Mobile Usage Patterns and Privacy Implications

Michael Mitchell, Ratnesh Patidar, Manik Saini, Parteek Singh, An-I Wang
Department of Computer Science
Florida State University
Tallahassee, FL, USA
{mitchell, patidar, saini, pasingh, awang}@cs.fsu.edu

Peter Reiher
Department of Computer Science
University of California, Los Angeles
Los Angeles, CA, USA
reiher@cs.ucla.edu

***Abstract*—Privacy is an important concern for mobile computing. Users might not understand the privacy implications of their actions and therefore not alter their behavior depending on where they move, when they do so, and who is in their surroundings. Since empirical data about the privacy behavior of users in mobile environments is limited, we conducted a survey study of ~600 users recruited from Florida State University and Craigslist. Major findings include: (1) People often exercise little caution preserving privacy in mobile computing environments; they perform similar computing tasks in public and private. (2) Privacy is orthogonal to trust; people tend to change their computing behavior more around people they know than strangers. (3) People underestimate the privacy threats of mobile apps, and comply with permission requests from apps more often than operating systems. (4) Users' understanding of privacy is different from that of the security community, suggesting opportunities for additional privacy studies.**

Keywords—mobile computing, security, privacy, human factors

I. INTRODUCTION

Privacy is perceived as a major concern for mobile computing due to a broad spectrum of threats, ranging from wireless eavesdropping to location tracking. While arguably we should provide mechanisms to allow the most privacy-conscious users to achieve the levels of privacy they desire, if these mechanisms incur too much cost for typical users, they are likely to be removed, disabled, or avoided. Thus, it is important to first understand users' needs and opinions of what privacy means to them, as well as what behavioral alterations we might expect from them as they move among different environments. This will ensure that the correct tools are created to protect what is important to the user. It will also help improve the acceptance rate of privacy enhancing tools, especially those that may provide protection at the expense of convenience or ease of use.

Unfortunately, we do not have a very clear sense of how users regard privacy issues in the mobile environment. The main reason is that the notion of privacy is subjective. People in the same objective circumstances may feel very differently about whether they are computing in private and whether the information being accessed should be treated as confidential. Although embedded mobile sensor technologies are constantly improving, it is not always feasible to automatically determine how people feel about their privacy status in a given situation, without being intrusive. Thus, prior privacy research based on automated mechanisms was largely confined to location tracking and sharing [9-13] (e.g., whether a user is willing to share location information [9], privacy implications of mobile ads [11], and trading privacy for useful services [12]). Although human subject studies can help us discern privacy situations that cannot be automatically determined, the process can be tedious due to the need to obtain IRB approvals and recruit human subjects. The relatively few studies that exist include: examining perceived risks of application permission requests [14], studies on specific locations [17], and assessing very limited social groups [18].

This paper presents a survey study based on a ~100-question questionnaire and provides a further step in obtaining a better picture of how today's users of mobile devices regard privacy issues. The results are based on ~600 users who were queried about their use of mobile devices, their attitudes toward privacy for different kinds of activities in various situations, and their awareness and understanding of existing tools to improve privacy. We present key elements of the results from the survey, focusing on points that shed light on which privacy scenarios different classes of real users consider important. We then suggest directions for building mobile privacy mechanisms and areas where more information from users would help determine how to provide them with the privacy they actually desire.

The primary goal of this survey study was to examine how mobile computing users feel about privacy. What does it mean to be private? Do people change their computing behavior when in the presence of other people? What types of people? Do people change their computing behavior in public? Does the perception of privacy differ by gender, age, ethnicity, device ownership, or technical background? How do we measure the perception of privacy based on differences between the numbers of hours spent using certain applications in public and private?

A secondary goal of this research was to better understand user behavior and general usage patterns—more specifically, to identify how, when, and where people use their mobile devices. What types of applications do people use the most? Does gender, ethnicity, age, income, choices of technology, or technical sophistication influence behavior?

