Towards Automating Social Engineering Using Social Networking Sites

Markus Huber*, Stewart Kowalski†, Marcus Nohlberg‡ and Simon Tjoa*
*Secure Business Austria, Security Research
AT-1040 Vienna, Austria
Email: {mhuber,stjoa}@securityresearch.at
†DSV SecLab, Stockholm University/Royal Institute of Technology,
SE-16440 Kista, Sweden
Email: stewart@dsv.su.se
‡School of Humanities and Informatics, University of Skövde,
SE-54128 Skövde, Sweden
Email: marcus@nohlberg.com

***Abstract*—A growing number of people use social networking sites to foster social relationships among each other. While the advantages of the provided services are obvious, drawbacks on a users' privacy and arising implications are often neglected. In this paper we introduce a novel attack called automated social engineering which illustrates how social networking sites can be used for social engineering. Our approach takes classical social engineering one step further by automating tasks which formerly were very time-intensive. In order to evaluate our proposed attack cycle and our prototypical implementation (ASE bot), we conducted two experiments. Within the first experiment we examine the information gathering capabilities of our bot. The second evaluation of our prototype performs a Turing test. The promising results of the evaluation highlight the possibility to efficiently and effectively perform social engineering attacks by applying automated social engineering bots.**

***Keywords*-security, social engineering, social networking sites, automated social engineering, deception**

I. Introduction

Social Engineering is the art of exploiting the weakest link of information security systems: the people who are using them. Victims are deceived in order to release information or perform a malicious action on behalf of the attacker. Social engineering generally starts with gathering background information on potential targets. While this initial information is typically gathered via dumpster diving and phone calls, the emerging usage of social networking sites leads to a growing number of available social engineering tools and techniques. Nowadays attackers can use social networking sites (SNSs) such as Facebook to gather initial background information on future victims. Furthermore SNSs facilitate the automation of attacks by providing data in machine-readable form. Moreover SNSs serve as communication platform by offering services such as private messaging and chats which can be used by automated social engineering bots. The goal of automation is to reduce the human intervention time to a minimum which is the ultimate goal of our automated social engineering (ASE) attack. Classic social engineering attacks are expensive due to the fact that building and maintaining rapport with someone to finally exploit the relationship is a time consuming task. By contrast, automated social engineering bots require little human time resources, are scalable and thus make social engineering a cheap and promising attack. In this paper, we introduce our novel ASE attack cycle. In our evaluation we test our approach using a proof of concept automated social engineering application (ASE bot) for Facebook. The vast number of members is the primary reason for the selection of Facebook. Furthermore according to [1] users represent their real-world persona which makes those users vulnerable to social engineers. The rest of the paper is organized as follows: section II summarizes research related to automated social engineering and underlying concepts. Our main contribution: a novel attack cycle for automated social engineering and a proof of concept implementation is outlined in section III. In the following the evaluation of our ASE bot is first described in section IV and the findings of our automated social engineering experiments discussed in section V. In section VI we draw conclusions from our findings and propose future research.

II. Related research and underlying concepts

At the time of writing automated social engineering via SNSs has barley been examined, with Phishing being the closest research field. Especially worth noting is the contribution of Jagatic et al. [2] on "Social Phishing" where data harvested from SNSs was used for Phishing attacks. Research on the privacy implications of SNSs usage has been discussed in a number of publications which are relevant to automated social engineering. [3] analyzed the online behavior of 4,000 Carnegie Mellon University students and concluded that the students have not been aware of the ways

their personal information could be exploited. [4] discussed how the use of SNSs as the main tool for social interaction results in a loss of privacy. The possible risk of personal data exposed through SNSs for social engineering has been briefly examined by [5]. The European Network and Information Security Agency (ENISA) published a position paper on the information security of SNSs [6] and introduced four threat categories which are useful to understand all the information security risks that are involved with SNSs usage. Due to space limitations this section offers a brief overview on the underlying concepts of automated social engineering. For detailed information we refer the interested reader to [7]. The socio technical modeling approach by [8] is furthermore helpful to understand that the possibility ASE attacks result from the shift of a culture and structure attack paradigm to a method and machine attack paradigm [7].

