Virtual Individual Servers as Privacy-Preserving Proxies for Mobile Devices

Ramón Cáceres
AT&T Labs
Florham Park, NJ, USA
ramon@research.att.com

Landon Cox
Duke University
Durham, NC, USA
lpcox@cs.duke.edu

Harold Lim
Duke University
Durham, NC, USA
harold@cs.duke.edu

Amre Shakimov
Duke University
Durham, NC, USA
shan@cs.duke.edu

Alexander Varshavsky
AT&T Labs
Florham Park, NJ, USA
varshavsky@research.att.com

ABSTRACT

People increasingly generate content on their mobile devices and upload it to third-party services such as Facebook and Google Latitude for sharing and backup purposes. Although these services are convenient and useful, their use has important privacy implications due to their centralized nature and their acquisitions of rights to user-contributed content. This paper argues that people's interests would be be better served by uploading their data to a machine that they themselves own and control. We term these machines Virtual Individual Servers (VISs) because our preferred instantiation is a virtual machine running in a highly-available utility computing infrastructure. By using VISs, people can better protect their privacy because they retain ownership of their data and remain in control over the software and policies that determine what data is shared with whom. This paper also describes a range of applications of VIS proxies. It then presents our initial implementation and evaluation of one of these applications, a decentralized framework for mobile social services based on VISs. Our experience so far suggests that building such applications on top of the VIS concept is feasible and desirable.

Categories and Subject Descriptors

C.0 [**General**]: System Architectures; D.4.6 [**Operating Systems**]: Security and Protection—*Access Controls*; D.4.7 [**Operating Systems**]: Organization and Design—*Distributed Systems*

General Terms

Design, Experimentation, Performance, Security

Keywords

Cloud computing, location based services, mobile devices, online social networks, privacy, utility computing, virtual machines

MobiHeld'09, August 17, 2009, Barcelona, Spain.

1. INTRODUCTION

Individuals generate a growing variety of personal data on their mobile devices, and they commonly upload this data to services owned by third parties. For example, a person may upload his photos to Facebook [8] and his geographic locations to Google Latitude [10]. These services are useful and extremely popular, but their use raises important privacy issues. One, they concentrate data for many users under a single administrative domain. This centralization introduces the possibility of large-scale privacy breaches from intentional or unintentional data disclosures [21]. Two, their terms of service typically grant the providers rights to users' data. These rights often include a license to display and distribute all content contributed by users in any way the provider sees fit [8][12]. As these services become more entwined with people's lives and gain access to more of their personal data, the threat of privacy violations will grow.

In this paper, we argue that individuals would be better served by having their mobile devices upload this data to machines owned by the individuals themselves. We term these machines Virtual Individual Servers (VISs) because our preferred instantiation is a virtual machine running in a utility computing infrastructure such as Amazon Elastic Compute Cloud (EC2) [1]. We believe that individuals will adopt virtualized utility computing for many of the same reasons that enterprises are adopting it: VISs unburden users from maintaining their own high-availability systems without forcing them to give up control of their data, software, and policies. In contrast to many free web services, paid utility computing providers such as EC2 do not take on rights to the content that users place on these services [1].

VISs can protect their owners' privacy in significant ways. One, VISs are resistant to large-scale privacy breaches because each runs in its own administrative domain. Two, a VIS gives its owner thorough control over the software and policies that determine what data is shared with whom. A major reason people upload personal data to web services is to share it with others. A VIS can run a complete server software stack, including web, application, and database servers. It can thus be configured to share data with other people and machines as its owner sees fit.

In addition to improving privacy, a VIS can serve as a proxy that greatly enhances the capabilities of a mobile device. Because it resides on stationary infrastructure, a VIS has much higher availability and is much less subject to processing, storage, communication, and energy limitations. It can thus offload a great deal of work from the mobile device while presenting highly available and high-performance services to the rest of the world.

The mobile computing research community has long recognized the need for stationary proxies to augment the capabilities of mobile devices. However, the question of proxy ownership has been largely unexplored. In this paper we propose that a mobile device and its data-serving proxy should be owned by the same person. Common ownership allows them to place a higher degree of trust in each other than in systems owned by third parties. The device can thus upload personal data exclusively and opportunistically to its VIS. The VIS can then respond to all third-party requests for the data according to the owner's privacy preferences, thus preserving the limited resources of the device for the direct needs of its owner.

