Lowering the Barriers to Large-Scale Mobile Crowdsensing

Yu Xiao
Carnegie Mellon University
Aalto University
yuxiao@cs.cmu.edu

Pieter Simoens
Carnegie Mellon University
Ghent University - iMinds
pieter.simoens@intec.ugent.be

Padmanabhan Pillai
Intel Labs
padmanabhan.s.pillai@intel.com

Kiryong Ha
Carnegie Mellon University
krha@cmu.edu

Mahadev Satyanarayanan
Carnegie Mellon University
satya@cs.cmu.edu

ABSTRACT

Mobile crowdsensing is becoming a vital technique for environment monitoring, infrastructure management, and social computing. However, deploying mobile crowdsensing applications in large-scale environments is not a trivial task. It creates a tremendous burden on application developers as well as mobile users. In this paper we try to reveal the barriers hampering the scale-up of mobile crowdsensing applications, and to offer our initial thoughts on the potential solutions to lowering the barriers.

1. INTRODUCTION

In recent years, there has been phenomenal growth in the richness and diversity of sensors on smartphones. It is now common to find two cameras, a GPS module, an accelerometer, a digital compass, a gyroscope and a light sensor in a single smartphone. And there is more to come! The rich information about the smartphone user's activity and environment provided by these sensors inspired the first wave of sensing applications that personalized user experience based on the sensed context. Now, a second wave of mobile sensing applications is gaining momentum. The focus has shifted from individual sensing towards *crowdsensing,* defined as "individuals with sensing and computing devices collectively sharing information to measure and map phenoma of common interest" [10]. Initially, crowdsensed inputs were analyzed offline, for example in the analysis of transportation activities in urban spaces [34], for the measurement of interperson similarity [14], or for mental and physical health assessment of elder people [23]. In more recent crowdsensing applications, the collected inputs are processed in real time. Examples include traffic monitoring [35, 36], public safety management [27], and collaborative searching [31].

A hypothetical use case serves to illustrate the potential benefits of crowdsensing using information-rich multimedia sensors and some potential pitfalls [25]. Imagine that a small child gets lost while watching a parade in the middle of a large city. The distraught parents, upon noticing their child is missing, immediately use their smartphone to initiate a search, providing sample images with their child's face. A crowdsensing search application tries to match these with the videos and images being captured by the many smartphone cameras in the crowd. Any potential matches are forwarded to the parents' phone, along with GPS location information. With thousands of electronic eyes applied to this problem, the child is quickly found, before she herself is even aware of being lost. For this use case, the large number of smartphone cameras in use makes it likely the child appears in one or more captured images; however, the crowdsensing search application itself can succeed only if a sufficiently large number of smartphone users participate.

More generally, there is a growing realization that *scale* is the key to the success of crowdsensing applications. Since individual users may go offline and individual sensor readings may be inaccurate or corrupted, the reliability and trustworthiness of crowdsensing applications scales more than proportionally with the number of users. Access to a vast user base is thus crucial. However, our survey of the literature shows that today's mobile crowdsensing applications using physical sensors like GPS have rarely been scaled up to more than 1,000 participants.

Table 1 shows a representative sample of crowdsensing studies. Much to our surprise, the number of participants is often omitted in the papers reporting these studies. When concrete numbers are provided, the crowd sizes are usually small. It is only with data sources that are easy to collect (e.g. from social networking applications such as Twitter) that larger crowds have been studied. The one notable exception is the work of Balan et al [2], discussed in Section 2.

What limits the scaling of crowdsensing applications? In this paper, we explore this issue and and propose an architectural solution. We then explore the merits of this architecture, and discuss potential implementation challenges.

2. OBSTACLES TO CROWD SCALING

Crowdsensing applications, including the ones that exist today and the emerging class of applications making use of richer multimedia sensors, face three major barriers to achieving the large crowd sizes critical to their success.

