

Lessons Learned from the NetSense Smartphone Study

Aaron Striegel, Shu Liu, Lei Meng,
Christian Poellabauer
Department of Computer Science and
Engineering
University of Notre Dame
striegel,slui6,lmeng,cpoellab@nd.edu

David Hachen, Omar Lizardo
Department of Sociology
University of Notre Dame
dhachen,olizardo@nd.edu

ABSTRACT

Over the past few years, smartphones have emerged as one of the most popular mechanisms for accessing content across the Internet driving considerable research to improve wireless performance. A key foundation for such research efforts is the proper understanding of user behavior. However, the gathering of live smartphone data at scale is often difficult and expensive. The focus of this paper is to explore the lessons learned from a two year study of two hundred smartphone users at the University of Notre Dame. In this paper, we offer commentary with regards to the entire process of the study covering aspects including funding considerations, technical architecture design, lessons learned, and recommendations for future efforts gathering live user data.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: [Wireless communication]

General Terms

Design, Human Factors, Measurement

Keywords

Smartphone, User Study, Wireless, Cellular Networks, WiFi

1. INTRODUCTION

Over the past few years, a dizzying array of wireless devices and technology have emerged that are fundamentally transforming how we as a society gather and react to information. Whether it is the hot feature on the latest smartphone, a debate on the size of the ‘best’ tablet, or predicting the latest ‘hot’ application or service, the veritable transformation to wireless as a pivotal delivery mechanism and platform has been nothing short of amazing. Further, wireless data consumption has been increasing at incredible rates with the most popular cited estimates slating traffic to double every year for the next ten years. In part, driven by

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HotPlanet'13, August 16, 2013, Hong Kong, China.

Copyright 2013 ACM 978-1-4503-2177-8/13/08 ...\$15.00.

the pervasive penetration of broadband, users now bring requests for speeds and similar rich content from the wired broadband world to their mobile device. Dubbed the *wireless data tsunami*, the dominant question for many wireless service providers has been how one can meet what appears to be an insatiable need for data to mobile devices.

To address the rising expanding needs of wireless users, critical research is necessary with regards to both improved performance as well as improved efficiency. Although researchers are typically quite comfortable operating on the performance side of research, there is a growing need for an improved understanding of users in a realistic setting, at a reasonable scale, and over a sufficient observation period. For wireless networks, particularly cellular networks, the gathering of such data can be difficult to manage and quite expensive to conduct. Most notably, there exist only a handful of studies with limited published results with respect to effective design or management of the instrumentation aspects of such a deployment in the research community.

Hence, the purpose of this paper is to share the experiences that we have gained managing the NetSense smartphone study, a study of roughly two hundred users now approaching the end of its two year study duration. Our study involved gathering the complete ‘digital’ smartphone view of over two hundred freshmen at the University of Notre Dame and tracking that cohort from the moment of arriving at campus in fall of 2011 through the completion of the study in the summer of 2013. As opposed to a purely technical contribution, our contributions in this paper are centered more on the creation, instrumentation, and management of the study itself, hoping in part to share aspects that went well, aspects that could be improved, and critical areas where other researchers could address gaps in our study. We focus on two critical aspects that include:

- *Resources necessary for the study:* For studies involving cellular data gathering, the key challenges are how one instruments the smart devices (relatively easy) and how one acquires the financial resources to provide appropriate user subsidies and industry cooperation (considerably more difficult). We comment in particular on our particular infrastructure for instrumentation, issues associated with requisite study approvals (human subjects) and note various aspects for improvement.
- *Challenges with study management:* Although the resource challenges to launching a large scale can be significant, the on-going management of the study can be equally challenging. Notably, the *users* themselves

present unique challenges with regards to on-going participation, retention, and on-going support.

Notably, we hope that our effort will be one of a host of efforts in part inspired by our own and other fledgling studies (Nokia Data Challenge [5], PhoneLab¹) attempting to gather live, large scale, finely instrumented user data. We hope that our candid discussions of our study will impart some wisdom to readers and also hope to guide aspiring study designers based on our efforts through a set of open research study areas and challenges for the community to explore.