Many mobile applications can be built on the base VIS concept. One example is mobile social services, which take advantage of the physical proximity of devices to their owners to enable a wide range of new social interactions, for example finding out when friends are nearby. We have designed a decentralized framework in which VISs share location information about their owners through self-organizing overlay networks, one overlay per social group. Experimental results using our prototype implementation indicate that our approach is a viable alternative to centralized mobile social services.

In short, this paper takes the position that using Virtual Individual Servers as trusted and resource-rich proxies for mobile devices is superior to the prevailing practice of uploading personal data to third-party servers directly from mobile devices. Section 2 discusses the advantages and disadvantages of VIS proxies. Section 3 describes a range of applications of VIS proxies. Sections 4 and 5 present our initial implementation and evaluation of our mobile social services framework based on VISs. Section 6 outlines related work and Section 7 discusses future work.

2. PROS AND CONS OF VIS PROXIES

2.1 VISs v. Centralized Third-Party Services

The use of Virtual Individual Servers to store and share personal data has a number of advantages over using centralized third-party services. The main one is improved privacy. As mentioned earlier, VISs give their owners more control over how to share their personal data. In addition, they are less prone to large-scale privacy breaches because each VIS runs in its own administrative domain.

A second advantage is flexibility. A VIS owner is free to add or remove functionality as he sees fit. Since the VIS is a system-level virtual machine that runs a complete operating system environment, the owner can install arbitrary software packages and exploit their full set of configuration options to implement desired functionality. In contrast, the user of a third-party service is limited to the features and configuration options chosen by the provider.

A third advantage is long-term availability. The encapsulation and portability properties of virtual machines allow a VIS owner to make backups of the complete VIS image and resume that image elsewhere. For example, if the VIS is hosted at a utility computing provider, the VIS owner can make periodic copies of the image to a personal disk residing at his home. He is free to resume that image on any suitable hardware that runs a compatible virtual machine monitor, for example at a different utility computing provider or even on a home computer. In contrast, users of a third-party web service are dependent on that provider's continued existence. If the provider shuts down its operations abruptly, as has happened with numerous free web services, then not only the service but all user content stored there may become permanently unavailable.

A related benefit of a VIS is that its owner can ensure that his data is stored in open formats. The owner can thus export that data from one platform to another as desired, while a third-party service may use proprietary formats and decline to make an export facility available to its users.

A fourth advantage is cost scalability. The VIS approach distributes the costs of hosting user-contributed content among all users. In contrast, a centralized third-party service concentrates these costs on its provider, while the ability of many popular services to recoup these costs remain unproven. For example, free services that host videos and photos for millions of users are commonly believed to be so far unprofitable because of the aggregate bandwidth costs of providing such services [13].

There are also disadvantages of using VISs instead of free third-party web services. One is the need for users to manage their own virtual machines. Web services completely offload management responsibilities from their users, and this ease of use contributes to their popularity. To mitigate this disadvantage of VISs, we envision a market of software and services to help users manage their VISs. Such managed virtual machine services are beginning to appear, so far aimed at enterprises [17]. Another disadvantage is the need for users to pay for the computing resources used by their VISs. The free nature of popular web services contributes to their popularity, but creates incentives for providers to share user data in ways that may diminish user privacy. To recoup the costs of operating a VIS, individuals may be able to run advertisements on their own VISs without divulging private information to third parties. It is also important to note that VIS owners can amortize these costs across a variety of VIS applications, some of which are described in Section 3. We feel it is important to explore alternatives to the prevailing third-party services as popular awareness of privacy issues grows and the cost of computing drops.

2.2 VISs v. Serving Data from Mobile Devices

The use of VISs to store and share personal information also has significant advantages over serving the information directly from mobile devices. Most of these advantages stem from the VISs residing on wired infrastructure instead of wireless devices. For example, VISs have higher availability because they do not suffer from frequent periods of disconnection due to being put to sleep in a pocket, or moving into an area without wireless connectivity. VISs also enjoy better network performance, both bandwidth and delay, due to their wired instead of wireless connectivity. They also have more processing and storage resources due their residing on server-class hardware instead of portable hardware. Finally, VISs do not have the energy constraints of devices running off batteries. These energy advantages not only give VISs higher availability, as mentioned above, but also reduce reluctance on the part of their owners to devote some of their resources to serving data to others.

