Area 3 (Security and Privacy)

Danny Yoo (dyoo@cs.wpi.edu)

Abstract

Imagine you are designing a Web-based system that enables sharing and communication between users who are as young as ten years of age. What are the security and privacy considerations of supporting users that young (in contrast to supporting adults of consenting age)? What techniques are available to help with implementing protections for these young users and, crucially, what necessary and important techniques do you feel are missing? Consider both technical and non-technical issues in your answer.

1 Introduction

Online social networks, like Facebook and MySpace, allow members to share themselves to the world, be it through gossip, photos or video, or through creative actions like writing stories, fan-fiction, or even video games. A key element to these web sites is the ability for users to share these *gifts* [1] with one another. By enabling gift-giving, web sites allow people to perform acts of self-affirmation, cooperation, and friendship.

Networks form in these social networks through relationships of topic and societal status. Children have been notable in their increasing use of online social networks. The Kaiser Foundation, for example, reports an up-swing in adolescent use of social networks: online social networking is the most popular computer activity for children [13]. This group, though, poses a particular complication for web sites that permit gift-giving. Gifts can reveal something about oneself, including real-life information like one's address or identity. Children, generally speaking, have an underdeveloped judgment and empathy, which can result in risky or hurtful actions. They are also susceptible to sexual predators, peer harassment (cyberbullying), and exposure to indecent material. Girls may be particularly susceptible in online environments. According to a study by Berson and Berson, a survey of adolescent girls reveals that they will willingly give personal information through forms and online questions, send pictures of themselves, and arrange offline meetings. In the US, 59% willingly give personal information through a form or online questionnaire, 23% disclose their picture on request, and 12.5% have gone on to face-to-face meetings [9].

1.1 The susceptibility of online communication

Part of the reason why children are so susceptible are due to their lack of developed judgment: they may not realize that certain gifts are dangerous to give to certain people. The nature of online web interactions, though, makes online web sites riskier than real life:

- Berson and Berson hypothesize that the moving imagery of multimedia web sites, like those of television advertisements, encourages impulsive behavior instead of reasoned action.
- Online communication can lack the kinds of sensory cues that mark risky behavior. For example, the audio and visual cues for communication with a stranger look just the same in email as does communication with a close friend.
- Finally, children may not realize the implications of digital data, that gifts can be re-shared online: once published, a gift is accessible, not just by one's immediate friends, but potentially with the world.

[1] I use the word "gift" to emphasize the potentially dynamic nature of the content: a gift may be an arbitrary program, for example.

1.2 Legislation

The potential threats and the real vulnerability of children demand some action to guard children. One way to address this is through legislation. Some laws take an extreme, and ban the use of social networking web sites altogether in schools and libraries, like the Safe School Act HB 5941 of Rhode Island [3] or the Deleting Online Predators Act of 2006 [2]. This approach is almost certainly ineffective: it depends on children to willingly limit their computer usage only to those schools and libraries. Children are rebellious, and web site access is ubiquitous; these laws account for neither.

COPPA

Other legislation aim at websites rather than children, and reduce the power of social networking web sites to collect information. The most relevant law, as of this writing, is the the Children's Online Privacy Protection Act (COPPA) [1], which is meant to limit the power of web sites from collecting information from children under the age of 13. It affirms that parents have the ultimate right over their child's information, and enforces a disclosure policy.

A COPPA-conformant web site that caters to children needs to:

- disclose the nature of what information is recorded
- get explicit parental consent to use the web site
- provide parental control to any information about their child, and to delete such data on request
- requires reasonable effort to guard that information
- restricts a web site from coercing information orthogonal to the provided web service

This law does have teeth; the FTC, for example, has applied it to Disney Playdom for a $3 million penalty [5], and Xanga for $1 million [7] for violating COPPA. At the very least, it forces social web sites to consider the data-collection policy of their users.

1.3 The threat

The considered threats (sexual predation, cyberbullying, exposure to obscenity) can be distilled to the following, that the act of gift-giving can allow:

- the sender to reveal their location or identity to another.
- the receiver to receive a hurtful or dangerous gift, a Trojan horse.

This threat model applies generally to people, but is especially acute when the participants include children. Although COPPA restricts the power of web sites from collecting personal information, it focuses only on the danger of sexual predation based on revealing identity and personal information.

However, the threat model includes more than the exposure of identity. The protection guaranteed by COPPA is insufficient because even if a web site is following COPPA to the letter of the law, it still says nothing about dangerous and hurtful actions. For example, a child can send bullying messages to another, and the web site will still be following COPPA, because the content doesn't necessarily contain personal information.

