

Feasibility of a Socially Aware Authentication Scheme

Andrew Dathan Frankel and Muthucumaru Maheswaran

School of Computer Science

McGill University

Montreal, QC H3A 2A7, Canada

{afrank10, maheswar}@cs.mcgill.ca

Abstract—Social interaction is already a proven component of informal identification: humans are naturally skilled at recognizing other people and are unlikely to be duped by impersonation. Based on this premise, a fourth-factor, *someone-you-know*, has already been proposed as an emergency authentication method. This paper explores leveraging a user’s preexisting social actions as a primary authentication tool, one that operates transparently and automatically without explicit user guidance. Specifically, we describe the feasibility of capturing a user’s local social context using short range wireless devices and evaluate the uniqueness of that context in comparison to that of possible aggressors.

Index Terms—biometrics, fourth-factor authentication, identification, mobile networks, recognition, security, social networks, technology social factors.

I. INTRODUCTION

Multi-factor authentication systems are premised on the difficulty of defeating multiple, independent security measures. A pickpocket may be able to steal someone’s bank card out of an unprotected wallet, while a snoop may be able to learn another’s PIN by glancing at an unguarded keypad. Yet compromising a bank account requires both thefts to occur within short order to the same victim.

Independent factors are essential to gaining additional security over a single checkpoint. Merely requiring two instances of the same factor, such as two separately issued bank cards, would not provide the same protecting as requiring a PIN, as many users would be inclined to keep both cards together for convenience. Requiring *something-you-know* in addition to *something-you-have* is the most commonly used form of two-factor authentication systems, which are generally considered robust against casual attacks.

Yet such systems are far from perfect for either high-value or casual use situations. High-value targets, such as the contents of a person’s bank account, are still at risk to old-fashioned physical assaults, where an aggressor may demand a victim surrender both tokens under duress. More frequently, the second factor may simply be ignored for the sake of speed or convenience. Many credit card companies now allow vendors to process small payments without collecting the cardholder’s signature [1].

Biometrics are a commonly suggested and much researched alternative. Unlike traditional items you have which can easily be lost or stolen, or things you know which may be forgotten or written down, a direct measurement of *someone-you-are* is

permanent, untransferable, and unforgettable. A wide variety of usable metrics have been proposed including DNA, face, fingerprint, hand geometry, and iris: each characteristic has its own advantages and tradeoffs with regards to accuracy, convenience for the user, speed of analysis, and cost [2]. Commercial security systems implementing each of the aforementioned biometrics are readily available.

A fourth factor, *somebody-you-know*, has been proposed by researchers from RSA laboratories. Their system of vouching is not intending to be used as an additional primary authentication factor. Instead, it provides a convenient alternative to traditional methods in emergency authentication situations, that is, the case where a user has forgotten his password or hardware token. It allows a trusted agent to temporarily substitute his personal confirmation of the user’s identity for one of the forgotten factors [3].

Yet, a fourth factor security scheme does not need to measure its user’s social interaction directly. Interacting wireless devices could stand in for the actual interaction of people if we assume such devices were convenient and popular enough to always accompany their owners. This is not an unreasonable assumption. People are already accustomed to keeping some valuable items such as wallets on their persons throughout the day. Every morning, many strap time-keeping machines onto their wrists without giving a moment’s thought to the prevalence and pervasion of technological tools. This paper presumes that a social authentication device could become as ubiquitous as a wallet, watch, or house key (or, more likely, social authentication capabilities could be built-in another device).

In fact, one wireless device has already reached such a stature. Cellular phones are immensely popular with over 3.3 billion subscriptions worldwide and over 100 percent market penetration in some countries (a statistical quirk caused by some individuals’ possession of multiple phones) [4]. Bluetooth capable phones are already fully equipped for short-range wireless communication. Such phones are increasingly the norm: at the start of 2008 nearly 80 percent of all phones sold were Bluetooth capable [5].

This paper discusses the feasibility of using such recorded social knowledge as a primary (non-emergency) factor. In particular, we expect the user’s social context, those people in reasonably close proximity, to typically contain friends:

people the user has frequently seen in the past. On the contrary, a random aggressor is unlikely to share many mutual acquaintances with his victim. If the recorded context suddenly and unexpectedly changes to a group of strangers, we can conclude the user may no longer be who he claims to be. Such a security scheme is said to be socially aware.