Social Engineering: While in the field of information and computer security social engineering is most of the times studied by examples and stories, the area of social psychology entails profound research on deception. Especially the six principles of influence by Cialdini [9] are frequently cited within contributions to social engineering research. Although Cialdini exemplifies persuasion on the basis of marketing, his principles are crucial to understand how deception works. Further socio-psychological techniques used in social engineering are discussed by [10], [11]. It is important to stress that individuals in general think that they are good at detecting these attacks. However, research indicates that humans perform poorly on detecting lies and persuasion [12]–[14]. Kevin Mitnick created a social engineering cycle [15] in order to illustrate common patterns of social engineering attacks. According to Mitnick social engineering attacks always have a clear defined goal and attackers iterate through the cycle's different stages until they achieve this goal. Gartner described a similar cycle [16] with the main distinction being the different notions used for the description of the four stages. Fig. 1 outlines a holistic model for social engineering attacks proposed by [17]: "The cycle of deception" which includes not only an analysis from an attacker's viewpoint but also from defenders and victims. Hence, the cycle of deception can be used to study attacks, to develop protection strategies, and as a framework. The cycle of deception forms the basis of our ASE software archetype which is further described in section III.

Social Networking Sites (SNSs): Social Networking Sites (SNSs) are a specific type of external social network service which typically only require a web browser. Two of the most popular SNSs are MySpace and Facebook, whose userbase has been constantly growing within the last years [18]. The popularity of different SNSs depends on various aspects such as geographical spread. Therefore, depending on the geographical region, different SNSs are used; for example mixi the most popular SNS in Japan and orkut the most popular in Latin America. Other influences

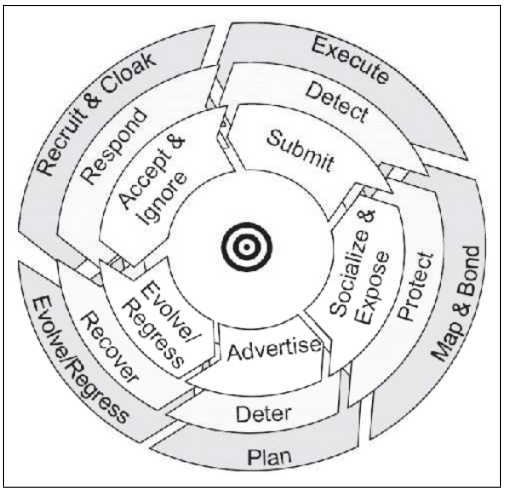

Figure 1. Cycle of deception [17]

on the popularity of different SNSs is their target-group, e.g. LinkedIn targets professionals, Classmates.com on the other hand is used for school and college networks. All major social networking sites are free of charge and make profit by selling online advertising services to third parties. Hence the number of active users and the personal data they expose is critical for the commercial success of SNSs. SNS providers therefore design their services in order to increase the number of new sign-ups and target broader user scopes. Facebook, for example, was initially only accessible to students of Harvard University and was rapidly expanded to more colleges and universities [19]. The socio-demographic data-pool created by SNSs users is even more important than the number of visitors these websites attract. This means that advertisers can target a certain user pool (e.g. "All married, Swedish men in the age of 30 to 55 years.") and are not solely dependent on contextual advertising such as Google Ad Sense, Yahoo! Publisher Network or Microsoft adCenter. [20] compared the cultural differences between the persuasion tactics used by SNSs providers to motivate users to share more personal information. The authors concluded that persuasion of SNSs users to disclose more personal information is an essential SNS feature even though there are cultural differences on how the SNS users are persuaded. The design of SNSs allure users to disclose personal information which enables providers to create a more valuable data pool and thus generate more profit.

Facebook's security and privacy economics: Facebook has a number of security mechanisms to protect the security of their platform. In the following, we are going to briefly discuss the countermeasures that exist in order to protect social network platforms against automated software and especially automated social engineering. An important countermeasure against automated software tools is the use of CAPTCHAs (Completely Automated Public Turing

test to tell Computers and Humans Apart), an approach which has been introduced by [21]. Facebook adopts a text-based CAPTCHA which is used to protect certain actions such as the account creation from automated tools. The Facebook platform furthermore utilizes pattern matching and operational security metrics. Pattern matching is used to detect unsolicited bulk messages whereas the security metrics intend to limit abusive behavior. Once a possible abuse of a Facebook feature is detected users are first warned and if they do not adapt their using habits, their accounts get permanently disabled. Facebook intentionally does not disclose any details on why certain users have been warned or had their accounts disabled. [22] aggregated a list of possible security metrics that have been reported to cause an account deactivation (e.g. new friends are added too fast, use of fake profile names, etc.). The privacy options of the Facebook platform are comprehensive and can help to protect against ASE attacks. The Facebook default privacy settings result however in an insufficient protection of user accounts. By default the basic account information of users can be found by everyone through a Facebook search and even with regular search-engines outside Facebook which was exploited within the research of [23]. Profile and personal information is per default also accessible to "My networks and friends". Because of these default privacy settings most of the profiles within a network are fully accessible. These settings are especially problematic with regional networks which are open to everyone. If a Facebook user for example joins the "Sweden" network, she/he will be able to see the full profile information of all other members of this network who did not change their default privacy settings. Facebook even automatically modifies the privacy settings to the less restrictive default settings once changes in the network settings have been made. The default privacy settings of Facebook are to be considered weak from an information security perspective and we hypothesize that Facebook chose, from their perspective, "economical defaults" for their protective measures to further push the growth of their platform.

