

Measuring wireless network usage with the Experience Sampling Method

Tristan Henderson
Department of Computer Science
Dartmouth College
Hanover, NH 03755, USA
tristan@cs.dartmouth.edu

Denise Anthony
Department of Sociology
Dartmouth College
Hanover, NH 03755, USA
denise.anthony@dartmouth.edu

David Kotz
Department of Computer Science
Dartmouth College
Hanover, NH 03755, USA
dfk@cs.dartmouth.edu

Abstract—Measuring wireless local area networks has proven useful for characterizing, modeling and provisioning these networks. These measurements are typically taken passively from a vantage point on the network itself. Client devices, or users, are never actively queried. These measurements can indicate *what* is happening on the network, but it can be difficult to infer *why* a particular behavior is occurring. In this paper we use the Experience Sampling Method (ESM) to study wireless network users. We monitored 29 users remotely for one week, and signaled them to fill out a questionnaire whenever interesting wireless behavior was observed. We find ESM to be a useful method for collecting data about wireless network usage that cannot be provided by network monitoring, and we present a list of recommendations for network researchers who wish to conduct an ESM study.

Researchers have conducted many wireless local area network (WLAN) measurement studies, collecting data to help characterize, model and provision these networks. Typically these studies take network-side measurements, counting the number of clients, the amount of traffic transmitted, client session durations and so forth. Such studies are useful for determining *what* is occurring on a network. They may not, however, explain *why* a particular behavior is observed. For instance, in one of our recent studies [13], we observe increased peer-to-peer filesharing on our campus WLAN, despite the presence of a higher-bandwidth wired network. Why are users choosing to use the slower wireless network for downloading such large files? Is it because the convenience of the wireless network outweighs its limited capacity? Or is it because the users are unaware of the wired network? Such questions cannot be answered by network-side measurement alone.

To understand why particular wireless network behavior occurs, it is useful to ask network users about their usage patterns. A one-off survey or questionnaire, however, may not be appropriate. If a user is asked about a wireless behavior some time after the event has occurred, they may have forgotten the conditions at the time of the event, or the particular reason for their actions. Instead, we need a way to non-intrusively record all aspects of a user's wireless network experience.

In this paper, we describe an Experience Sampling Method (ESM) study of a campus-wide WLAN. ESM is a psychology method for studying user experience. We compare the data obtained through ESM to data from network monitoring.

We introduce the Experience Sampling Method in the next section, and describe previous ESM studies and wireless network measurement studies. Section II describes our experimental setup and methodology, and Section III presents some brief results. In Section IV we list some of the lessons that we learnt in the course of conducting the study, and Section V concludes.

I. BACKGROUND AND RELATED WORK

Diary methods are a group of research tools and methods used in psychology for “documenting the particulars of life.” [6]. The distinguishing feature of a diary method is that participants self-report their own ongoing experiences. This enables the recording of events and experiences in a more natural context than a formal interview. Diary methods can be divided into three categories [22]: interval-contingent, where participants report at regular intervals; signal-contingent, where participants report when they receive a signal; and event-contingent, where participants report whenever a defined event occurs.

Within the broader category of diary methods, the Experience Sampling Method (ESM) [18] has emerged as a popular method for evaluating user experiences and situations. In an ESM study, also referred to in the literature as Ecological Momentary Assessment (EMA), participants fill out a questionnaire several times a day. The questionnaire asks about the participant's current activities, conditions and feelings. A typical ESM study is signal-contingent, involving seven signals a day over seven days. ESM has been used to evaluate a wide variety of experiences, such as job satisfaction amongst teachers [5], happiness amongst American youths [9], and the quality of experience of new parents before and after childbirth [10].

Most relevant to our work are ESM studies that look specifically at communications or technology. Kubey et al. [17] present a survey of communications-related ESM studies, such as television viewing, and communication patterns within families and marriages. Consolvo and Walker [8] use ESM to evaluate user experiences with the Intel Personal Server ubiquitous computing device. Gaggioli et al. use ESM to analyze quality of experience and presence in virtual environments [12].

Technology can also be used to improve ESM methodology. Instead of using a notebook or paper questionnaires to record responses, Palen and Salzman ask their participants to record voicemail messages upon receiving an alert [19]. Another common electronic ESM tool is to use a PDA for questionnaires [4].