2. NETSENSE SMARTPHONE STUDY

The NetSense smartphone study was originally conceived as an offshoot of a larger effort stemming from the Wireless Institute out of the University of Notre Dame. The basic premise was relatively simple but broad in scope, how could we instrument the spectrum surrounding the campus / local area with considerations for not only the technical side of the wireless spectrum but possibly instrumenting the social side impacting the network as well. Over time and several unsuccessful grant proposals later, the vision of NetSense was born in full drawing inspiration in part from work on network diffusion from the social sciences literature [2] as well as the original work from the MIT Reality Mining group exploring social interactions via phone-based agents [7, 3]. The NetSense project was funded at the start through the National Science Foundation SoCS (Socially-oriented Computing Systems) program.

The NetSense project was born out of the core question of: *How do the always-on networks afforded by wireless impact how we make and keep friends?* Specifically, how does the digital world (Facebook, SMS, e-mail, phone call) impact how we make (tie creation) and keep (tie persistence) friends? Could we in part leverage the digital devices present (circa 2010) to unobtrusively instrument our digital world?

Notably, the university environment itself affords a unique opportunity to study such aspects with freshmen in particular being an especially appealing study group. At the core, college represents a significant period of change but more particularly for the purposes of social constructs, freshmen are much more likely to be forming new ties (individuals met at college) in addition to reducing or breaking old ties (ties from high school). In contrast to state universities which would tend to draw from a more localized population in the area, Notre Dame draws from across nearly the entire country effectively offering a nearly blank slate for students to meet and establish new friendships from.² In tandem with the rich bevy of social interactions (including location and proximity), a host of technical data becomes available for the purposes of network optimization affording opportunities to study technical performance as well as the impacts of social constructs on usage and performance.

While instrumentation of smart devices is a natural fit for networking researchers, the selection of study participants, the approval with respect to human subjects research, the acquisition of funding specifically for user studies, and the requisite industry partnerships are relatively foreign topics.

¹www.phone-lab.org

²Students are unlikely to know roommates or classmates prior to entering ND due to the geographic recruiting diversity.

For the purposes of the original NetSense study, many of the resulting efforts (technical, sociological, management) were undertaken in parallel during the first year after which funding was acquired. We comment briefly on the ‘softer’ aspects of the study in terms of study preparation and mechanics before we focus on the specific technical aspects of the study.

2.1 Study Participation and Approval

With the launch of the study, a key design decision emerged with respect to instrumentation of the user population, namely could one instrument the actual ‘phone’ of the users in the study or should each user receive an auxiliary device for appropriate instrumentation? In the case of the auxiliary device, funding for the auxiliary devices is considerably simpler as music players or digital assistants (PDAs) would incur only the fixed cost of the device themselves (one-time purchase cost and on-going repairs) and would need only access to WiFi at the cost of being slightly less than optimal for logging. In contrast, having the presence of an ‘agent’ and / or owning the smartphone introduces considerable obstacles with respect to higher subject costs (whole or partial subsidization of service) and on-going support (device outages must be repaired immediately). For even a limited subject pool of only two hundred users, the costs for complete subsidization (phone plus monthly service) could easily eclipse nearly \$400,000 for unlimited data / texting plans (circa 2010). Through a partnership with a carrier (Sprint) and arrangements with the university, we were able to provide the devices and service free for the study participants settling on the Nexus S 4G.

The rationale for ownership and complete subsidization was two-fold. First and most critically, a key thrust of the study was to explore the impact of off-line (face-to-face) interactions as well as on-line (digital) interactions. While all digital interactions were available once a device was rooted, the detection of face-to-face interactions necessitated a fine-grained approach to detecting the proximity of users. For the detection of such distances roughly analogous to face-to-face interactions (less than 2.5 meters), existing techniques such as WiFi triangulation and GPS were not nearly accurate enough nor power-sensitive enough so as not to overwhelm the typical phone useable duration [6]. In contrast, our own early explorations of Bluetooth showed that Bluetooth could indeed be tuned to be accurate for proximity with exceptional accuracy in the 0 to 5 m range, perfect for the purposes of the study. Moreover, such data could be gleaned without pairing, requiring only that the device be discoverable in a Bluetooth sense.

