Information Revelation and Privacy in Online Social Networks (The Facebook case)

Pre-proceedings version. ACM Workshop on Privacy in the Electronic Society (WPES), 2005

Ralph Gross
Data Privacy Laboratory
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
rgross@cs.cmu.edu

Alessandro Acquisti
H. John Heinz III
School of Public Policy and Management
Carnegie Mellon University
Pittsburgh, PA 15213
acquisti@andrew.cmu.edu

ABSTRACT

Participation in social networking sites has dramatically increased in recent years. Services such as Friendster, Tribe, or the Facebook allow millions of individuals to create online profiles and share personal information with vast networks of friends - and, often, unknown numbers of strangers. In this paper we study patterns of information revelation in online social networks and their privacy implications. We analyze the online behavior of more than 4,000 Carnegie Mellon University students who have joined a popular social networking site catered to colleges. We evaluate the amount of information they disclose and study their usage of the site's privacy settings. We highlight potential attacks on various aspects of their privacy, and we show that only a minimal percentage of users changes the highly permeable privacy preferences.

Categories and Subject Descriptors

K.4.1 [**Computer and Society**]: Public Policy Issues—*Privacy*

General Terms

Human Factors

Keywords

Facebok, Online privacy, information revelation, social networking sites

1. EVOLUTION OF ONLINE NETWORKING

In recent years online social networking has moved from niche phenomenon to mass adoption. Although the concept dates back to the 1960s (with University of Illinois Plato computer-based education tool, see [16]), viral growth and commercial interest only arose well after the advent of the Internet.[1] The rapid increase in participation in very recent years has been accompanied by a progressive diversification and sophistication of purposes and usage patterns across a multitude of different sites. The Social Software Weblog[2] now groups hundreds of social networking sites in nine categories, including business, common interests, dating, face-to-face facilitation, friends, pets, and photos.

While boundaries are blurred, most online networking sites share a core of features: through the site an individual offers a "profile" - a representation of their sel[ves] (and, often, of their own social networks) - to others to peruse, with the intention of contacting or being contacted by others, to meet new friends or dates (Friendster,[3] Orkut[4]), find new jobs (LinkedIn[5]), receive or provide recommendations (Tribe[6]), and much more.

It is not unusual for successful social networking sites to experience periods of viral growth with participation expanding at rates topping 20% a month. Liu and Maes estimate in [18] that "well over a million self-descriptive personal profiles are available across different web-based social networks" in the United States, and Leonard, already in 2004, reported in [16] that world-wide "[s]even million people have accounts on Friendster. [...] Two million are registered to MySpace. A whopping 16 million are supposed to have registered on Tickle for a chance to take a personality test."

The success of these sites has attracted the attention of the media (e.g., [23], [3], [16], [4], [26]) and researchers. The latter have often built upon the existing literature on social network theory (e.g., [20], [21], [11], [12], [32]) to discuss

WPES'05, November 7, 2005, Alexandria, Virginia, USA.

[1] One of the first networking sites, SixDegrees.com, was launched in 1997 but shut down in 2000 after "struggling to find a purpose for [its] concept" [5].
[2] `Http://www.socialsoftware.weblogsinc.com/` .
[3] `Http://www.friendster.com/` .
[4] `Http://www.orkut.com/` .
[5] `Http://www.linkedin.com/` .
[6] `Http://www.tribe.net/` .

its online incarnations. In particular, [7] discusses issues of trust and intimacy in online networking; [9] and [8] focus on participants' strategic representation of their selves to others; and [18] focus on harvesting online social network profiles to obtain a distributed recommender system.

In this paper, we focus on patterns of personal information revelation and privacy implications associated with online networking. Not only are the participation rates to online social networking staggering among certain demographics; so, also, are the amount and type of information participants freely reveal. Category-based representations of a person's broad interests are a recurrent feature across most networking sites [18]. Such categories may include indications of a person's literary or entertainment interests, as well as political and sexual ones. In addition, personally identified or identifiable data (as well as contact information) are often provided, together with intimate portraits of a person's social or inner life.

Such apparent openness to reveal personal information to vast networks of loosely defined acquaintances and complete strangers calls for attention. We investigate information revelation behavior in online networking using actual field data about the usage and the inferred privacy preferences of more than 4,000 users of a site catered to college students, the Facebook.[7] Our results provide a preliminary but detailed picture of personal information revelation and privacy concerns (or lack thereof) in the wild, rather than as discerned through surveys and laboratory experiments.

The remainder of this paper is organized as follows. We first elaborate on information revelation issues in online social networking in Section 2. Next, we present the results of our data gathering in Section 3. Then, we discuss their implications in terms of users attitudes and privacy risks in Section 4. Finally, we summarize our findings and conclude in Section 5.

2. INFORMATION REVELATION AND ONLINE SOCIAL NETWORKING

While social networking sites share the basic purpose of online interaction and communication, specific goals and patterns of usage vary significantly across different services. The most common model is based on the presentation of the participant's profile and the visualization of her network of relations to others - such is the case of Friendster. This model can stretch towards different directions. In matchmaking sites, like Match.com[8] or Nerve[9] and Salon[10] Personals, the profile is critical and the network of relations is absent. In diary/online journal sites like LiveJournal,[11] profiles become secondary, networks may or may not be visible, while participants' online journal entries take a central role. Online social networking thus can morph into online classified in one direction and blogging in another.

Patterns of personal information revelation are, therefore, quite variable.