Several researchers have measured and characterized wireless LANs, mainly in academic settings. Academic WLANs that have been measured include those at Stanford [21], Georgia Tech [14], Dartmouth [13], [16], UNC [7] and the University of Saskatchewan [20]. These studies have all taken place on the wired side of the access points, using packet sniffers, SNMP polling, syslog records, or a combination of these techniques. The non-academic WLANs that have been studied include a conference WLAN [2], and the WLAN at a corporate research campus [3], both of which were also measured using SNMP or packet sniffing.

More recently, WLANs have been measured from the wireless side of the access point; that is, measuring the wireless medium itself. Wireless-side measurement is complex due to the unreliable nature of the wireless medium, hardware and software incompatibilities, and so forth. As such, most wireless-side studies are of a smaller scale than wired-side studies. Yeo et al. measure a single access point at the University of Maryland [23], while various other studies have taken place in controlled wireless conditions [1], [11], [15].

II. METHODOLOGY

In this section we describe our experimental setup. Our study differs slightly from most ESM studies in that it is both signal-contingent and event-contingent. Participants are sent signals to record details about their general wireless network experience, but in addition, we use our network monitoring infrastructure to detect interesting events and send event-contingent signals.

A. ESM setup

We recruited participants using a website and college bulletin boards. The website used a registration form to preselect participants (to make sure, for instance, that all participants owned a wireless device). We interviewed all of the respondents and selected 30 participants: 15 male and 15 female. All of the participants were undergraduate students. To attract students and encourage participation, we offered them \$100 each to participate in the study. One participant dropped out during the study; this participant's data is excluded from the results.

We issued each participant with a pager and a questionnaire notebook. The pager was a Motorola Bravo numeric pager operating in the 406-512MHz band. To programmatically send pages to the participants, we used a SMTP-to-page gateway operated by the pager service provider. The notebook contained questionnaires with questions about communications usage (wired, wireless and other communications devices), their current activities, location and so forth. We chose to use a notebook for several reasons. First, our questionnaires

were extensive, containing 25 questions, some of which had several parts and open-ended questions, and thus would be impractical to fill out using a PDA. Second, we considered using a PDA equipped with an 802.11 network adapter, as this would have allowed us to use the same device for signalling participants, for participants to fill out questionnaires, and for the questionnaire responses to be transmitted to a central data-collection server. We rejected this solution, however, as participants could not be signalled when they were out of range of an 802.11 network (for example, if they were off campus). Moreover, participants could not be signalled if the 802.11 network was having trouble, which was one of the aspects that we were interested in monitoring. Rather than require participants to carry two electronic devices (a PDA and a pager), which would only increase the probability of a device malfunction during the study, we chose to use a pager and a notebook.

Participants provided "conflict times" at which they did not want to be paged, for example, when they were asleep or in an examination. Using this information, we drew up a schedule of *signal-contingent alerts*. At each signal-contingent alert time, we sent participants a page containing a numeric code. This code indicated the page in the notebook to be filled out (in case a participant missed an alert). Each participant received seven signal-contingent alerts per day, and each alert was scheduled at least 45 minutes after the preceding alert, to prevent the alerts from being too intrusive.

B. Wireless monitoring setup

We monitored our campus WLAN using our existing measurement infrastructure. This uses four techniques to trace WLAN usage: syslog and DHCP records, SNMP polls and network sniffers (see [13] for further details). We record syslog messages from every AP on campus every time a user associates, disassociates, authenticates, deauthenticates or roams. Every five minutes, we poll each AP using SNMP to collect detailed client statistics including bytes, packets and errors transmitted; signal strength and quality. We have installed 18 Linux boxes in various switchrooms around campus, connected to the wired side of the APs, and use tcpdump to capture headers of packets sent to and from the APs.

In addition to the measurement mechanisms described in [13], we obtained logs from our central campus e-mail server. This server writes a timestamped record to a log whenever a user logs on to the e-mail server (logout times are not recorded). Our system administrators filtered these logs by username, to provide us with records for our 30 participants.

To detect interesting wireless events, we used the syslog and SNMP data since these are collected in real-time.¹ We detected the following events (the data source is listed in brackets):

- [syslog] Excessive associations and disassociations in a short time period ("ping-ponging" between APs). We define ping-ponging as a median interarrival between a client's last five associations of less than 10 seconds.