The first obstacle is the heterogeneity of sensing hardware and mobile platforms. In today's mobile device market, there are at least three popular software platforms, includ-

Reference	Mobile Platform	Application Category	Crowd Size	Input
Zhou et al. [35] (Mobisys 2012)	Android	Transportation	unknown	cell tower ID, audio signal accelerometer
Tiramisu [36] (CHI 2011)	iOS	Transportation	28	GPS
SignalGuru [12] (MobiSys 2011)	iOS	Transportation	13	video frames
Balan et al. [2](Mobisys 2011)	Car GPS	Transportation	15000	GPS
Mathur et al. [16] (Mobisys 2010)	Car GPS	Transportation	500	GPS
Niu et al. [20] (Com.geo 2011)	Blackberry	Transportation	unknown	GPS
Bao et al. [3] (Mobisys 2010)	Symbian, iPod	Social Application	unknown	video
Wirz et al. [30] (SCI 2011)	Android	Social Application	unknown	GPS
CrowdSearch [31] (Mobisys 2010)	iOS	Search	unknown	image
GeoLife [34] (WWW 2009)	GPS phones	User Behavior Study	107	GPS
SoundSense [15] (Mobisys 2009)	iOS	User Behavior Study	unknown	audio stream
#EpicPlay [28] (CHI 2012)	Twitter	Social Application	unknown	tweets
Wakamiya et al. [29] (ICUIMC 2012)	Twitter	User Behavior Study	39898	tweets
Fujisaka et al. [9](ICUIMC 2012)	Twitter	User Behavior Study	8139	tweets
CrowdSearcher [5] (WWW 2012)	Facebook	Search	137	text

Table 1: Representative Sample of Crowd-sensing Applications

ing Android, iOS and Windows 8. Applications written for any of these can not be run on the others. Even different versions of a particular platform are sometimes incompatible, due to changes in hardware or evolution of the software APIs. Furthermore, the Apps model in vogue today, along with the relatively low processing power of mobile devices, has encouraged smaller, stand-alone applications, and discouraged the development of external libraries, middleware, and virtualization techniques to bridge the differences between platforms. There is no sign that a single platform will dominate this fragmented market in near future. For true ubiquity, application developers need to write, test, support, and maintain versions of their applications for all of these platforms. In a sense, the complexity of the crowdsensing application space grows with cross product of the number of platforms and the number of applications.

This issue of heterogeneity is underscored by the experience of Balan et al. [2], who conducted one of the largest crowdsensing studies to date. It took them six months to deploy one version of their GPS-based crowdsensing application on 15,000 taxis in Singapore, mainly due to the heterogeneity of the on-car GPS devices provided by different vendors. Web-based applications implemented in HTML5 are sometimes put forward as a "write once, run everywhere" alternative. Unfortunately, the HTML5 sensor APIs available on mobile browsers are still quite limited and the support for different sensors varies from browser to browser and from platform to platform. While geolocation tracking using GPS is widely supported in the up-to-date mobile browsers, accessing the microphone and video cameras is not possible in most cases [17]. In practice, due to the potential incompatibility between browsers (e.g. the inconsistent support for audio and video codecs), developers still have to customize their code for different browsers in some cases.

The second obstacle is the burden today's crowdsensing applications place on users. Today, each user must install a separate proprietary application for every crowdsensed experiment in which s/he wishes to participate. As a result, the deployment of a single crowdsensing application is limited by the rate at which users adopt and install it on their devices. It can take weeks or months for a newly introduced application to reach the critcial mass of participants needed

for it to be useful. Rapid, large-scale deployment, as in the lost child usage scenario above, is impossible with an install-based deployment model. Users also have to be tolerant of the processing, memory, and battery life these applications consume. Because today's mobile operating systems are designed to shield applications from each other, each application is meant to be self-contained and does not share information with others. In addition, some sensors, such as cameras, need to be exclusively locked before use. Participating in more than one crowdsensing application at a time is therefore not easy, even if a user is positively inclined.