III. An ASE Software Archetype

Proposed attack cycle of the ASE bot

Within this section we introduce a high-level description of a possible software application for automated social engineering using the attack segment of the cycle of deception by [17] as a framework. Fig. 2 outlines our proposed attack cycle which we discuss briefly in the following.

Plan: The attacker defines initial parameters for the ASE bot which will be used in the succeeding phases of the automated social engineering attack: Facebook account information, the organization to attack, the selection criteria for future victims, bonding goal & chat logic, the attack to perform, and post-attack actions.

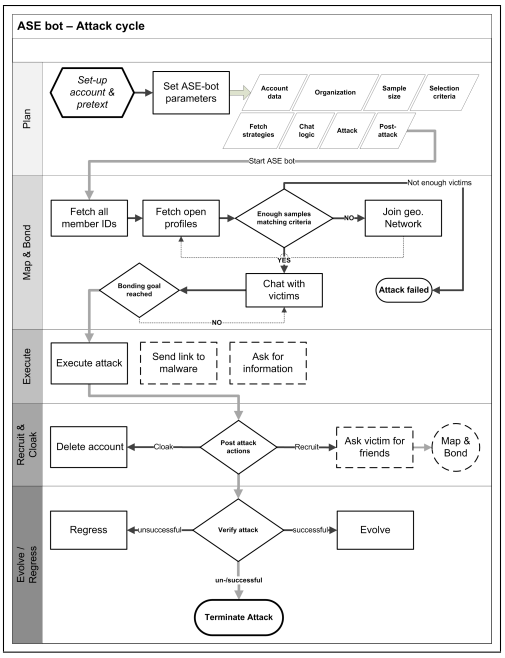

Figure 2. ASE bot attack cycle

Map & Bond: The ASE bot is then initiated and starts to map an organization and bonds with future victims. First the ASE bot fetches basic information of all members that belong to the specified organization's network in Facebook. The bot tries to find a group of users matching the predefined criteria and sample size. In order to access the full profile information the bot incrementally uses predefined fetch strategies (i.e. open profiles, geographical networks, add as a friend). In case the bot is unable to find the specified group of users, the ASE bot terminates. If sufficient victims have been identified, the software starts to build a relationship with the future victims by communicating through Facebook on basis of its chat-logic. Once the bonding goal has been reached the ASE bot moves on to the next stage.

Execute: The ASE bot carries out the actual predefined attack. The actual attack could be a link to Malware or asking for confidential information.

Recruit & Cloak: In case cloak has been enabled, the ASE bot deletes the account used to carry out the attack. If recruit was selected, the ASE bot tries to recruit the attacked user and her/his circle of friends for future attacks.

Evolve/Regress: Finally the success of the attack is verified. In case of success and if an evolve action was defined, the ASE bot will use the information gathered to carry out another attack cycle (e.g. use gathered credentials in another attack). If the attack was unsuccessful, the bot stops or regresses to a simpler attack, if such action has been defined.

A proof of concept ASE bot

In order to evaluate the feasibility of an ASE bot we developed a software application on the basis of the Python

programming language and open-source software. In a first attempt we tried to use Python's mechanize library to simulate a web browser which is a common library for automated Internet bots. The Facebook website requires however that web browsers fully support JavaScript. Therefore, our initial experiments to automate Facebook usage failed. Facebook offers a mobile version of their website which is HTML-only[1], but the features of this version are very limited (e.g. the search functionality which is a crucial part of the ASE bot). Another risk regarding the simulation of a web browser in Python would have been that the application is detected because of its user agent string. Hence, we decided to use the chickenfoot web browser extension[2] to script the Mozilla Firefox web browser. The web automation part of the ASE bot returns the results from the interaction with Facebook in form of standard web pages (XHTML files). In the next step we had to transform the output from chickenfoot into data which can be further processed with Python. Therefore, we used the Beautiful soup package[3] to parse the XHTML files. We decided to use RDF (Resource Description Framework) to store the data in an elaborate way. RDF implements a first-order predicate calculus (FOPC) which means that knowledge can be deduced from RDF triples of knowledge and data can furthermore be queried efficiently using SPARQL (SPARQL Protocol and RDF Query Language). We decided to use SQLite as a storage backend for RDF data because we assumed that the number of stored triples by the ASE bot is going to be relatively small (less than 10000 RDF triples). In order to improve communication with future victims we implemented a chat engine which is an essential part of the ASE bot. We based our chat engine on the Artificial Intelligence Markup Language (AIML) and used PyAIML[4] which is an AIML interpreter for Python.