¹While the sniffers collect data in real-time, logs are stored on the sniffers and uploaded to our analysis server in off-peak periods.

- [SNMP] A client associates with an AP that is heavily-loaded with clients. We define such a “busy” AP as one with more than 16 associated clients.
- [SNMP] A client associates with an AP that is not reachable via SNMP. An AP that cannot be reached by SNMP, but is still sending syslog messages, may be having network or operational problems. We define such an “unreachable” AP as one that has not been polled successfully for one hour.
- [SNMP] A client associates with an AP that is heavily-loaded with traffic. We define such a “high traffic” AP as one where the mean inbound or outbound bytes per second over the last two SNMP polls is greater than 250Kb/s.
- [SNMP] A client associates with an AP that is experiencing high levels of errors or retransmits. We define such an “error-prone” AP as one with an error/retransmit level of $\geq 10\%$.
- [syslog] A client associates with an AP after a long period of inactivity, defined as two hours. We ignored the first association of the day, as this might be the participant waking up.

Whenever one of the above events occurred, an *event-contingent alert* was triggered, which sent a page to the relevant participant. These pages contained a different code, indicating that the participant should fill out an event-contingent survey. This survey was shorter, containing questions that were solely about the participant’s wireless experience.

Each participant was sent up to three event-contingent alerts a day. Event alerts were only sent if the participant had not received a page in the last 45 minutes.

The study took place over one week, starting on a Sunday.

III. RESULTS

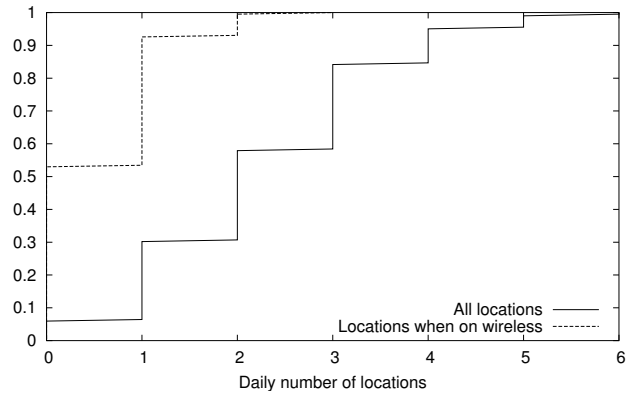
In this section we present some results from our study. The study generated an enormous amount of data; with 29 users filling out up to 10 questionnaires a day over seven days, there were over 270,000 individual data points from the questionnaires alone, in addition to the data from pre-study and post-study interviews, and the network monitoring. As such, in this paper we only provide selected statistics, and concentrate on those results that could not be obtained through network-monitoring.

A. User location

In previous measurement studies we have used syslog records to determine the approximate location of a user at a given point in time. We use a client’s association records, and assume that the client is located near to this AP. This location information may not always be accurate, however, due to different client driver behavior, or ping-ponging between APs, which can make it difficult to determine an accurate location. Asking users to self-report their location may solve these problems.

Comparing the locations determined via syslog with the participants’ self-reported locations at the time of a page shows that syslog accurately determined the building in which a participant is located 55% of the time. This proportion of

Fig. 1. Number of locations at which a participant is signalled, compared to the number of locations where a participant is signalled and connected to the wireless network. Distribution across locations per day.



correctly-determined locations is surprisingly low. It means that our raw syslog data may require some processing or additional context-specific information before conclusions about user locations or user mobility can be made.

The ESM questionnaires provide a useful mechanism for understanding errors in location-determination. Where there is a discrepancy in locations, syslog reported a building that is adjacent to the self-reported location 5% of the time. Errors in the self-reporting, such as a blank entry in the location field of the questionnaire, account for 13.4% of the discrepancies. The largest cause of discrepancies is where we believe the participant leaves their laptop at home, but still powered up and connected to the wireless network. We distinguish this behavior by looking for a long wireless session, where a client is observed associated with the same AP, but the corresponding participant’s self-reported locations change more than once. These situations account for 23.2% of the discrepancies.