The third obstacle, which primarily affects future applications, is the increasing network bandwidth demands of emerging crowdsensing applications. Table 1 shows that the GPS data has been the most widely used sensing information in the existing crowdsensing applications. However, looking ahead, we envision growing use of data-rich, multimedia sensing information like video [1] in emerging applications such as augmented reality, or the video-based lost child locator discussed above. These applications not only demand far more computing power, but also far more network bandwidth to send data to the cloud infrastructure. Based on data rate analysis of 80 videos on YouTube captured from a first-person viewpoint, each participant in a video-based crowdsensing application will upload between 0.6 Mbps (360p resolution) and 5.6 Mbps (1080p resolution). With many users, such an application can easily overwhelm link capacity in regional networks and into datacenters. For example, Verizon recently introduced state-of-the-art 100 Gbps links in their metro networks [21], yet these are only capable of supporting 1080p streams from just 18000 users. A broadly-deployed application with 1 million users will require 1–2 Tbps, 200x the total upload bandwidth of all YouTube contributors [33] today. An application model where each device sends data to centralized servers (as is typical today) cannot scale to support data-rich sensors. Ensuring the scalability of crowdsensing with data-rich sensors requires rethinking application and cloud architectures to acquire, process, and aggregate such data efficiently.

Ultimately, all three of these obstacles are ramifications of the deployment model in vogue today, where participation in each crowdsensing activity requires a separate application
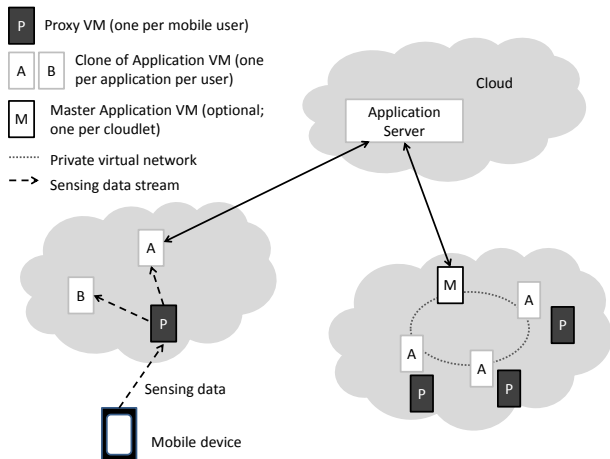
Figure 1: System Architecture.

that must be installed and run on user devices, and directly communicates to central servers. To overcome these obstacles, we must rethink the structure and deployment model used in crowdsensing applications.

3. PROPOSED SOLUTION

We propose a crowdsensing deployment model built around 3 core design principles:

- separation of data collection and sharing from application-specific logic.

- removal of application installation on smartphones from the critical path of application deployment.

- decentralization of processing, and data aggregation near the source of data.

These design principles address the obstacles discussed above. By construction, our proposed solution overcomes the key barriers to scaling up crowdsensing applications.

3.1 System Architecture

The 3-tier system architecture of our deployment model is illustrated in Fig. 1. The first layer is composed of mobile devices, whose roles are essentially reduced to that of (multi-input) sensors forwarding captured data to *proxy VMs* in the second layer. The second layer comprises of distributed cloud infrastructure deployed close to the mobile devices, typically in the access or aggregation network of Wi-Fi or cellular network providers. The concept of distributed cloud infrastructure here is akin to the concept of *cloudlet* presented in [26]. In practice, this can be a private cloud owned by a business or community, or a small data center such as Myoonet's Micro data center [18] that is deployed by a cloud operator. For the sake of simplicity, we will refer to this distributed cloud infrastructure as cloudlets in the remainder of this paper.

Each proxy VM is associated with a single mobile device, and is kept physically close to the mobile device through VM migration to other cloudlets or public clouds. This ensures network resources to transfer data from the mobile device is minimized. The proxy VM handles all the requests for sensor data on behalf of the mobile device. On the mobile device,

a single application is responsible for collecting sensor data and communicating it to the proxy VM. This application can be either implemented as a native application, or -if a good mobile browser is available on the device- as a HTML5 web application. The proxy VM is essentially an extension of the mobile device into the cloudlet, and can perform custom data preprocessing, e.g., to enforce privacy settings or handle quirks of the mobile platform, and enforce user preferences on data sharing and crowd participation. From here, data is forwarded to one or more *application VMs* also running on the cloudlet infrastructure.