Social networking web sites can be designed to resist these threats. It can passively shape the structure of the social network itself, or actively enforce constraints on gifts to restrict what content can be sent and received, both with automated techniques as well as parental oversight.

2 Techniques

2.1 Limiting the power of introduction

In traditional social networks, users can "friend" those who they either know through some public identification like an email address, or transitively send friend requests through the contacts of their existing friends. Some systems, like Facebook, can take this to an extreme, by allowing users to go through multiple levels of the friendship relation, as if friendship and trust were mathematically transitive. This kind of unrestricted befriending allows even relative strangers to become part of one's social network.

For a network that caters to children, this kind of unrestricted introduction is too powerful. One technique to protect children limits the social network to the pre-existing personal relationships that exist in a child's offline life. An example of this is the text-messaging network of cell phones, which uses the contacts on a phone address book, and whose network structure matches real-world acquaintances. The coupling to the address book means that a child need to know their contact in real life (or at the very least, have the contact's phone number) before being able to send SMS messages or photos.

2.1.1 Friend codes

The general idea of the technique of Friend Codes is to attach some unique identification, a *Friend Code* to a user. A Friend Code is a short, randomly-generated key associated to a user: if a code is known, the social network will allow communication with the owner of the code. Crucially, the social network does not provide any tools to share that Friend Code directly on that network. This forces the user to use a pre-existing, out-of-band channel to establish introductions.

A popular implementation of this technique is the "Friends Code" service that connects uses of the Nintendo game systems [4]. Nintendo's main audience targets adolescents, and its network reflects the needs for maintaining a safe gaming environment, to the frustration of some of its older users.

Some children channel that frustration to undermine the system of friend codes, by publishing the friend code to a public forum. When the key is associated to some information that's considered private, such as a phone number, it's usually clear to a child that leaking one's phone number is a bad decision. However, for Friend Codes, public sharing occurs in practice: it's somewhat common for sophisticated users to use public video gaming messaging forums to exchange codes with one another.

What is missing, and what would make this technique more effective, is an identification and introduction system that can't be shared to mass audiences. A hardware approach, such as a dongle or electronic id card, may be designed that implements an identity that makes in-person introduction easy, and online introduction difficult or impossible. Alternatively, a code that can only be used once to create a single introduction can serve a similar purpose and be resistant to being published to a public forum, at the cost of making it more inconvenient to use.

2.1.2 Deriving the social network from parents

Another approach to shape the social network for a child uses the existing social network structure of the child's parent. The child social network Togetherville [6] pulls relationship information from Facebook. Within Togetherville, the children of parents can only interact with the children of their parents friends. This assumes, of course, that their parents are part of a social network, and this may be too restrictive of a constraint.

2.2 Limiting the content of gifts

Along with shaping the social network, a web site can actively enforce constraints on the contents of gifts. These gifts can be categorized as:

- text messages
- images

- sounds
- videos
- computer programs

The content of a gift can be atomic, or be composed of individual chunks. Each these have the potential to encode controversial material, and so a web site that guards that limit the expressivity of gifts needs to focus on both the (1) atomic primitives and the (2) means of composing those primitives.

2.2.1 Primitives

Many child-centric web sites focus on collecting a blessed whitelist of primitive values. This collection can be manually curated by adults and automatically curated with software tools [12] [11]. The blessed set of primitives may also be a static or dynamic collection: users may upload new primitive values, which then need to go through a validation process before they are added to the whitelist.

- Text messages can be considered atomic data if they're are pre-selected from a collection of phrases. Togetherville, for example, takes this approach in maintaining a sanitized set of phrases that users can select to send messages.

  Focus can be placed on individual words as well as entire phrases, as is often done in message forums which replace swear words with grawlixes.

- Images, sounds, and music, may contain objectionable content.

  These values may also contain metadata content. Images, for example, often contain geographic details, not only in the image itself, but as geolocation data in metadata. Web sites should scrub metadata out of these values to prevent personal information from being exposed.

2.2.2 URL resources

A special class of textual content are URLs that indirectly reference external content on the Web. It might seem reasonable to treat URLs just as any other atomic value, and to have a curator validate the content before adding it to the whitelist. However, URLs are "soft" references, and the referent can be changed without notice to the referrer, allowing a malicious user to maneuver a bad primitive value into the whitelist.