II. EXPERIMENTAL METHODOLOGY

Subject recruitment: We initially recruited survey participants from the Florida State University (FSU) campus. We solicited participation flyers posted on campus and mass emails to university departments. Over 6 weeks (2/1/2013 –

3/15/2013), these efforts resulted in 292 student responses, nearly all (252) from the mass emailing.

We later decided to expand the survey to determine if our results could be generalized further. We solicited participants through the volunteer section of Craigslist of the ten most populated U.S. cities. Over 6 weeks (6/1/2013 – 7/15/2013), 303 responses were collected from this part of the survey.

We allocated ~$1,000 for prizes for participation. With the goal of sufficient motivation for participation, without excessive motivation to cheat, we decided to offer a chance to win one of 66 $15 Amazon.com gift cards.

Mobile usage questionnaire: Participants were asked to answer ~100 questions through a web interface [15]. The questionnaire started by asking about demographic information such as gender, ethnicity, expertise, device ownership, and background knowledge of privacy-enhancing tools such as encryption.

The questionnaire then asked about the frequency of performing 43 mobile activities in seven categories: entertainment (e.g., listen to music), communication (e.g., access emails), productivity (e.g., calendar), tools (e.g., reviews), financial (e.g., online banking), administration (e.g., configure network), and personal (e.g., health monitoring). The user answered whether an activity is performed hourly, daily, weekly, monthly, or never. To estimate the number of accesses per month during waking hours, we summarized the per-user frequency for a given activity within a month with the following formula: (#hourly*16*30) + (#daily*30) + (#weekly*(30/7)) + (#monthly). As a sanity check, a prior study showed that users on average access their mobile phone 150 times per day [8], and we achieved similar results.

For each activity, we also asked about the frequency of performing it either in a public setting (defined as with anyone else present) or a private setting (no one else present). These definitions ensure that participants used the same meanings of public/private settings to estimate the frequency of activities.

III. Demographics and Market Share

Survey demographics: Our subject pool reflects the general population where the surveys were conducted.

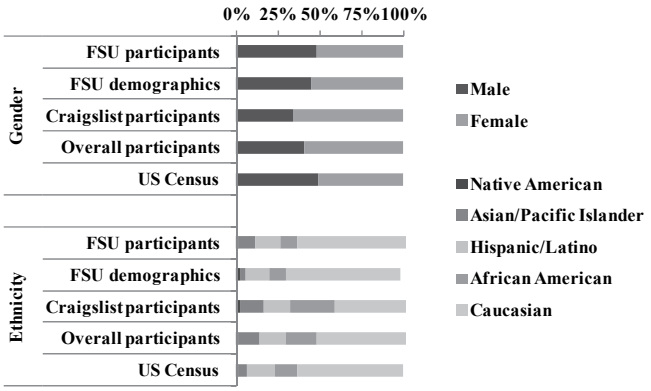

Fig. 1: Demographics comparisons for our survey participants [1, 2, 3, 4].

FSU survey participants: For the FSU survey, the 292 participants had a median age of 22, with an average of 6 years of computing experience. The gender split of the participants was within 3% of the FSU demographics, with slightly more male participation (Fig. 1).

Academic/education background was not quite as characteristic of FSU. Our survey had greater participation from computer science (CS)/engineering (by 28%) and undecided/other (by 9%), and correspondingly lower-than-expected participation by literature/language/social science (by 15%) and business-related students (by 8%). This may be due, in part, to the reasonably tech-savvy target audience of the survey, as well as greater access to survey recruitment e-mails and flyers for CS/engineering students.

The ethnicity of the participants largely reflected FSU demographics, except that there were more Asian/Pacific Islander participants (by 7%), possibly due to more participants with a CS/engineering background (Fig. 1).

Although these participants may not reflect the findings for the general population, mobile computing nevertheless has the deepest penetration among this age group [19], and this large user base's perception of privacy is worthy of study.

U.S. survey participants: For the U.S. Craigslist survey, the 303 participants had a median age of 27, also with 6 years of computing experience. The gender split was within 8% of U.S. demographics, with more female participation than expected (Fig. 1). The ethnicities of survey participants also reflected U.S. demographics. However, the U.S. Craigslist survey had higher minority participation rates (Fig. 1).