Embedding an identification system into a such a socially aware device does not reduce the system to yet another form of something-you-have. A successful theft of a social authentication device only temporally compromises the system: the new social context should reveal the theft and allow the security system to disable the device. Such security checks can be made continuously without direct user input, thus reducing the need for burdensome active authentication.

This poses the essential compromise that must be made when implementing a socially aware security (SAS) scheme: balancing convenience with security. Attempting to identify thefts too quickly could cause excessive false alerts when the rightful owner simply meets new people. For simplicity, we have only evaluated our simulation from a windowed perspective that evaluates each day separately. A practical implementation should record and evaluate all of its available data in realtime.

II. PRIOR WORK

Measuring social interaction using bluetooth devices has been previously done. The Reality Mining group within the Massachusetts Institute of Technology Media Lab has been recording and analyzing data on human behavior using Bluetooth phones loaded with software that stores information about both the phone’s environment and its usage. In particular, the group has published a data set collected by 100 volunteers equipped with Bluetooth cell phones over the course of the 2004-2005 academic year [6].

We have used this data set to analysis the feasibility of constructing a socially aware security scheme. The data set is awash with information that could be used to construct a full behavioristic model, such as phone numbers for calls placed and received, nearest cell tower locations, keypad activity and time-spans for each the described. Although any practical implementation of cellular phone anti-theft system will surely benefit from leveraging all of this available data, we have limited our analysis to what can be learned from sightings of other Bluetooth devices in the vicinity (~ 10 m). These responding devices include both the phones issued by the study to the other subjects, as well as privately owned cell phones, computers, and other Bluetooth devices. This it remains applicable to any short-range wireless device and in any environment where electronics are commonplace.

Although the Reality Mining data set is convenient for preliminary investigations into the feasibility of constructing a SAS scheme, it is not particularly well suited to security research. In particular, one of our underlying assumptions is that unrelated people are unlikely to share the company of many mutual acquaintances. Yet, the study participants were expressly chosen to maximize the density of subjects within

a limited geographical area. Thus, 75 of the participants are students, faculty, or staff of the MIT Media Lab and the other 25 are students from the adjacent Sloan business school. This limitation only increases the difficulty of successfully distinguishing expected from anomalous behavior, thus we feel we can present our findings as a worst-case example.

Nathan Eagle wrote his PhD thesis [7] discussing a variety of findings in the data set. Although security was not his focus (most of his discussions focus on individual user’s behaviors rather than comparing different users) many of his techniques are applicable to our work.

Location information can be collected by cell phones in three different ways. The first two, GPS and cell tower visibility, are widely used but are unfortunately limited in resolution and to places with adequate reception. Eagle proposes a third technique: static bluetooth devices. Many of the places with limited cell tower reception are filled with numerous unmoving bluetooth capable devices (such as office buildings equipped with desktop computers). Proximity to these devices can be used in addition to or instead of tower based location information. [pg. 43-46]

Intuitively, some people’s lives are less predictable than others. Eagle has applied the concept of Shannon entropy (Equation 1) from information theory to quantify the randomness of life. In particular, he considered the entropy associated with location and daily schedules. He found that freshman undergraduates in the study had the most random schedules, followed by graduate students, then by the MIT faculty and staff who typically had low entropies. [pg. 68-71]

$$H(x) = \sum_{i=1}^n p(i) \log_2 \frac{1}{p(i)} \quad (1)$$

Such low entropy subjects are predictable. Using Bayes’ rule conditioned on the time of day and any available location information, Eagle was able to predict whether a given user will see a given subject within the hour with accuracies of up to 90%. [pg. 78]

Although Bayesian and Markovian probability models can be successful for limited sets of specific behaviors, they are less applicable in general. Eagle developed a more adaptable technique by decomposing an individual’s behavioral data set into “a set of vectors of characteristic behaviors named eigenbehaviors” [pg. 89]. These vectors can be used for prediction: 75% of low entropy users’ location behaviors (where they are at given times of the day) can be described by one eigenbehavior. They can also be used to quantify how similar two users’ behaviors are. If the eigenbehaviors are calculated from an entire group’s behaviors and taken as a basis, then each user’s unique behavior can be represented as vector in this eigenspace. The difference between their behaviors is then the Euclidean distance between these vectors. The same approach can compare a given user to a larger group of his peers. If the same user is compared to multiple groups, then attempting to minimize the distance provides a method of classifying the user into one group versus another. Using 6 eigenbehaviors,