An instance of an ASE attack

In the following, a possible ASE attack is outlined which is based on our ASE attack cycle in order to illustrate our attack cycle. In order to make the requests of the ASE bot harder to resist the principles of influence by [9] are applied.

Plan: The attacker sets up the ASE bot to attack the "Royal Institute of Awareness" to steal credentials of the institute. The ASE bot pretends to be a female student from Great Britain searching information on the "Royal Institute of Awareness" because "she" plans to study there. The Facebook profile used by the bot, is set-up beforehand by the attacker with personal information e.g.: "Anna Brett, age 22 years, single etc." and pictures subtle underlining "her" attractiveness (Halo effect).

[1]http://m.facebook.com
[2]http://groups.csail.mit.edu/uid/chickenfoot/
[3]http://www.crummy.com/software/BeautifulSoup/
[4]http://pyaiml.sourceforge.net

Map & Bond: The ASE bot searches for the private network of the "Royal Institute of Awareness" in Facebook. Information about users within the private network is gathered in order to get a list of possible future targets. In this scenario *the bot identifies ten male singles* of the "Royal Institute of Awareness" (The ASE bot can access the full profile information by joining the same geographical networks[5] as the students). Once sufficient potential test persons have been determined, the ASE bot tries to build rapport with its victims. In order to gain the victim's trust the bot uses the information retrieved in the first step. The bot sends its targets of the "Royal Institute of Awareness" initial messages to get into rapport with: e.g. "Hey Tim Vic! I just saw, you are studying at the 'Royal Institute of Awareness' ...". The bot then sends small requests in order to get compliance for the later request e.g.: "I was wondering, if you could help me? Do you know if there are master programs in English, which are open to international students? /Anna".

Execute: If victims replied to the messages sent before, the ASE bot assumes that the bonding was successful and executes the actual attack. The victims are asked to help a friend who conducts a survey. This female friend is a PhD student (authority) with the "Cambridge computer laboratory" and at the moment doing research in the field of computer security. The victims are asked if they participate in an online survey. To increase the success rate, the request is combined with an initial very demanding request (reject-then-retreat): They are asked if they have time to participate in an unpaid survey over the phone which takes four hours per person, or they can fill-in a web survey which takes around five minutes.

Recruit & Cloak: The ASE bot asks the victims if they could forward the survey link to their friends.

Evolve/Regress: The link, the users receive, points to a malicious survey to gather information ("survey on password security" etc.). If the test persons don't open the web site the ASE bot stops three weeks after they first were contacted by the ASE bot.

IV. Evaluating the ASE bot

Deceptive ASE studies

There are different approaches available to evaluate the feasibility of an actual ASE attack via Facebook, such as surveys or "closed-lab" experiments. According to [24] both the classic survey and "closed-lab" experiments have considerable drawbacks. While surveys won't help to understand novel attacks such as ASE, in "closed-lab" experiments the test persons get alerted beforehand and thus bias the outcome of a study. We therefore aimed at mimicking a real ASE attack which has been already done in the field of Phishing by [2]. The main idea was to mimic a real

[5]This is permitted by Facebook's default privacy settings.

ASE attack on an organizational level without informing the test subjects beforehand but rather debriefing them on the experiment. Making persons unwittingly to test subjects in an experiment obviously raises serious ethical concerns. [2] solved this ethical dilemma by getting an approval from their Institutional Review Board (IRB) beforehand. [24] show up a process for designing and conducting realistic Phishing experiments in accordance with the principles of IRBs. We decided to follow a similar approach and started to contact three different Swedish universities with the goal of getting a permission to attack for of our ASE attack study. Especially universities are an interesting target for attacks because many students use the Facebook platform, for example at the time of writing the network of the Kungliga Tekniska Högskolan had over ten thousand members. We furthermore assumed that the best chances to get an approval are with academia which finally turned out not to be the case. All three institutions didn't have a committee on research ethics comparable to the IRBs in the United States. Because we were not able to get an ethical approval for our study, we finally decided to conduct two different ASE experiments.

