An analysis of various tools, methods and systems to generate fake accounts for social media

A Thesis Proposal Presented
by

Avanish Pathak

to the faculty of
College of Computer and Information Science

in Partial Fulfillment of the Requirements
for the Degree of

Master of Science in Information Assurance

Northeastern University
Boston, Massachusetts

December 2014

Contents

Abstract

Fake accounts on Online Social Networks (OSNs) have become a basic resource used in various kinds of online attacks. While some of these attacks are annoying but innocuous, other attacks are more serious and can wreak havoc online. Popular OSNs and webmail providers have adopted many security measures to halt the mass creation of fake accounts. However, their security measures are often rendered ineffective by the many tools available on underground marketplaces that allow unscrupulous individuals to cheaply acquire fake accounts in bulk.

The purpose of this paper is to study the mechanisms used by modern account creation programs and their overall effectiveness. This study analyzes the different ways in which these tools create fake accounts and how they manage to circumvent existing security measures. It also helps to get an insight into what websites do in order to handle fake accounts, both during the account sign-up process, as well as and after the fake accounts have been created. Tests that reveal the number of accounts that can be fabricated prior to an OSN's countermeasures and their longevity due to the inability of the OSN's detection mechanisms are presented. This study highlights whether major websites are following security best practices to mitigate fake account creation, and if existing security countermeasures are effective.

Chapter 1

Introduction

Fake accounts on Online Social Networks (OSNs) and webmail providers have become a basic resource used in various kinds of online attacks. Some of these attacks are annoying but innocuous, for example using fake accounts to generate "likes" on Facebook, follows on Twitter, and views on YouTube [40]. Other attacks are more serious, including using fake accounts to influence trending topics [16], conduct Search Engine Optimization (SEO) [42], spread spam advertisements, and trick users into installing malware. In the most extreme cases, attackers have been able to wreak havoc online by using fake accounts to spread false political content (astroturfing) [36], censor speech [52], and even spread false rumors that contribute to panic during real-world disasters [5]

Popular OSNs and webmail providers are aware of the problems created by fake accounts, and have adopted many security measures designed to halt the mass creation of fake accounts. CAPTCHAs [10], email address verification and phone number verification are some of the user-facing techniques that hinder mass account creation. Similarly, many providers limit the number of accounts that can be created from an individual IP address, or even ban certain IP addresses from making accounts all together (*e.g.* Tor exit nodes [45]). Advanced techniques that leverage machine learning have also been presented in the research literature [4].

Unfortunately, the security measures put in place by web services to prevent mass account creation are often ineffective: there are many online marketplaces that allow unscrupulous individuals to cheaply acquire fake accounts in bulk. In the past, Internet Relay Chat (IRC) rooms were used to connect people with stockpiles of accounts to buyers looking to leverage those accounts for attacks. However, nowadays blackhat forums and websites that operate out in the open serve as the marketplace between buyers and sellers. Recent studies have shown that Google, Hotmail, Facebook, Twitter, *etc.* accounts are widely available on these sites with prices ranging from $15-$180 for a 1000 accounts, depending on the target site, and the type of account [9] (*e.g.* Do they have many followers? Are they Phone Verified Accounts (PVA)? *etc.*). The fact that these marketplaces exist and remain profitable, is a testament to the massive demand for their product, *i.e.* fake accounts.

Although it is known that marketplaces for fake accounts exist, it remains unclear where all of these accounts come from. Some studies have found that the marketplaces are fueled by crowdsourced workers who manually create fake accounts in exchange for small monetary payments [38, 57]. Other studies have found that some marketplaces sell access to real people's accounts based on credentials that have been stolen through phishing, malware infection, or social engineering [34, 47]. These techniques circumvent existing security systems by relying on human beings to solve CAPTCHAs and respond to verification emails.

However, by far the most economical way to create fake accounts in bulk is to use a program or a script to automate the sign-up process. In their simplest form, these tools simply automate the process of filling out and submitting the account sign-up forms on various websites. Figure 1.1 shows an example of an actual tool called FacebookDevil that streamlines the process of signing up for Facebook accounts; the tool can create accounts as quickly as the user can solve the CAPTCHAs. As we will see later,

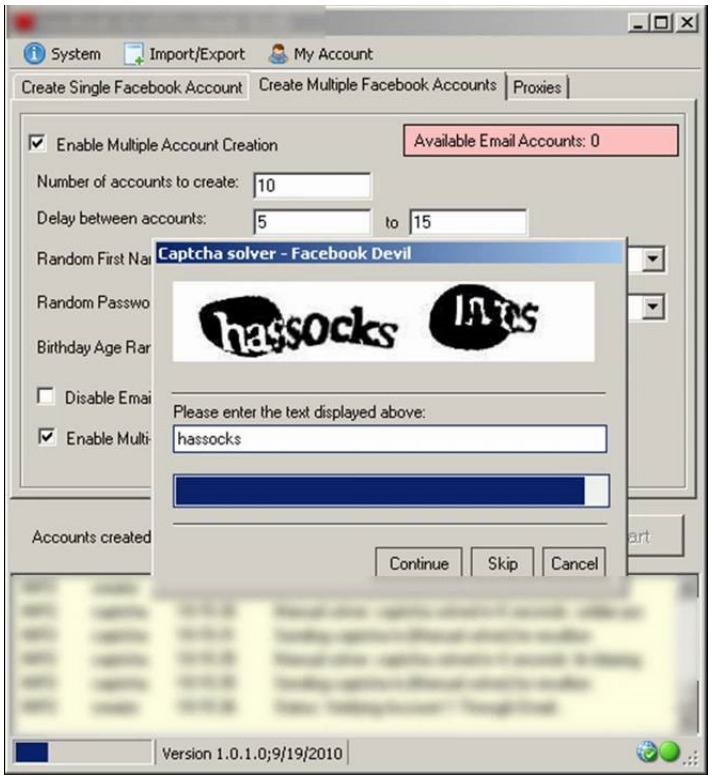

Figure 1.1: FacebookDevil: An easy to use, automated account creation tool.

some tools include more advanced features like automatic handling of verification emails, CAPTCHA solving [37], and support for proxies. By using a tool, a single attacker can, in theory, create hundreds or thousands of accounts per day at no (or little) cost. In contrast, accounts created by crowdsourced labor are potentially more expensive, while accounts stolen from real people may be reclaimed by the owner, and are therefore less reliable.

The goal of this study is to understand the mechanisms used by, and the overall effectiveness of, modern account creation programs. A deeper look inside the functionality of these applications will help us analyze the different ways in which these tools create fake accounts and how they manage to circumvent existing security measures. We will also gain knowledge as to what websites do in order to handle fake accounts, both during the account sign-up process, as well as and after the fake accounts have been created. This will help us understand whether major websites are following security best practices to mitigate fake account creation, and if existing security countermeasures are effective.

To conduct this study, we first identify five popular account creation tools from blackhat forums. These tools cover seven popular websites, including OSNs (Facebook, Twitter, *etc.*) and webmail providers (Hotmail). Next, we stress-tested each tool in a controlled environment, in order to determine how successful it was at creating fake accounts, and what low-level mechanisms the tool used to complete the sign-up process. We leverage network traces collected by Wireshark, as well as decompilation tools, to study the inner workings of the five tools. These stress tests allow us to gauge the effectiveness of each tool and also highlight the countermeasures taken by the target websites in dealing with our sign-up attempts.

Our analysis shows that different websites have very different defense strategies against automated account creation attempts. For example, the tools we tested were able to create up to 40 Twitter accounts on a single IP address within the duration of half an hour. Almost none of the account creation attempts were presented with challenges (*i.e.* CAPTCHAs), and the vast majority of the Twitter accounts we

created were not banned after 24 days[1]. This finding helps explain why Twitter is one of the largest victims of fake accounts [15]. On the other hand, eBay only allowed us to create 3 accounts per IP address. In this case, eBay has a clear interest in preventing e-commerce fraud, and thus they implement stringent countermeasures against fake accounts.

Another interesting finding from our study is that the account creation tools we surveyed are updated frequently. Three of the five tools are only available via paid subscriptions, and thus the developers of these tools are financially incentivized to update the tools whenever the target websites implement new security measures. The remaining two tools are free (one was open-source), and they are also well maintained, with a thriving community of contributors.

The takeaway of our study is twofold. First, many popular, widely available, and free/affordable account creation tools exist on the Web. Although we find that these tools are quite successful at making fake accounts, they are also relatively unsophisticated. The success of these tools, and the fact that they do not need to implement complicated evasive techniques, suggests that popular websites can and should implement stricter security measures against fake accounts.

[1]For ethical reasons, all accounts created in our experiments were deleted or disabled after our study was completed. This helps to minimize the impact of our experiments on the target websites.

Chapter 2

Background and Related Work

We begin our study by discussing background information related to fake accounts on major websites. *First*, we discuss the black market for fakes, including what is known about its size and economics. *Second*, we examine the various ways in which fake accounts can be abused to harm real users, and motivate why dealing with fake accounts is a serious concern on the modern Web. *Third*, we briefly survey existing techniques that websites can deploy in order to hinder the creation of fake accounts or suspend them after-the-fact. We will examine whether these security measures are actually used by real websites, and whether they are deployed correctly, in § 5.