TABLE I
DISTRIBUTION OF ENCOUNTERS BY DURATION

	< 1 min.	≥ 1 min.	All
# of Encounters	145181	145081	285490
Percentage	49.2	50.8	

Eagle was able to correctly classify 92% of the Reality Mining subjects into their self-identified groups. He also suggests that eigenbehaviors would be useful as a biometric to “detect incidents that are far from the user’s behavior space” [pg. 101]. [pg. 89-101]

A. G. Milkas et al. have also explored the behavior of participants in the same Reality Mining Project in the interest of extracting useful knowledge from social context recordings [8]. Their definition of friends, two people whose phones come within proximity on ten or more separate days, is particularly useful to our work. Some of their relevant discoveries are:

- The longest period with no encounters reported is 4 hours and 24 minutes
- Individuals’ encounter rates are predictable, changing by more than 5 encounters between consecutive hour slots less than 7% of the time
- Most pairs of people (71%) encounter on only one day. Less than 7% of pairs are friends (those who encounter on 10 or more days)
- Encounters between friends account for two-thirds of all encounters

III. FEASIBILITY ASSESSMENT

Bluetooth encounters can casually be partitioned into two types based on their durations: short, near-instantaneous sightings, as would be expected from two people walking past one another on the street, and longer engagements corresponding to actual interaction or shared activities. The Reality Mining data set supports this assumption. In fact, the two types of encounters occur in nearly equal proportions: 49.2% of encounters were less than one minute in duration, while 50.8% lasted one minute or longer (Table I). We expect the longer “engagement” type encounter to more accurately represent an individual’s social context and have limited the rest of our discussion to only this type of encounters. We thus modify Milkas’ definition of a friend to someone with whom the given user has engaged on 10 or more days.

This definition is not intended to correspond with the colloquial definition of friendship. If a user were to purchase coffee every morning at the same cafe from the same barista, he may not consider his server to be a friend in the traditional sense, but he would be labeled as such in our model. Similarly, two people who happen to ride the train to work each morning at the same time would be included. Such regular social events, although hardly representative of friendship, are still inherently valuable in detecting and predicting a person’s daily habits.

It is also worth emphasizing that, although we have been using the terms bluetooth device, person, and friend interchangeably, not every responding device corresponds to a unique

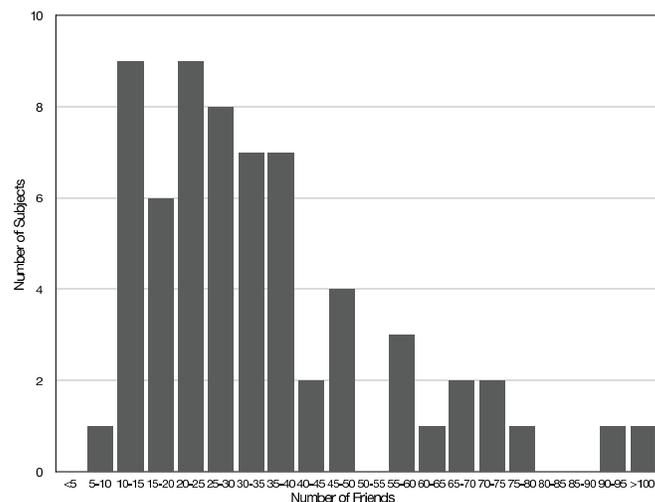


Fig. 1. A histogram showing the distribution of subjects based on how many friends they formed during the course of the entire study.

person. Some of the bluetooth devices are stationary, such as desktop computers. Likewise, some people may carry multiple devices. We do not believe such details are problematic from a security standpoint.

As would be expected from a large scale study of human beings, the Reality Mining data set is incomplete. Some data was lost to storage corruption, more to phones simply being turned off or accidentally damaged. Fortunately, the original researchers also logged when logs have coverage, allowing missing time to be excluded. During the nine months in which the study was conducted with 97 participants, there is log coverage for at least part of 10,522 days. Of these, 9,716 days (92%) were logged by 70 distinct subjects. The other 27 did not participate in the study for at least 60 days each.

By the end the study, the subjects made a total of 2,859 friendship pairs. On average, they each made 29 friends. Of the original 97, 13 subjects made fewer than 10 friends. The other 84 accounted for 98% of the recorded friendships. A histogram showing the distribution of these 84 according to their number of friends is included in Figure 1. Some of these 13 outliers merely did not participate in the study long enough to form many friendships. We expect the others to have frequently forgotten their phones or were prone to some other sort of unreliable behavior.