First, the pretense of identifiability changes across different types of sites. The use of real names to (re)present an account profile to the rest of the online community may be *encouraged* (through technical specifications, registration requirements, or social norms) in college websites like the Facebook, that aspire to connect participants' profiles to their public identities. The use of real names may be tolerated but filtered in dating/connecting sites like Friendster, that create a thin shield of weak pseudonymity between the public identity of a person and her online persona by making only the first name of a participant visible to others, and not her last name. Or, the use of real names and personal contact information could be openly *discouraged*, as in pseudonymous-based dating websites like Match.com, that attempt to protect the public identity of a person by making its linkage to the online persona more difficult. However, notwithstanding the different approaches to identifiability, most sites encourage the publication of personal and identifiable personal photos (such as clear shots of a person's face).

Second, the type of information revealed or elicited often orbits around hobbies and interests, but can stride from there in different directions. These include: semi-public information such as current and previous schools and employers (as in Friendster); private information such as drinking and drug habits and sexual preferences and orientation (as in Nerve Personals); and open-ended entries (as in LiveJournal).

Third, visibility of information is highly variable. In certain sites (especially the ostensibly pseudonymous ones) any member may view any other member's profile. On weaker-pseudonym sites, access to personal information may be limited to participants that are part of the direct or extended network of the profile owner. Such visibility tuning controls become even more refined on sites which make no pretense of pseudonymity, like the Facebook.

And yet, across different sites, anecdotal evidence suggests that participants are happy to disclose as much information as possible to as many people as possible. It is not unusual to find profiles on sites like Friendster or Salon Personals that list their owners' personal email addresses (or link to their personal websites), in violation of the recommendation or requirements of the hosting service itself. In the next subsection, we resort to the theory of social networks to frame the analysis of such behavior, which we then investigate empirically in Section 3.

2.1 Social Network Theory and Privacy

The relation between privacy and a person's social network is multi-faceted. In certain occasions we want information about ourselves to be known only by a small circle of close friends, and not by strangers. In other instances, we are willing to reveal personal information to anonymous strangers, but not to those who know us better.

Social network theorists have discussed the relevance of relations of different depth and strength in a person's social network (see [11], [12]) and the importance of so-called weak ties in the flow of information across different nodes in a network. Network theory has also been used to explore how distant nodes can get interconnected through relatively few random ties (e.g., [20], [21], [32]). The privacy relevance of these arguments has recently been highlighted by Strahilevitz in [27].

Strahilevitz has proposed applying formal social network theory as a tool for aiding interpretation of privacy in legal cases. He suggests basing conclusions regarding privacy "on

[7] `Http://www.facebook.com/` .
[8] `Http://www.match.com/` .
[9] `Http://personals.nerve.com/` .
[10] `Http://personals.salon.com/` .
[11] `Http://www.livejournal.com/` .

what the parties should have expected to follow the initial disclosure of information by someone other than the defendant" (*op cit*, p. 57). In other words, the consideration of how information is expected to flow from node to node in somebody's social network should also inform that person's expectations for privacy of information revealed in the network.

However, the application of social network theory to the study of information revelation (and, implicitly, privacy choices) in online social networks highlights significant differences between the offline and the online scenarios.

First, offline social networks are made of ties that can only be loosely categorized as weak or strong ties, but in reality are extremely diverse in terms of how close and intimate a subject perceives a relation to be. Online social networks, on the other side, often reduce these nuanced connections to simplistic binary relations: "Friend or not" [8]. Observing online social networks, Danah Boyd notes that "there is no way to determine what metric was used or what the role or weight of the relationship is. While some people are willing to indicate anyone as Friends, and others stick to a conservative definition, most users tend to list anyone who they know and do not actively dislike. This often means that people are indicated as Friends even though the user does not particularly know or trust the person" [8] (p. 2).

Second, while the number of strong ties that a person may maintain on a social networking site may not be significantly increased by online networking technology, Donath and Boyd note that "the number of weak ties one can form and maintain may be able to increase substantially, because the type of communication that can be done more cheaply and easily with new technology is well suited for these ties" [9] (p. 80).

Third, while an offline social network may include up to a dozen of intimate or significant ties and 1000 to 1700 "acquaintances" or "interactions" (see [9] and [27]), an online social networks can list hundreds of direct "friends" and include hundreds of thousands of additional friends within just three degrees of separation from a subject.

This implies online social networks are both vaster and have more weaker ties, on average, than offline social networks. In other words, thousands of users may be classified as friends of friends of an individual and become able to access her personal information, while, at the same time, the threshold to qualify as friend on somebody's network is low. This may make the online social network only an imaginary (or, to borrow Anderson's terminology, an *imagined*) community (see [2]). Hence, trust in and within online social networks may be assigned differently and have a different meaning than in their offline counterparts. Online social networks are also more levelled, in that the same information is provided to larger amounts of friends connected to the subject through ties of different strength. And here lies a paradox. While privacy may be considered conducive to and necessary for intimacy (for [10], intimacy resides in selectively revealing private information to certain individuals, but not to others), trust may decrease within an online social network. At the same time, a new form of intimacy becomes widespread: the sharing of personal information with large and potential unknown numbers of friends and strangers altogether. The ability to meaningfully interact with others is mildly augmented, while the ability of others to access the person is significantly enlarged. It remains to be investigated how similar or different are the mental models people apply to personal information revelation within a traditional network of friends compared to those that are applied in an online network.