2.1 The Market for Fake Accounts

This year several social networking sites addressed the issue of fake accounts on their respective platforms. Facebook reported 8.7% of their accounts are fake [43] whereas Twitter claims 10% of accounts are fake [56]. The fact that the two biggest OSNs are struggling with over a 100 million fake accounts each does not provide any sort of comfort to other platforms. Other OSNs have not made any public statements about the issue of fake accounts and how they plan to deal with it.

There are many underground websites that specialize in selling fake accounts on major websites in bulk. These accounts are so widely used and common that they are extremely affordable. Table 2.1 presents the average selling price of fake accounts for common OSNs being advertised on `buyacc.com` [8] in Fall 2014. The numbers reveal that basic bulk accounts are quite affordable. Phone Verified Accounts (PVA) are more expensive; our investigation in § refsec:analysis reveals that popular automated account creation tools do not have the ability to phone verify accounts, which explains why PVA accounts are more expensive.

Platform	**Amount**	**Type**	**Price**
Twitter	1000	Basic / PVA	$35/ $ 500
Facebook	1000	Basic / PVA	$50 / $600
Wordpress	1000	Basic	$25
Hotmail	1000	Basic / PVA	$14 / $120
Pinterest	1000	Basic	$70

Table 2.1: Prices for fake accounts on major websites advertised on buyacc.com.

Academic studies have uncovered crowdsourced marketplaces where you can pay real users to create fake accounts. Motoyama *et al.* looked at these "dirty jobs" on Freelancer [38], while Wang *et al.* examined market dynamics the two largest "crowdturfing" sites in China [57]. These studies reveal that millions of dollars is being spent each month on fake accounts and social spam.

The natural question that arises from these revelations is: *what are all these fake accounts being used for*? In the following sections, we explain the various nefarious uses for fake accounts.

2.2 Spam and Scams

One of the biggest reasons fake accounts are bought is to promote social media spam. In the first half of 2013 itself, social spam has increased by 355% [20]. To no one's surprise, social spammers have made a thriving business of up to $200 million on Facebook alone [2]. Some studies even suggest that the volume of junk posts are now more prevalent than real posts on social media [39].

Social spam campaigns can have a variety of objectives. The most obvious uses are promoting shady e-commerce sites, foreign pharmaceuticals, surveys, and scams [3, 22, 48, 59], *i.e.* the same kinds of content found in email spam. Social spam may also be used to spread malicious social applications that leverage the graph structure of OSNs propagate from friend to friend [28, 44]. To give an idea of the scope of this problem: 8% of 25 million URLs that are posted to Twitter point to sites that are known for phishing, scams, or malware. Unfortunately it has been shown that 90% of the visitors click on these malicious links before they are blacklisted by OSNs [25].

Another spam phenomenon that is unique to social networks is the manipulation of trending topics. Trending topics are highlighted by many OSNs, and receive many clicks and views. Thus, attackers often use fake accounts to try and create their own trending topics [58], or inject spam content into existing trending topics [26, 50].

Fake Reviews and Promotion. "Word-of-mouth marketing truly is the world's best-known marketing secret", is a common truth known among advertisers [35]. Before purchasing a product many prospective buyers browse through online reviews, in order to make an informed purchasing decision. The importance of online reviews has given rise to a market for fake reviews that promote one's own products, or slander a competitor's products [38]. Researchers have identified fake reviews on sites like Amazon, Yelp, and Tripadvisor [29–31, 41]. Studies have also shown that there hundreds websites that offer to promote products, services, and even individuals by inflating "like" and follower counts on social media [47].

2.3 Online Propaganda

Perhaps more disturbing than spam is the use of fake accounts for political ends. Astroturfing is the process of hiding a sponsored message, post or product by an organization or an individual by making it appear as if it originates from public opinion and grass-root participants. The term crowdturfing applies when astroturfing is implemented through online crowdsourcing with the help of fake profiles and accounts.

Recent evidence suggests that cases where fake accounts are used to manipulate political information are on the rise. During the Free Tibet Movement in 2013–2014, there was a large volume of protest on OSNs, especially Twitter. To counter this, many users began posting pro-China messages. However, a deeper investigation discovered that the tweets that were spreading this content were not generated by legitimate accounts [21]. In this case, fake accounts had been used to alter public opinion and start a viral spread of news and facts.

Russia too has been implicated in using fake Facebook accounts to stifle supporters of Ukraine. In this case, the content posted by fake accounts painted a picture of a nation that had gone out of control and was in need of Russia's help [24].

Even the US Military is known to develop software to manipulate social media in order to spread pro-American propaganda. A Californian company, Centcom, was awarded a $2.76 million contract for software that would allow an officer to control to control at least ten accounts or identities on OSNs. Realistic and convincing backgrounds and histories for these accounts had been promised by Centcom as part of the contract [11].

2.4 Countermeasures

In order to deal with the problems discussed above, many websites employ countermeasures designed to stop these attacks at their source, *i.e.* prevent the mass creation of fake accounts. Below, we discuss various common techniques designed to limit the creation of fake accounts, as well as their shortcomings.

- **Email Verification:** Most websites require users to supply an email address in order to sign-up for an account. The website then sends an activation email to the supplied address which contains a random link that the user must click in order to prove the ownership of the associated address. In many cases, limited or no features are offered to users that have not verified their email address. Attackers have figured out ways to circumvent email verification requirements by making dynamic email addresses on the fly. Free services offered by websites such as `fakemailgenerator.com` [18] and `mailhazard.com` [32] can be used to create dynamic, disposable email accounts. Once the registration details are entered into a target website, an automated tool can wait patiently in the inbox of the disposable email account and click the verification link as soon as it is received.

- **CAPTCHA Resolution:** The Completely Automated Public Turing test to tell Computers and Humans Apart (CAPTCHA) is a test that is used by developers to determine whether the user is human. CAPTCHA challenges have become extremely common around the Web, which has forced attackers to develop two ways around this type of challenge. They can either require a human to manually input the value of the CAPTCHA or utilize Optical Character Recognition methods to automatically determine the value of the CAPTCHA. Popular CAPTCHA solving websites such as DeCaptcha and DeathByCaptcha offer paid services to resolve CAPTCHA via APIs [37]. The rates for CAPTCHA solving on the Web are as low as $7 for 5000 CAPTCHAs [14].

- **Text and Call Verification:** Similar to email verification, some websites require that users supply a phone number while creating an account. The website can then text a random code to the user, or call them and use an automated process to read the code aloud. Users must enter the code into the website before their account will be created. Although it is slightly more expensive for attackers, phone verification challenges can be bypassed [53]. Services like `numberproxy.com` offer API access to virtual and disposable phone numbers which can be interfaced with to receive verification codes via text or calls.

- **Time and IP Address-based Restrictions:** Many websites rate limit the number of accounts that can be created from each IP address over time. Once the rate limit is met, the website may increase the number of challenges shown to the user, or entirely block them from making more accounts. Attackers can circumvent IP address-based limits by switching to fresh IP addresses obtained from proxy and VPN services. Many of these proxy and VPN services give out IP addresses for free, although faster, more reliable services are typically available for a nominal fee.

- **Challenge Questions and Puzzles:** Some websites implement their own modified CAPTCHAs in the form of challenge questions and puzzles. These sometimes take the form of asking the user to solve a math problem, listen to an audio clip and answer a question about it, or identify objects in an image. However, as we show in § 4, modern account creation tools circumvent these challenges by creating crowdsourced databases of answers that are updated by and accessible to all users of the tools.

- **Crowdsourced Flagging and Reporting.** Social media websites often include functionality that allows normal users to flag or report content/accounts that are spammy or abusive. This mechanism is essentially a crowdsourced mechanism for identifying fake accounts. There are two challenges with these systems: first, attackers can abuse them by flagging legitimate content. Thus, flagged content is typically validated a second time by trained moderators, in order to identify

spurious reports. Second, users may ignore fake accounts and spam rather than flag them. Thus, the success of crowdsourced detectors relies on the incentives offered to users to encourage them to diligently participate.

As we will show, major websites tend to implement a combination of the above techniques to combat fake accounts. Unfortunately, we will also demonstrate that many websites do not have strict enough policies, which allow automated tools to create many fake accounts even in the presence of security countermeasures.

Chapter 3

Methodology

In this section we present an overview of the methodology we will use in this study. *First*, we discuss how we gathered account creation tools. Since there are many underground scripts and paid tools that aim to create fake accounts, we select a subset of tools that 1) are popular (*i.e.* have many users and thus are high impact) and 2) target a variety of websites. We gathered these tools and information about them from various underground forums. *Second*, we present a high-level overview of how we evaluated each tool by using it to create accounts, and verifying that these accounts were not banned after several days. *Third* and finally, we discuss the ethics of our study and methodology.

3.1 Identifying Account Creation Tools

The first step in this study was to identify websites that were being actively targeted by account creation tools. We started searching on well-known black hat forums where tools are often bought, sold, traded, and advertised. Blackhatworld [6] (a popular underground SEO forum) and TheBot [49] are examples of underground forums with active threads and discussions. The discussions of tools on these forums influenced our choice of target websites. Specifically, we chose to target: Twitter, Facebook, Pinterest, LinkedIn, Wordpress, eBay, and Hotmail. All of these sites were well represented in the forums (*i.e.* many people were selling accounts from, and tools targeting, these sites) and they are extremely popular with web users in general [1].