Unfortunately, numerous recent security exploits centered around Bluetooth required user intervention to periodically re-enable the device to be Bluetooth discoverable³. The notion that a student would be willing to periodically refresh a device to be discoverable (even with prompting) every few minutes over the entirety of the day was quickly ruled out. Although the newest release of Android (Honeycomb) supported the notion of an infinitely discoverable Bluetooth, Honeycomb was a variant only for tablets and the next release (Ice Cream Sandwich) was still considerably in the

³Other devices possessed the ability to make Bluetooth permanently discoverable such as with the MIT Reality Mining study but neither Android or iOS at the time supported such a configuration in the base release (summer 2011).

future. Fortunately, a popular ROM for the community (Cyanogenmod) incorporated early features into its active build allowing us to leverage the infinite Bluetooth discoverability at the cost of installing a custom ROM on each handset in the study. Thus, we felt reasonably certain that while an agent would not be perceived as unusual, requiring a participant to jailbreak and install a customized ROM with questionable support outside of certain models (Nexus S as we used), would be viewed as a significant impediment to participation.

Second, by directly owning and subsidizing the device ourselves, we could better offer the ability to quickly hotswap devices, minimizing the time that a student might be out with a broken phone. Working with Sprint, we were able to secure a pool of cold standby devices in addition to priority support to effectively roll over a number from a broken phone to a new phone, conducting repairs potentially on the order of less than an hour. Moreover, we felt the homogeneity of devices in the pool would allow for the quicker identification of problems and consistency with respect to performance across the Notre Dame campus (3G, WiFi) and beyond (WiMax).

With the ownership model resolved, two twin challenges then emerged, namely what sort of data could we gather and depending upon what data we could gather, what would be acceptable for the local Institutional Review Board (IRB)? As the data gathering involved an a convenient but coercible pool (college students) and further involved data of a private nature (all communications), appropriate cautions and considerations needed to be weighed with regards to data gathering. At the foundation, basic environmental data with regards to wireless network performance (WiFi APs visible, WiFi signal strength, battery level) presented little challenge but the translation of said data into more useful aspects (location, application usage, application consumption) in tandem with social aspects presented a rich but sensitive data pool. Through discussions with our IRB, the compromise with regards to data was to restrict data solely to where, when, and with whom communications occur but not the actual contents of the communications themselves. Thus, for a text message, we would know the recipients of the text message, where it was sent, when it was sent, and the length but not the actual content itself. The same restrictions applied for all communication data gleaned from the phone and from other mediums (Facebook).

Student participants were targeted from a set of four dormitories (two male, two female) each having roughly fifty new freshmen arrive distributed roughly equally amongst various majors (intended) and other demographic aspects. Two alternative dormitories (one male, one female) were selected to reach the full cadre of two hundred students. A unique challenge to our institution was the fact that room assignments were not made until July thus increasing the likelihood that a student would have already purchased a new smartphone before the beginning of the school year. Students were sent an invitation letter soon after initial dormitory assignments in the targeted forms and were asked to browse full consent forms (via our website) denoting all forms of gathered data and the purposes of the study. As the purpose of our study did not involve subterfuge for the purposes of intervention but rather focused on unobtrusive monitoring, we had the flexibility to give detailed examples

of all gathered data and to expand in detail on the terms of the study.

As noted earlier, the terms of the study were relatively straightforward. In exchange for study participation and complete monitoring rights on the provided phone for the course of the study, the participant would receive a new Nexus S 4G smartphone upon arrival at campus with complementary cellular service for the study duration including a plan for 200 anytime minutes (derived through consultation with Sprint⁴), unlimited mobile-to-mobile minutes, unlimited texting, and unlimited data. Students were also expected to complete detailed entrance surveys (personality surveys, current social network), friend our study participant, complete on-going and relatively infrequent surveys (2-3 questions, multiple choice answers provided via a quiz app), and use the phone as their primary phone. Students would fax or scan their consent form to the faculty members (if under 18, parental consent required) and would receive their phone on the day of move-in.