We see one situation where VISs are at a disadvantage with respect to mobile devices: in locations where there is no access to the wired infrastructure. In that situation, mobile devices can still offer service to other mobile devices in the immediate vicinity even if the devices have lost connectivity to the wired infrastructure, and therefore lost connectivity to their VISs. For example, the devices may share information via an ad hoc network of local wireless links. However, the use of VISs does not preclude devices from operating in this ad hoc manner when necessary.

Overall, we feel that the advantages of VISs outweigh their disadvantages. We therefore believe that it is worthwhile to explore their use further.

3. APPLICATIONS OF VIS PROXIES

This section describes a range of applications of VISs as privacy-preserving proxies for mobile devices. It focuses on two: mobile social services and participatory sensing.

3.1 Mobile Social Services

Mobile social services use a continuous stream of location information from participants to coordinate social interactions such as notifying users when people of interest are nearby, delivering location-bound virtual sticky notes, and forwarding location-scoped queries to live mobile users. Unfortunately, all services of which we are aware suffer from the same drawback of concentrating users' sensitive location information under a single administrative domain. We believe that VIS proxies can enable alternative, distributed approaches for building mobile social services that provide stronger privacy protections than the dominant centralized architecture.

Under such a scheme, each administrative domain consists of a set of mutually trusting VISs and mobile devices. Secure communication among VISs and devices under the same domain is built on shared cryptographic state that is distributed out of band. Devices are only responsible for uploading their location to their proxy as resources allow, and all inter-domain interactions occur between proxies. Placing the burden of storing and serving location information on VISs rather than mobile devices has two advantages. First, it saves devices' energy, storage, and compute resources. Second, serving location information from VISs avoids the complexities of building a highly-available service from intermittently-connected mobile devices and desktop PCs.

We can apply key-establishment techniques such as LoKey [16] to provide mutual authentication between domains. LoKey distributes shared keys between a pair of users by taking advantage of two features that are inherent to mobile social services: 1) access to the closed Short Message Service (SMS) network via users' mobile phones, and 2) a pre-established social connection such as one user's knowledge of the other's mobile phone number. Keys can be initially established between two user's mobile phones over SMS, and then relayed to their proxy VISs.

This approach to establishing secure communication is attractive for two reasons. First, it avoids the Sybil attack by binding identities to mobile phone numbers. Mobile phone numbers are expensive to acquire, which severely limits how much of the identity space an attacker can control. Second, reducing the problem of establishing cryptographic state out of band to distributing mobile phone numbers takes advantage of existing practices among friends and colleagues. Within a social network, users already commonly share their phone numbers, and even if a user's number is not present in a phone's address book, it may be accessible via an online social networking website like Facebook or LinkedIn.

In the absence of widespread collusion, partitioning responsibility for location information among many domains reduces the likelihood of a large-scale privacy breach. However, proxies must also self-organize such that mobile social applications can efficiently access location information from multiple domains.

We have begun exploring this question through location-based extensions to Vis-à-Vis [19], our decentralized framework for online social networking. Vis-à-Vis organizes a social network into groups of users with similar attributes and interests, plus a well-known meta group used to advertise other groups. Groups of VISs are organized into distributed hash tables to provide efficient, fault-tolerant lookup and routing. We support location-based operations by adding skip graphs [2], which are well suited to storing location data because they support range queries of a key space. Figure 1 shows three location-enabled groups in Vis-à-Vis.

Locations within the skip graph are represented using Z-order space-filling curves. These curves map two-dimensional coordinates to a one-dimensional space of *Z-values*, while preserving the geographic locality of nearby points. Z-values also support arbitrary granularities so that participants can represent their locations as, for example, a street address, a neighborhood, or a city. These features allow applications to query a social group for the identities of users within a geographic range, while allowing users to publish different granularities of location information to different groups for privacy reasons. A VIS can also lie about its owner's location on his behalf. More detail on our VIS-based framework for mobile social services is available here: [20].