The second step was to identify specific tools for further analysis. Ideally, we want to choose the most popular tools, since this will give us a representative sample of the tools that are being used in the wild. Unfortunately, it is difficult to say which tools are most popular, since people buy these tools anonymously; the real transaction activity between buyers and sellers happens via private messages or third-party websites.

Fortunately, many tools have active forum threads dedicated to them. On one hand, sellers and developers are incentivized to advertise their tools, and thus promote active discussion and vouches of their software. On the other hand, these threads also serve as support forums for the users of these tools. For example, the "[GET] FACreator - Fast Web 2.0 Account Creator" thread on blackhatworld.com has 433 posts and is active since 4/3/2014; the "YouTube AIO V4 BETA -Subscribe, Like, Comment, Friend Add, Video View & MORE!" thread on thebot.net has 2,749 posts and is active since 5/15/2011. Thus, we leverage thread length and activity as a proxy to identify popular account creation tools.

The third step in our methodology was actually acquiring the tools we identified as popular. This meant downloading software directly from blackhat forums, as well as shady websites promoting account creation tools. To ensure the safety of our systems, we scanned all the tools that we downloaded for viruses [46] and verified the file hashes against the most authentic versions we could find. Many files raised virus alarms, but they all appeared to be false positives.

In total, we acquired 12 tools, including seven that require payment. However, as one might expect

Tool Name	Version	Downloaded From	Cost	Target Platform(s)
Twitter Account Creator Bot	2.0.0.6	Dedicated Site	Free / Open Source	Twitter
FB Mass Account Generator	4.0.0	Dedicated Site	Paid Subscription	Facebook
PinMass	4.0	Dedicated Site	Paid Subscription	Pinterest
FACreator	1.0	Underground Forum	Paid Subscription	LinkedIn, Hotmail
Account Creator Extreme	4.2	Underground Forum	Free	Wordpress, eBay

Table 3.1: The five tools we evaluated in our experiments.

when dealing with software from the underground, several of the tools were broken, non-functional, or abandoned by the developer. After some initial testing, we ended up selecting the five tools (three paid and two free) shown in Table 3.1 for further study.

3.2 Evaluation Methodology

We now describe the methodology we use to test the capabilities of each of the tools in Table 3.1. Recall that the goal of our study is twofold: 1) we want to understand the capabilities and success rate of account creation tools, 2) we want to examine the security mechanisms currently deployed by websites against these tools, and evaluate their efficacy. The only way to accomplish these goals is to actually execute the tools in Table 3.1 and have them create as many fake accounts as possible. Clearly this experiment is true to the intended use of these tools, and by pushing them to the limit, we hope to invoke and measure the countermeasures employed by the target websites.

In total, we ran 21 tests with the selected tools. As shown in Table 3.1, the five tools target seven websites; we ran tests against each website on four separate days to take temporal variations into account. All experiments were conducted during October 2014. During each test, all network traffic to and from each tool was recorded using Wireshark. We discuss how each tool was parameterized in § refsec:tools. Each tool was allowed to run until it was forced to stop by the target website, *i.e.* security mechanisms on the remote server prevented any further account creation. The login details for all created accounts were stored and used in later validation tests, to observe whether the created accounts were deleted over time (see § 5.2).

For our tests, we configured each tool to send traffic via a VPN service and web proxies alternately. We also made sure to use free and paid VPN services to study the different factors that affect the life of a fake account. VPNbook's free VPN service [55] was used for several tests. We observed that the accounts created using the European certificate bundle were immediately suspended by Twitter. However, the fake accounts that were created using HMA Pro VPN [27] were not banned during the duration of our study. Elite Proxy [17] also provides free web proxies that could be configured for usage with most tools.

Our use of free IP addresses from these services represents a methodological tradeoff: on one hand, real attackers often leverage these services to obfuscate their traffic, so this makes our tests more realistic. On the other hand, real attackers may have used these IP addresses recently to create accounts on the same websites we are investigating (in other words, the IP addresses may be tainted). Thus, the target websites may already have limited or banned account creation from these IP addresses. As we show in § 5.2, other attackers have definitely used the same IP addresses as us, but it turns out this has minimal impact on the results of our experiments.

After our final round of account creation experiments on October 23, 2014, the accounts were kept alive for one final week so we could observe if they were banned by the target websites. As of October 31, 2014, all of the accounts were deleted or disabled in order to clean up the residue of our experiments.

3.3 Ethics

There are two ethical issues concerning our experimental methodology. *First*, we paid for three of the five tools evaluated in this study. The total cost that was spent on the tools for the entire duration of the study was <$200, which is a negligible contribution to the underground market ecosystem. Conversely, although our money is unlikely to have a significant impact on any dubious enterprises, the knowledge we gain from examining these tools will have significant positive benefits for websites that are trying to defend themselves against these attacks.

Second, by testing these tools to create fake accounts on websites, we are violating the terms of use on those sites. In total, less than <700 accounts were made in the duration of the study, and <200 accounts were made on any individual website. These numbers are negligible compared to the millions of daily active users on the websites we evaluated, and the stress or harm caused to the respective servers was minimal. Furthermore, our accounts were totally inactive: they never generated content or interacted with other accounts. Thus, we believe that the knowledge gained from this study outweigh the negligible costs that we imposed on the target websites. Finally, all the accounts that were created during this study were promptly deleted as soon as the experiments were completed.

Chapter 4

Tools

In this section, we examine the features and characteristics of the five tools highlighted in Table 3.1. We discuss the advertised capabilities of each tools, the websites they target, and how we configured each tool for the experiments conducted in § 5.

4.1 Twitter Account Creator Bot

Tool: Twitter Account Creator Bot 2.0.0.6
Obtained: `http://sourceforge.net/projects/kipesoft-acb/`
Active Since: 10/18/2014
Platforms: Windows, Linux, Mac
Cost: Free / Open Source
Features: Automated Following
Challenge Bypass: CAPTCHA Input, Email Verification, Proxy Support

Kipesoft Inc. is a famous "company" in the underground, known for creating open source automation tools. The account generating bots built by them are available to download for free or to compile from source. The bot projects are licensed under the Creative Commons v3 License, include rich documentation, and developer participation is encouraged. Twitter Account Creator Bot incorporates an automated update engine that keeps the tool up-to-date when new patches are rolled out. The Sourceforge page for Twitter Account Creator Bot shows that it benefits from active community contributions, including Pull Requests for changes and features requests, as well as tickets to report bugs. According to Sourceforge, Twitter Account Creator Bot gets downloaded 90 - 200 times per week.

Some of the features that the tool offers are:

- Automated email verification: Leverages disposable, temporary email address providers such as `fakemailgenerator.com` [18] and `mailhazard.com` [32] to verify accounts by automatically clicking on the verification links in the confirmation emails.
- Automatic updates with the latest patches
- Browser configuration: The automated browser embedded in the tool can be configured to change the User-Agent string, set proxies or clear the cache, history, temp and form data.
- Twitter API support for added customization

Besides a clean user interface, Twitter Account Creator Bot offers a visual view of the current tasks that are being performed by the tool. Figure 4.1 show a screenshot of Twitter Account Creator Bot after an account has been created and verified. Since the Twitter account creation is a step by step procedure,

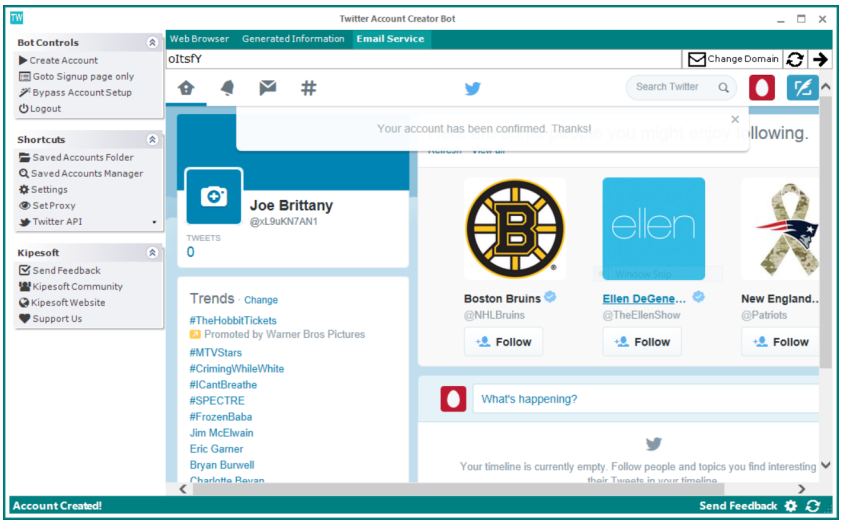

Figure 4.1: An example of Twitter Account Creator Bot successfully making a Twitter account.

from registering account to following trending profiles, a user can choose to stay or abandon the creation process at any time. This allows an added flexibility to the level of fakes one wishes to create.

In § 5, we ran experiments with Twitter Account Creator Bot by leaving all the settings at their default configuration. No additional inputs were entered by us. The generated fake accounts are stored directly in the "Generated Information" tab and can be exported to a text file. In deeper tests, we setup web proxies to test how the tool handles proxy based connections. The tool worked efficiently and outputted fake accounts as expected.