2.2 Study Launch

While we had hoped (perhaps naively) to capture nearly all of the incoming freshmen for each of the four dormitories, the response rate necessitated an expansion to two additional dormitories and finally a limited broadcast to additional incoming freshmen (via Facebook) for the remaining few students. A wait list was formed of roughly 10 students to allow for students electing not to participate in the study or dropping out shortly after joining the study. Successful enrollment required the receipt of the consent form several weeks before arrival to campus.

Generally, most freshmen at Notre Dame arrive over the period of roughly three days just before the start of classes. A few students arrived earlier (band, international) students which provided an early test of the enrollment process. A large set of times (roughly 8-10 hours per day) were set aside for each of the primary arrival days allowing students to pick up their phone as time allowed. A faculty member was present at nearly all times which was actually quite informative for the purposes of ascertaining rationales for joining the study, the most common of which: *'I will be glad when their texting is off of my plan⁵.'*

At enrollment time, a checklist was used for the set of graduate students walking each student through the enrollment process. Enrollment steps including setting up the campus WiFi (noting the requirement that Bluetooth and WiFi must always be on), setting up the students respective e-mail accounts, and answering any immediate questions as to usage of the phone. Terms for broken and lost phone were specifically acknowledged by the student as well as the requirement to complete the two entrance surveys (each requiring roughly 20-30 minutes) for the purposes of a baseline assessment. Users were also specifically reminded that only instances of communication are recorded, not the contents of the communications themselves. Terms were also reiterated with respect to leaving the study (no penalty provided

⁴The actual minutes were more than enough with only a limited number of students incurring overages in the first few months (< 3%).

⁵For most students, their cellular service was part of a family service plan with most plans (circa Fall 2011) still having fees for various levels of texting versus the typical default unlimited texting of today.

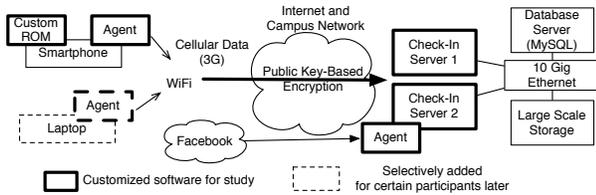


Figure 1: NetSense Infrastructure

phone is returned) and that students could elect to leave and have all data removed from consideration in the study.

Selective support was offered for the importing of contacts but was typically deferred noting that the ‘cloud’ meant that this would be the last time one would need to enter contacts. Notably, number porting was not offered due in part to the cost constraints associated with porting two hundred numbers to Sprint and early terminating existing lines. In retrospect, while this decision significantly streamlined the distribution process, it was noted as the primary reason for not participating in the study. We also noticed a subset of our users who clearly kept their old phone exhibiting poor usage rates relative to the typical user in our study. Additional commentary on this subject is offered at the end of the paper.

2.3 Technical Architecture

Once distributed, the data gathering process was done unobtrusively by virtue of an agent pre-installed on each of the phones. The agent (which executed at the user level) executed a series of periodic data gathering instruments including but not limited to: network environment (WiFi APs detected with signal strength, name, MAC address), proximity of other users (Bluetooth devices detected, name, signal strength), device state (battery, charging, screen on / off), application usage (list of installed applications, application consumption), device usage (cellular usage, WiFi usage), location, communications (SMS, e-mail, phone call), as well as various other aspects (contact book, etc.). Data was sampled at frequencies ranging from a few minutes (proximity of other users, data consumption, WiFi state) to a few hours (application installation) to each day (contacts) to on-demand (screen on / off, battery level change).

Data was temporarily spooled in a local SQLite database and then reported to one of two check-in servers through an encrypted channel (see Figure 1). Two check-in servers were used to provide for fault tolerance in the event of failure. In the event that neither server is available, data is spooled on the device and then relayed when a connection is found or a pre-defined limit is reached. Data from social networking sites (Facebook) is gathered via a separate process at one of the check-in servers and is not gathered directly on the phone. The agent itself underwent two major revisions to add in additional data gathering mechanisms (security stance - screen lock or not) and performance improvements. A limited laptop agent was also deployed to twenty users to gauge data loads on the laptop versus the smartphone of the user drawing in part on a separate security study also running concurrent with the NetSense study.