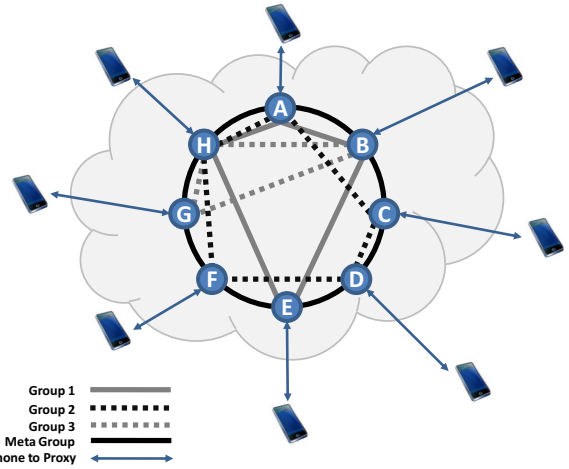

Figure 1: Example Vis-à-Vis network with eight mobile phones, their corresponding Virtual Individual Server proxies, and three social groups.

3.2 Participatory Sensing

The combination of widespread mobile phone adoption, improvements in mobile phone hardware, and the advent of flat-fee mobile data plans have enabled participatory sensing [15], a new sensing paradigm in which users' mobile phones collect and contribute data from the sensors available on the phones (e.g., GPS, accelerometer, camera, microphone). The data collected from a large population of users can then be used for a variety of applications, including understanding environmental impact [15] and urban planning [5].

Involvement in a participatory sensing project generally requires installation on a mobile device of a custom client application that collects sensor data and transfers it to a data collection server managed by a third party. There are two concerns with this mode of user involvement. First, a user has little control over the raw data that is being collected by the third-party applications running on his device. Second, the number of applications that need to run on a mobile device grows with the number of sensing projects a user participates in. This load may have a significant negative effect on the battery life and other performance of the user's device.

To mitigate these concerns, we propose a new scheme for involvement in participatory sensing projects. In our scheme, a mobile phone uploads sensitive raw data only once and only to its VIS, regardless of the number of projects the user is involved in. The VIS then applies the user's privacy preferences on the raw data and interacts with the participatory sensing servers according to the user's specifications.

3.3 Other Applications of VISs

There are many other possible applications of Virtual Individual Servers, not all of them exclusively related to mobile devices. Here we briefly outline five:

Personal Location Server: A number of recent papers dealing with location privacy assume the presence of trusted proxies between mobile devices and third-party location-based services [9, 11, 14]. VISs can fill that role.

Personal Backup and Synchronization Server: A VIS can be used for privacy-preserving backup and synchronization of personal information (e.g., calendar entries, address book entries, to-do lists) that is generated and consumed on both mobile devices and personal computers.

Personal Web Server: A VIS can play the role of a general-purpose but privacy-preserving web server, for example for sharing blogs, photos, and videos generated and consumed on both mobile devices and personal computers.

Personal Email Server: A VIS can also play the role of a privacy-preserving email server to clients running on both mobile devices and personal computers, thus freeing people from their current reliance on third-party email services.

Incoming Connection Manager: A VIS can serve as an incoming connection manager for its associated mobile device. In this role, the VIS aggregates notifications and updates from third parties and propagates them to the mobile device only when it is convenient and energy-efficient for the device.

4. INITIAL IMPLEMENTATION

To show the viability of the VIS concept, we built a prototype of the privacy-preserving mobile social services architecture described in Section 3.1. Recall that this Vis-à-Vis architecture organizes VISs into peer-to-peer overlay networks corresponding to social groups.

Our Vis-à-Vis prototype uses Pastry [18] to provide basic distributed hash table (DHT) functionality. It also uses Scribe [6] to provide multicast functionality on top of Pastry DHTs, but only in groups whose configuration options require multicast. Similarly, location-based functionality is simply another configuration option when creating a Vis-à-Vis group. Communication within a location-based group uses Pastry, with the addition of skip-graph data structures to implement location-based operations.