Application VMs perform data processing steps specific to each crowdsensing application. Each application VM hosts a single crowdsensing application, which is not customized to any particular mobile platform. Generally, for each crowdsensing activity, one application VM is assigned to each participant, making it easy to migrate a user's proxy VM together with the associated application VMs, preserving any hard state they may contain. If an application does not need to maintain hard state for each user, then a single application VM can be shared by all users on a particular cloudlet.

The application VMs for each sensing service are deployed by a coordinating entity on the highest layer in our architecture, typically by the application server running on the centralized cloud infrastructure. In practice, when many application VMs are run on each cloudlet, the application server can initiate a master application VM (MAVM) on each cloudlet and delegate management and data aggregation tasks. The MAVM will coordinate, clone, and configure the application VMs on the cloudlet, and aggregate data within the cloudlet before forwarding results. Depending on the application, the MAVMs on multiple cloudlets may form a peer-to-peer overlay network / tree to scalably aggregate data to the central application server.

Our deployment model is predicated on two assumptions. First, this architecture depends on distributed cloud infrastructure near the user. The vision of executing customized VMs on nearby infrastructure has been articulated many times, e.g. in [8] and [11]. It has also be argued that offloading to nearby computing infrastructure (cyber foraging) is needed for compute-intensive and latency-sensitive mobile applications [4] for the purpose of energy savings and latency reduction. Our concept of distributed cloud infrastructure to host proxy and application VMs fits perfectly in this vision.

Second, our approach assumes a standard API exists for the data transfer between the proxy VM and the associated application VMs. However, we argue this is a much easier task to accomplish than having to write an individual application for each mobile platform (and possibly for each individual version of the mobile platform). Indeed, the output of scalar sensors can be represented as a few integers (e.g. GPS coordinates, temperature value, ...), and standards for multimedia data (e.g., video formats) already exist. Combining such data with standardized XML format descriptions, one can establish a standard for communication between proxy VMs and application VMs. In fact, several programming frameworks for crowdsensing applications have proposed solutions to abstract sensing information [32, 24] and task description [22]. These programming frameworks can be leveraged in our model as well.
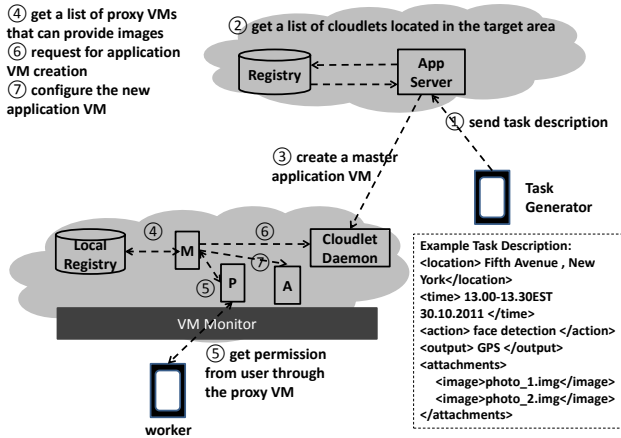
④ get a list of proxy VMs that can provide images
⑥ request for application VM creation
⑦ configure the new application VM

② get a list of cloudlets located in the target area

Registry   App Server

① send task description

③ create a master application VM

Task Generator

Local Registry   ④   M   ⑥   Cloudlet Daemon

⑤   P   ⑦   A

VM Monitor

⑤ get permission from user through the proxy VM

worker

Example Task Description:
<location> Fifth Avenue , New York</location>
<time> 13.00-13.30EST 30.10.2011 </time>
<action> face detection </action>
GPS
<attachments>
    <image>photo_1.img</image>
    <image>photo_2.img</image>
</attachments>

Figure 2: Workflow of crowd bootstrap.