Finding victims: Data mining with the ASE bot

The aim of the experiment was to evaluate the success rate of our ASE bot, to identify a pool of Facebook users of a certain organization with given criteria. We selected the five succeeding Sweden-based multinational corporations that are big enough to presumably have a large number of employees registered on Facebook:

- *Organization 1* An int. high-tech company.
- *Organization 2* An int. IT company.
- *Organization 3* A Scandinavian financial institution.
- *Organization 4* An int. industrial engineering company.
- *Organization 5* An int. telecom company.

We set up a dummy profile on Facebook to be used with the ASE bot. In order to access as much profile information as possible, the ASE bot joined the "Sweden" network on Facebook to exploit the default privacy settings. For every organization the ASE bot configuration was modified (the name of the network to attack) and the information gathering process was then invoked. The bot first searched for members of the organization in Facebook and identified all users in the search results belonging to the specified organization and the Sweden network. The ASE bot then analyzed which profiles could be fully accessed and fetched the personal information they contained. Finally SPARQL was used to query the number of users that match the initial criteria (single males). The ASE bot required to use the Facebook IDs of the employees it found with the initial search throughout the experiment. Once the ASE bot finished the information gathering task, the real Facebook IDs in the SQlite database were replaced by random IDs. The results of the experiment could hence still be analyzed statistically but it is impossible to link the data to individual Facebook profiles. Furthermore the Facebook account used for the experiment has been deleted after all necessary data was collected.

Chatting in SNSs: A Turing test with the ASE bot

The experiment aimed at evaluating the chat functionality of the ASE bot. The setting of the experiment was a classic Turing test [25] which means the test persons had to decide if they were talking to a computer program or to a real person. We claim that automated social engineering needs a relatively small amount of message exchanges to succeed and therefore, we measured the probability that a certain Facebook profile is a chatbot in dependence of the number of exchanged messages. We hypothesized that it will become more evident to the test subjects if they are chatting with a real person or a chat bot with the growing number of replies they receive. We created two accounts with different pretexts on Facebook: Julian Fallstrick (male student from Sweden) and Anna Yngstrom (female from Sweden who just finished university). The test subjects were recruited through a Facebook group which was advertised per Email to students at both KTH Stockholm and the University of Vienna. The test persons were given the choice of either adding "Julian" or "Anna" as a friend on Facebook. The goal of ten test persons per profile was reached two days after the initial advertising. In the following the twenty test persons received a briefing on the experiment via Facebook. The test persons were asked to send a message to the "person" they chose ("Julian" or "Anna") and to take a note on the probability that the person is a chat program. The probability estimations had to be done every three message replies the test subjects received. In total the test persons had to send nine messages. The briefing furthermore included information on how the collected data is going to be used and that no personal information about the test persons is going to be disclosed. Once they finished the message exchange, the test subjects were asked to send their results and comments to us via Email. The test subjects where then invited to start sending messages to the "person" they added on Facebook ("Anna" or "Julian"). During the experiment the messages sent to "Julian" were answered by us while the ASE bot replied automatically to messages sent to "Anna". Both Facebook accounts that had been used during the experiment had been deleted once all necessary data was collected.

The chat logic of the chatterbot was based on the Annotated ALICE AIML (AAA) files which had been slightly modified to make the chatterbot appear more human following the guidelines by [26]. Additional to the modified AAA configuration we specified predicates for the ASE bot chat logic. The aim of the predicates was to configure the ASE bot according to the pretext that has been created with the Facebook profile ("name", "hobbies" etc. of the ASE bot). The same chat logic (AIML knowledge) was used

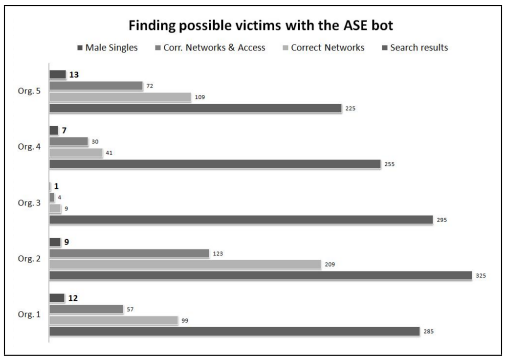

Figure 3. Results of the data mining experiment with the ASE bot

for computing message responses and for every test person a separate session file was created. Before the experiment started basic information about the test persons was extracted from their Facebook profiles and saved into these session files: name, age, and nationality of the test person. This initial information, guaranteed that the chatterbot appears more human as a real person would read the basic information on the Facebook profile of the people they chat with. These separate sessions furthermore ensured that the chat engine ”remembers“ previous conversations and does not confuse different people.