Although Craigslist has its own bias in terms of user demographics, these users represented a broader age spectrum. Interestingly, the findings for the FSU population are similar to the findings for the Craigslist population. Thus, unless noted, the results reflect the combined 595 responses.

Device market share: The smartphone ownership of our survey participants reflected the U.S. market share [5] (3% more iPhone owners and 2% fewer Android phone owners). In addition, tablet ownership of the survey participants reflected the U.S. market share of tablet ownership [6]. Participants in the survey had fewer Android tablets (by 7%) and more iPad/non-Android tablets (by 4%). The demographics of our participants' laptop ownership are less aligned with published data on market share [7]. There were significantly fewer Windows users (by 28%), and significantly more Apple (by 21%) and Linux users (by 7%).

Device ownership: We wanted to see if device ownership played a role in mobile privacy impressions, so we also split users into groups based on the brand of the mobile device they use. More often, men, tech-savvy users (defined as participants who work or major in computer science or related fields), and minorities own Android devices (by 10-20%). In addition, more often, men, tech-savvy users, and minorities own Windows laptops (by 9-19%). Tech-savvy users own Android phones, Android tablets, and laptops running Windows or Linux more frequently (by 9-36%). African Americans own iPhones less frequently (by 20%).

Brand homogeneity: Overall, participants tended to own multiple devices from the same manufacturer. iPhone owners more frequently own an Apple laptop or tablet (by 15-28%) compared to non-iPhone owners and Android owners more frequently own an Android tablet (by 15%). The Apple trend was more pronounced in the FSU data set where iPhone owners even more frequently own Apple laptops and tablets (by 15-40%).

IV. Privacy Results

All results are reported at 95% confidence intervals. Unless otherwise stated, the confidence intervals are within 9% of the mean.

Who makes us change our behavior? Participants were asked whose presence makes them change their computing behavior. Fig. 2 shows that people are most likely to change behavior around their parents, boss, friends, and significant others and least likely to change behavior around subordinates, foreign strangers, roommates, and the technically savvy. A significant number (>10%) do not care who is around and never change their behavior. Women and men behave similarly, but women alter their behavior more often when they are around their parents (by 11%). Tech-savvy users change their behavior more often around their roommates and others who they believe to also be tech-savvy (by 11%).

User hardware preference has the most significant influence over behavioral change around specific people: Apple laptop owners tend to change their behavior around their parents, significant others, friends, and siblings (by 9-16%) more than around other device owners. No significant differences in behavior changes were observed across ethnic groups.

One element of privacy concerns is whether one worries about consequences if certain information is given away. The consequences can be how people perceive you, how people with influence and authority can hold the information against you, etc. Trust of strangers may reflect that the availability of such privacy information to them would be inconsequential. This attitude may also reflect complacency toward the potential privacy threats that strangers can pose. More technically savvy people seem to be more aware of such threats when around their technically savvy peers.

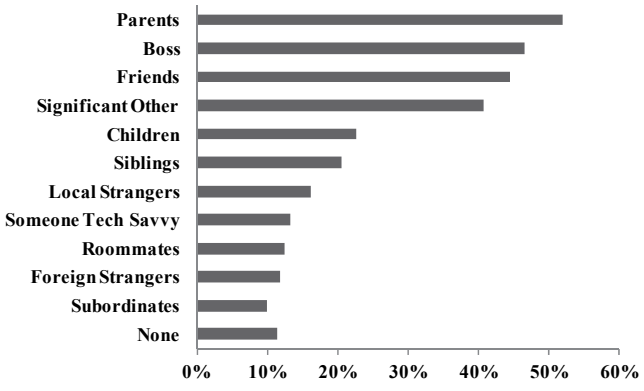

Fig. 2: Percentage of participants who change their computing behavior around certain people.

What do we do in public? In private? Fig. 3 shows the top ten most frequently performed activities in public (when anyone is present) and in private (otherwise). Texting, emailing, web browsing, social networking, and listening to music are the top five. Among them, people text equally often in public and in private, while people prefer to email, browse the web, social network, and listen to music in private (by up to 31%). Overall, people largely engage in the same kinds of activities in public and private environments.