Web sites that allow URLs in gifts can limit the set of allowable URLs to those from trusted domains. Another approach is to keep a fingerprint of the content, such as an md5 checksum, and check that the linked content has not changed before presenting the content to the user. To avoid race conditions, it would be best to have this validation happen on the client side of a browser, but the same-origin policy in web browsers restricts client-side JavaScript from being able to perform low-level actions, like getting an image's bitmap for remote URLs.

The safest approach is to have the web side keep its own copy of the content, to guarantee its immutability.

2.2.3 Composition

Although the primitives that make up a gift can be validated, that doesn't necessarily imply the acceptability of a composition. As an example, a computer program can create an image out of simple shapes like overlapping circles and lines. For those with sufficient imagination, these simple composition operators pose no limitations. Trying to control what one can express by limiting the power of computation is a difficult task.

With regards to text messages, some social networks attempt to control message content by forcing the user to construct sentences with templates. The ScuttlePad social network, for example, requires all messages to be written using a sentence-templating tool. The tool enforces the sentence structure in such a way as to present an illusion of freedom.

Detecting the patterns of obscenity through composition seems to be a much more difficult problem to tackle with automated techniques. For rich media, such as computer programs, manual curation is the most common technique.

2.2.4 Information extraction and education

One technique that can be applied to free-form text is information extraction [10], which seeks to discover structured data such as postal and email addresses, and phone numbers. Rather than just redact this data, a web site can proactively educate a child that they're doing something potentially harmful. Just as in HTML form validation, a prompt or some visual cue can indicate that the user's in the process of exposing their personal information.

2.3 Playgrounds

A metaphor that's missing from most social networks is that of the park or the playground. In these environments, people establish safety through the act of mutual surveillance [8]: adults watch over the children, and just as importantly, children know that they are being monitored, and that they can easily get the attention of an adult if they feel threatened. In contrast, in a traditional social network, the atmosphere is, strangely enough, a solitary experience. It's of a child with a toy who is plays with it alone, who doesn't share it with others during the process of creation. Parents participate only when children publish their gifts. In most child-centric social networks, the observation power is asymmetric and asynchronous, with parents able to look down, but with children unable to see the online presence of each other, or their vanguards.

One thing that child-centric social networks may do is to present such a playground spatial and temporal metaphor, to create spaces of mutual surveillance that constructs a safety-in-numbers environment. In it, parents and children should be able to observe each other's actions in real time, where tools to create gifts allow for collaboration. Such an environment would allow parents to help educate children on the appropriateness of their actions, and allow children to police each other.

References

[1] The Children's Online Privacy Protection Act. `www.coppa.org/`.

[2] The Deleting Online Predators Act of 2006. `en.wikipedia.org/wiki/Deleting_Online_Predators_Act_of_2006`.

[3] HB 5941: The Safe School Act. `www.rilin.state.ri.us/BillText11/HouseText11/H5941Aaa.pdf`.

[4] Nintendo friends codes. `www.nintendo.com/consumer/wfc/en_na/ds/gameSupportFriendCodes.jsp`.

[5] Operators of Online "Virtual Worlds" to Pay $3 Million to Settle FTC Charges That They Illegally Collected and Disclosed Children's Personal Information. `www.ftc.gov/opa/2011/05/playdom.shtm`.

[6] Togetherville. `togetherville.com/`.

[7] Xanga.com to Pay $1 Million for Violating Childrens Online Privacy Protection Rule. `www.ftc.gov/opa/2006/09/xanga.shtm`.

[8] Anders Albrechtslund. Online Social Networking as Participatory Surveillance. *First Monday*, 13(3-3), 2008.

[9] Illene R. Berson and Michael J. Berson. Challenging Online Behaviors of Youth: Findings from a Comparative Analysis of Young People in the United States and New Zealand. *Social Science Computer Review*, 23(1), 2005.

[10] Chia-Hui Chang, Mohammed Kayed, Moheb Ramzy Girgis, and Khaled F. Shaalan. A Survey of Web Information Extraction Systems. *IEEE Transactions on Knowledge and Data Engineering*, 2006.

[11] Christian Jansohn, Adrian Ulges, and Thomas M. Breuel. Detecting Pornographic Video Content by Combining Image Features with Motion Information. *Proceedings of the 17th ACM international conference on multimedia*, 2009.

[12] Henry A. Rowley, Yushi Jing, and Shumeet Baluja. Large Scale Image-Based Adult Content Filtering. *International Confererence on Coputer Vision Theory and Applications*, 2006.

[13] Kaiser Family Foundation Study. Generation M^2: Media in the Lives of 8– to 18-Year-Olds. `www.kff.org/entmedia/upload/8010.pdf`.