resources to perform tasks in conjunction in an opportunistic manner.

However, merging the resources of multiple devices to work together according to the social relationships of their users is a tough challenge as it requires a common understanding of the context of each device. This consists of multiple properties and configurations that influence the runtime behavior of the




Fig. 2: (a) No interaction available, (b) Interaction based on context and (c) Interaction on social profile.

device, e.g., nearby other devices, activities currently taking place, type of device, temperature, CPU usage, memory, bandwidth, sensing frequency, number of apps, etc.

In addition, since smart devices function in multiple roles, understanding what really constitutes context becomes difficult as different types of contexts need to be modeled depending on the type of task. Context information is critical in multi device setups, e.g., for the formation of collaboration groups, and for negotiating responsibilities between the devices. Any misunderstandings in context can thus be extremely counterproductive for the collaboration between devices. While IoP introduces a new level of context awareness for devices, the formation of pervasive communities remain challenging as devices need a deep understanding about their context and owner’s preferences to engage into device to device collaborations. In this section, we highlight these challenges and potential solutions.

A. Community context modeling

Modeling the context of a device is a complex task as it is influenced by many parameters [18], e.g., type of communication, type of device, device’s temperature, number of apps in background, etc. It has been demonstrated that a single device is unable to characterize its own context as it is an exhaustive and overwhelming task [3]. One potential solution to overcome this problem is to rely on a community of devices and then model the device’s context by using passive aggregated sampling. The EMCO (Evidence-aware Mobile Computational Offloading) framework is a clear example of this approach [3]. While IoP introduces extra complexity for context modeling as aggregated social data also needs to be collected, it reduces the complexity of decision making over a context as a device is equipped with different level of context awareness, whose computation complexity gets higher based on the type of social interactivity that a particular user has with other users. This means that the computational cost of granting context awareness to devices to form a pervasive community is higher at the first encounter, and reduces over time.

B. Community formation

A device needs to be aware about its context, and other devices’ context to dynamically evaluate the impact of sharing resources in a collaborative manner. A collaboration between devices can be productive or counterproductive when sharing resources, e.g., bandwidth, processing, etc. It is productive when the computational cost (measured in terms of energy or performance) is reduced, otherwise, it is counterproductive. Usually, collaborative tasks are goal oriented, which suggests that devices share the same goal. However, in the IoP, as long as the collaboration is meaningful from users perspective, devices have to interact to fulfill the owner’s goals. This suggests that devices can merge resources without evaluating whether their own computational cost is reduced or not. For instance, a user with many devices can select the device with the highest battery life, e.g., smartphone, to help his other devices with lower energetic sources, e.g., smartwatch, smart glasses. Another example is a group of friends, in which the user with the lowest battery life can rely on his friends’ devices to extend his battery life.

Naturally, a collaborative model, where devices are exploited without achieving gains in the long term is doomed to collapse, e.g., selfishness. Thus, incentive mechanisms where devices are awarded fairly for their help is encourage. For instance, HyMobi (Social-aware Hybrid Mobile Offloading) [4] is a framework that allows a smartphone to offload computational tasks to other mobile devices in proximity. By using HyMobi, devices that received the computational task for processing are paid using a digital currency, and the payment is proportional to the amount of energy wasted by the devices when processing the task [19].

C. Opportunistic community

While devices can be aware about explicit collaborations to improve energy and performance, devices can also be opportunistically utilized in periods of no interactivity with others (idle times). Devices in fixed locations, e.g., smart television, smart refrigerators, etc., and personal devices, e.g., mobile devices, smart watches, and personal computers, can be used in these periods to process data as a computing cluster (Micro

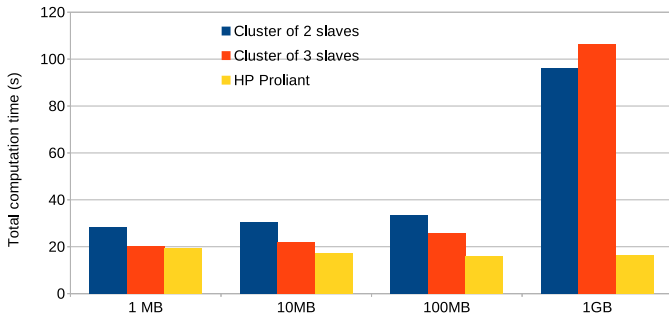


Fig. 3: Execution speed of the Mandelbrot set task on a smartphone cluster and a server.




Fig. 4: Total energy consumption the Mandelbrot set task on a smartphone cluster and a server.

data centers). Example uses include air quality monitoring inside a building, level of humidity of an environment, network conditions of the local area, or discovery of users within a particular place. Smartphones can be a powerful computing platform; our initial results in the UbiSpark Project¹ show promise in this area (See Section VI).

VI. FEASIBILITY STUDY

In this section, we explore the potential gains of merging multiple devices using latest solutions. In addition, we envision a more sophisticated scenario that involves a large number of multi-device configurations.