One use for tracking the location of network users is to estimate or model their mobility. Using AP associations, however, only allows us to track user mobility while they are associated with the network. While this may be appropriate for laptop users, who are nomadic in their network usage, newer wireless devices such as VoIP phones tend to be always-on, and so have different mobility characteristics. To better anticipate the nature of these devices’ mobility, we may want to track users’ overall physical mobility, i.e., track all of the locations that they visit, rather than just those where they are using the WLAN.

Figure 1 shows the total number of different locations at which a participant self-reported each day, compared to the number of locations at which they self-reported themselves and were also connected to the wireless network. The median total number of locations is 2, whereas the median number of wireless locations is 0. By only observing users on our campus while they are connected to the WLAN, we are clearly missing out on a large part of their mobility.

B. Device usage

By monitoring an 802.11 wireless network, we can observe the behavior of the clients on that network. By analyzing TCP flows, it may also be possible to determine the device type of each client, which we have done in a previous study [13]. We are limited, however, to wireless devices that use the 802.11 PHY and MAC layers. A given user may be carrying a variety of other wireless devices, such as cellphones or Bluetooth devices. It is useful to understand how these multiple devices are used, to influence the design of future hybrid network devices, such as phones that can operate on both 802.11 and UMTS networks.

We asked participants about all of their communications, both electronic and non-electronic, and the devices that they used, at the time of an alert. Laptops were the most commonly-used devices, in use at 28.8% of alerts. The television was the next most popular device, in use at 10.3% of alerts. Desktops and public terminals were used at 1.9% and 2.9% of alerts respectively. Other devices in use included MP3 players (1.3%), cellphones (1.0%) and analog phones (0.9%).

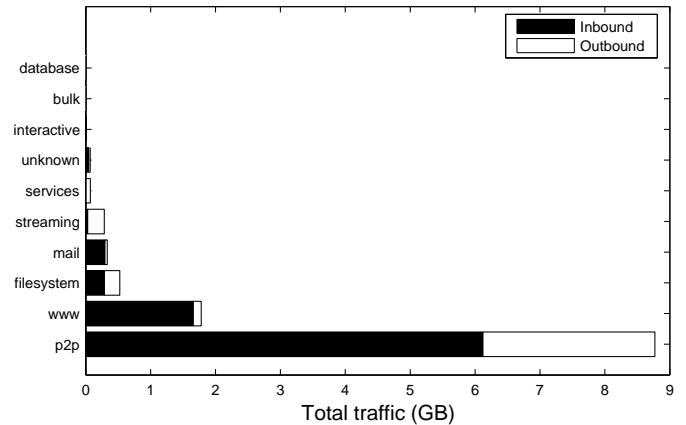
Participants were often using more than one device. The television, in particular, was used in conjunction with other devices; 48.3% of the events that involved television usage also involved the use of another device. 23.7% of the events that involved laptop usage also involved another device; of these multi-device events, 34.1% involved television and laptops, and 12.7% involved MP3 player and laptops.

Dartmouth has gradually migrated from an analog PSTN telephony system to a VoIP system over the past year. Wireless VoIP devices (Cisco 7920 and Vocera badges) are deployed on campus, and students are provided with Cisco Softphone VoIP software for use on their laptops. We have found very little wireless VoIP usage so far, however [13]. This was confirmed in our ESM study; only one participant used VoIP. We also found very little cellphone usage, and participants reported using a cellphone in the hour prior to an alert at only 2.1% of the alerts. There is limited cellphone coverage on our campus, due to its somewhat isolated location. As such, students tend to prefer using e-mail communication over telephony; participants reported sending e-mail in the 30 minutes before an alert at 15.5% of alerts, with an average of 4.7 e-mail messages sent at these times. Our low observed VoIP usage might therefore be due to low telephone usage in general.

C. Network usage

Our campus measurement studies have concentrated on the wireless portion of the campus LAN. To understand trends in wireless network usage, it may be useful to examine usage on the wired network as well, to determine those behaviors that are wireless-specific. Measuring our wired network is difficult, due to the lack of convenient central points to monitor traffic, and the large amounts of data involved. Moreover, both our wired and wireless networks do not require authentication, and so it is impossible to determine which wired and wireless MAC addresses represent the same client machine. Asking

Fig. 2. Total participant traffic (GB), by TCP or UDP protocol.



participants to self-report their application usage means that we can examine their usage of both wired and wireless networks.

Examining DHCP records for participants' wired and wireless MAC addresses shows that of our 29 participants, 13 connected to the wired network at some point during the week. All participants connected to the wireless network.