2.2 Privacy Implications

Privacy implications associated with online social networking depend on the level of identifiability of the information provided, its possible recipients, and its possible uses. Even social networking websites that do not openly expose their users' identities may provide enough information to identify the profile's owner. This may happen, for example, through face re-identification [13]. Liu and Maes estimate in [18] a 15% overlap in 2 of the major social networking sites they studied. Since users often re-use the same or similar photos across different sites, an identified face can be used to identify a pseudonym profile with the same or similar face on another site. Similar re-identifications are possible through demographic data, but also through category-based representations of interests that reveal unique or rare overlaps of hobbies or tastes. We note that information revelation can work in two ways: by allowing another party to identify a pseudonymous profile through previous knowledge of a subject's characteristics or traits; or by allowing another party to infer previously unknown characteristics or traits about a subject identified on a certain site. We present evaluations of the probabilities of success of these attacks on users of a specific networking site in Section 4.

To whom may identifiable information be made available? First of all, of course, the hosting site, that may use and extend the information (both knowingly and unknowingly revealed by the participant) in different ways (below we discuss extracts from the privacy policy of a social networking site that are relevant to this discussion). Obviously, the information is available within the network itself, whose extension in time (that is, data durability) and space (that is, membership extension) may not be fully known or knowable by the participant. Finally, the easiness of joining and extending one's network, and the lack of basic security measures (such as SSL logins) at most networking sites make it easy for third parties (from hackers to government agencies) to access participants data without the site's direct collaboration (already in 2003, LiveJournal used to receive at least five reports of ID hijacking per day, [23]).

How can that information be used? It depends on the information actually provided - which may, in certain cases, be very extensive and intimate. Risks range from identity theft to online and physical stalking; from embarrassment to price discrimination and blackmailing. Yet, there are some who believe that social networking sites can also offer the solution to online privacy problems. In an interview, Tribe.net CEO Mark Pincus noted that "[s]ocial networking has the potential to create an intelligent order in the current chaos by letting you manage how public you make yourself and why and who can contact you." [4]. We test this position in Section 4.

While privacy may be at risk in social networking sites, information is willingly provided. Different factors are likely to drive information revelation in online social networks. The list includes signalling (as discussed in [9]), because the perceived benefit of selectively revealing data to strangers may appear larger than the perceived costs of possible privacy invasions; peer pressure and herding behavior; relaxed

attitudes towards (or lack of interest in) personal privacy; incomplete information (about the possible privacy implications of information revelation); faith in the networking service or trust in its members; myopic evaluation of privacy risks (see [1]); or also the service's own user interface, that may drive the unchallenged acceptance of permeable default privacy settings.

We do not attempt to ascertain the relative impact of different drivers in this paper. However, in the following sections we present data on actual behavioral patterns of information revelation and inferred privacy attitudes in a college-targeted networking site. This investigation offers a starting point for subsequent analysis of the motivations behind observed behaviors.

3. THE FACEBOOK.COM

Many users of social networking sites are of college age [8], and recent ventures have started explicitly catering to the college crowd and, in some cases, to specific colleges (e.g., the Facebook.com, but also Universitysingles.ca, quad5.com, CampusNetwork.com, iVentster.com, and others).

College-oriented social networking sites provide opportunities to combine online and face-to-face interactions within an ostensibly bounded domain. This makes them different from traditional networking sites: they are communities based "on a shared real space" [26]. This combination may explain the explosive growth of some of these services (according to [26], the Facebook has spread "to 573 campuses and 2.4 million users. [...] [I]t typically attracts 80 percent of a school's undergraduate population as well as a smattering of graduate students, faculty members, and recent alumni.") Also because of this, college-oriented networks offer a wealth of personal data of potentially great value to external observers (as reported by [6], for example, the Pentagon manages a database of 16-to-25-year-old US youth data, containing around 30 million records, and continuously merged with other data for focused marketing).

Since many of these sites require a college's email account for a participant to be admitted to the online social network of that college, expectations of validity of certain personal information provided by others on the network may increase. Together with the apparent sharing of a physical environment with other members of the network, that expectation may increase the sense of trust and intimacy across the online community. And yet, since these services can be easily accessed by outsiders (see Section 4) and since members can hardly control the expansion of their own network (often, a member's network increases also through the activity of other members), such communities turn out to be more *imagined* than real, and privacy expectations may not be matched by privacy reality.

The characteristics mentioned above make college-oriented networking sites intriguing candidates for our study of information revelation and privacy preferences. In the rest of this paper we analyze data gathered from the network of Carnegie Mellon University (CMU) students enlisted on one of such sites, the Facebook.

The Facebook has gained huge adoption within the CMU student community but is present with similar success at many other colleges nationwide. It validates CMU-specific network accounts by requiring the use of CMU email addresses for registration and login. Its interface grants participants very granular control on the searchability and visibility of their personal information (by friend or location, by type of user, and by type of data). The default settings, however, are set to make the participants profile *searchable* by anybody else in any school in the Facebook network, and make its actual content *visible* to any other user at the same college or at another college in the same physical location.[12]

The Facebook is straightforward about the usage it plans for the participants' personal information: at the time of this writing, its privacy policy [30] reports that the site *will* collect additional information about its users (for instance, from instant messaging), not originated from the use of the service itself. The policy also reports that participants' information may include information that the participant has not knowingly provided (for example, her IP address), and that personal data may be shared with third parties.

3.1 Access Tools

In June 2005, we separately searched for all "female" and all "male" profiles for CMU Facebook members using the website's advanced search feature and extracted their profile IDs. Using these IDs we then downloaded a total of 4540 profiles - virtually the entire CMU Facebook population at the time of the study.