4.2 FB Mass Account Generator

Tool: FB Mass Account Generator, 4.0.0
Obtained: `http://www.latestautomationbots.com/facebook-mass-account-creator/`
Active Since: 5/6/2014
Platforms: Windows
Cost: $7 per month
Features: Automated Liking, Following, Unfollowing, and Commenting
Challenge Bypass: Email Verification

FB Mass Account Generator specialized in the creation of realistic fake accounts on Facebook. This tools requires certain input parameters to function, including: a license key (since this is a subscription-based tool), a default password for the generated accounts, sex of the accounts, and a disposable email provider (the tool includes built-in support for several). The tool includes an embedded web browser that allows users to see the various steps and progress made by the automation engine.

Some of the features of this tool are:

- Automated email verification
- Random name generation: As shown in Figure 4.2, the tool automatically fetches random realistic names for generated accounts from a web service.
- Email handler customization: Choose between different disposable email services

Unlike most other tools, this tool does not randomly generate names from a name list. It navigates to the website `onerandomename.com` and copies realistic, random names from the site. We contacted

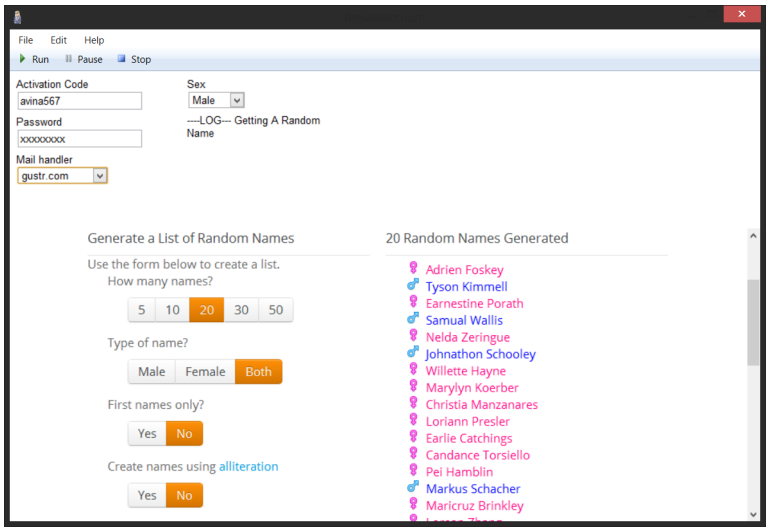

Figure 4.2: FB Mass Account Creator fetches random names for created accounts from a web service.

the author of FB Mass Account Generator and asked why the tool went to such great lengths to acquire names. As explained by the author, Facebook has a huge database of real names, which it leverages to identify names that have been constructed purely at random. Hence, the tool relies on One Random Name to give the tool an unpredictable yet not totally random name. The tool also goes to great lengths to fill out the "about me" page of the newly created fake account by getting a random quote from `onerandom.com`. As we show in § 5.2, verification tests demonstrate that the accounts generated by FB Mass Account Generator are not suspended by Facebook's security systems.

For the experiments in § 5, we made the following required configuration changes to FB Mass Account Creator. The tool required us to enter our activation code which is obtained by buying the subscription from the vendor. We also entered a password that was used for all fake accounts that we created. We alternated between creating "Male" and "Female" fake accounts in our experiments. Although FB Mass Account Creator comes with support for several disposable email providers, Facebook rejected account creation attempts when it was supplied with email addresses from most of these domains. The sole exception was `gustr.com`: Facebook accepted email addresses from this domain, and thus we used it for all of our experiments. Although FB Mass Account Creator does not include support for proxy servers, we were able to test it successfully with VPNbook and create fake accounts as expected.

4.3 PinMass

Tool: PinMass 4.0.0
Obtained: `http://www.latestautomationbots.com/pinterest-mass-account-creator`
Created on: 5/1/2014
Platforms: Windows
Cost: $7 per month
Features: Automatic Following and Unfollowing
Challenge Bypass: Email Verification

PinMass is a very sophisticated account creation tool for Pinterest that requires no configuration (other than selecting a desired number of accounts) and starts with a single click. There was no human interaction necessary during the entire course of the account creation process. PinMass not only creates

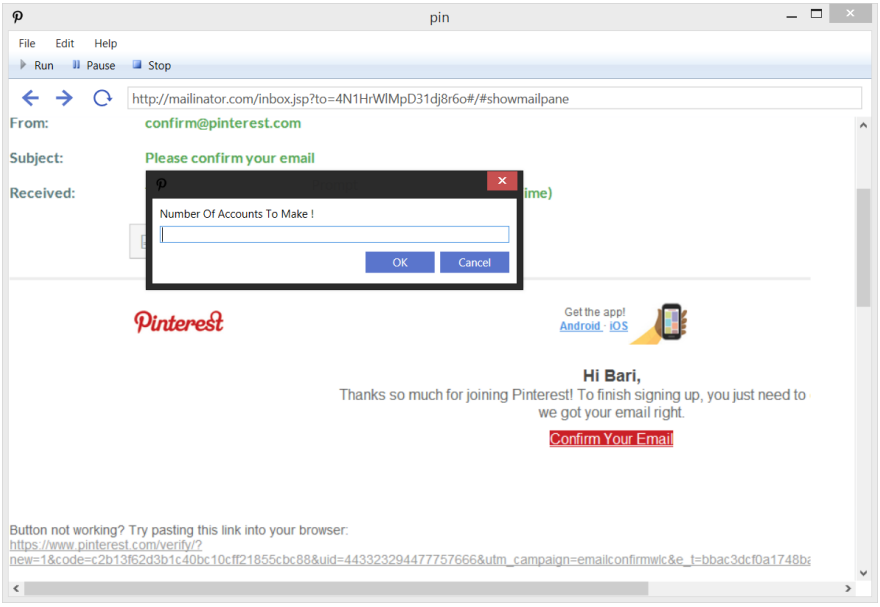

Figure 4.3: The only configuration needed with PinMass was inputting the number of accounts we wished to create.

accounts but also locates and follows a few pins randomly once each account is created, to add to the credibility of the accounts. The authors of this tool recommended that other tools be used to periodically generate some activity on the created accounts, lest they be suspended by Pinterest due to suspicious inactivity.

PinMass creates email verified accounts. It relies on disposable email addresses provided by `mailinator.com` [33]. PinMass waits for an email on the inbox page and automatically clicks on the verification link. After the creation of email verified accounts, PinMass safely outputs the credentials to an "accounts.txt" text file. This is thoughtfully done in the username:password format to ensure compatibility with other tools that require Pinterest account credentials to function.

It is interesting to note that Pinterest does not appear to deploy any challenges beyond email verification to hinder mass account creation. In our tests with PinMass, we were never served CAPTCHAs, nor did we observe IP address-based rate limiting. As we discuss in § 5.2, all of the accounts created by PinMass worked flawlessly for re-pinning and commenting.

4.4 FACreator

Tool: FACreator 1.0
Obtained: `http://www.blackhatworld.com/blackhat-seo/black-hat-seo-tools/663405-get-facreator-fast-web-2-0-account-creator-100-supported-websites.html`
Active Since: 4/3/2014
Platforms: Windows
Cost: $12 per month
Features: Creates Accounts on Many Websites, Automated Content Posting
Challenge Bypass: Captcha Resolution, Proxy Support, Security Question Bypass, Email Verification

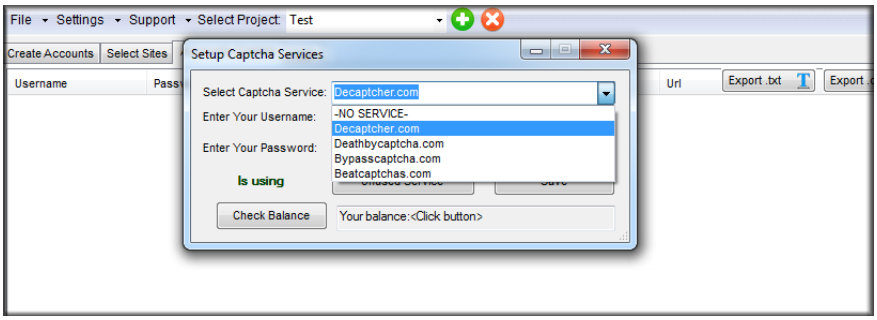

Figure 4.4: FACreator can automatically solve CAPTCHAs by relying on paid CAPTCHA-solving services.

Unlike the other tools we have studied so far, FACreator is a multi-account creation engine. FACreator can create accounts on over 100 target websites, which are selected and implemented by the developer based on public demand. To use FACreator, the user first chooses a target website from ”select site” section. Once the target is set, a new profile creation “campaign” can be created, at which point the tool is ready to begin creating accounts. FACreator randomly generates data to use in registration forms based on rules that are hard-coded for each respective target website.

Besides the standard account creation process FACreator also comes with may additional features:

- Automated email verification
- Solving challenge questions by database lookups
- Automated CAPTCHA solving
- Proxy support
- Multiple profiles
- Posts and comment submission
- Automatic updates with the latest patches

For $12 a month, FACreator offers many advanced features not available in other tools. Besides the standard automatic email verification we have seen so far, it solves security challenge questions that are presented in during the account creation process as well. It stores a database of challenge questions it may encounter and their answers obtained from a crowdsourced effort. For example, a website may ask the user to solve a math question as part of a challenge during account creation. The tool checks if this particular question was encountered before by querying the database. It then asks the user to manually enter the answer and stores the question along with the answer in a global database for future use. The next time the tool is challenged with the same question, it will perform a lookup and solve the challenge automatically.