Smartphones of today have enough power to rival super-computers of the past. We have conducted an experiment running a distributed Mandelbrot Set computation task on the UbiSpark Prototype developed at University of Helsinki¹, using smartphones, and compared execution speed and total energy (taking into account execution time). We found that up to 100 MB task size, a smartphone cluster of 3 slaves has similar performance to the server, and since the execution time scales linearly, that two clusters, or 8 Samsung Galaxy S 4 smartphones (1 master and 3 slaves each), would achieve better performance than a HP Proliant server about the same age (see Figure 3). Our energy efficiency results in Figure 4 show that as the size of the Mandelbrot Set result is increased, energy consumption grows, and that up to 100 MB task size, two smartphone clusters would use 12%–18% of the energy that the server spends for the same task. With a 1000 MB task size the WiFi communication channel is saturated and performance decreases; this shows that there are limits to how much data can be processed by smartphone clusters at a time.

In summary, it is feasible to carry out heavy computation on smartphones, up to a point. It is also much more energy efficient than using server hardware. We can leverage the power of computing on smart devices in the pervasive communities in the Internet of People, letting nearby devices form computing clusters to execute tasks on their owners' behalf.

¹<http://ubispark.cs.helsinki.fi/>

VII. DISCUSSION

In addition to the challenges and solutions discussed, forming device communities in IoP requires further investigation, particularly on the topics outlined below.

A. Identity Management of Users and Social Roles

A key enabling mechanism for the IoP are identity management and the social profiles of users, which allow on one hand users to access community networks and services and on the other hand to make devices act on behalf of the users as their *cyber me*. The functional and non-functional requirements need to be analyzed in order to build a scalable identity management architecture. In addition, the identity management must be able to detect free-riders that only utilize resources or services without contributing. This issue can also be overcome using incentive mechanisms, such as those in P2P networks, to avoid selfish peers. Some of the above issues are covered in our recent article [20] where coalitions of trusted devices are formed using an advanced voting architecture. Furthermore, users may take on different social roles and exhibit different behaviors in different communities.

B. Partial Sharing of Data and Data Access

Certainly, the context of devices is a key component to consider when merging computational resources of smart devices in proximity. Usually, the context of devices is modeled from passive data as it provides better characterization of specific situations, e.g., battery life characterization via crowdsensing. However, in the IoP, the user plays a critical role in the management of the resources of devices. Consequently, active data, e.g., activity, gesture and audio recognition among others, has to be considered within the context of devices for providing better control and awareness over actions performed autonomously by devices. For instance, a user can be reluctant to share information even between relatives from time to time. In this case, specific gestures or voice commands can be included within the context in real-time to prevent sharing. Together with identity management, mechanisms to access data and certain social profiles are to be developed.

In the IoP, devices act on behalf of a user or social entity. However, user preferences on social interactivity can change drastically on the fly. Thus, means are required for dynamically managing the communication and synchronization of devices.

This relates to issues of responsibility. Given the social, interactive and often multi-device nature of the IoP, challenges are raised regarding which user preferences have priority, and who is responsible when failures inevitably occur.

Moreover, technology is increasingly the subject of legal and regulatory attention; for instance, data protection regulation means that certain rights, responsibilities and obligations can flow with the exchange of personal data, and many legal considerations will have direct implications for the technical infrastructure [21]. These concerns will affect technology adoption and acceptance. Therefore, it is important that the technical community also consider the broader legal and social context in order for the IoP to flourish.

D. Naming and Addressing Challenges

So far, we have assumed unlimited visibility among IoP devices. However, addressing issues can be efficient show-stoppers for exchanging information and tasks. Therefore, a generic, robust, scalable and future-proof naming and addressing scheme needs to be devised, which even allows to incorporate relationship and context information.

E. A General Manifesto

Certainly, IoP is ruled primarily by policies and guidelines that define the actions that devices can perform when interacting between each other. However, recent technological advances toward the virtualization of physical devices introduce new challenges and opportunities to define social profiles and the behavior of devices as discussed in [6]. For instance, a physical device can be virtualized in the cloud from its collected data, such that the virtual device can emulate the behavior of the physical one. In this case, the physical device can delegate interactions to the virtual one that can be categorized as risky or potential threats.

VIII. CONCLUSIONS

We have outlined key challenges with community formation in the Internet of People (IoP), as well as provided potential solutions. We can define typical contexts via crowdsourcing, use offloading to take advantage of nearby resources, as well as bring computation directly to smart devices as in the UbiSpark Project. We suggest using an incentive model based on social relationships as well as digital currency as in the HyMobi system. Our results show that computing speeds similar to servers can be achieved by the combined effort of several smart devices, making pervasive communities possible in terms of computing resources.

ACKNOWLEDGMENTS

Key insights about the paper were discussed in the Dagstuhl Seminar 17412 on Internet of People in October 8-11, 2017. This work was partially funded by the Academy of Finland projects 297741, 296139, 295913, and 303825.