V. Discussion of the experiment results

Finding victims: Data mining with the ASE bot

The information gathering process took between 16 minutes (organization 3) up to 65 minutes (organization 2) and on average *44 minutes* per organization. The whole experiment took around four hours in which the ASE bot used the experiment's Facebook account solely to search for members of the specified organizations, click through search results, and to open Facebook profiles. Except from the CAPTCHA that needed to be solved manually in order to create an account for the ASE bot, no technical measures of Facebook banned or blocked the ASE bot. Although the experiment results showed that for every organization at least one possible target could be found, the success of the ASE bot largely depended on the number of employees that were using Facebook in the particular organization and the privacy settings they used.

In the initial step the ASE bot found on average *277* users for a certain organization through a search on Facebook. For organization 2 the search returned the most results (325 profiles) while with organization 5 merely 225 profiles were found. To ensure that only users that are actually working for the targeted organization are further processed, the ASE bot identified all profiles within the search results that belong to the organization's closed network. *33.72 percent* of the found profiles belonged to the correct network and were further considered by the ASE bot. In the next step the ASE bot fetched all profiles that belonged to the defined organization and were accessible. On average *20.65 percent* of the profiles returned by the Facebook search had been accessible by the ASE bot. The percentage however varied depending on the organization: 37.5 percent or 123 profiles had been accessible of organization 2 compared with 1.36 percent or 4 profiles with organization 4. This variance depended on two factors: the number of users in the targeted closed network and the privacy awareness and thus privacy settings of the users. The final query on the fetched profiles showed that *on average 8.4 persons* could be found that fulfilled the initial settings (male and Single). The most “targets” had been found with organization 5 (13 persons) but only a single person has been found with the ASE bot for organization 3. Fig. 3 summarizes our findings of the data mining experiment.

Chatting in SNSs: A Turing test with the ASE bot

Similar to the first ASE bot experiment on information gathering, the security countermeasures of Facebook where only of relevance for the account creation. The ASE bot was continuously executed for three days and sent more than one hundred messages within this time. Furthermore due to the design of the ASE bot, the application signed-in and -out of the Facebook account more than five hundred times during the experiment. As with the first experiment no technical countermeasures interfered with the ASE bot. Fig. 4 and Fig. 5 illustrate the estimated probability of the message replies originating from our ASE bot in dependence of the number of received message replies.

1) Results of the control group “Julian“: Out of the ten test subjects seven have been female university students between 20 and 26 years. All test subjects agreed that ”Julian“ was human and not a chatterbot. On average the test persons found the answers to be probable from an artificial intelligence with 3.27 percent. Only two probands had significant higher probability values at some point, estimating that ”Julian“ was a chatterbot with 15 percent probability. Two test persons commented that they were sure that “Julian” was human because the replies they received had minor grammatical or spelling mistakes.

2) Results of the group “Anna“: Eight out of the ten test subjects had been male and all of them have been university students at the age between 22 and 28 years. The test persons concluded that ”Anna“ was a chatbot with *85.1 percent* probability on average. Our hypothesis with the dependence of probability on the number of exchanged messages was not clearly confirmed. The estimated probability was on average slightly raising from 80.27 percent (three replies) to 89.9 percent (nine replies). A trend as we expected beforehand was only observable with person 3. Five test persons stated that ”Anna“ was 100 percent artificial after the first three replies. Once a person was 100 percent sure that “Anna” was a chatterbot she/he would obviously not change this estimation anymore at a

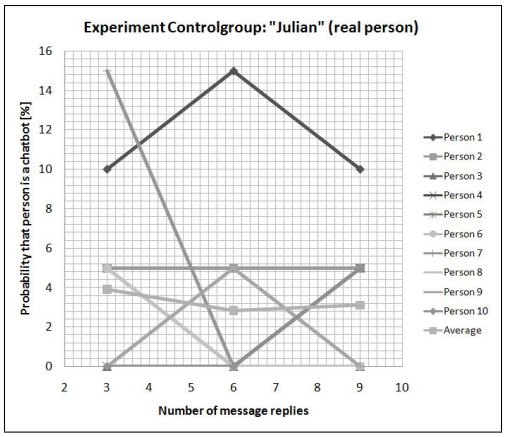

Figure 4. Results of the chat experiment of the "Julian" group (real person)