3.2 Crowd Bootstrap

Let's revisit the lost child scenario from Section 1 to see how a crowdsensing task can be rapidly bootstrapped using our deployment model. As shown in Fig. 2, the process of crowd bootstrap can be summarized in the following seven steps.

1. The task generator (here, it is the parents' smartphone) constructs and sends a task description to the application server, typically located in the public cloud. The actual format of this can be application-specific, but is shown as an XML snippet here. The critical information includes type of search (face detection), the sample images, and a location area to scope the search. How this description is constructed and communicated is also left to the application, e.g., with a front-end app on the phone, or through a web-form on the application server.

2. The application server parses the task description, and consults a global registry for a list of cloudlets that are located in the target area.

3. The application server contacts the cloudlet daemon on each target-area cloudlet, and requests a MAVM instance be created. It forwards the VM disk image and memory snapshot to launch the MAVM. In practice, techniques employing demand paging or VM synthesis can minimize overheads of launching the MAVM.

4. The MAVM on each target cloudlet uses a cloudlet-local registry to discover proxy VMs connected to devices that can provide the desired sensor data (here, videos and images).

5. The MAVM requests participation from the mobile users through the proxy VMs. Depending on user-defined policies, the proxy may require explicit permission from the user, or the proxy VM can automatically join crowds on behalf of the user when particular criteria are met (e.g., share video when in a public space, but not audio).

6. Once permission is granted, the MAVM will request the cloudlet daemon to create application VMs. In

practice, these can simply be clones of the MAVM, operating in a different mode.

7. The MAVM configures the networking setup of the application VMs, while the the proxy VM will add the new application VM to the subscriber list.

When the above steps are finished, the proxy VMs will start forwarding images and videos to the application VMs, which will apply face detection and forward potential matches through the MAVM and application server to the parents' smartphone. We believe our architecture has the potential to bootstrap large crowds in just a matter of minutes, making this on-demand crowdsensing use case possible.

3.3 Benefits of Our Design

Our deployment model is architected to support scalable, efficient data sharing between multiple applications and users, while reducing the burden on application developers and end users. It scales up crowdsensing tasks by making it easier to access data from a larger pool of diverse smartphones, allow users to simultaneously particpate in multiple applications, and support rich, high-data-rate sensors at global scale.

Separating the process of data collection and sharing from application-specific processing, our system lets developers focus on the latter, rather than porting their application to a myriad of mobile platforms and understanding the idiosynchrasies of each. In fact, our deployment model increases the choices in programming languages, as the application is self-contained in its application VM and does not have to meet specific compatibility constraints for mobile platforms. Similarly, the developer is free to use a variety of programming models to distribute computation and aggregate results, and not forced to use a one-size-fits-all paradigm. Deploying VMs to users boils down to rapid cloning of the application VMs on cloudlets, regardless of the mobile devices.

In our architecture the personal data of the users is stored and processed on their own proxy VMs. According to [7] this approach provides a higher degree of privacy if compared to the traditional approach of storing and processing the data using centralized third party services. Our framework also allows flexibility in partitioning work between the proxy VM and mobile device. For example, supporting multiple applications with differing fidelity or resolution requirements simultaneously will entail some amount of preprocessing; this can be done in the proxy VM, mobile device, or a combination of both depending on hardware capabilities, processing overheads, and energy availability.

Our approach reduces the burden on users and their mobile devices for participating crowdsensing. First, instead of installing individual apps on their devices for each crowdsensing application, users only need to install one app that allows users to participate in different crowdsensing applications. Users can join a crowd by simply granting permission, and if willing, can direct their proxies to automatically participate in some forms of crowdsensing. When a user leaves a crowd, the application VM is simply destroyed, and does not require additional attention from the user. Second, demands on the mobile device can also be reduced, as processing is offloaded to the cloud, and only a single copy of the sensor data is uploaded even when participating in multiple applications. The potential reduction in data transmission helps save energy for users' mobile devices.