Figure 2 shows the overall application types that we saw on our tcpdump wireless sniffers. We use well-known port numbers to identify applications, and in addition we parse data payloads to look for some of the more popular peer-to-peer filesharing programs (DirectConnect, BitTorrent, Kazaa and so forth; see [13] for further methodology details). It is clear from Figure 2 that such peer-to-peer filesharing accounted for the majority of the traffic generated by our study participants (73.0%). At first glance this is not reflected in our questionnaires, where participants report that they are using a peer-to-peer application at only 2.6% of the alerts, and only six participants report using a peer-to-peer application at all. Further traffic analysis indicated that most of the peer-to-peer traffic (96.4%) was generated by four participants, one of whom never reported this usage. The timing of the alerts and the location of our packet sniffers, however, meant that we did not capture all peer-to-peer usage, with two users reporting peer-to-peer usage off-campus. ESM can thus complement packet sniffers; ESM can indicate application usage in areas that are difficult to monitor, and packet sniffing can be used to verify the accuracy of ESM reports.

The ESM questionnaires are also useful for determining the purpose of network traffic. For instance, if participants were using a web browser at the time of an alert, we asked them whether they were surfing for work-related purposes, pleasure, or both. When participants were not at home, they surfed for work 56.9% of the time and pleasure 54.2% of the time. When at home, however, the level of pleasure-related surfing rose significantly to 74.1%, while participants surfed for work at 30.2% of alerts.

IV. LESSONS FROM THE STUDY

This was the first ESM study to be conducted by the authors, and one of the first ESM studies to involve the interactive use of a wireless network. We have learnt several lessons from this initial study, and we list some of these for other researchers who may wish to carry out an ESM study.

- *Choose the time of the study carefully.*
ESM studies typically last for one week, as it is difficult to ensure participation for longer periods. When holding a one-week long study, it is important to choose an appropriate week. We chose a week during the summer term, as this was convenient for the researchers and faculty members who were conducting the study. Unfortunately there are fewer students on our campus during this term, which meant that there was a smaller pool of students to recruit from, and that there were fewer users on the network. In addition, some of our participants had examinations during the week of the study, and they commented that they did not use their wireless devices as much as usual because they were studying for their examinations.
- *Allow sufficient time between participant selection and the start of the study.*
The timing of participant selection for a sociology study is a complicated procedure. If the time interval between participant selection and the start of a study is too long, then the participants may forget to participate, or forget their instructions, by the time the study begins. In our study, the users were selected on Friday, and the study began on a Sunday. We found that this was insufficient time to prepare the study, due to the amount of data that needed to be collected from all of the participants before the study could begin. For instance, several participants provided us with incorrect MAC addresses, that required verification. Several participants gave us too many or overlapping conflict times, so that it was impossible to schedule seven signal-contingent alerts.
- *Consider all event-contingent alerts.*
Our study resulted in very few event alerts being sent to participants. This was due to a number of factors. Holding the study during the summer term meant that there were fewer users on the network, and this may have led to less problems of network congestion or loss. Although we asked our participants to minimize their list of conflict times, it turns out that there was insufficient free time in each day to send participants three event-contingent alerts each day, in addition to the seven scheduled signal-contingent alerts. This was because of our requirement that each alert be sent at least 45 minutes apart. With hindsight, we should have scheduled the signal-contingent alerts at least 45 minutes apart, but sent the event-contingent alerts whenever we observed one of the event-triggering network behaviors.
- *Provide a sufficiently large time window for responses.*
Our questionnaires asked participants about their activities in the five minutes previous to receiving a page. Some users often mentioned in their responses that they had

recently disconnected from the wireless network, but over five minutes ago. We may have been able to obtain more data if we had asked participants about their wireless activity over a longer timescale, such as the last 30 minutes.