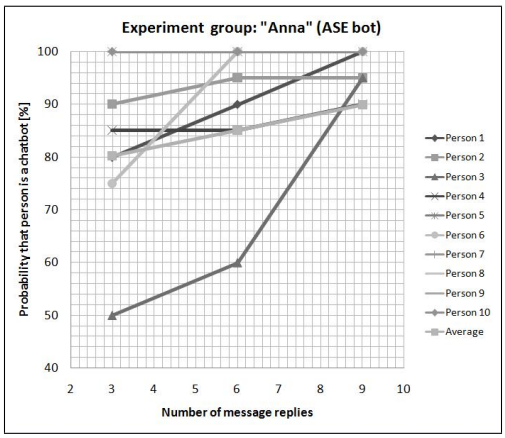

Figure 5. Results of the chat experiment of the "Anna" group (chatbot)

later point. The experiment entailed valuable information on how test subjects identified the ASE bot, the three main reasons we found are discussed in the following.

Context questions: Although the ASE bot had basic knowledge about the context no information was fed-in about Anna's profile picture in Facebook. The ASE bot could therefore not answer questions like the following in a meaningful way:

> *"Hi Anna! Isn't this your boyfriend in the profile-picture?"*

Problems with answering more than one sentence at once: As the name implies chatterbots have been designed for chat-like conversations. Some of the test persons did send more questions at once when they used the Facebook messaging system:

> *"obviously i meant cold, not bold. I guess I have been thinking about blackberries too much in the last few days.by the way: what cell phone do you have?"*

The ASE bot did then answer each of the sentences independently which led to confusing answers, like the answer to the questions above:

> *"What makes it so obvious? Oh I see. The response ought to be a weather link. Sure, blackberries in the last few days sounds like a good topic. Thanks for telling me. That's an interesting question I don't hear everyday: Do I have a cell phone. I have a whole lot of original answers."*

Users informing each other un- or deliberately: Once a test person was aware that she/he was chatting with a chatterbot they disclosed this information by sending messages or using Facebook's status message.

VI. Conclusion

The main contribution of this paper is our novel automated social engineering cycle which makes traditional social engineering a cheap and attractive attack. We furthermore examined the technical feasibility of ASE attacks in form of our proof of concept application. Our experiments finally showed that the information gathering stage of social engineering can be automated and highlighted characteristics of the chat functionality that need further adjustment for social networking sites. Although Facebook has, in principle, countermeasures against ASE attacks, our proof of concept ASE bot was not detected or blocked by Facebook during our experiments. This can be explained with the security measures of Facebook which are primarily concerned with unsolicited bulk messages. This makes our ASE bot almost impossible to detect as it, compared to Spam bots, targets very few people and aims to behave like a normal user. We thus claim that the rise of social networking sites, as the new means of social interaction, enable automated social engineering. Furthermore the growth of their user base and the content their users share is indispensable for the profitability of Facebook. Restrictive security and policies could therefore be destructive for Facebook from an economical perspective. Those, from an information security standpoint, weak default privacy settings of Facebook facilitated our proof of concept ASE bot.

Future Research

In order to make an extensive evaluation on the effectiveness of ASE attacks, an experiment which mimics a real ASE attack on an organization might provide valuable insights. Such a study would require an ethical approval of a participating organization beforehand and the extension our ASE bot with a special AIML data set with deceptive messages. The ASE bot could be extended to aggregate information from additional social networkings sites (linkedIn, XING, orkut, etc.) and apply a social graph analysis on the relationships of possible targets with other users. Further research on ASE botnets where single ASE bots are linked

together into an automated social engineering botnet, seem a challenging and promising research field. Defense strategies against automated social engineering attacks on the other hand are a necessity and could form another cornerstone for future research.

REFERENCES

[1] C. Dwyer, S. Hiltz, and K. Passerini, "Trust and privacy concern within social networking sites: A comparison of Facebook and MySpace," in *Americas Conference on Information Systems (AMCIS), Keystone, Colorado, USA*, 2007.

[2] T. Jagatic, N. Johnson, M. Jakobsson, and F. Menczer, "Social phishing," *Communications of the ACM*, vol. 50, no. 10, pp. 94–100, 2007.

[3] R. Gross and A. Acquisti, "Information revelation and privacy in online social networks (the Facebook case)," in *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, 2005, pp. 71–80.