Lastly, our design performs processing and data aggregation close to the data sources. This brings two benefits: 1) it reduces traffic on wide-area networks; 2) it reduces network latency by avoiding long-distance data transmission through the backbone networks. This makes it possible to scale up crowdsensing with high-data-rate sensors. VM migration can ensure that processing remains close to data source even as users move around.

4. CHALLENGES

There are several technical hurdles in the path of a real-world deployment of our proposed architecture. We discuss these below.

4.1 Virtualization Overhead

Leveraging virtualization allows us to create a flexible platform in a multi-party setting where user privacy, scalability and isolation between crowdsensing applications are key requirements. These advantages come at the price of both VM creation overhead and the need for more advanced inter-VM communication management. In our design, a new clone of the application VM is instantiated for each user joining the crowd. Ideally, this new VM should start as fast as possible with minimal cost on resources. In practice, when a VM Monitor starts a new VM, it must first reserve all of the memory resources needed for the VM. This constraint prevents rapid creation of multiple VMs concurrently.

One way to solve this problem is to reduce the number of running VMs by replacing the per-user application VMs with one multiplexing application VM on each cloudlet. However, this will introduce the complexity of process migration in mobile scenario when any hard state contained in the application VMs must be preserved. An alternative way is to reduce the overhead of VM creation through advanced cloning mechanisms. There are several efforts that try to reduce the memory copy overhead by cloning the memory from running VMs. SnowFlock [13] proposes to fetch memory on demand while cloning VMs. It manages to clone 32 clones in 32 different hosts within one second by combining on-demand fetching with TCP multicasting for network scalability. Kaleidoscope [6] takes this one step further by discriminating VM memory state into semantically related regions to achieve prefetching and efficient transmitting.

An additional challenge is configuration and performance of inter-VM communication. The performance of inter-VM communication is relatively low compared to inter-process communication. When the system workload on the cloudlet increases, this may result in delayed transmission of sensor data between proxy and application VMs. Note that this low performance is due to inefficient CPU scheduling of the host, as the physical network interface is not touched by inter-VM traffic.

4.2 Migration-induced Reconfiguration

Physical mobility of a device may trigger the migration of the proxy VM and the associated application VMs that are not stateless. Consequently, the IP address of the mobile device as well as those of the VMs may change. To maintain established connections between mobile device and proxy VM, as well as between proxy VMs and application VMs, automated advanced network reconfiguration is needed. This potentially includes network addressing, NAT settings and firewall setup in VMs. Due to this overhead, IP-based so-

lutions may not provide adequate performance in our envisaged scenarios. Non IP-based solutions such as the Host Identify Protocol [19] have been designed from scratch with these limitations in mind, but these protocols still need to be evaluated in real networks. The deployment of these is unlikely to be easy, given the fact that today's Internet is built almost exclusively on the TCP/IP stack.

4.3 Standardization of Sensing Interfaces

Sensor data is distributed from the proxy VMs to the application VMs through a publish-subscribe mechanism. Standard sensor data descriptions are needed to realize communication between proxy VMs and application VMs of various developers. As discussed in Section 3.1, some efforts [32, 24] have been invested on developing such interfaces, however, unfortunately so far no consensus has been made yet.

Another challenge lies in the fact that different crowdsensing applications might be built on the same sensor data, but require a different format or sample rate. However, the sensor data collected from the devices provided by different vendors may not be able to always provide the data in the right format or at the right sample rate.

There is a trade-off to be studied on whether the conversion from the original sensor data to the requested output format(s) must be done on the mobile device, the proxy VM or inside the application VM itself. At first sight, running inside the application VM is the most logical choice, as it removes as much logic as possible from the mobile device and the proxy VM. However, this results in a lack of synchronization and a potential waste of resources. For example, what if all currently running application VMs only need camera frames at 10 fps, while the mobile device emits at a standard 30 fps? In this case, it would make sense to put downsampling application logic on the mobile device, and to put logic in the proxy VM that can configure the sensor capturing on the mobile device. When a new application VM is deployed needing 15 fps, the proxy VM may instruct the mobile device to increase its frame rate accordingly. Support for device-level configuration may vary significantly by platform and specific sensor hardware, so proxy VMs need to be designed to abstract away such differences.