Each VIS runs an Apache Tomcat server in addition to the core DHT-based software. We deployed Java Server Pages (JSPs) and Servlets in the Tomcat server to implement external interfaces and their underlying logic. For example, the JSPs present web forms that a person can use to create, join, and leave a group, as well as advertise and search for information.

When deploying a VIS we encapsulate all the above software along with the requisite Java JDK inside a self-contained virtual machine image. We have successfully run VISs in a variety of virtualized computing environments, including Amazon EC2, Emulab, PlaceLab, a cluster of virtualized machines at AT&T Labs, and an experimental utility computing facility at Duke University.

We also created a mobile application for users to interact with our location-based groups. Figure 2 shows a screen sample from this application running on a Nokia N95 phone. Using the Google Maps API, the application allows a user to define a rectangular region on a map. These regions are used for two purposes: to set the precision of location updates sent to the mobile device's VIS, and to define the range of search queries sent to the VIS.

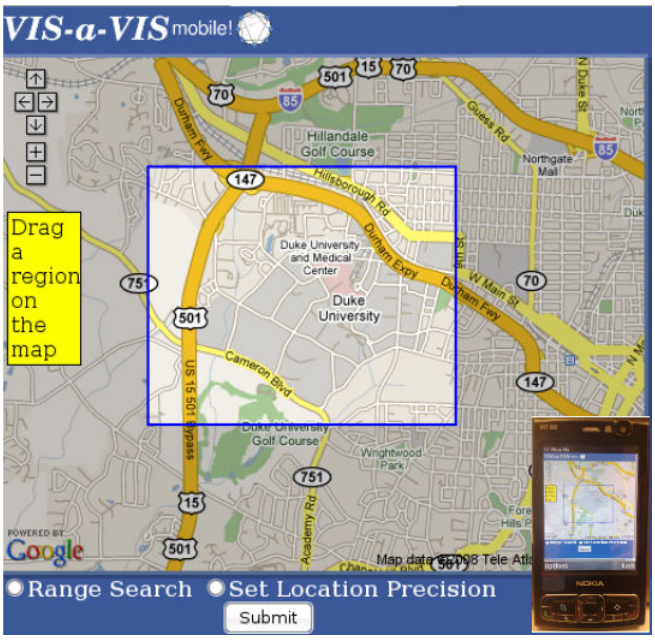

Figure 2: Mobile app running on a Nokia N95 phone.

5. INITIAL EVALUATION

We deployed our location-based Vis-à-Vis prototype on Emulab [23], a testbed that provides resources for experiments on distributed systems. In Emulab all resources reside at the same geographic site and users can request a set of virtual machines on which to run experiments.

Among other aspects of the system, we have explored how the size of a social group impacts performance. To capture the effects of group size, we evaluated our prototype's primitive operations, specifically the location update and range-search operations. Many other operations are a combination of these two. In this section we present experimental results for location updates.

To measure the latency of the location-update operation, we ran our experiment 25 times on a group of size n, where n is 20, 40, 60, 80 and 100. We first created a group of size n from n randomly chosen nodes, and then measured the latency for the $n + 1^{th}$ (randomly chosen) node to successfully insert itself to the skip graph. One such insertion is needed for every location-update operation. These group sizes are meaningful because our prior work on characterizing groups within Facebook found an average group size of close to 250 [19].

Figure 3(a) shows the results of our experiment. The mean latency of the insert operation grows slowly between 20 and 100 nodes. This result confirms our choice of distributed hash tables and skip graphs as our core data structures.

In order to characterize the expected number of location updates our system needs to handle, we investigated the amount of movement a user exhibits in a typical day. To do so, we provided each of 10 students living in the Duke University Smart Home with a Nokia N95 mobile phone, and recorded their location throughout the day for two weeks [20].