Thus, we have filtered the data set in the following discussion to those subjects who participated for at least 60 days and formed at least 10 friendships. There are 64 such subjects. We have recorded the number of days each subject participated in the study, and the total number of friends he made in that time. We then measured the average amount of time he spent per day in the company of a friend (where spending 5 minutes with two friends simultaneously adds 10 minutes to the total) and the standard deviation of the means.

Finally, we calculated the entropy of the subject’s friendships. Entropy in this case is defined as the amount of time spent with a given friend divided by the time spent with any friend. Thus, this entropy calculation measures the randomness

TABLE II
AVERAGES OF SUBJECTS' ENCOUNTERS OF FRIENDS AND STRANGERS

	Selves	Stolen by Friends	Stolen by Strangers
# of Days Participated	111.0		
# of Friends	35.0		
Time spent with friends (hours) per day	13.0	2.1	0.8
Standard deviation of time spent with friends (hours) per day	14.0	3.7	1.5
Entropy of friends (bits)	3.7		
Time spent with strangers (hours) per day	16.5	8.9	8.0
Standard Deviation of time spent with strangers (hours) per day	16.9	12.7	11.0
Entropy of Strangers (bits)	5.0		
Percent of total social time spent with friends	63	16	6.2

of which friends the user frequents. For example, if someone were to divide his time equally between two friends, the entropy would be 1 bit, the same as for a fair coin flip. Conversely, spending 75 percent of one's time with one friend and 25 percent with the other would reduce the entropy to .81 bits. Spending equal amounts time with 35 friends (the average number of friends in the study) would yield an entropy of 5.1 bits.

Entropy calculations are dependent on the choice of a probability model and ours is the most simplistic one possible. A more sophisticated model might be conditioned on the location, time of day, day of the week, and lengths of previous engagements. More accurately predicting how long a meeting with a given friend will last would decrease the entropy calculated from the same data.

The entropy calculations can be used to evaluate the accuracy of the social model or can be leveraged directly as metric to detect unexpected behavior. Subject's friend entropies averaged to 3.7 bits, much less than the value of 5.1 which would correspond to engagements of random lengths. Thus, not surprisingly, subjects spent more time with some of their friends than with others. Although we have only considered friendship as a binary state, you either are or are not friends with someone, a more accurate model could evaluate how close the user is to different friends and weigh their engagements accordingly. Similarly, we expect a user's engagements of strangers to show more randomness than of their friends. This is supported by the data. An unexplained decrease in this randomness could be used to detect a theft. This would correspond to the thief's interaction with his own friends, who would be classified as strangers by the stolen device. More importantly, quantifying the randomness in a user's life is invaluable when deciding whether a given scenario represents a theft or merely an unusual, but innocent event. For example, a low entropy user, someone who associates with only a few friends, interacting with a large group of strangers should draw more suspicion than a gregarious person doing the same.

We measured the amount of time a given subject spends with strangers (anyone engaged by a subject who is not a friend), found the ratio of time spent with friends to strangers, and calculated the entropy of which strangers are engaged by the subject.

We then simulated what a theft of the social authentication

device would do to the same amount of time spent with friends and strangers respectively. For each of the subjects, say Alice, we calculated how much time a thief Terri would spend with Alice's friends and correspondingly with her strangers. We have separately considered the cases where Terri is one of Alice's friends and where Terri is a stranger to Alice. Many thieves did not spend any time with their victims' friends. The entropy is undefined for such cases, thus we have not included its average value. These findings are summarized in Table II.

The most striking difference is the absolute amount of time the subjects spent with their friends compared to the amount of time a thief happened to do the same. Even if one of the victim's "friends" were to steal his device, the rightful owner still typically spends over 600% more time with his friends. This leaves a wide margin for a social context based theft-detection algorithm to work. Comparing time spent with friends instead of strangers is similarly promising. While typical users are with friends for 63% of the total time they spend with other people, thieves only do so for 16% and 6.2% when the theft is a false-friend or a stranger respectively. Figures 2, 3, and 4 emphasize this disparity by displaying the increasingly dramatic right-skew of the three distributions.

Despite this large difference in social context, a simple model based solely on these statistics is unlikely to be successful in detecting thefts. For all of the time-based values, the standard deviation is as large as the measurement itself. Stated simply: although a typical user spends around 13 hours each day in the company of friends, he is also quite likely to spend none, or to spend twice as many. Individuals have a wide variety of habits and hobbies; the social interaction of multiple individuals are even more varied. Some of this variance may be a result of limited Bluetooth market penetration. We expect that a contemporary study would detect far more Bluetooth devices in the wild than was the case in 2005. A more sophisticated model might also be able to ignore expectedly erratic encounters. Including location and time information, for example, would allow the system to consider encounters made in a subway car or on a bus at 9 am differently than those made in the user's home at midnight.