3.2 Demographics

The majority of users of the Facebook at CMU are undergraduate students (3345 or 73.7% of all profiles; see Table 1). This corresponds to 62.1% of the total undergraduate population at CMU [31]. Graduate students, staff and faculty are represented to a much lesser extent (6.3%, 1.3%, and 1.5% of the CMU population, respectively). The majority of users is male (60.4% vs. 39.2%). Table 2 shows the gender distribution for the different user categories. The strong dominance of undergraduate users is also reflected in the user age distribution shown in Figure 1. The vast majority of users (95.6%) falls in the 18-24 age bracket. Overall the average age is 21.04 years.

[12] At the time of writing, the geography feature which generates networks based on physical location is by default not available to undergraduate students. However, the status of a profile can easily be changed to e.g. "graduate student" for which the feature is accessible.

Table 1: Distribution of CMU Facebook profiles for different user categories. The majority of users are undergraduate students. The table lists the percentage of the CMU population (for each category) that are users of the Facebook (if available).

	# Profiles	% of Facebook Profiles	% of CMU Population
Undergraduate Students	3345	74.6	62.1
Alumni	853	18.8	-
Graduate Students	270	5.9	6.3
Staff	35	0.8	1.3
Faculty	17	0.4	1.5

Table 2: Gender distribution for different user categories.

		# Profiles	% of Category	% of CMU Population
Overall	Male	2742	60.4	-
	Female	1781	39.2	-
Undergraduate Students	Male	2025	60.5	62.0
	Female	1320	39.5	62.3
Alumni	Male	484	56.7	-
	Female	369	43.3	-
Graduate Students	Male	191	70.7	6.3
	Female	79	29.3	6.3
Staff	Male	23	65.7	-
	Female	12	34.3	-
Faculty	Male	17	100	3.4
	Female	0	0.0	0.0

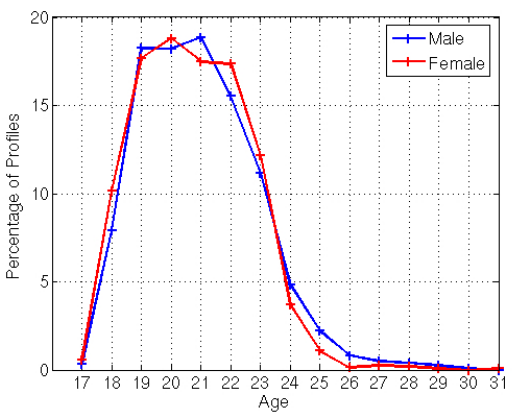

Figure 1: Age distribution of Facebook profiles at CMU. The majority of users (95.6%) falls into the 18-24 age bracket.

3.3 Types and Amount of Information Disclosed

The Facebook offers users the ability to disclose a large and varied amount of personal information. We evaluated to which extent users at CMU provide personal information. Figure 2 shows the percentages of CMU profiles that disclose different categories of information.

In general, CMU users of the Facebook provide an astonishing amount of information: 90.8% of profiles contain an image, 87.8% of users reveal their birth date, 39.9% list a phone number (including 28.8% of profiles that contain a cellphone number), and 50.8% list their current residence. The majority of users also disclose their dating preferences (male or female), current relationship status (single, married, or in a relationship), political views (from "very liberal" to "very conservative"), and various interests (including music, books, and movies). A large percentage of users (62.9%) that list a relationship status other than single even identify their partner by name and/or link to their Facebook profile.

Note that, as further discussed below in Section 3.4, Facebook profiles tend to be fully identified with each participant's real first and last names, both of which are used as the profile's name. In other words, whoever views a profile is also able to connect the real first and last name of a person to the personal information provided - that may include birthday or current residence.

Across most categories, the amount of information revealed by female and male users is very similar. A notable exception is the phone number, disclosed by substantially more male than female users (47.1% vs. 28.9%). *Single* male users tend to report their phone numbers in even higher frequencies, thereby possibly *signalling* their elevated interest in making a maximum amount of contact information easily available.

Additional types of information disclosed by Facebook users (such as the membership of one's own network of friends at the home college or elsewhere, last login information, class schedule, and others) are discussed in the rest of this paper.

3.4 Data Validity and Data Identifiability

The terms of service of the site encourage users to only

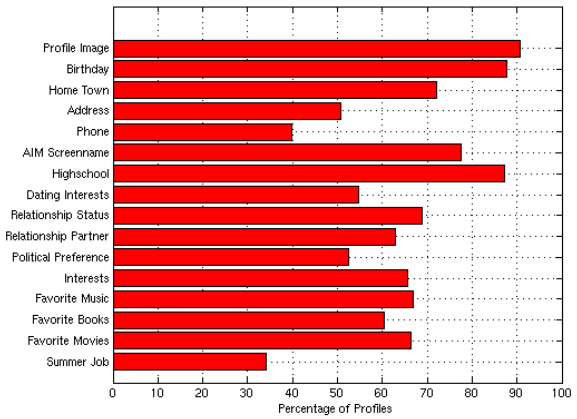

Figure 2: Percentages of CMU profiles revealing various types of personal information.

publish profiles that directly relate to them and not to other entities, people or fictional characters. In addition, in order to sign up with the Facebook a valid email address of one of the more than 500 academic institutions that the site covers has to be provided. This requirement, along with the site's mission of organizing the real life social networks of their members, provides incentives for users to only publish accurate information.

We tested how valid the published data appears to be. In addition, we studied how identifiable or granular the provided data is.

In general, determining the accuracy of the information provided by users on the Facebook (or any other social networking website) is nontrivial for all but selected individual cases. We therefore restrict our validity evaluation to the measurement of the manually determined *perceived* accuracy of information on a randomly selected subset of 100 profiles.