5. CONCLUSIONS

This paper has argued that the existing deployment model for crowdsensing applications does not support either efficient crowd scaling over heterogeneous mobile platforms or the data sharing between crowdsensing applications. While VM-based cloudlets have been widely studied and utilized for computation offloading, we explore the potential uses of VM-based cloudlets for lowering the barriers to scaling up crowdsensing applications. Our solution leverages the existing programming frameworks for crowdsensing applications. There are still several challenges that must be addressed before this kind of deployment model can be adopted, we are implementing the deployment platform with specific focus on the research challenges discussed in this paper.

6. ACKNOWLEDGMENTS

7. REFERENCES

[1] BAHL, P., PHILIPOSE, M., AND ZHONG, L. Vision: cloud-powered sight for all: showing the cloud what you see. In *Proc. of MCS* (2012).

[2] BALAN, R. K., NGUYEN, K. X., AND JIANG, L. Real-time trip information service for a large taxi fleet. In *Proc. of MobiSys* (2011).

[3] BAO, X., AND ROY CHOUDHURY, R. Movi: mobile phone based video highlights via collaborative sensing. In *Proc. of MobiSys* (2010).

[4] BONOMI, F., MILITO, R., ZHU, J., AND ADDEPALLI, S. Fog computing and its role in the internet of things. In *Proc of MCC* (2012), pp. 13–16.

[5] BOZZON, A., BRAMBILLA, M., AND CERI, S. Answering search queries with crowdsearcher. In *Proc. of WWW* (2012).

[6] BRYANT, R., TUMANOV, A., IRZAK, O., SCANNELL, A., JOSHI, K., HILTUNEN, M., LAGAR-CAVILLA, A., AND DE LARA, E. Kaleidoscope: cloud micro-elasticity via vm state coloring. In *Proc. of EuroSys* (2011).

[7] CÁCERES, R., COX, L., LIM, H., SHAKIMOV, A., AND VARSHAVSKY, A. Virtual individual servers as privacy-preserving proxies for mobile devices. In *Proc of MobiHeld* (2009), pp. 37–42.

[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 *Proc. of MobiSys* (2010).

[9] FUJISAKA, T., LEE, R., AND SUMIYA, K. Discovery of user behavior patterns from geo-tagged micro-blogs. In *Proc. of ICUIMC* (2010).

[10] GANTI, R. K., YE, F., AND LEI, H. Mobile crowdsensing: current state and future challenges. *IEEE Communications Magazine 49*, 11 (2011).

[11] KOSTA, S., AUCINAS, A., HUI, P., MORTIER, R., AND ZHANG, X. Thinkair: Dynamic resource allocation and parallel execution in the cloud for mobile code offloading. In *Proc of INFOCOM* (2012).

[12] KOUKOUMIDIS, E., PEH, L.-S., AND MARTONOSI, M. R. Signalguru: leveraging mobile phones for collaborative traffic signal schedule advisory. In *Proc. of MobiSys* (2011).

[13] LAGAR-CAVILLA, H. A., WHITNEY, J. A., SCANNELL, A. M., PATCHIN, P., RUMBLE, S. M., DE LARA, E., BRUDNO, M., AND SATYANARAYANAN, M. Snowflock: rapid virtual machine cloning for cloud computing. In *Proc. of EuroSys* (2009).

[14] LANE, N. D., XU, Y., LU, H., HU, S., CHOUDHURY, T., CAMPBELL, A. T., AND ZHAO, F. Enabling large-scale human activity inference on smartphones using community similarity networks (csn). In *Proc. of UbiComp* (2011).

[15] LU, H., PAN, W., LANE, N. D., CHOUDHURY, T., AND CAMPBELL, A. T. Soundsense: scalable sound sensing for people-centric applications on mobile phones. In *Proc. of MobiSys* (2009).

[16] MATHUR, S., JIN, T., KASTURIRANGAN, N., CHANDRASEKARAN, J., XUE, W., GRUTESER, M., AND TRAPPE, W. Parknet: drive-by sensing of road-side parking statistics. In *Proc. of MobiSys* (2010).