To avoid unnecessary location updates, our systems sends a location update to the server only when the user's location changes by a threshold number of meters from the previously reported location. Figure 3(b) plots the cumulative distribution function of the number of location updates per day when the threshold is set to

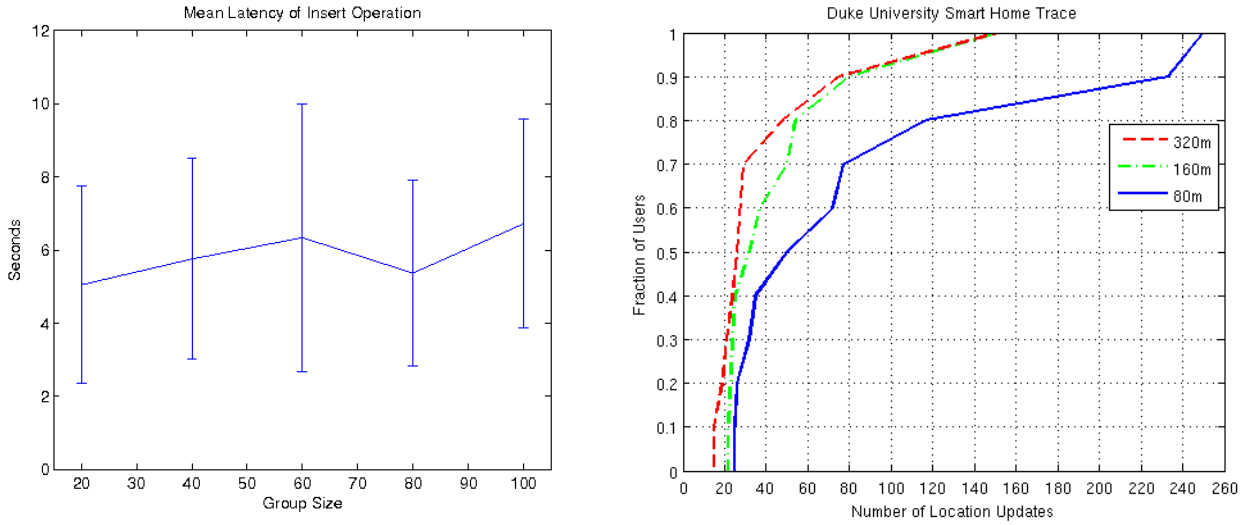

(a) Insert operation latency

(b) Human mobility characterization

Figure 3: (a) Mean latency of inserting a node with a given location into a group's distributed skip graph. Error bars show the standard deviation. (b) Number of location updates per day for different location granularities, based on a human mobility trace.

80m, 160m and 320m. The results show that the expected number of location updates is relatively small. For instance, when the movement threshold is set to 80m, the median number of updates per day is less than 50. This suggests that a 6-second latency for a location update operation should not significantly affect the user experience.

6. RELATED WORK

In this section, we first contrast the concept of a VIS proxy with a mobile personal server and with cyber foraging. Then, we contrast our Vis-à-Vis framework for mobile social services with other recently proposed decentralized architectures for online social networking.

The Personal Server [22] is a mobile device that carries a user's data and acts as a server to give the user access to her data through I/O and computing resources found in the currently local environment. In contrast, a VIS is a stationary proxy for any mobile device that a user might choose to carry with her. The VIS lives in the cloud and can serve data regardless of the availability of its associated mobile devices.

Cyber foraging refers to opportunistically using computing resources in the currently local environment to help with computations that would otherwise be carried out on a mobile device [3]. In contrast, a VIS is a privacy-preserving platform that lives somewhere in the cloud and is owned by the same person as the mobile device.

Other recent work [4, 7] has also proposed decentralized architectures for online social networking to improve privacy. Our Vis-à-Vis approach has availability, manageability, and scalability advantages due to our novel use of personal virtual machines running in a professionally managed utility computing platform. In addition, online social networking is only one of many possible applications of VISs.

7. CONCLUSIONS AND FUTURE WORK

In this paper, we have proposed using Virtual Individual Servers as privacy-preserving proxies for mobile devices. We discussed advantages and disadvantages of this approach and argued that the advantages outweigh the disadvantages. We also described a number of applications that can benefit from this approach, and presented our initial implementation and evaluation of a decentralized framework for mobile social services based on VIS proxies. Our experience so far suggests that building such applications on top of the VIS concept is feasible and desirable.