After reaching the check-in server, data is spooled temporarily before being parsed and stored in a relational database (MySQL). Critically, none of the devices talk directly to the

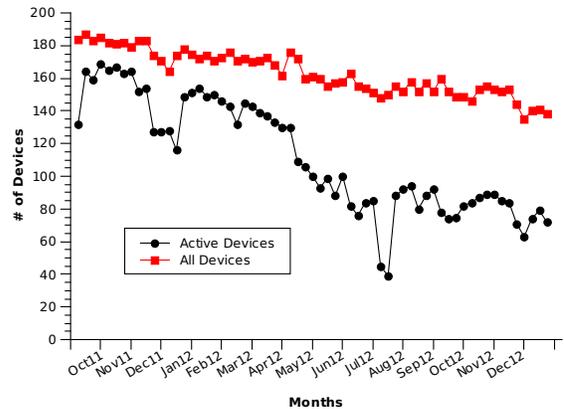


Figure 2: Active User and Overall Enrollment in Study

database nor is the database containing the data made available on the public Internet for querying, instead requiring a proxy for all accesses / queries. To the extent that we can, data is immediately anonymized with mappings saved offline in an entirely separate storage location until needed for the purposes of analytics. Data access is strictly limited to researchers in the study though agreements have been pursued where data can be reasonably anonymized. The ability of researchers to shared such data also represents a topic of future research that we comment on at the end of the paper. The integrity of the data is verified through multiple health scripts with various reports passed on to undergraduate liaisons and study managers for the purposes on following up with non-compliant participants.

Due to the limited battery life of the devices, special attention was paid to the energy impact of the devices. As originally constructed, the basic agent (Bluetooth infinitely discoverable, queried every three minutes, environmental monitoring, no GPS) was developed to last for over a day and a half with low usage, a single day with nominal use. As the study evolved, several instrument points were gradually increased in granularity based on user charge patterns but still maintaining a typical phone lifetime of twenty four hours.

3. STUDY MANAGEMENT

As the study involves actual users, a key challenge is how to manage the users, namely unexpected issues, attrition, and non-compliance. We begin first with Figure 2 which shows the overall study population over the lifetime of the study starting out at a peak of nearly two hundred students and slowly decreasing over time. A data point in the graph represents cases where a minimal corpus of data was successfully gathered from the student during that particular timeframe (single week). A few key aspects are noteworthy for commentary on the figure.

First, although we expressed significant worry with going over and over-running the budget, we were perhaps too cautious and elected to stay at or under two hundred devices on the study despite having multiple students on the wait list. At the launch of the study, the budgetary figures seemed quite onerous though in retrospect, such attrition could have been better exploited. Over time, the number

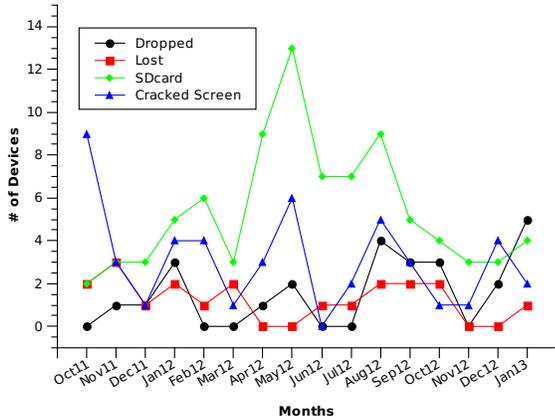


Figure 3: Causes of Broken Phones

gradually shrinks with peak dips occurring either during break (students do not use the phone or break the phone while away). Major changes tended to occur as students signed up for international study (ND has significant international study opportunities) and after the holidays (students receive a new smartphone for Christmas). Gradual attrition occurred most frequently for those who still kept their old cellular phone line with the most common reason for leaving the study that ‘the phone was too difficult to use.’

We believe a significant reason for keeping the dropout rate as high as it was for the first year of the study was due to a modification of the lost / broken phone costs. Originally, students were to be responsible for the full repair costs of any cracked screens (\$200) or lost phones (\$500). Unfortunately, issues with cracked screens occurred far more frequently due in part to the significant areas of concrete in a campus environment (versus carpet in a home). One other notable occurrence was the use of the phone as an alarm clock and the distance to the floor from the loft far exceeded the tolerances of the glass screen as used on the Nexus S. Following internal discussions with our group, we elected to subsidize the student breakages switching to a tiered scheme with a broken screen requiring a student cost of \$50 (subsidy of \$150) and the lost phones being supplemented from our spare pool. Figure 3 captures the frequency of various phone issues throughout the study.