As shown in Figure 4.4, FACreator also supports CAPTCHA-solving services such as De-Captcher [12] and DeathByCaptcha [13] by leveraging their APIs. The user selects a service and enters the username and password for their account on the corresponding service. When FACreator encounters CAPTCHAs during account creation, the images are automatically forwarded to and solved by the selected service.

FACreator allows the user to create multiple “profiles” for running different and various campaigns. This is handy for attackers who would like to create X accounts on a particular website, pause, and then create Y accounts on another website. The tool can also be configured to submit content via HTTP POST requests to websites, *e.g.* status updates, tweets, *etc*. Built-in support for mediaWiki based websites such as Wikipedia are also included. New websites get added as modules each week and are automatically updated within the tool.

FACreator required significant configuration before we were able to evaluate it. The tools requires that email addresses and passwords be supplied manually in a series of text fields; the tool will not load this information from a file. These email addresses are used to create accounts and deal with email verification. Unlike the other tools we have examined thus far, FACreator does not have automated support for disposable email services. Instead, FACreator comes bundled with a separate tool for creation of fake Hotmail and Outlook email accounts. We used the email creation tool to create Hotmail accounts, which we then used as input for an account creation campaign on LinkedIn. We tested FACreator using it's built-in support for web proxies, and successfully created accounts on both target websites. Unfortunately we did not get a chance to test the challenge question and answer database as it was filled with questions in a foreign language.

4.5 Account Creator Extreme

Tool: Account Creator Extreme 4.2
Obtained: `https://www.blackhatspot.com/Thread-GET-Account-Creator-Extreme-4-2-25-Supported-Websites`
Active Since: 3/17/2011
Platforms: Windows
Cost: Free
Features: Creates Accounts on Many Websites, Bot maker
Challenge Bypass: Proxy Support, Captcha Resolution

Account Creator Extreme is a Swiss Army Knife for generating fake accounts. Like FACreator, it has support for creating accounts on multiple websites. However, FACreator must be updated by the developer at the source code-level to support new websites; in contrast, Account Creator Extreme relies on XML configuration files that specify the necessary information to create accounts on websites. Account Creator Extreme includes a GUI tool that allows users to build new XML configurations, by: 1) browsing to a target website, 2) selecting the URL of the account registration form, 3) selecting all the pertinent form fields on the registration page and choosing a data generation algorithm to fill each one, and finally 4) selecting the button that submits the form. These XML files can then be shared with other users of the tool. Figure 4.5 shows a screenshot of Account Creator Extreme, including the tiles that represent XML configurations for different targets.

The downside to Account Creator Extreme's generalized approach is that it does not include many of the more advanced challenge-bypass features supported by other tools. The tool does automatically identify CAPTCHAs and displays them to the user, so the user can solve them.

Account Creator Extreme requires minimal configuration in order to create accounts. All the fields in the profile information are randomized, but can also be altered manually. Account Creator Extreme relies on `AirMail.com`, a disposable email service to confirm verification emails. However, this is not fully automated: it requires the user's manual clicks. A user may also choose to enter email addresses of his or her own choice for verification emails. We were able to successfully create accounts on Wordpress using the tool's default configuration options. To evaluate Account Creator Extreme's XML configuration building tools, we trained a new script for creating accounts on eBay. As shown in § 5, we were also successful in automating the creation of eBay accounts with Account Creator Extreme.

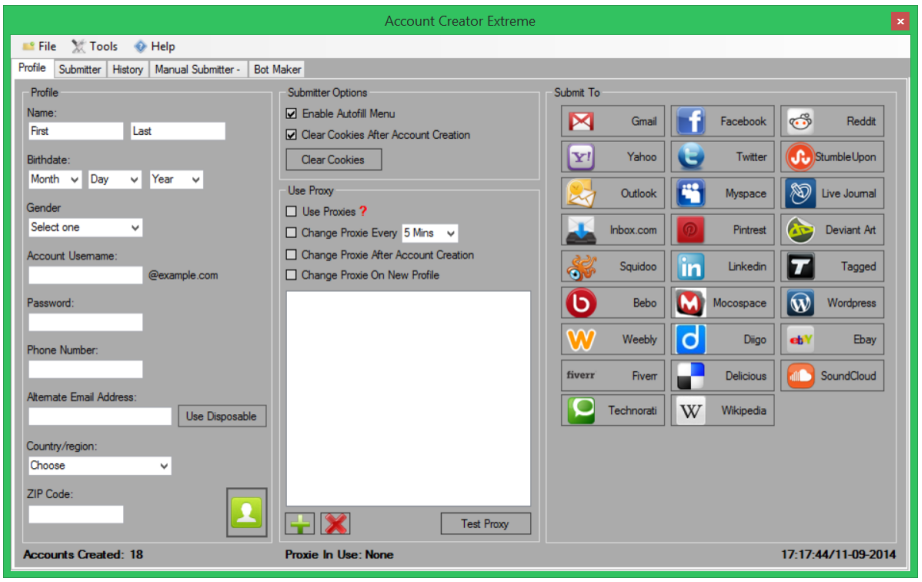

Figure 4.5: Account Creator Extreme enables easy fake account generation on over a dozen websites by leveraging crowdsourced, XML configuration files.

Chapter 5

Analysis

Now that we are familiar with the tools of the trade, our next step is to analyze their effectiveness at creating fake accounts. *First*, we present the results of our stress tests, where we measured the number of accounts that each tool could make from a single IP address. *Second*, we examine whether the accounts we created were banned or suspended over time, which would indicate that the target website noticed and responded to our account creation. *Third*, we examine the security countermeasures implemented by the target websites in order to halt mass account creation, and discuss their effectiveness or ineffectiveness. Finally, we briefly analyze whether the tools we evaluated are trustworthy, *i.e.* do they leak account credentials or other information surreptitiously back to the tool creator?

5.1 Creating Accounts

The first step in our analysis was to stress test the account generation capabilities of each tool. As discussed in § 3.2, we used five tools to create accounts on seven target websites. This was done in an effort to study the effectiveness of security countermeasures implemented by these websites to thwart automated account generation. We tested each target website four times; each of the attempts was made one week apart (from 10/2/2014 to 10/23/2014) and from a different IP address. In order to maintain the integrity of the tests and not raise suspicion, a VPN was used to ensure that the visible IP address for each attempt was never the same, but within the same /24 IP address range. Each tool was configured as described in § 4, and was allowed to run until it encountered an error that prevented it from making additional accounts. We discuss these errors in more detail below. Each individual test completed in <2 hours, and in some cases completed within minutes.

Figure 5.1 shows the total number of accounts we were able to create on each social media website before the corresponding tool was halted (refer to Table 3.1 to identify which tool was used for each of these websites). Our results show that these sites greatly differ in the strictness of their countermeasures against mass account creation. Ebay allowed us to create at most 3 accounts before banning our IP address for the day: no challenges were encountered but a simple message was displayed to let us know that the daily limit has been exceeded and to try again tomorrow. Facebook, Pinterest, and to a lesser extent LinkedIn, implement less strict policies that halted the automated tools after 10–20 accounts were created. Twitter implements a much less restrictive policy, allowing a single IP address to create >40 accounts in a single day. Shockingly, Wordpress does not seem to implement any meaningful strategies to block mass account creation: we manually shut down Account Creator Extreme after 50–60 Wordpress accounts were created, rather than allowing the tool continue to create accounts indefinitely. Hotmail accounts have not been included in the stress test, since we needed to manually solve a CAPTCHA to create each account.

In summary, the results in Figure 5.1 are somewhat disheartening. Many of the sites allow tens of accounts to be created from a single IP address. Given that attackers can acquire a large number of

IP addresses from free proxy and VPN services, as well as rent IP addresses from botnets, a dedicated attacker would be able to create thousands of accounts relatively easily. In § refsec:countermeasures, we discuss the security countermeasures we observed these websites implementing, and identify shortcomings that allowed us to create so many accounts.