3.4.1 Profile Names

We manually categorized the names given on Facebook profiles as being one of the the following:

1. Real Name
   Name appears to be real.

2. Partial Name
   Only a first name is given.

3. Fake Name
   Obviously fake name.

Table 3 shows the results of the evaluation. We found 89% of all names to be realistic and likely the true names for the users (for example, can be matched to the visible CMU email address provided as login), with only 8% of names obviously fake. The percentage of people that choose to only disclose their first name was very small: 3%.

In other words, the vast majority of Facebook users seem to provide their fully identifiable names, although they are not forced to do so by the site itself.

As comparison, 98.5% of the profiles that include a birthday actually report the *fully identified* birth date (day, month, and year), although, again, users are not forced to provide

Table 3: Categorization of name quality of a random subset of 100 profile names from the Facebook. The vast majority of names appear to be real names with only a very small percentage of partial or obviously fake names.

Category	Percentage Facebook Profiles
Real Name	89%
Partial Name	3%
Fake Name	8%

the complete information (the remaining 1.5% of users reported only the month or the month and day but not the year of birth). Assessing the validity of birth dates is not trivial. However, in certain instances we observed friends posting birthday wishes in the comments section of the profile of a user on the day that had been reported by the user as her birthday. In addition, the incentives to provide a fake birth date (rather than not providing one at all, which is permitted by the system) would be unclear.

3.4.2 Identifiability of Images on Profile

The vast majority of profiles contain an image (90.8%, see Section 3.3). While there is no explicit requirement to provide a facial image, the majority of users do so. In order to assess the quality of the images provided we manually labelled them into one of four categories:

1. Identifiable
   Image quality is good enough to enable person recognition.

2. Semi-Identifiable
   The profile image shows a person, but due to the image composition or face pose the person is not directly recognizable. Other aspects however (e.g. hair color, body shape, etc.) are visible.

3. Group Image
   The image contains more than one face and no other profile information (e.g. gender) can be used to identify the user in the image.

4. Joke Image
   Images clearly not related to a person (e.g. cartoon or celebrity image).

Table 4 shows the results of labelling the profile images into the four categories. In the majority of profiles the images are suitable for direct identification (61%). Overall, 80% of images contain at least some information useful for identification. Only a small subset of 12% of all images are clearly not related to the profile user. We repeated the same evaluation using 100 randomly chosen images from Friendster, where the profile name is only the first name of the member (which makes Friendster profiles not as identifiable as Facebook ones). Here the percentage of "joke images" is much higher (23%) and the percentage of images suitable for direct identification lower (55%).[13]

[13]We note that Friendster's profiles used to be populated by numerous fake and/or humorous profiles, also called "Fakesters" (see [8]). Friendster management tried to elim-

Table 4: Categorization of user identifiability based on manual evaluation of a randomly selected subset of 100 images from both Facebook and Friendster profiles. Images provided on Facebook profiles are in the majority of cases suitable for direct identification (61%). The percentage of images obviously unrelated to a person ("joke image") is much lower for Facebook images in comparison to images on Friendster profiles (12% vs. 23%).

Category	Percentage Facebook Profiles	Percentage Friendster Profiles
Identifiable	61%	55%
Semi-Identifiable	19%	15%
Group Image	8%	6%
Joke Image	12%	23%

3.4.3 Friends Networks

The Facebook helps in organizing a real-life social network online. Since Facebook users interact with many of the other users directly in real-life, often on a daily basis, the network of friends may function as profile fact checker, potentially triggering questions about obviously erroneous information. Facebook users typically maintain a very large network of friends. On average, CMU Facebook users list 78.2 friends at CMU and 54.9 friends at other schools. 76.6% of users have 25 or more CMU friends, whereas 68.6% of profiles show 25 or more non-CMU friends. See Figure 3 for histogram plots of the distribution of sizes of the networks for friends at CMU and elsewhere. This represents some effort, since adding a friend requires explicit confirmation.

3.5 Data Visibility and Privacy Preferences

For any user of the Facebook, other users fall into four different categories: friends, friends of friends, non-friend users at the same institution and non-friend users at a different institution.[14] By default, everyone on the Facebook appears in searches of everyone else, independent of the searchers institutional affiliation. In search results the users' full names (partial searches for e.g. first names are possible) appear along with the profile image, the academic institution that the user is attending, and the users' status there. The Facebook reinforces this default settings by labelling it "recommended" on the privacy preference page. Also by default the full profile (including contact information) is visible to everyone else at the same institution.

Prior research in HCI has shown that users tend to not change default settings [19]. This makes the choice of default settings by website operators very important. On the other hand, the site provides users a very granular and relatively sophisticated interface to control the searchability and visibility of their profiles. Undergrad users, for example, can make their profiles searchable only to other undergrad users, or only users who are friends, or users who are friends of friends, or users at the same institution - or combinations of the above constraints. In addition, *visibility* of the entire profile can be similarly controlled. Granular control on contact information is also provided.

inate fake profiles and succeeded in significantly reducing their number, but not completely extirpating them from the network. Based on our manual calculations, the share of fake Friendster profiles is currently comparable to the share of fake Facebook profiles reported above.

[14]The Facebook recently introduced a new relationship category based on user location, e.g. Pittsburgh, which we did not consider in this study.

Sociological theories of privacy have noted how an individual may selectively disclose personal information to others in order to establish different degrees of trust and intimacy with them (see [10]). In light of these theories, we tested how much CMU Facebook users take advantage of the ability the site provides to manage their presentation of sel[ves]. By creating accounts at different institutions, and by using accounts with varying degree of interconnectedness with the rest of the CMU network, we were able to infer how individual users within the CMU network were selecting their own privacy preference.