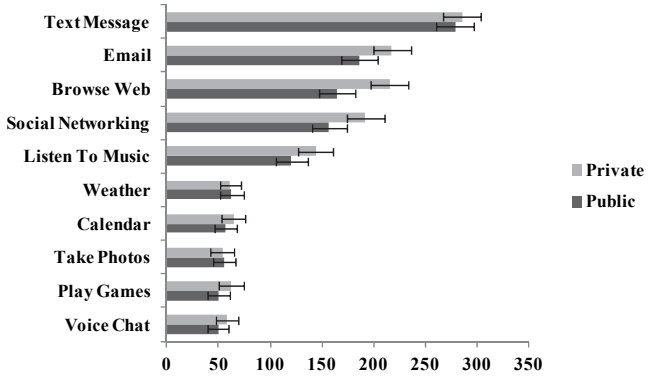

Fig. 3: Tasks most frequently performed in public and private in total number of accesses per month.

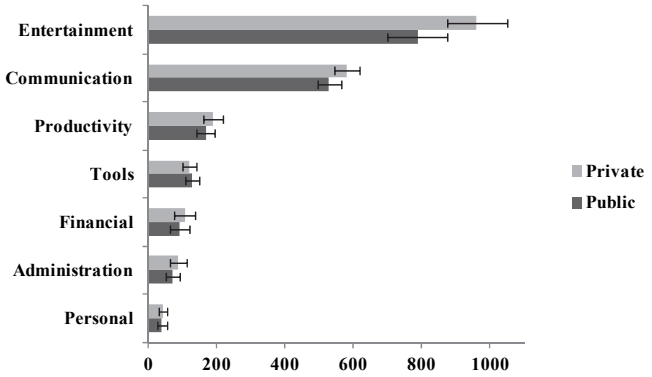

Fig. 4: Frequency of public and private mobile tasks, by task category in total number of accesses per month.

Fig. 4 shows a similar trend for categories of mobile tasks performed. The most commonly performed task categories are entertainment (44%) and communication (28%), with a significant margin over all other mobile tasks. Various task categories are performed more frequently in private (by up to 20%) than those performed in public, but people largely perform the same types of activities in both settings. However, the frequency does not tell the whole story; for example, entertainment tasks may present a low privacy risk, while financial tasks may have significant privacy implications.

Thus, we also classified mobile computing tasks based on the risk level of the exposed information. High-risk tasks involve information that, if exposed, could be used in identity theft [15], such as financial transactions or online banking. Low-risk tasks, such as playing games or watching videos, have low personal information exposure. The remaining tasks, such as web browsing and social networking, are classified as

medium risk, since they can involve risk ranging from low to high. Surprisingly, for high-risk tasks, people behave the same way in public and private (Fig 5), while people perform lower-risk tasks more frequently in private settings. Thus, privacy may also mean not wanting to be disturbed; performing high-risk tasks such as online shopping may benefit from consulting others in public settings.

Overall, the behavioral differences in public and private environments across genders, technical backgrounds, and ethnicities are not statistically significant, with a few exceptions. We found that women use social networking more frequently than men both in public (by 40%) and private (by 32%), but otherwise behave similarly to men. Tech-savvy users are more likely to access emails than less sophisticated users both in public (by 63%) and private (by 24%), but otherwise behave similarly as well. No significant behavioral differences were found among different ethnic groups.

What do we do when OSs and apps ask for permission? Users seem to trust and comply with their apps more often than with their operating systems (OSs). Fig. 6 shows that 38% of users always comply with OS permission requests, but 61% always comply with such requests from mobile apps. Since apps are written by parties of presumably variable trustworthiness and OSs are written by one well-known party, this result is counterintuitive, and raises interesting questions. Why are people 23% more likely to always agree with a permission request from an app than from their OS? Do the kinds and frequency of permission requests play a role? What makes the mobile app more trustworthy? With the access most apps have to personal information on a device, this behavior is risky. Perhaps people underestimate the power of apps and the privacy implications. If so, mechanisms to proactively protect user privacy are more important to add to mobile devices. The kinds and frequency of permission requests might also play a role.