IV. CONCLUSION

SAS presents the opportunity to add fourth-factor security to a wide range of wireless devices. In fact, the increasing

miniaturization and mass manufacturing of computing systems brings the potential to embed wireless chips in traditionally non-interactive objects. A typical person is already accustomed to carrying several high-value items on his person throughout the day such as credit cards, keys, a watch, a cell phone, etc.. Each of these objects could implement a SASS that is especially attuned to the proximity of the other objects. Thus, a credit card could disable itself if divorced from its owner's wallet. Likewise, a cell phone may only function outside of its owner's home if a house key is also present.

Of course, the best implementation of such a system would consider all of the objects typically in a user's "device cloud" or "personal area network." The social security model is well suited to aggregation with additional biometric models. Some of devices are inherently more "intelligent" than others, that is, they have access to more data by design. Cell phones have access to calling habits and GPS or cell-tower based location information. Credit card companies already monitor individual's spending habits for unusual changes that could indicate a theft. A wristwatch equipped with a thermometer could track its wearer's body temperature. All of these devices could then share wireless tokens that not only include announce their presence, but also quantify their confidence in their carrier's identity. Individually, each of these biometric models would not be particularly useful for theft detection, but their combination would be difficult to dupe. The additional, non-social factors, also prevent a bootstrapping identity theft vector where a mugger could defeat a naive SAS simply by taking all of his victim's possessions.

Strictly fourth-factor based security is unlikely to be practical as a stand-alone security component. Although users spend a great amount of time on average in the company of friends, the variance of that amount of time is also very high. Additional biometrics would decrease the frequency of false alerts when the rightful owner of a device merely does something unexpected. Yet occasional such errors should not be considered a fatal flaw in the system. A practical implementation would simply have to ask the user for some alternative form of authentication. For example, a cell phone might require a previously chosen PIN, as some models already offer as an optional security feature. However, few users choose to complete the chore of manually activating and deactivating such a security feature before and after using their phones. Instead of requiring that PIN each time the phone is used, a socially aware cell phone would only ask when a potential theft is detected. Even a relatively eager implementation, which asks for explicit re-authorization at the slightest possibility of a theft, would require less interaction than a manual system. This would greatly improve the convenience of such security features and thus increase the actual overall security of the system. Such an implementation reunites fourth-factor security with the notion of emergency authentication, only in the reverse of its original proposal.

Shared emergency authentication information can be used to augment the security of lesser devices that lack keypads or direct input channels. For example, consider the case of a

ordinary person traveling on a family vacation who wishes to use his credit card in another country. Currently, unexpected international charges are considered highly suspect by many credit card companies, who often wish to confirm them with the cardholder. SAS presents a variety of options the credit card could use to confirm its owner's identity, which may be selected as warranted by the price of the attempted transaction. If the traveler is accompanied by his wife and children, their presence, combined with possession of the credit card and cell phone (two factors: something-you-have and someone-you-know), may be sufficient to authorize a small purchase. If the traveler is alone, or simply if more re-assurance is desired, the system could leverage the biometric information available to a wristwatch (something-you-are). For a larger value, where transparent security is no longer expected, the card could signal its paired phone and request the entry of a PIN (something-you-know). Most dramatically, the card could query the phone for its contact number and provide it to the credit card company, who can then immediately call the card holder and confirm his identity and travel plans directly, even if the company only had his home phone number on file. The fourth factor element can be used as an enabling technology, facilitating the use of the traditional three factors where they are most appropriate.

We hope to continue exploring this concept by developing a software application which volunteers will be able to install on their own Bluetooth mobile phones. This will allow us to collect data from a large number of test subjects spread over a greater geographic area than those of the Reality Mining study. More importantly, it will provide a platform to implement a realtime, proof-of-concept of the socially aware security model.