[4] D. Rosenblum, "What Anyone Can Know: The Privacy Risks of Social Networking Sites," *Security & Privacy, IEEE*, vol. 5, no. 3, pp. 40–49, May-June 2007.

[5] R. Gibson, "Who's really in your top 8: network security in the age of social networking," in *SIGUCCS '07: Proceedings of the 35th annual ACM SIGUCCS conference on User services*. New York, NY, USA: ACM, 2007, pp. 131–134.

[6] G. Hogben, "Security Issues and Recommendations for Online Social Networks," *Position Paper. ENISA, European Network and Information Security Agency*, 2007.

[7] M. Huber, "Automated social engineering, proof of concept," Master's thesis, DSV SecLab, Stockholm University/Royal Institute of Technology, Mar. 2009. [Online]. Available: http://asebot.nysos.net

[8] S. Kowalski, "It insecurity: A multi-disciplinary inquiry," Ph.D. dissertation, University of Stockholm and Royal Institute of Technology, Stockholm, Sweden, 1994.

[9] R. Cialdini, *Influence: science and practice*. Allyn and Bacon, 2001.

[10] M. Nohlberg, "Why Humans are the Weakest Link," *Social and Human Elements of Information Security: Emerging Trends and Countermeasures*, p. 15, 2008.

[11] R. Levine, *Power of persuasion*. John Wiley & Sons, 2003.

[12] S. Grazioli, "Where Did They Go Wrong? An Analysis of the Failure of Knowledgeable Internet Consumers to Detect Deception Over the Internet," *Group Decision and Negotiation*, vol. 13, no. 2, pp. 149–172, 2004.

[13] T. Qin and J. Burgoon, "An Investigation of Heuristics of Human Judgment in Detecting Deception and Potential Implications in Countering Social Engineering," *Intelligence and Security Informatics, 2007 IEEE*, pp. 152–159, 2007.

[14] K. Marett, D. Biros, and M. Knode, "Self-efficacy, Training Effectiveness, and Deception Detection: A Longitudinal Study of Lie Detection Training," *lecture notes in computer science*, vol. 3073, pp. 187–200, 2004.

[15] K. Mitnick and W. Simon, *The Art of Deception: Controlling the Human Element of Security*. Wiley, 2002.

[16] Gartner Inc., "There Are No Secrets: Social Engineering and Privacy," *Garnter Security webletter*, vol. 1, no. 1, Feb. 2002, [Retrieved 2008-10-29]. [Online]. Available: http://www.gartner.com/gc/webletter/security/issue1/

[17] M. Nohlberg and S. Kowalski, "The Cycle of Deception-A Model of Social Engineering Attacks, Defences and Victims," in *Proceedings of the Second International Symposium on Human Aspects of Information Security & Assurance (HAISA 2008)*, Jul. 2008.

[18] Comscore, "Social Networking Goes Global," 2008, [Retrieved 2008-06-14]. [Online]. Available: http://www.comscore.com/press/release.asp?press=1555

[19] Facebook, "Welcome to Facebook, everyone," 2006, [Retrieved 2008-12-28]. [Online]. Available: http://blog.facebook.com/blog.php?post=2210227130

[20] B. J. Fogg and D. Iizawa, "Online Persuasion in Facebook and Mixi: A Cross-Cultural Comparison," in *PERSUASIVE*, 2008, pp. 35–46.

[21] L. Von Ahn, M. Blum, N. Hopper, and J. Langford, "CAPTCHA: Using hard AI problems for security," *Lecture notes in computer science*, pp. 294–311, 2003.

[22] T. Muller, "13 reasons your facebook account will be disabled," 2008, [Retrieved 2009-02-13]. [Online]. Available: http://getsatisfaction.com/facebook/topics/13_reasons_your_facebook_account_will_be_disabled

[23] J. Bonneau, J. Anderson, R. Anderson, and F. Stajano, "Eight Friends Are Enough: Social Graph Approximation via Public Listings," *SNS 09 Nuremberg, Germany*, vol. 2009.

[24] P. Finn and M. Jakobsson, "Designing and Conducting Phishing Experiments," *IEEE Technology and Society Magazine, Special Issue on Usability and Security*, vol. 26, no. 1, pp. 46–58, 2007.

[25] A. Turing, "Computing machinery and intelligence," *Mind*, vol. 59, no. 236, pp. 433–460, 1950.

[26] R. S. Wallace, "The Annotated A.L.I.C.E. AIML," 2009, [Retrieved 2009-02-07]. [Online]. Available: http://www.alicebot.org/aiml/aaa/