Notably, the most frequent issue with the phones arose with the on-board flash / SDCard. The phones possessed a hard-soldered 16GB flash chip that represented the most common outage for the phone. The net result of the flash issues was a non-trivial increase in our phone standby pool in order to tolerate the month plus backlog for phone repairs. While our phone selection was limited to only the Nexus S line (to ensure Cyanogenmod would work), the prevalence of flash failures under SQLite loads would imply that a user replaceable flash is an essential component for any study. Unfortunately, the general trend seems to be away from user replaceable SD cards which is an item to monitor for prospective studies.

Finally, while user attrition is difficult to avoid, user compliance can also present significant challenges. While we could largely ensure that users were compliant respect to

WiFi and Bluetooth being on by virtue of our agent, the extent to which users always employed the phone as their primary communication device waned considerably over the course of the study. In addition to the base enrollment, Figure 2 also contained the number of ‘active’ (compliant) users where an ‘active’ user is defined as having the phone on, using the phone for a reasonable portion of the week, and completing quizzes in a timely manner. The exact thresholds varied over but we eventually finalized on a threshold of the phone gathering data for roughly 50% of the time per day on average for the week as the trigger for further investigation. However, it is noteworthy that not all decreased activity is non-compliance but rather may be family travels such as associated with spring, winter, or summer break periods (ex. June, July, August 2012).

Unfortunately, as the devices age, there is a noticeable decay in the percentage of active users most frequently decaying just after a break period where users do not drop but divert much of their usage to a secondary device. Due in part to shortages of devices from flash repairs, we were quite aggressive in recovering devices from consistently poor users who had clearly switched to a different device for their primary usage or whom were non-compliant in completing surveys / quizzes. Each case was discussed on a one-on-one basis between a faculty member and the study participant to avoid false positives and an amicable exit from the study if desired.

Although devices can be reclaimed, the completion of quizzes and other non-phone instruments that require time investments decay as well. Figure 4 plots the 24 hour and 48 hour response compliance rates for the multiple choice quizzes. In part, we believe the decrease in response rates can be attributed to the students simply being busier (sophomore versus freshmen year) and a reduced desire for compliance (typically decreased in sophomores). While repeated avoidance of survey / quiz completion can lead to removal from the study, the lack of a fine-grained incentive mechanism (or penalty) clearly shows its weakness in the latter part of the study. A relatively simple modification to the quiz agent such as posited in [4] which proposed the use of visualized feedback to improve user association with the study would likely improve compliance and we intend to adopt similar techniques (including the observations noted in [1]) for subsequent efforts.

4. CONCLUSIONS / FUTURE REMARKS

While the NetSense study has provided a rich corpus of data that will likely keep our group busy for the next few years, we note several shortcomings and opportunities that other researchers could perhaps build on as well as several research challenges.

- *Subsidy vs. ownership:* While ownership of the devices and making the devices fully subsidized (free to the student) satisfied the needs at the time, we believe a subsidy-based model is much more viable, effectively freeing the researchers from the chore of device maintenance. However, the use of a subsidy model creates other issues with regards to how to tie the level of subsidy to compliance (reward for data, reward for survey, etc.), how to subsidize but stay compliant for relevant tax / grant regulations, and the emergence of significant device heterogeneity (two year old handset, iOS support).