REFERENCES

- [1] E. Roseman. (2008) No need to sign on the dotted line. [Online]. Available: <http://www.thestar.com/Business/article/243082>

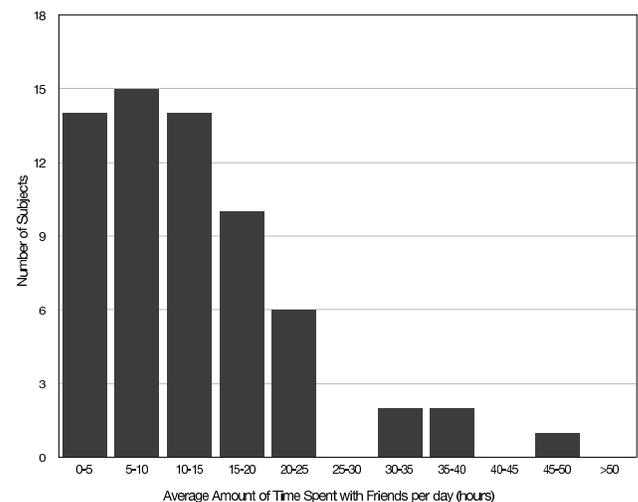


Fig. 2. A histogram showing the distribution of subjects based on how much total time they spent with their own friends.

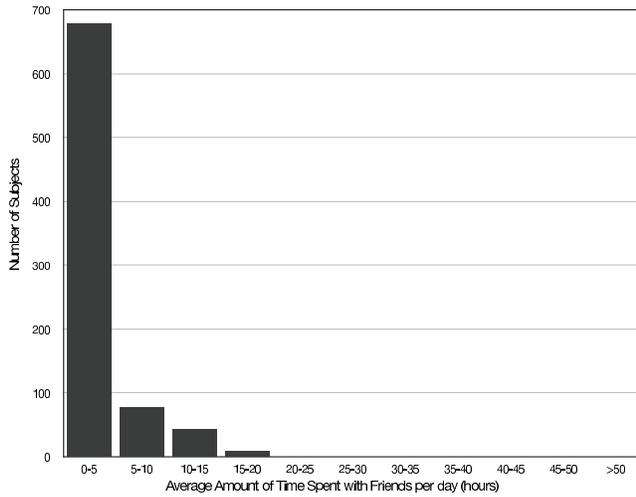


Fig. 3. A histogram showing the distribution of thief-victim pairs based on how much total time the thief spent with the victim's friends. In this case, the thieves were limited to be friends to the victim.

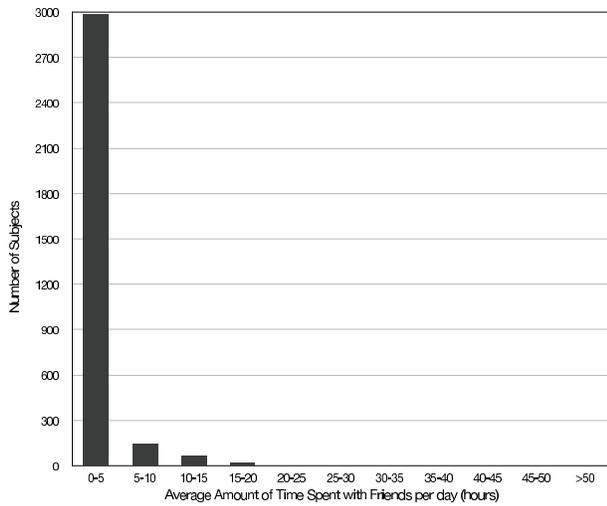


Fig. 4. A histogram showing the distribution of thief-victim pairs based on how much total time the thief spent with the victim's friends. In this case, the thieves were limited to be strangers (non-friends) to the victim.

- [2] A. K. Jain, A. Ross, and S. Prabhakar, "An introduction to biometric recognition," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 14, Jan. 2004.
- [3] J. Brainard, A. Juels, R. L. Rivest, M. Szydlo, and M. Yung, "Fourth-factor authentication: Somebody you know," *ACM CCS*, pp. 168-78, Oct./Nov. 2006.
- [4] Reuters. (2008) Global cellphone penetration reaches 50 pct. [Online]. Available: <http://tinyurl.com/3cmmfr>
- [5] M. Hachman. (2008) Study: Few phone buyers purchase accessories. [Online]. Available: <http://www.pcmag.com/article2/0,1895,2321441,00.asp>
- [6] (2008) Reality mining dataset. [Online]. Available: <http://reality.media.mit.edu/dataset.php>
- [7] N. N. Eagle, "Machine perception and learning of complex social systems," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, May 2005. [Online]. Available: <http://reality.media.mit.edu/pdfs/thesis.pdf>
- [8] A. G. Miklas, K. K. Gollu, K. K. Chan, S. Saroiu, K. P. Gummadi, and E. de Lara, "Exploiting social interactions in mobile systems," *UbiComp*, 2007.