No significant statistical differences for operating system and app compliance were observed across gender, ethnicity, or device ownership.

Usage of privacy-enhancing tools. Subjects were asked about their use of privacy-enhancing software tools, specifically encryption and password vaults. Fig. 7 shows that 44% of participants responded that they had used encryption, 31% never encrypt, and 25% were not sure.

Men were found 21% more likely to use encryption than women. Women were 25% more likely to be unsure about whether they use encryption (35% total). To see if this gender gap is caused by the pool of computer science and engineering students in our FSU sample, we also compared the usage patterns between men and women with and without technical backgrounds. Fig. 7 shows that the differences in awareness of encryption are even more pronounced for people without computer science or related backgrounds.

Minorities encrypt more often (by 12%) than Caucasians. In particular, Asians encrypt more often at 12% ±11% above the mean, and Caucasians are the least likely, at 22% below the mean. Tech-savvy people encrypt more often (by 31%) than less sophisticated users. Device ownership plays a minor role: Android users encrypt more often (by 7% ± 6%) than non-Android users, and iPhone and Apple laptop owners encrypt less often (by 7% ± 4%).

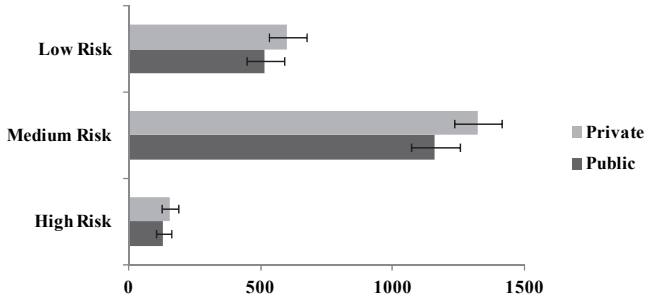

Fig. 5: Task frequency organized by risk level of information exposure in total number of accesses per month.

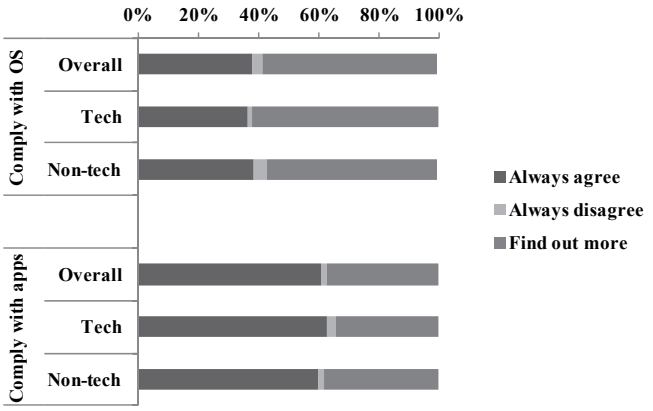

Fig. 6. Compliance levels for permission requests from operating systems and mobile apps for subjects with technical and non-technical backgrounds.

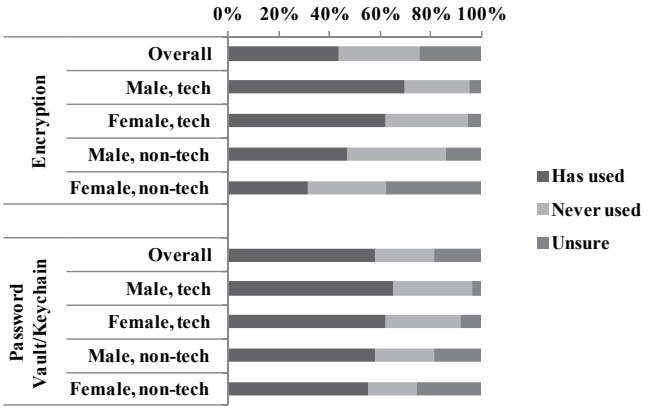

Fig. 7: Encryption and password vault/keychain usage patterns by gender and technical background level.