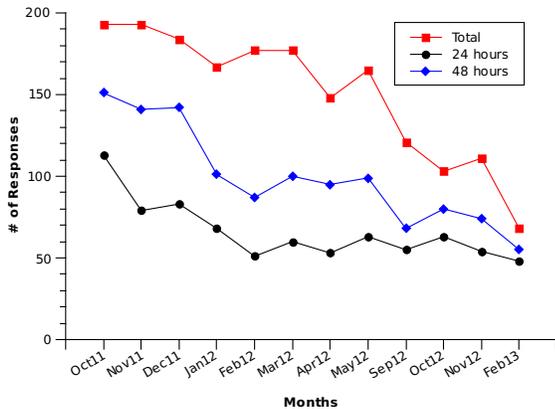


Figure 4: Average 24 / 48 hr Response Rate for Quizzes

- *Scale:* Although two hundred users was quite expensive to instrument, we quickly felt the limitations of the overall pool, particularly with regards to proximity detection. Two hundred students was only 10% of the class and overall, only 2.5% of the ND student body effectively blinding us to a large portion of the population. A critical question would be what a ‘good’ level would need to be and if perhaps shorter but larger populations are a better use of resources for technical aspects.
- *Device Breadth:* While the smartphone is the nexus of the digital communication world, it is only one piece of the puzzle. Although we selectively instrumented a subset of laptops, one must be careful to consider the picture of the ‘whole’ user, capturing not only the phone but also the tablet and laptop.
- *Beyond the Device:* Although the device is a natural place to instrument, performance is an issue that has roots in both the device as well as the wireless base station. A critical challenge looking forward, particularly with heterogeneous cellular networks (aka small cells) is how one can instrument the user side as well as the wireless infrastructure side of interactions where providers are considerably more guarded in such information sharing.
- *Data sharing:* Data anonymization quite is quite difficult for the dataset given the auxiliary social and personal information gleaned by the study across various modalities. While we can envision partitioning several aspects (application usage, WiFi performance, anonymized proximity for DTNs), there is not a robust manner by which to easily share broad data swaths outside of institutional agreements. We intend to share (with student permission) as much data as possible with insitutions such as CRAWDAD starting hopefully in the summer of 2013.

From a forward looking perspective, significant research and institutional explorations need to be conducted with respect to data sharing. The PhoneLab effort (noted earlier) represents an interesting initial foray

into this space though any research that links social or behavioral beyond foundational technical explorations will likely be difficult to navigate.

Acknowledgments

Beyond the contributions of the core faculty and graduate students driving the effort, numerous students and faculty members have played a role. We would like to thank Sandy Pentland from the MIT Reality Mining group for his invaluable insight while we were beginning our study, Trish Smithwick and Kara Worman from Sprint for their invaluable support, Christine Broadbent for assisting with the day-to-day operations of managing the study, Margaret Pickard for serving as the undergraduate contact, numerous participants in the ERWiN REU site, and finally Dirk Van Bruggen, Chris Miller, Nikhil Yadhav, and Salvador Aguinaga for their technical expertise in development.

This material is based upon work supported by the National Science Foundation under Grant No. IIS-0968529. Further support was provided through the University of Notre Dame and Sprint.

5. REFERENCES

- [1] F. B. Abdesslem, I. Parris, and T. Henderson. Reliable online social network data collection. In *Computational Social Networks*, pages 183–210, 2012.
- [2] N. Christakis and J. Fowler. Social contagion theory: Examining dynamic social networks and human behavior. In *Arxiv Preprint arXiv:1109.5235*, 2011.
- [3] N. Eagle and A. Pentland. Social serendipity: Mobilizing social software. *Pervasive Computing*, 4(2):28–34, 2005.
- [4] G. Hsieh, I. Li, A. Dey, J. Forlizzi, and S. E. Hudson. Using visualizations to increase compliance in experience sampling. In *Proc. of Ubicomp*, pages 164–167, New York, NY, USA, 2008. ACM.
- [5] J. K. Laurila, D. Gatica-Perez, I. Aad, J. Blom, O. Bornet, T.-M.-T. Do, O. Dousse, J. Eberle, and M. Miettinen. The mobile data challenge: Big data for mobile computing research. In *Proc. of Nokia Mobile Data Challenge Workshop*, 2012.
- [6] S. Liu and A. Striegel. Accurate extraction of face-to-face proximity using smartphones and Bluetooth. In *Proc. of WiMAN*, pages 1–5, 2011.
- [7] A. P. Nathan Eagle and D. Lazer. Inferring social network structure using mobile phone data. *Proc. of the National Academy of Sciences (PNAS)*, 106(36):15274–15278, September 2009.