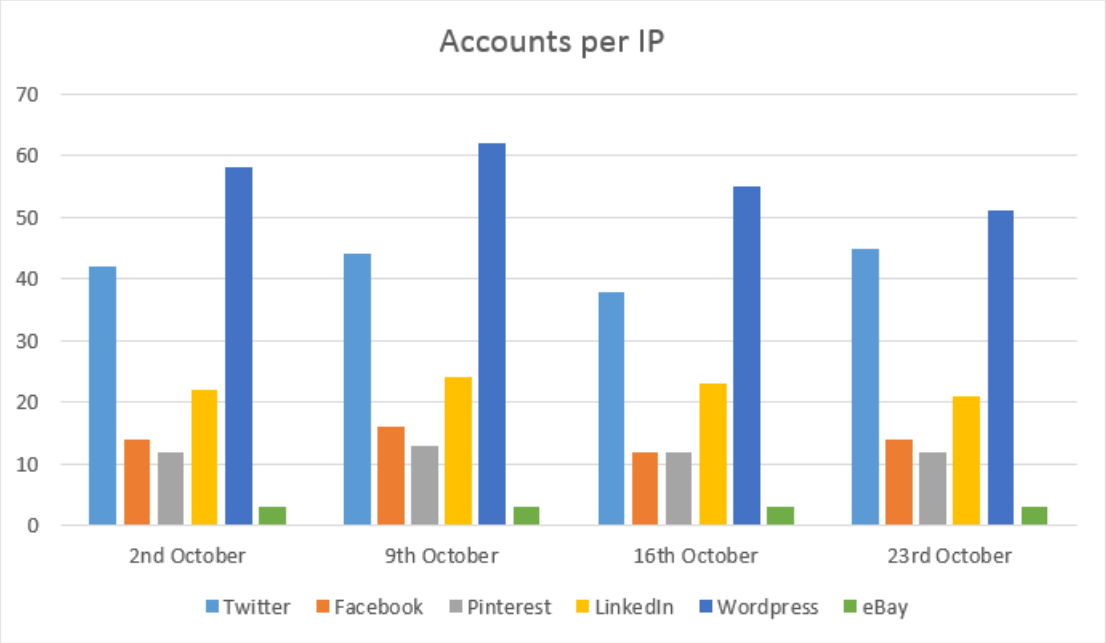

Figure 5.1: Experimental results showing the number of accounts we could create on various websites using a single IP address.

5.2 Verification

Figure 5.1 demonstrates that existing tools for automated account creation are shockingly effective. However, simply showing that account can be created is not the whole story: it is possible that the target websites may quickly identify and suspend these accounts using a batch process that runs periodically.

To determine if this was the case, we monitored all of our created accounts over a period of 24 days to see if they were suspended or banned. All of our measurements were passive, *i.e.* we just logged-in to each account to see if it was still alive. We made no attempt to make the accounts appear "active" or "real" by setting profile pictures, generating content, interacting with other users, *etc*. In practice, real attackers would almost certainly employ additional methods to make their fake accounts look real and avoid being banned over time.

Figure 5.2 shows how many of our created accounts remained active over time on various websites. In general, the lines are essentially flat, meaning that the vast majority of accounts were not banned or suspended during the 24-day observation period. The one notable exception is for Twitter accounts between the 6 and 14 day period. The banned accounts were all created using a single IP address from VPNBook using the European Certificate Bundle; thus, we assume that Twitter's systems flagged this IP address for suspicious activity and suspended all accounts recently created from that IP address. This suggests that other attackers may also have been leveraging this specific IP address to create Twitter accounts, since we used IP addresses from the same VPN to create all of our other Twitter accounts, but very few of them were suspended.

The takeaway from Figure 5.2 is that the websites we examined do not take aggressive steps to suspend fake accounts after-the-fact. We suspect that these websites may only ban fake accounts after

they have been flagged by other users, or once they begin to engage in attacks like spamming.

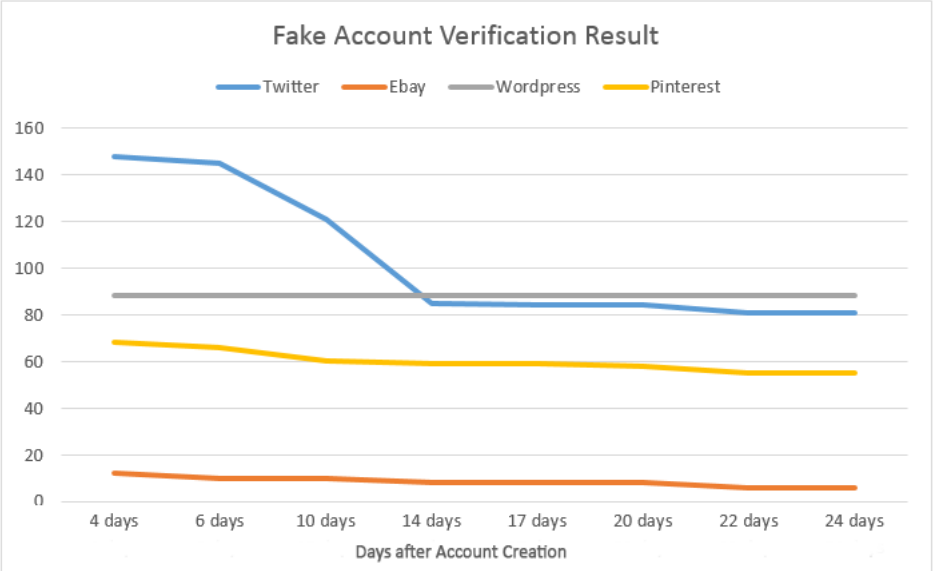

Figure 5.2: Experimental results showing how many of our created accounts were banned or suspended by various websites.

5.3 Security Countermeasures

The next issue that we examine are the countermeasures to mass account creation employed by the websites we investigated. As we discussed in § 5.1, these sites appear to vary in terms of how strictly they enforce security measures. We now discuss the countermeasures employed by each site.

Twitter. Twitter Account Creator Bot makes accounts after a few simple clicks. The tool allows users to manually click the verification links for new fake accounts by automatically loading the inbox of the accounts corresponding disposable email address. We created an average of 40 accounts before Twitter's anti-bot challenge response system commenced: in this case, a phone verification challenge was added if we wanted to create further accounts. Since the tool does not support automatically handling phone verification requests, there was no way around the challenge except to change the IP address. However, the IP address as seen by Twitter's system can be easily changed using the tool's inbuilt proxy tool. After changing the proxy one can easily create new accounts with the tool again.

Facebook. FB Mass Account Creator automatically verifies newly created accounts with the help of confirmation emails received by disposable email addresses. It created an average of 14 accounts without any challenges from Facebook. Once this limit was reached, Facebook presented us with a account creation warning banner. The banner explains how Facebook is doing everything it can to prevent the use of automated tools from registering fake accounts on their platform. We had to manually click the continue button where we were presented a CAPTCHA challenge to create accounts. After about 5 CAPTCHA account creations, Facebook presented us with a phone verification challenge. Since FB Mass Account Creator does not include mechanisms to handle phone verification, the tool could not proceed. Although FB Mass Account Creator does not support proxies, and attacker could alter their IP address using a tunnel or VPN and continue making accounts.

Pinterest. The process PinMass uses to create accounts never required any kind of interaction from us besides the initial setup. It uses the Mailinator disposable email service to verify fake accounts.

Pinterest did not require us to solve a single challenge or CAPTCHA to create accounts. However, if the accounts were not email verified it stopped us after an average of 12 accounts and demanded email verification (as shown in Figure 5.1). Since this security challenge is already bypassed with pinMass's inbuilt feature we could virtually create endless accounts.

Hotmail. FACreator's email creator tool has no web view or inbuilt browser. The user interface is optimized for user's keyboard input. The only interaction that the tool needs is CAPTCHA solving. We were required to supply a CAPTCHA either manually or by configuring a CAPTCHA solving service like DeathByCaptcha or DeCapther. Since our IP address was never blacklisted from creating new accounts by Hotmail the proxy feature of the tool was never used.

LinkedIn. FACreator comes bundled with more than 100 supported websites. Each week new support for sites are added or removed: LinkedIn was one of the sites the tool briefly supported from September, 2014 to November, 2014. We set up the tool and initiated the campaign (as explained in § 4.4). We created approximately 5 accounts before we were asked to solve a CAPTCHA challenge. If we chose not to manually solve the CAPTCHA, we could configure the CAPTCHA solving service to continue or use a web proxy to switch IP addresses, and thus bypass the challenges altogether for another 5 accounts.

Wordpress. Account Creator Extreme comes with many features and utilities to bypass most types of security challenges (refer § 4.5). To our amazement, Wordpress did not require us to solve any challenge irrespective of the number of accounts we created. We created 60 accounts as part of each stress test and eventually gave up due to exhaustion.

eBay. Given that eBay is a marketplace with high concerns for forgery and fraud on their platform, the method by which they handle automated attempts at account creation are outstanding. eBay enforces a strict daily constraint for the registration process, which is limited to 3 accounts per Ip address per day. There is no negotiating with this system: no CAPTCHAs, email, or phone verification. The only way to bypass this restriction is for the attacker to switch IP addresses.

In general, these observations confirm that each website uses different mechanisms to prevent mass account creation. eBay has the strongest policy overall, while several of the other sites leverage a *graduated response* mechanism, where the difficulty of challenges presented to users gradually ramp up as more accounts are created from an IP. We discuss the implications of these findings and provide advice for website administrators in § 6.3.

5.4 Phoning Home

One potential concern when using tools from the underground is that they may include secret features that harm the user. For example, it is possible for an account creation tool to "phone home," *i.e.* secretly send the credentials for generated accounts back to the tool creator. The tool creator could then sell the accounts or otherwise profit from them.

To determine if the tools in our study phoned home, or otherwise leaked data over the network, we recorded the activity of each tool using Wireshark. Figure 5.3 shows an example Wireshark trace from FB Mass Account Creator, highlighting a web request for a quotation to be inserted into a new accounts profile. Although it is possible that a tool author could obfuscate a leak by encrypting the traffic, we would still be able to see the outgoing packets in Wireshark; any packets being sent to strange destinations (*i.e.* destinations not related to the website we are making accounts on) are a potential red-flag that something is amiss.