Password vaults/keychain usage trends were similar to encryption. Fig. 7 shows that 58% of users have used password vaults, and 23% never use them. Password vaults may be slightly less confusing than encryption, but still 19% were unsure about their use. Similar to encryption results, men were 10% more likely to use password vaults than women; women were more likely to be unsure by 12%. Non-technical users were 19% more likely to be unsure than technical users if

they use password vaults/key chains. Ethnic group and device ownership did not have a statistically significant influence on password vaults/key chain usage.

How do we connect to WiFi: WiFi use is nearly ubiquitous, but how privacy-aware are people when they connect? Subjects were asked about the types of networks they connect to using their mobile device. Eighty-one percent of the participants responded that they use public WiFi without security; only 10% required at least password protection to connect. Eight percent of respondents never use public WiFi of any kind. Tech-savvy users were 11% more likely to only use protected WiFi than non-tech users. Device ownership was also significant, iPhone owners were 10% more likely to use open WiFi networks than were non-owners. Little difference was observed for WiFi usage between genders and ethnic groups. These results suggest we should not rely on WiFi encryption to protect data in transit, since most users will happily use an unprotected network.

What kind of apps do we use? Questionnaire subjects were asked about their usage of different types of mobile apps. As shown in Fig. 8, 88% use market apps, 79% use web apps, and 37% use non-market apps. Men and women use market apps and web apps similarly, but men are more likely to use non-market apps (by 12%). For respondents with both technical and non-technical backgrounds, market app use and web app use are similar; however, technical users are more likely to use non-market apps (by 13%) than other users. In terms of ethnicity, minorities are more likely to use non-market apps (by 15%) than the ethnic majority. In particular, African Americans are 21% more likely to use non-market apps.

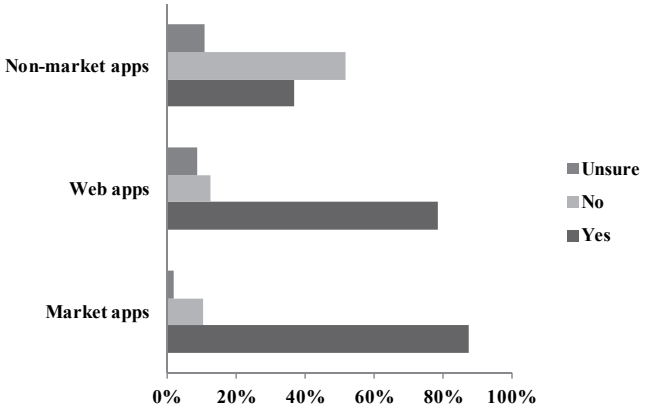

Fig. 8: Percentage of subjects who use market, web, and non-market apps.

V. Other Findings

Where do we compute? As shown in Fig. 9, the top computing locations are at home, in class, in the library, while waiting in line, and in restaurants. Locations least likely to be used for computing are parks, while exercising, and in the washroom. Men and women behave similarly for most locations, though women are more likely to compute at a park, restaurants, and while waiting in line (by 7-10%). Technical users more often compute in the classroom and at the office (by 14-18%). Non-technical users more often compute in restaurants, while exercising, and while waiting in line (by 9-14%).

Device ownership is perhaps the most interesting split. iPhone owners are more likely to compute in restaurants, airports/bus/train stations, public transportation, while exercising, in class, and while waiting in line (by 13-16%). Android device owners are more likely to compute at a park, office, and washroom (by 10-15%). No significant differences in computing locations were found between ethnic groups.

In the FSU survey, the effects of device ownership are even more pronounced. iPhone owners are more likely to compute in restaurants, on a bus/train/airplane, at airports/bus stops/train stations, while exercising, and waiting in line (by 19-26%).

Implications of Apple ownership: Compared to Android owners, Apple users more frequently use their iPhones, iPads, and Apple laptops in public locations (by 13-16%). One plausible explanation is Apple's greater choices of apps. Another possibility is their use as a status symbol.