- [1] Y. Lee, Y. Ju, C. Min, S. Kang, I. Hwang, and J. Song, "Comon: Cooperative ambience monitoring platform with continuity and benefit awareness," in *Proceedings of the ACM International Conference on Mobile Systems, Applications, and Services (MobiSys 2012)*, (Low Wood Bay, Lake District, United Kingdom), June 25–29 2012.
- [2] J. Leppänen, M. Pelkonen, H. Guo, S. Hemminki, P. Nurmi, and S. Tarkoma, "Collaborative and Energy-Efficient Speech Monitoring on Smart Devices," *Computer*, vol. 49, no. 12, pp. 22–30, 2016.
- [3] H. Flores, P. Hui, P. Nurmi, E. Lagerspetz, S. Tarkoma, J. Manner, V. Kostakos, Y. Li, and X. Su, "Evidence-aware Mobile Computational Offloading," *IEEE Transactions on Mobile Computing*, 2017.
- [4] H. Flores, R. Sharma, D. Ferreira, V. Kostakos, J. Manner, S. Tarkoma, P. Hui, and Y. Li, "Social-aware hybrid mobile offloading," *Pervasive and Mobile Computing*, vol. 36, pp. 25–43, 2017.
- [5] M. Conti, A. Passarella, and S. K. Das, "The Internet of People (IoP): A new wave in pervasive mobile computing," *Pervasive and Mobile Computing*, vol. 41, pp. 1–27, 2017.
- [6] J. Miranda, N. Mäkitalo, J. Garcia-Alonso, J. Berrocal, T. Mikkonen, C. Canal, and J. M. Murillo, "From the Internet of Things to the Internet of People," *IEEE Internet Computing*, vol. 19, no. 2, pp. 40–47, 2015.
- [7] J. Guillen *et al.*, "People as a Service: A Mobile-centric Model for Providing Collective Sociological Profiles," *IEEE software*, vol. 31, no. 2, pp. 48–53, 2014.
- [8] L. A. Cutillo, R. Molva, and T. Strufe, "Safebook: Feasibility of transitive cooperation for privacy on a decentralized social network," in *World of Wireless, Mobile and Multimedia Networks*, 2009, pp. 1–6.
- [9] A. Passarella and M. Conti, "Analysis of individual pair and aggregate intercontact times in heterogeneous opportunistic networks," *IEEE Transactions on Mobile Computing*, vol. 12, no. 12, 2013.
- [10] X. Mao *et al.*, "Energy-efficient opportunistic routing in wireless sensor networks," *IEEE transactions on parallel and distributed systems*, vol. 22, no. 11, pp. 1934–1942, 2011.
- [11] C. Boldrini, M. Conti, and A. Passarella, "The stability region of the delay in Pareto opportunistic networks," *IEEE Transactions on Mobile Computing*, vol. 14, no. 1, pp. 180–193, 2015.
- [12] Y. Li, T. Wu, P. Hui, D. Jin, and S. Chen, "Social-aware D2D communications: Qualitative insights and quantitative analysis," *IEEE Communications Magazine*, vol. 52, no. 6, pp. 150–158, 2014.
- [13] B. Han, P. Hui, and A. Srinivasan, "Mobile data offloading in metropolitan area networks," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 14, no. 4, pp. 28–30, 2011.
- [14] A. K. Dey, "Understanding and using context," *Personal and ubiquitous computing*, vol. 5, no. 1, pp. 4–7, 2001.
- [15] A. Beach *et al.*, "Fusing mobile, sensor, and social data to fully enable context-aware computing," in *Proceedings of the Eleventh Workshop on Mobile Computing Systems & Applications (HotMobile 2010)*, (Annapolis, MD, USA), February 22 - 23, 2010.
- [16] N. Mäkitalo *et al.*, "Social Devices: Collaborative Co-located Interactions in a Mobile Cloud," in *Proceedings of the ACM International Conference on Mobile and Ubiquitous Multimedia (MUM 2012)*, (New York, NY, USA), December 04–06, 2012.
- [17] M. Seufert, G. Darzanos, I. Papafili, R. Łapacz, V. Burger, and T. Hoßfeld, "Socially-aware traffic management," in *Socioinformatics-The Social Impact of Interactions between Humans and IT*. Springer, 2014, pp. 25–43.
- [18] H. Flores and S. Srirama, "Mobile code offloading: should it be a local decision or global inference?" in *Proceedings of the ACM International Conference on Mobile systems, applications, and services (MobiSys 2013)*, (Taipei, Taiwan), June 25–28, 2013.
- [19] S. Hosio *et al.*, "Monetary assessment of battery life on smartphones," in *Proceedings of the ACM Annual International Conference on Human Factors in Computing Systems (CHI 2016)*, (San Jose, CA, USA), May 7–12, 2016.
- [20] N. Mäkitalo, A. Ometov, J. Kannisto, S. Andreev, Y. Koucheryavy, and T. Mikkonen, "Safe, Secure Executions at the Network Edge: Coordinating Cloud, Edge, and Fog Computing," *IEEE Software*, vol. 35, no. 1, pp. 30–37, 2018.
- [21] J. Singh, T. Pasquier, J. Bacon, J. Powles, R. Diaconu, and D. Eyers, "Policy-driven Middleware for a Legally-Compliant Internet of Things," in *Proceedings of the ACM 17th International Middleware Conference (Middleware 2016)*, (Trento, Italy), December 12–16, 2016.