We did not observe any of the five tools we analyzed passing information back to the developer. To confirm this observation, we reverse engineered the tools using commonly available disassembly and debugging tools. Based on analysis of the source code, it does not appear likely that these tools contain hidden backdoors or "phone home" routines. This makes sense, given that one of the tools we examined was open-source (so it would be difficult to hide malicious code), and three tools required paid subscriptions (the developers would lose all their customers if they stole from them). These results

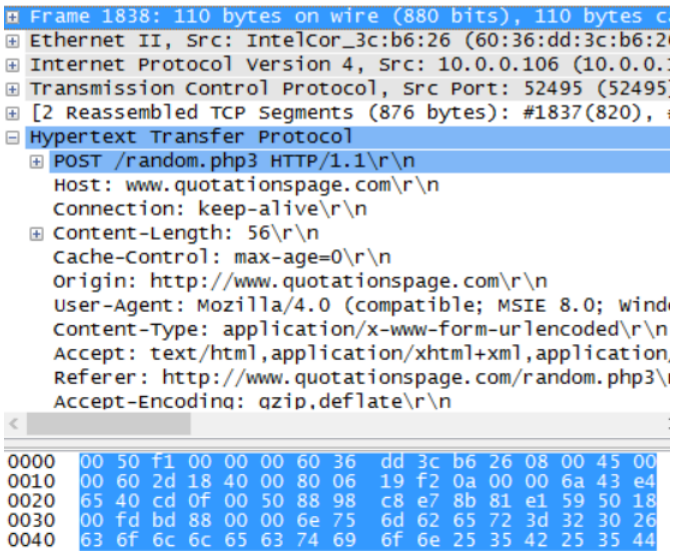

Figure 5.3: Wireshark dump of packets from FB Mass Account Creator.

point to the maturity of the market for account creation tools: these tools we examined demonstrate a high-level of professionalism, despite their origins in the underground.

Chapter 6

Discussion

Our experimental results in § 5 indicate that account automation tools are quite effective at their task, while conversely, many websites fail to implement strict security measures against mass account creation. In this section we take a broader view and discuss the implications of our findings for website operators.

6.1 Effectiveness of Existing Countermeasures

The tools we tested are readily available online and were surprisingly effective. We would expect that popular tools are blocked first by website operators, since they are so common. However, we see that many websites are still greatly vulnerable to these tools.

As seen from the results of the stress test (Figure 5.1) Twitter allowed us to create an average of 40 accounts on a single IP address with absolutely no challenge thrown. The duration it took for the bot to create these 40 accounts were approximately 30 minutes. Similarly, Wordpress allowed us to create as many accounts as we wanted. We even posted nonsensical content onto the blogs during the creation process. On the other hand, eBay allowed us to create only 3 accounts per IP address, before it blocked the creation of more.

These results suggest that websites can and should be more aggressive in limiting the number of account creations per IP address. Obviously, websites do need to allow some slack in their sign-up process to allow multiple users behind a single Network Address Translation (NAT) device. However, eBay manages to be successful and have relatively strict sign-up policies. Growth-obsessed social media sites would be wise to also adopt stricter policies.

Once in a while, websites make changes that prevent account creation tools from functioning correctly. We observed this happening several times during the course of our experiments. Paid developers are the fastest to patch their tools and release an update. Freeware and open-source developers on the other hand are not very prompt at addressing incorrect functionality immediately. Since the tool would understandingly be a part-time project for them, they may take some time before rolling out a fix for their tool. However, in both cases, changes to websites rarely render account creation tools inoperative for long, meaning that stronger security measures are needed in order to stop mass account creation.

Finally, we note that there are many additional techniques for detecting fake accounts, but these techniques only work after the fake has initiated some kind of action. These methods include:

- A sudden increase in friends, followers, likes *etc.*
- Minimal or no interaction on the profile page
- Same profile pictures on multiple accounts, or images available via Google Image Search
- Content containing spam or ads

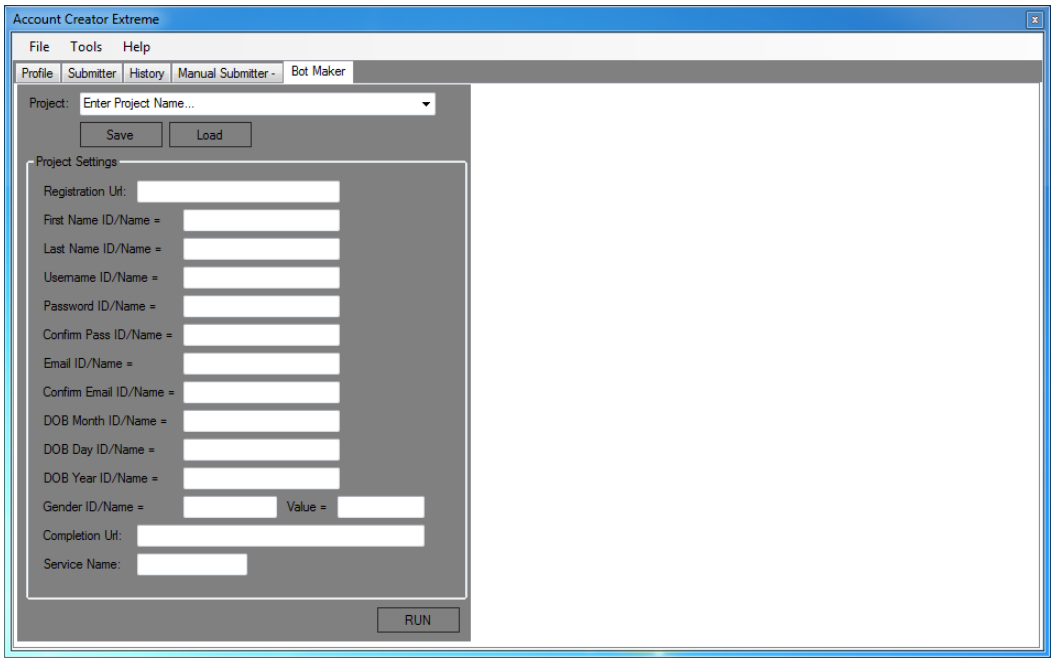

Figure 6.1: Account Creator Extreme allows users to quickly create bots with the help of simple drag and drop of in-browser HTML elements.

- Real users asking a suspicious user if they know them
- Interconnected fakes in their list of friends

Since our fake accounts did not generate any activity on the target sites, it is unclear if these techniques are being used to identify fakes. However, given that these methods are only effective after an attack has already begun, website operators should still invest in improved preemptive countermeasures against fake accounts.

6.2 Next-Generation Attack Tools

The tools we analyzed in this study incorporate many features that can bypass security challenges they face. However, since there is always space for improvement and creativity, we brainstormed ideas that could be the future of account creation tools.

There have been several cases of copied identities on OSNs [7]. In our analysis, we already see tools like FB Mass Account Generator moving away from using random information to fill out sign-up forms, towards more realistic sources of profile information. Continuing this trend, we believe that next generation account creation tools will copy data from several OSN profiles and shuffle the details as a way to create unique user profiles. Given that OSN users have surprisingly little control over the privacy of their data [51], this gives tools an opportunity to scrape OSN data and remix it for other purposes. The resulting fake accounts will have attributes from a number of different profiles, making it extremely hard to catch or detect even by the real user.

Tool authors have gone out of their way to help their customers make their own bots and share them with the community in a crowdsourced effort. As seen in Figure 6.1, Account Creator Extreme provides a bot maker. This bot maker helps users create automated scripts for websites it does not have in its database yet. Once they are created they can be exported, shared and imported by others.

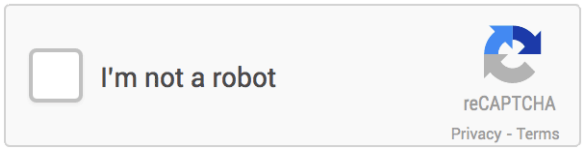

Figure 6.2: No CAPTCHA reCAPTCHA: A simple checkbox tick distinguishes automated tries.

We believe that future account creation tools will be multi-account generators instead of single targeted platform bots. These tools will require a step-by-step setup process where users can configure solvers for various challenges they are prepared to handle. This includes existing mechanisms that we observe in current tools like support for CAPTCHA solving services and disposable email address providers, as well as mechanisms we have not observed in the wild like virtual phone number services for handling phone verification and proxy/VPN services that have APIs for requesting new external IP addresses. It is only a matter of time before tools with this level of automation emerge, which will force websites to fundamentally reevaluate their countermeasures against mass account creation.

6.3 Improving Countermeasures

It is clear from our results that many websites are not implementing best practices when it comes to blocking mass account creation. One obvious countermeasure is that websites should reduce the number of accounts that can be created per IP address per day. Restricting the number of accounts that can be made per IP address per day would naturally put more pressure onto the free proxy and VPN providers (*i.e.* by forcing more attackers to use more IP addresses), which would in turn allow service providers to identify these misused IP addresses more readily. Therefore, this would reduce the number of fake accounts that are generated in the long term.

Websites can also lookout for characteristics that help identify fake accounts from real ones. For example, the email addresses that were generated by the tools for creation of fake accounts were so unreasonably long and complex that a normal person could never remember them. Websites could check the length and entropy of email addresses to identify suspicious signups. Furthermore, we were surprised that many websites accepted email addresses from disposable email providers. It is not an unreasonably difficult task to identify and ban such domains from registering new accounts on websites altogether.