[17] MOBILE HTML5. http://mobilehtml5.org, 2013.

[18] MYOONET. Unique scalable data centers. http://www.myoonet.com/unique.html, 2011.

[19] NIKANDER, P., GURTOV, A., AND HENDERSON, T. Host identity protocol (hip): Connectivity, mobility, multi-homing, security, and privacy over ipv4 and ipv6 networks. *Communications Surveys Tutorials, IEEE 12*, 2 (2010), 186 –204.

[20] NIU, Z., LI, S., AND POUSAEID, N. Road extraction using smart phones gps. In *Proc. of COM.geo* (2011).

[21] PC WORLD. http://www.pcworld.com/article/255519/verizon_to_offer_100g_links_resilient_mesh_on_optical_networks.html, 2012.

[22] RA, M.-R., LIU, B., LA PORTA, T. F., AND GOVINDAN, R. Medusa: a programming framework for crowd-sensing applications. In *Proc. of MobiSys* (2012).

[23] RABBI, M., ALI, S., CHOUDHURY, T., AND BERKE, E. Passive and in-situ assessment of mental and physical well-being using mobile sensors. In *Proc. of UbiComp* (2011).

[24] RAVINDRANATH, L., THIAGARAJAN, A., BALAKRISHNAN, H., AND MADDEN, S. Code in the air: simplifying sensing and coordination tasks on smartphones. In *Proc. of HotMobile* (2012).

[25] SATYANARAYANAN, M. Mobile computing: the next decade. In *Proc of MCS* (July 2010).

[26] SATYANARAYANAN, M., BAHL, P., CACERES, R., AND DAVIES, N. The case for vm-based cloudlets in mobile computing. *IEEE Pervasive Computing 8*, 4 (Oct. 2009).

[27] SHAH, S., BAO, F., LU, C.-T., AND CHEN, I.-R. Crowdsafe: crowd sourcing of crime incidents and safe routing on mobile devices. In *Proc. of ACM SIGSPATIAL GIS* (2011).

[28] TANG, A., AND BORING, S. #epicplay: crowd-sourcing sports video highlights. In *Proc. of CHI* (2012).

[29] WAKAMIYA, S., LEE, R., AND SUMIYA, K. Crowd-sourced urban life monitoring: urban area characterization based crowd behavioral patterns from twitter. In *Proc. of ICUIMC* (2012).

[30] WIRZ, M., STROHRMANN, C., PATSCHEIDER, R., HILTI, F., GAHR, B., HESS, F., ROGGEN, D., AND TRÖSTER, G. Real-time detection and recommendation of thermal spots by sensing collective behaviors in paragliding. In *Proc. of SCI* (2011).

[31] YAN, T., KUMAR, V., AND GANESAN, D. Crowdsearch: exploiting crowds for accurate real-time image search on mobile phones. In *Proc. of MobiSys* (2010).

[32] YE, F., GANTI, R., DIMAGHANI, R., GRUENEBERG, K., AND CALO, S. Meca: mobile edge capture and analysis middleware for social sensing applications. In *Proc. of WWW Companion* (2012).

[33] YOUTUBE. http://www.youtube.com/t/press_statistics, 2012.

[34] ZHENG, Y., ZHANG, L., XIE, X., AND MA, W.-Y. Mining interesting locations and travel sequences from gps trajectories. In *Proc. of WWW* (2009).

[35] ZHOU, P., ZHENG, Y., AND LI, M. How long to wait?: predicting bus arrival time with mobile phone based participatory sensing. In *Proc. of MobiSys* (2012).

[36] ZIMMERMAN, J., TOMASIC, A., GARROD, C., YOO, D., HIRUNCHAROENVATE, C., AZIZ, R., THIRUVENGADAM, N. R., HUANG, Y., AND STEINFELD, A. Field trial of tiramisu: crowd-sourcing bus arrival times to spur co-design. In *Proc. of CHI* (2011).