We also found that Apple device owners tend to use their devices for most social mobile computing tasks: texting, e-mailing, and social networking more than owners of other devices. Apple device owners are more likely to e-mail both in public (by 27%) and in private (by 19%), send text messages in public (by 19%) and in private (by 22%), and use social networking in public (by 63%) and in private (by 35%).

As previously discussed, Apple users use their device in more public places. Apple device owners also have less regard for WiFi security. Eighty-six percent of iPhone owners use open, public WiFi without security, 6% above average. And Apple device users are less likely to use encryption (by 7%).

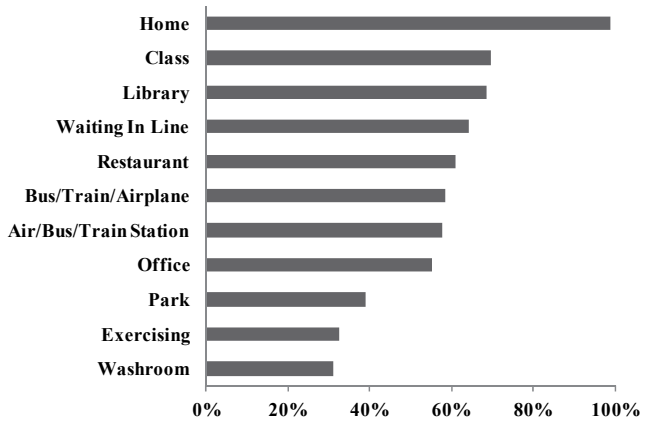

Fig. 9: Percentage of subjects who regularly compute in each location.

VI. Lessons from this Survey

This survey speaks to user attitudes towards privacy, not necessarily actual behavior. However, experience shows that user attitudes are critical in determining whether a privacy or security measure is widely used; thus, in some senses such attitudes are just as important to a privacy mechanism's success as the technical details of how it works. Certainly designers of mobile computing privacy mechanisms should keep this point in mind when going about their work.

Implications of privacy on systems: One finding is that mobile users are far more concerned about protecting their privacy from those who know them well. The presence of parents, for example, seemed to be viewed as a threat to users' privacy more than twice as often as strangers. Users may be more likely to accept privacy preserving mechanisms designed to protect against family and friends than against random eavesdropping by strangers.

Privacy, trust, and anonymity: The results suggest that trust and privacy are largely orthogonal. People we trust the most are also the ones whose loss of trust we most fear. This result is consistent with studies of teenagers' behavior on social networking sites [20], but is not necessarily what many academic researchers think about when they address how to achieve privacy in mobile computing systems. Since we should provide users with the privacy they want and will actually use, researchers must ensure that their goals align with users' real privacy desires.

Users' relative indifference to privacy threats posed by strangers might suggest a perception of anonymity, leading to a false sense of security. Frequently, computer users have been behind the curve on what malicious parties can do with information they obtain, and this may be another such case. The perception of anonymity is probably true only if strangers are not interested in us. If they are, and we do not protect ourselves, we may suffer serious consequences. Researchers and developers are clearly interested in protecting users from such threats, but our results suggest that, in today's world, only transparent and simple protection mechanisms against such threats will succeed, since users may be unwilling to take actions that seem inconvenient.

Privacy vs. solitude: Privacy may also mean the desire to be alone (e.g., when playing games, watching videos). Thus, future studies may need to distinguish between achieving privacy for security versus for user solitude.

Mobile apps' privacy implications underestimated? Why do the respondents trust applications more than OSs? Clearly, the consequences of mistakenly permitting an OS to take action may cause greater harm. However, it appears users are unaware of how much information a modern mobile app can access and of the resulting potential privacy implications.

Overall: Various aspects of the survey, such as the places people feel comfortable computing, the kinds of applications used in public, and the disregard for the presence of other people when using mobile devices, suggest that many users are not concerned about preserving the privacy of their computing in mobile environments. Yet such risks are real. The obvious question is whether the majority of users are unaware of the risks, or are reasonably aware and simply do not care about them. Unfortunately, our study did not include questions that allow us to provide insight on this point. This point is crucial. If users ultimately care little about preserving their privacy from such risks, only the cheapest, most transparent, least intrusive privacy enhancing mechanisms will succeed.