- *Be flexible about user preferences.*
In our initial testing, we assumed that participants would have somewhat static conflict times, so that it would be possible to predetermine all of the times for the signal-contingent alerts. It turned out that students have more unpredictable schedules than we had assumed, and many participants contacted us during the week of the study to modify their conflict times. Some students, in fact, seemed to think that it was reasonable to modify their conflict times with only five minutes warning! We had to quickly modify the program that sent out the signal-contingent alerts to adapt to changing schedules. For future studies, however, we recommend instituting a rule that participants must provide sufficient notice (e.g., 12 hours or 24 hours) before rescheduling their conflict times.
- *Clearly explain the technical parameters of the study.*
The study was described to participants as a study about their communication patterns and wireless network usage. This non-technical description was not specific enough. During our study, one participant contacted us to say that they had left campus and gone home (to another state). They thought that they could still participate in the study, because they were using their own WLAN at home. We had to explain that we could only monitor network behavior on our campus WLAN.
- *Automatically prescreen participants.*
In our study the only automated prescreening was through the use of a webpage, where prospective participants had to fill out contact details and state whether they owned a wireless laptop. Additional prescreening took place in face-to-face interviews. It might, however, be preferable to increase the level of automated prescreening. For instance, the webpage could ask for a user's MAC address, and the network usage of this MAC address could be monitored prior to interviewing the prospective participant. This could be used to ensure participants with specific home locations or usage patterns.
- *Use a two-way signalling channel.*
The biggest practical problem with our study was the radio pagers used to signal the participants. We chose to use pagers as they provided a simple interface for participants, and as they would still function even when out of range of our campus WLAN. Unfortunately our pagers were one-way devices; they could receive pages, but the device itself had no means of transmitting any data. Thus there was no way to receive feedback to determine whether a message had been successfully sent. In practice we found that some participants did not receive all of their pages, or that their pages were queued by the SMTP gateway and received during a conflict time. For future studies we recommend the use of a two-way pager, so that participants can acknowledge receipt of a signal.

- *Maintain an anonymous support channel.*

In any study that involves the collection of data from human subjects, it is vital to maintain participants' privacy. In an academic environment, this may be a requirement of the Institutional Review Board (IRB) or equivalent. To maintain privacy, we assigned anonymous identifiers to each participant and used these instead of participants' names in the signalling program. The identifier to name mapping was kept in a separate location to the researcher that operated the signalling program, so that a security breach would not result in the revelation of any private information. There may, however, be situations where participants need to contact researchers, and by doing so they may reveal their identities. Several participants had support issues, such as pagers failing to work, or needed to contact us to reschedule their conflict times. Providing an online bulletin board or mailing list, where participants can communicate using anonymous identifiers, is a useful means for maintaining participant anonymity.

V. CONCLUSIONS AND FUTURE WORK

In this paper we present some results from an Experience Sampling Method study of a campus-wide wireless network. Our study is novel in that it is the first to integrate network monitoring into an event-contingent study, so that events can be triggered by observed network behavior.

We find that an ESM study is a useful method for augmenting the data collected from a traditional network monitoring study. Conducting an ESM study, however, is a complex process, given the interactions with people and the large number of variables that can cause things to go wrong. By presenting the lessons that we have learnt from our study, we hope that other researchers can benefit from our experiences.

In future work we intend to improve our analysis methods, in particular allowing us to better correlate data from ESM questionnaires and network monitoring. We also intend to carry out further ESM studies, using the lessons that we have learnt from this study.

VI. ACKNOWLEDGEMENTS

This work was a project of the Center for Mobile Computing at Dartmouth College, and was supported by McKinsey & Company, the Cisco Systems University Research Program and NSF Infrastructure Award number EIA-9802068. The authors thank Linda Lomelino '06 for her help with screening participants and processing the questionnaire data, and WaveComm of West Lebanon, NH, for providing the pagers. We also thank Minkyong Kim and the anonymous reviewers for their helpful comments on earlier drafts of the paper.