In the future, we plan to investigate ways for users to specify the privacy preferences and policies that will be enforced by their VISs. We hope to find a balance between usability and expressiveness. On the one hand, giving too many options to users can prove counterproductive. On the other hand, overly restricting the options may inhibit users from specifying their needs. In addition, we plan to further explore the scalability of our decentralized approach to mobile social services. Using simulation, we hope to show that our Vis-à-Vis architecture works for very large groups such as a community of interest around a geographic region or an item of popular culture. Finally, we plan to directly compare the performance and usability of decentralized and centralized approaches to location-based services.

8. REFERENCES

[1] Amazon Elastic Compute Cloud (EC2). `http://aws.amazon.com/ec2/`.

[2] J. Aspnes and G. Shah. Skip graphs. *ACM Transactions on Algorithms, 3(4)*. November 2007.

[3] R. Balan, J. Flinn, M. Satyanarayanan, S. Sinnamohideen, and H. Yang. The case for cyber foraging. *Proc. of 10th ACM SIGOPS European Workshop*. 2002.

[4] S. Buchegger and A. Datta. A case for P2P infrastructure for social networks: Opportunities and challenges. *Proc. of 6th International Conference on Wireless On-demand Network Systems and Services (WONS)*. February 2009.

[5] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava. Participatory sensing. *Proc. of the 1st Workshop on World-Sensor-Web: Mobile Device Centric Sensory Networks and Applications (WSW)*. 2006.

[6] M. Castro, P. Druschel, A.-M. Kermarrec, and A. Rowstron. Scalable application-level anycast for highly dynamic groups. In *Proc. of NGC*. 2003.

[7] L. A. Cutillo, R. Molva, and T. Strufe. Privacy-preserving social networking through decentralization. *Proc. of 6th International Conference on Wireless On-demand Network Systems and Services (WONS)*. February 2009.

[8] Facebook. `http://www.facebook.com`.

[9] B. Gedik, L. Liu, and G. Tech. Location privacy in mobile systems: A personalized anonymization model. *Proc. of 25th International Conference on Distributed Computing Systems (ICDCS)*. 2005.

[10] Google Latitude. `http://www.google.com/latitude`.

[11] B. Hoh, M. Gruteser, H. Xiong, and A. Alrabady. Preserving privacy in GPS traces via uncertainty-aware path cloaking. *Proc. of 14th International Conference on Computer and Communications Security (CCS)*. 2007.

[12] LinkedIn. `http://www.linkedin.com`.

[13] F. Manjoo. Do you think bandwidth grows on trees? *Slate*. April 2009.

[14] J. T. Meyerowitz and R. R. Choudhury. Realtime location privacy via mobility prediction: Creating confusion at crossroads. *Proc. of 10th Workshop on Mobile Computing Systems and Applications (HotMobile)*. 2009.

[15] M. Mun, S. Reddy, K. Shilton, N. Yau, P. Boda, J. Burke, D. Estrin, M. Hansen, E. Howard, and R. West. PEIR: the personal environmental impact report as a platform for participatory sensing systems research. *Proc. of 7th International Conference on Mobile Systems, Applications, and Services (MobiSys)*. 2009.

[16] A. J. Nicholson, I. Smith, J. Hughes, and B. D. Noble. LoKey: Leveraging the SMS network in decentralized, end-to-end trust establishment. *Proc. of Pervasive*. 2006.

[17] pair Virtual Private Servers (pairVPS). `http://www.pair.com/services/vps/`.

[18] A. I. T. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. *Proc. of Middleware*. 2001.

[19] A. Shakimov, H. Lim, L. P. Cox, and R. Cáceres. Vis-à-Vis: Online social networking via virtual individual servers. Duke University Technical Report TR-2008-05. October 2008.

[20] A. Shakimov, H. Lim, L. P. Cox, and R. Cáceres. Mobile social services via virtual individual servers. Duke University Technical Report TR-2009-01. January 2009.

[21] L. Story and B. Stone. Facebook retreats on online tracking. The New York Times. November 2007.

[22] R. Want, T. Pering, G. Danneels, M. Kumar, M. Sundar, and J. Light. The personal server: Changing the way we think about ubiquitous computing. *Proc. of 4th International Conference on Ubiquitous Computing (Ubicomp)*. 2002.

[23] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb, and A. Joglekar. An integrated experimental environment for distributed systems and networks. *Proc. of OSDI*. 2002.