3.5.1 Profile Searchability

We first measured the percentage of users that changed the search default setting away from being searchable to everyone on the Facebook to only being searchable to CMU users. We generated a list of profile IDs currently in use at CMU and compared it with a list of profile IDs visible from a different academic institution. We found that only 1.2% of users (18 female, 45 male) made use of this privacy setting.

3.5.2 Profile Visibility

We then evaluated the number of CMU users that changed profile visibility by restricting access to CMU users. We used the list of profile IDs currently in use at CMU and evaluated which percentage of profiles were fully accessible to an unconnected user (not friend or friend of friend of any profile). Only 3 profiles (0.06%) in total did not fall into this category.

3.5.3 Facebook Data Access

We can conclude that only a vanishingly small number of users change the (permissive) default privacy preferences. In general, fully identifiable information such as personal image and first and last name is available to anybody registered at any Facebook member network. Since the Facebook boasts a 80% average participation rate among undergraduate students at the hundreds of US institutions it covers, and since around 61% of our CMU subset provides identifiable face images, it is relatively easy for anybody to gain access to these data, and cheap to store a nation-wide database of fully identified students and their IDs. In other words, information suitable for creating a brief digital dossier consisting of name, college affiliation, status and a profile image can be accessed for the vast majority of Facebook users by anyone on the website. (To demonstrate this we downloaded and identified the same information for a total of 9673 users at Harvard University.)

Additional personal data - such as political and sexual ori-

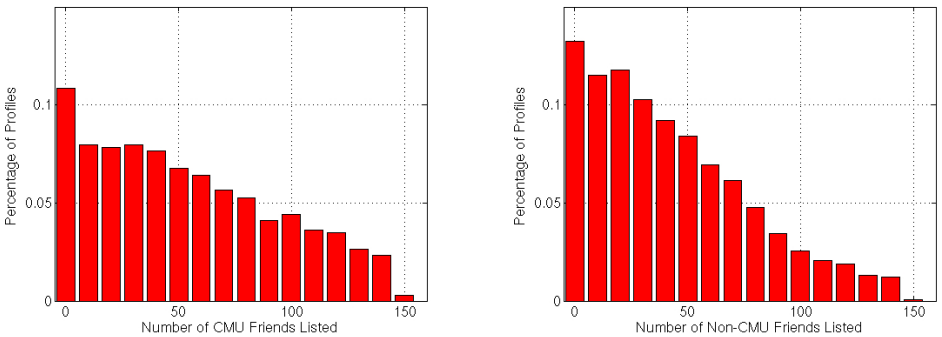

(a) Network of CMU friends (b) Network of Non-CMU friends

Figure 3: Histogram of the size of networks for both CMU friends (a) and non-CMU friends (b). Users maintain large networks of friends with the average user having 78.2 friends at CMU and 54.9 friends elsewhere.

entation, residence address, telephone number, class schedule, etc. - are made available by the majority of users to anybody else at the same institution, leaving such data accessible to any subject able to obtain even temporary control of an institution's single email address.

4. PRIVACY IMPLICATIONS

It would appear that the population of Facebook users we have studied is, by large, quite oblivious, unconcerned, or just pragmatic about their personal privacy. Personal data is generously provided and limiting privacy preferences are sparingly used. Due to the variety and richness of personal information disclosed in Facebook profiles, their visibility, their public linkages to the members' real identities, and the scope of the network, users may put themselves at risk for a variety of attacks on their physical and online persona. Some of these risks are common also in other online social networks, while some are specific to the Facebook. In this section we outline a number of different attacks and quantify the number of users susceptible based on the data we extracted. See Table 5 for an overview.

4.1 Stalking

Using the information available on profiles on the Facebook a potential adversary (with an account at the same academic institution) can determine the likely physical location of the user for large portions of the day. Facebook profiles include information about residence location, class schedule, and location of last login. A students' life during college is mostly dominated by class attendance. Therefore, knowledge of both the residence and a few classes that the student is currently attending would help a potential stalker to determine the users whereabouts. In the CMU population 860 profiles fall into our definition of this category (280 female, 580 male), in that they disclose both their current residence and at least 2 classes they are attending. Since our study was conducted outside of the semester (when many students might have deleted class information from their profiles) we speculate this number to be even higher during the semester.

A much larger percentage of users is susceptible to a form of cyber-stalking using the AOL instant messenger (AIM). Unlike other messengers, AIM allows users to add "buddies" to their list without knowledge of or confirmation from the buddy being added. Once on the buddy list the adversary can track when the user is online. In the CMU population 77.7% of all profiles list an AIM screen name for a total of more than 3400 users.

4.2 Re-identification

Data re-identification typically deals with the linkage of datasets *without* explicit identifiers such as name and address to datasets *with* explicit identifiers through common attributes [25]. Examples include the linkage of hospital discharge data to voter registration lists, that allows to re-identify sensitive medical information [28].

4.2.1 Demographics re-identification

It has been shown previously that a large portion of the US population can be re-identified using a combination of 5-digit ZIP code, gender, and date of birth [29]. The vast majority of CMU users disclose both their full birthdate (day and year) and gender on their profiles (88.8%). For 44.3% of users (total of 1676) the combination of birthdate and gender is unique within CMU. In addition, 50.8% list their current residence, for which ZIP codes can be easily obtained. Overall, 45.8% of users list birthday, gender, and current residence. An adversary with access to the CMU section of the Facebook could therefore link a comparatively large number of users to outside, de-identified data sources such as e.g. hospital discharge data.