VII. ACKNOWLEDGMENTS

All interaction with human subjects was approved by the Florida State University IRB Human Subjects Committee, approval number 2013.10175. This work is sponsored by NSF CNS-1065127. Opinions, findings, and conclusions or recommendations expressed in this document do not necessarily reflect the views of the NSF, FSU, UCLA, or the U.S. government.

REFERENCES

[1] U.S. Census Bureau,. "State & county quickfacts 2007," http://quickfacts.census.gov, 2014.

[2] Florida State University. "About Florida State University: student body," http://www.fsu.edu/about/students.html, 2014.

[3] Craigslist.org. "Craigslist: classifieds for jobs, apartments, personals, for sale, services, community, and events," http://craigslist.org, 2014.

[4] Alexa, Inc. "Craigslist.org site info," http://www.alexa.com/siteinfo/craigslist.org, 2014.

[5] Smith A. "Pew Internet Research," *Smartphone Ownership – 2013 Update*. http://www.pewinternet.org/files/old-media//Files/Reports/2013/PIP_Smartphone_adoption_2013_PDF.pdf, 2014.

[6] King P. Strategy Analytics. "Android Dominates the Tablet Market in 2013," *Q2*. http://www.strategyanalytics.com/default.aspx?mod=pressreleaseviewer&a0=5403, 2014.

[7] NetMarketShare. "Desktop Operating System Market Share," http://www.netmarketshare.com/operating-system-market-share.aspx?qprid=10&qpcustomd=0, 2014.

[8] Sullivan B. *How the Smartphone Killed the Three-day Weekend*. http://www.cnbc.com/id/100765600, 2014.

[9] Lederer S, Mankoff J, Dey AK. "Who wants to know what when? privacy preference determinants in ubiquitous computing," In *CHI,* 2003.

[10] Wiese J, Kelley PG, Cranor LF, Dabbish L, Hong JI, Zimmerman J. "Are you close with me? are you nearby? Investigating social groups, closeness, and willingness to share," In *UbiComp*, 2011.

[11] Kelley PG, Benisch M, Cranor LF, Sadeh N. "When are users comfortable sharing locations with advertisers?" In *CHI*, 2011.

[12] Barkhuus L, Dey A. "Location-based services for mobile telephony: a study of users' privacy concerns," In *INTERACT*, 2003.

[13] Barkhuus L. "Privacy in location-based services, concern vs. coolness. In *Workshop on Location System Privacy and Control*, 2004.

[14] Felt AP, Egelman S, Wagner D. "I've got 99 problems, but vibration ain't one: a survey of smartphone users' concerns," *In SPSM*, 2012.

[15] Consolvo S, Jung J, Greenstein B, Powledge P, Maganis G, Avrahami D, "The wifi privacy ticker: improving awareness & control of personal information exposure on wifi," In *Ubicomp*, 2010.

[16] Google, Inc. "Google docs—online documents, spreadsheets, presentations, surveys, file storage and more," http://docs.google.com, 2014.

[17] Kientz JA, Choe EK, Truong KN. "Texting from the toilet: mobile computing use and acceptance in private and public restrooms. Knowledge Media Design Institute," University of Toronto. Technical Report KMD-13-1, 2013.

[18] Walton M, Donner J. "Public access, private mobile: the interplay of shared access and the mobile Internet for teenagers in Cape Town." Global Impact Study Research Report Series, 2012.

[19] Neilson.com. "Survey new U.S. smartphone growth by age and income," http://www.nielsen.com/us/en/newswire/2012/survey-new-u-s-smartphone-growth-by-age-and-income.html, 2014.

[20] Boyd D. "Why youth (heart) social network sites: The role of networked publics in teenage social life," MacArthur Foundation Series on Digital Learning–Youth, Identity, and Digital Media Volume, 2007.