REFERENCES

- [1] A. Aguiar and J. Klaue. Bi-directional WLAN channel measurements in different mobility scenarios. In *Proceedings of the Vehicular Technology Conference (VTC Spring)*, Milan, Italy, May 2004.
- [2] A. Balachandran, G. M. Voelker, P. Bahl, and P. V. Rangan. Characterizing user behavior and network performance in a public wireless LAN. In *Proceedings of the 2002 ACM SIGMETRICS Conference*, pages 195–205, Marina Del Rey, CA, June 2002.
- [3] M. Balazinska and P. Castro. Characterizing Mobility and Network Usage in a Corporate Wireless Local-Area Network. In *Proceedings of the First International Conference on Mobile Systems, Applications, and Services (MobiSys)*, pages 303–316, San Francisco, CA, May 2003.
- [4] L. F. Barrett and D. J. Barrett. An introduction to computerized experience sampling in psychology. *Social Science Computer Review*, 19(2):175–185, Summer 2001.
- [5] A. Bishay. Teacher motivation and job satisfaction: A study employing the experience sampling method. *Journal of Undergraduate Science*, 3:147–154, Fall 1996.
- [6] N. Bolger, A. Davis, and E. Rafaeli. Diary methods: Capturing life as it lived. *Annual Review of Psychology*, 54:579–616, Feb. 2003.
- [7] F. Chinchilla, M. Lindsey, and M. Papadopouli. Analysis of wireless information locality and association patterns in a campus. In *Proceedings of the 23rd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 906–917, Hong Kong, China, Mar. 2004.
- [8] S. Consolvo and M. Walker. Using the experience sampling method to evaluate ubicomp applications. *IEEE Pervasive Computing*, 2(2):24–31, Apr.-June 2003.
- [9] M. Csikszentmihalyi and J. Hunter. Happiness in everyday life: The uses of experience sampling. *Journal of Happiness Studies*, 4(1):185–199, 2003.
- [10] A. D. Fave and F. Massimini. Parenthood and the quality of experience in daily life: A longitudinal study. *Social Indicators Research*, 67(1-2):75–106, June 2004.
- [11] G. Gaertner and V. Cahill. Understanding link quality in 802.11 mobile ad hoc networks. *IEEE Internet Computing*, 8(1):55–60, Jan/Feb 2004.
- [12] A. Gaggioli, M. Bassi, and A. D. Fave. Quality of experience in virtual environments. In G. Riva, F. Davide, and W. A. Ijsselstein, editors, *Being There: Concepts, effects and measurement of user presence in synthetic environments*, pages 121–135. Ios Press, Amsterdam, The Netherlands, 2003.
- [13] T. Henderson, D. Kotz, and I. Abyzov. The changing usage of a mature campus-wide wireless network. In *Proceedings of the Tenth Annual International Conference on Mobile Computing and Networking (MobiCom)*, pages 187–201, Philadelphia, PA, Sept. 2004.
- [14] R. Hutchins and E. W. Zegura. Measurements from a campus wireless network. In *Proceedings of the IEEE International Conference on Communications (ICC)*, volume 5, pages 3161–3167, New York, NY, Apr. 2002.
- [15] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala. Sniffing out the correct physical layer capture model in 802.11b. In *Proceedings of the 12th IEEE International Conference on Network Protocols (ICNP)*, pages 252–261, Berlin, Germany, Oct. 2004.
- [16] D. Kotz and K. Essien. Analysis of a campus-wide wireless network. *Wireless Networks*, 11:115–133, 2005. Previously appeared in *MobiCom 2002* and as Dartmouth TR2002-432.
- [17] R. Kubey, R. Larson, and M. Csikszentmihalyi. Experience sampling method applications to communication research questions. *Journal of Communication*, 46(2):99–120, June 1996.
- [18] R. Larson and M. Csikszentmihalyi. The experience sampling method. *New Directions for Methodology of Social and Behavioral Science*, 15:41–56, 1983.
- [19] L. Palen and M. Salzman. Voice-mail diary studies for naturalistic data capture under mobile conditions. In *Proceedings of the ACM Conference on Computer Supported Collaborative Work*, pages 87–95, New Orleans, LA, Nov. 2002.
- [20] D. Schwab and R. Bunt. Characterising the use of a campus wireless network. In *Proceedings of the 23rd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 862–870, Hong Kong, China, Mar. 2004.
- [21] D. Tang and M. Baker. Analysis of a local-area wireless network. In *Proceedings of the Sixth Annual International Conference on Mobile Computing and Networking (MobiCom)*, pages 1–10, Boston, MA, Aug. 2000.
- [22] L. Wheeler and H. T. Reis. Self-recording of everyday life events: Origins, types and uses. *Journal of Personality*, 59(3):339–354, Sept. 1991.
- [23] J. Ye, M. Youssef, and A. Agrawala. A framework for wireless LAN monitoring and its applications. In *Proceedings of the Third ACM Workshop on Wireless Security (WiSe'04)*, pages 70–79, Philadelphia, PA, Oct. 2004.