4.2.2 Face Re-Identification

In a related study we were able to correctly link facial images from Friendster profiles without explicit identifiers with images obtained from fully identified CMU web pages using a commercial face recognizer [13]. The field of automatic face recognition has advanced tremendously over the last decade and is now offering a number of commercial solutions which have been shown to perform well across a wide range of imaging conditions [14, 17, 24]. As shown in Section 3.4 a large number of profiles contain high quality images. At CMU more than 2500 profiles fall in this category [15]. Potential de-identified data sources include other social networking sites (e.g. Friendster) or dating sites (e.g. Match.com) that typically host anonymous profiles.

[15]In fact, 90.8% of profiles have images, out of which 61% are estimated to be of sufficient quality for re-identification.

Table 5: Overview of the privacy risks and number of CMU profiles susceptible to it.

Risk	# CMU Facebook Profiles	% CMU Facebook Profiles
Real-World Stalking	280 (Female) 580 (Male)	15.7 (Female) 21.2 (Male)
Online Stalking	3528	77.7
Demographics Re-Identification	1676	44.3
Face Re-Identification	2515 (estimated)	55.4

4.2.3 Social Security Numbers and Identity Theft

An additional re-identification risk lies in making birthdate, hometown, current residence, and current phone number publicly available at the same time. This information can be used to estimate a person's social security number and exposes her to identity theft.

The first three digits of a social security number reveal where that number was created (specifically, the digits are determined by the ZIP code of the mailing address shown on the application for a social security number). The next two digits are group identifiers, which are assigned according to a peculiar but predictable temporal order. The last four digits are progressive serial numbers.[16]

When a person's hometown is known, the window of the first three digits of her SNN can be identified with probability decreasing with the home state's populousness. When that person's birthday is also known, and an attacker has access to SSNs of other people with the same birthdate in the same state as the target (for example obtained from the SSN death index or from stolen SSNs), it is possible to pin down a window of values in which the two middle digits are likely to fall. The last four digits (often used in unprotected logins and as passwords) can be retrieved through social engineering. Since the vast majority of the Facebook profiles we studied not only include birthday and hometown information, but also current phone number and residence (often used for verification purposes by financial institutions and other credit agencies), users are exposing themselves to substantial risks of identity theft.

4.3 Building a Digital Dossier

The privacy implications of revealing personal and sensitive information (such as sexual orientation and political views) may extend beyond their immediate impact, which can be limited. Given the low and decreasing costs of storing digital information, it is possible to continuously monitor the evolution of the network and its users' profiles, thereby building a digital dossier for its participants. College students, even if currently not concerned about the visibility of their personal information, may become so as they enter sensitive and delicate jobs a few years from now - when the data currently mined could still be available.

4.4 Fragile Privacy Protection

One might speculate that the *perceived* privacy protection of making personal information available only to members of a campus community may increase Facebook users' willingness to reveal personal information. However, the mechanisms protecting this social network can be circumvented. Adding to this the recognition that users have little control on the composition of their own networks (because often a member's friend can introduce strangers into that member's network), one may conclude that the personal information users are revealing even on sites with access control and managed search capabilities effectively becomes *public* data.

4.4.1 Fake Email Address

The Facebook verifies users as legitimate members of a campus community by sending a confirmation email containing a link with a seemingly randomly generated nine digit code to the (campus) email address provided during registration. Since the process of signing up and receiving the confirmation email only takes minutes, an adversary simply needs to gain access to the campus network for a very short period of time. This can be achieved in a number of well-known ways, e.g. by attempting to remotely access a hacked or virus-infected machine on the network or physically accessing a networked machine in e.g. the library, etc.

4.4.2 Manipulating Users

Social engineering is a well-known practice in computer security to obtain confidential information by manipulating legitimate users [22]. Implementation of this practice on the Facebook is very simple: just ask to be added as someone's friend. The surprisingly high success rate of this practice was recently demonstrated by a Facebook user who, using an automatic script, contacted 250,000 users of the Facebook across the country and asked to be added as their friend. According to [15], 75,000 users accepted: thirty percent of Facebook users are willing to make all of their profile information available to a random stranger and his network of friends.

4.4.3 Advanced Search Features

While not directly linked to from the site, the Facebook makes the advanced search page of *any* college available to *anyone* in the network. Using this page various profile information can be searched for, e.g. relationship status, phone number, sexual preferences, political views and (college) residence. By keeping track of the profile IDs returned in the different searches a significant portion of the previously inaccessible information can be reconstructed.

5. CONCLUSIONS

Online social networks are both vaster and looser than their offline counterparts. It is possible for somebody's profile to be connected to hundreds of peers directly, and thousands of others through the network's ties. Many individuals in a person's online extended network would hardly be defined as actual friends by that person; in fact many may be complete strangers. And yet, personal and often sensitive information is freely and publicly provided.

[16]See `http://www.ssa.gov/foia/stateweb.html` and `http://policy.ssa.gov/poms.nsf/lnx/0100201030`.

In our study of more than 4,000 CMU users of the Facebook we have quantified individuals' willingness to provide large amounts of personal information in an online social network, and we have shown how unconcerned its users appear to privacy risks: while personal data is generously provided, limiting privacy preferences are hardly used; only a small number of members change the default privacy preferences, which are set to maximize the visibility of users profiles. Based on the information they provide online, users expose themselves to various physical and cyber risks, and make it extremely easy for third parties to create digital dossiers of their behavior.

These risks are not unique to the Facebook. However, the Facebook's public linkages between an individual profile and the real identity of its owner, and the Facebook's perceived connection to a physical and ostensibly bounded community (the campus), make Facebook users a particularly interesting population for our research.

Our study quantifies patterns of information revelation and infers usage of privacy settings from actual field data, rather than from surveys or laboratory experiments. Still, the relative importance of the different drivers influencing Facebook users' information revelation behavior has to be quantified. Our evidence is compatible with a number of different hypotheses. In fact, many simultaneous factors are likely to play a role. Some evidence is compatible with a signalling hypothesis (see Section 3.3): users may be pragmatically publishing personal information because the benefits they expect from public disclosure surpass its perceived costs. Yet, our evidence is also compatible with an interface design explanation, such as the acceptance (and possibly ignorance) of the default, permeable settings (see Section 3.5). Peer pressure and herding behavior may also be influencing factors, and so also myopic privacy attitudes (see [1]) and the sense of protection offered by the (perceived) bounds of a campus community. Clarifying the role of these different factors is part of our continuing research agenda.

Acknowledgements

We would like to thank Anne Zimmerman and Bradley Malin for first bringing the Facebook to our attention. We would also like to thank danah boyd, Lorrie Cranor, Julia Gideon, Charis Kaskiris, Steven Frank, Bart Nabbe, Mike Shamos, Irina Shklovski, and four anonymous referees for comments. This work was supported in part by the Data Privacy Lab in the School of Computer Science and by the Berkman Fund at Carnegie Mellon University.

6. REFERENCES

[1] A. Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the ACM Conference on Electronic Commerce (EC '04)*, pages 21–29, 2004.

[2] B. Anderson. *Imagined Communities: Reflections on the Origin and Spread of Nationalism*. Verso, London and New York, revised edition, 1991.

[3] S. Arrison. Is Friendster the new TIA? *TechCentralStation*, January 7, 2004.

[4] J. Black. The perils and promise of online schmoozing. *BusinessWeek Online*, February 20, 2004.

[5] J. Brown. Six degrees to nowhere. *Salon.com*, September 21, 1998.

[6] D. Cave. 16 to 25? Pentagon has your number, and more. *The New York Times*, June 24, 2005.

[7] d. boyd. Reflections on friendster, trust and intimacy. In *Intimate (Ubiquitous) Computing Workshop - Ubicomp 2003, October 12-15, Seattle, Washington, USA*, 2003.

[8] d. boyd. Friendster and publicly articulated social networking. In *Conference on Human Factors and Computing Systems (CHI 2004), April 24-29, Vienna, Austria*, 2004.

[9] J. Donath and d. boyd. Public displays of connection. *BT Technology Journal*, 22:71–82, 2004.

[10] S. Gerstein. Intimacy and privacy. In F. D. Schoeman, editor, *Philosophical Dimensions of Privacy: An Anthology*. Cambridge University Press, Cambridge, UK, 1984.

[11] M. Granovetter. The strength of weak ties. *American Journal of Sociology*, 78:1360–1380, 1973.

[12] M. Granovetter. The strength of weak ties: A network theory revisited. *Sociological Theory*, 1:201–233, 1983.

[13] R. Gross. Re-identifying facial images. Technical report, Carnegie Mellon University, Institute for Software Research International, 2005. In preparation.

[14] R. Gross, J. Shi, and J. Cohn. Quo vadis face recognition? In *Third Workshop on Empirical Evaluation Methods in Computer Vision*, 2001.

[15] K. Jump. A new kind of fame. *The Columbian Missourian*, September 1, 2005.

[16] A. Leonard. You are who you know. *Salon.com*, June 15, 2004.

[17] S. Li and A. Jain, editors. *Handbook of Face Recognition*. Springer Verlag, 2005.

[18] H. Liu and P. Maes. Interestmap: Harvesting social network profiles for recommendations. In *Beyond Personalization - IUI 2005, January 9, San Diego, California, USA*, 2005.

[19] W. Mackay. Triggers and barriers to customizing software. In *Proceedings of CHI'91*, pages 153–160. ACM Press, 1991.

[20] S. Milgram. The small world problem. *Psychology Today*, 6:62–67, 1967.

[21] S. Milgram. The familiar stranger: An aspect of urban anonymity. In S. Milgram, J. Sabini, and M. Silver, editors, *The Individual in a Social World: Essays and Experiments*. Addison-Wesley, Reading, MA, 1977.

[22] K. Mitnick, W. Simon, and S. Wozniak. *The art of deception: controlling the human element of security*. John Wiley & Sons, 2002.

[23] A. Newitz. Defenses lacking at social network sites. *SecurityFocus*, December 31, 2003.

[24] P. Phillips, P. Flynn, T. Scruggs, K. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and J. Worek. Overview of the face recognition grand challenge. In *IEEE Conference on Computer Vision and Pattern Recognition, June 20-25, San Diego, California, USA*, 2005.

[25] P. Samarati and L. Sweeney. Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and cell suppression. Technical report, SRI International, 1998.

[26] I. Sege. Where everybody knows your name.

Boston.com, April 27, 2005.

[27] L. J. Strahilevitz. A social networks theory of privacy. The Law School, University of Chicago, John M. Olin Law & Economics Working Paper No. 230 (2D Series), December 2004.

[28] L. Sweeney. k-Anonymity: a model for protecting privacy. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, 10(5):557–570, 2002.

[29] L. Sweeney. Uniqueness of simple demographics in the U.S. populaiton. Technical report, Carnegie Mellon University, Laboratory for International Data Privacy, 2004.

[30] The Facebook. Privacy policy. http://facebook.com/policy.php, August 2005.

[31] University Planning. Carnegie Mellon Factbook 2005. Carnegie Mellon University, February 2005.

[32] D. Watts. *Six Degrees: The Science of a Connected Age*. W.W.Norton & Company, 2003.