It is extremely important to keep up to date with the latest technologies. New tools and libraries are constantly launched to help prevent the generation of fake accounts. Recently, Google Inc. launched the “no CAPTCHA reCAPTCHA”, an upgrade to the popular reCAPTCHA [23]. This upgrade proves to have a higher rate of confidence in determining whether the user is human or an automated bot. reCAPTCHA is available for free, has API support and is incredibly easy to incorporate. As seen in Figure 6.2, the selling point for reCAPTCHA is that the test comprises of a simple click on a checkbox. However, the inner-workings of this test are not so simple: it is a sophisticated tool that relies on a risk engine, click timing, client’s IP address and analyzes cursor movement to determine if the click was automated or not [54]. It is unclear at this point how long it will take attacks to reverse-engineer this new reCAPTCHA, and whether it can be defeated through automated means.

Chapter 7

Conclusion

Major websites provide critical functionality to billions of Internet users every day. However, some users will always try to abuse these websites and exploit their resources for personal or commercial gain. The tools we examined give us a cogent understanding of how easy it is to fabricate fake accounts on these services. These fake accounts make their way onto underground marketplaces where they can be cheaply purchased, and used to launch attacks like spam, political censorship, and blackhat SEO. The wide availability of account creation tools is proof that miscreants will find a mechanism to bypass any countermeasure put forward by websites.

We were particularly surprised by the lax security measures we observed in use by OSNs. These sites claim to be doing everything they can to deter the mass account creation; however, by scrutinizing their security systems, we see that in fact none exhibit stringent regulations, and at least one (Wordpress) appears to be doing nothing. These findings cause us to ponder the competing motivations that surround social media sites. On one hand, OSNs want to protect their users from attacks fueled by fake accounts. On the other hand, these are public companies and the stock price is influenced by the growth rate of the userbase [19]. Thus, there is some incentive for OSNs to be intentionally lax when fighting mass fakes. It is our hope that the results of this study will encourage major websites to reevaluate and strengthen their countermeasures against the mass creation of fake accounts.

Bibliography

[1] Alexa.com. Alexa - top sites in united states. http://www.alexa.com/topsites/countries/US.

[2] C. Arthur. Facebook spammers make $200m just posting links researchers say. http://www.theguardian.com/technology/2013/aug/28/facebook-spam-202-million-italian-research/.

[3] F. Benevenuto, G. Magno, T. Rodrigues, and V. Almeida. Detecting spammers on twitter. In *Proc. of CEAS*, 2010.

[4] A. Beutel, W. Xu, V. Guruswami, C. Palow, and C. Faloutsos. Copycatch: Stopping group attacks by spotting lockstep behavior in social networks. In *Proc. of WWW*, 2013.

[5] N. Bilton. Friends, and influence, for sale online. http://bits.blogs.nytimes.com/2014/04/20/friends-and-influence-for-sale-online/.

[6] Blackhatworld.com. Social networking sites. http://www.blackhatworld.com/blackhat-seo/f32-social-networking-sites/.

[7] B. Bosker. I was just friended by myself on facebook (and it only gets weirder from there). http://www.huffingtonpost.com/2012/11/09/fake-facebook_n_2104102.html.

[8] Buyaccs.com. Buyaccs.com, bulk accounts with instant delivery after payment. https://buyaccs.com/en/.

[9] buyfbpva.over blog.com. Buy fb pva accounts - buyfbpva.over-blog.com. http://buyfbpva.over-blog.com/buy-fb-pva-accounts.html.

[10] Captcha.net. The official captcha site. http://www.captcha.net/.

[11] I. Cobain. Revealed: Us spy operation that manipulates social media. http://www.theguardian.com/technology/2011/mar/17/us-spy-operation-social-networks.

[12] De-captcher.com. Short text image recognition, ocr, recognition software. http://de-captcher.com/.

[13] Deathbycaptcha.com. Death by captcha — best and cheapest captcha service! http://www.deathbycaptcha.com.

[14] Deathbycaptcha.com. Order captchas — death by captcha. http://www.deathbycaptcha.com/user/order.

[15] J. Ding. Twitter underground economy still going strong — barracuda labs. https://barracudalabs.com/2013/07/twitter-underground-economy-still-going-strong/.

[16] J. Elder. Inside a twitter robot factory. http://online.wsj.com/articles/SB10001424052702304607104579212122084821400.

[17] Eliteproxyswitcher.com. Elite proxy switcher - professional proxy software. http://www.eliteproxyswitcher.com/.

[18] Fakemailgenerator.com. Fake mail generator - free temporary email addresses. http://www.fakemailgenerator.com/.

[19] S. Feigerman. Twitter stock drops 10% on concerns about user growth in q3. http://mashable.com/2014/10/27/twitter-q3-earnings-2014/.

[20] L. Franceschi-Bicchierai. Social media spam increased 355% in first half of 2013 [study]. http://mashable.com/2013/09/30/social-media-spam-study/.

[21] Freetibet.org. Free tibet exposes hashtag chinaspam on twitter. http://freetibet.org/news-media/na/free-tibet-exposes-chinaspam-twitter.

[22] H. Gao, J. Hu, C. Wilson, Z. Li, Y. Chen, and B. Y. Zhao. Detecting and characterizing social spam campaigns. In *Proc. of IMC*, 2010.

[23] Google.com. recaptcha. https://www.google.com/recaptcha/intro/index.html.

[24] P. Gregory. Inside putin's campaign of social media trolling and faked ukrainian crimes. http://www.theguardian.com/technology/2011/mar/17/us-spy-operation-social-networks.

[25] C. Grier, K. Thomas, V. Paxson, and M. Zhang. @spam: the underground on 140 characters or less. In *Proceedings of the 17th ACM Conference on Computer and Communications Security (CCS)*, 2010.

[26] C. Grier, K. Thomas, V. Paxson, and M. Zhang. @spam: the underground on 140 characters or less. In *Proc. of CCS*, 2010.

[27] Hidemyass.com. Vpn secure virtual private network service hide my ass! https://www.hidemyass.com/pricing.

[28] T.-K. Huang, M. S. Rahman, H. V. Madhyastha, M. Faloutsos, and B. Ribeiro. An analysis of socware cascades in online social networks. In *Proc. of WWW*, 2013.

[29] N. Jindal and B. Li. Opinion spam and analysis. In *Proc. of WSDM*, 2008.

[30] N. Jindal, B. Liu, and E.-P. Lim. Finding unusual review patterns using unexpected rules. In *Proc. of CIKM*, 2010.

[31] E.-P. Lim et al. Detecting product review spammers using rating behaviors. In *Proc. of CIKM*, 2010.

[32] Mailhazard.com. Mailhazard.com : Anonymous and disposable email for you ! http://mailhazard.com/.

[33] Mailinator.com. Mailinator. http://www.mailinator.com/.

[34] R. McMillan. 1.5 million stolen facebook ids up for sale. http://www.infoworld.com/article/2626660/hacking/1-5-million-stolen-facebook-ids-up-for-sale.html.

[35] I. Misner. Word-of-mouth: The world's best-known marketing secret. http://www.entrepreneur.com/article/53188.

[36] G. Monbiot. The need to protect the internet from 'astroturfing' grows ever more urgent — george monbiot. http://www.theguardian.com/environment/georgemonbiot/2011/feb/23/need-to-protect-internet-from-astroturfing.

[37] M. Motoyama, K. Levchenko, C. Kanich, D. McCoy, G. M. Voelker, and S. Savage. Re: Captchas – understanding captcha-solving from an economic context. In *Proc. of USENIX Security*, 2010.

[38] M. Motoyama, D. McCoy, K. Levchenko, S. Savage, and G. M. Voelker. Dirty jobs: The role of freelance labor in web service abuse. In *Proc. of Usenix Security*, 2011.

[39] M. Neal. Inside the underground spam machine that's overrunning social media — motherboard. http://motherboard.vice.com/blog/inside-the-underground-spam-machine-thats-overrunning-social-media.

[40] M. Online. Celebrities buying 'bogus' facebook likes from 'click farms'. http://www.dailymail.co.uk/news/article-2534488/Celebrities-businesses-government-departments-buying-bogus-Facebook-likes-click-farms.html.

[41] M. Ott, Y. Choi, C. Cardie, and J. T. Hancock. Finding deceptive opinion spam by any stretch of the imagination. In *Proc. of ACL*, 2011.

[42] A. Press. Fake facebook, twitter and youtube clicks are big business. http://nypost.com/2014/01/05/fake-facebook-twitter-and-youtube-clicks-are-big-business/.

[43] E. Protalinski. Facebook estimates between 5.5% and 11.2% of accounts are fake. http://thenextweb.com/facebook/2014/02/03/facebook-estimates-5-5-11-2-accounts-fake/.

[44] M. S. Rahman, T.-K. Huang, H. V. Madhyastha, and M. Faloutsos. Efficient and scalable socware detection in online social networks. In *Proc. of USENIX Security*, 2012.

[45] O. V. Scanners and P. Scans. Tor exit nodes mapped and located — hackertarget.com. http://hackertarget.com/tor-exit-node-visualization/.

[46] Security. [buy] twitter account creator [online forum comment]. Retrieved from http://www.ubotstudio.com/forum/index.php?/topic/10547-buy-twitter-account-creator/.

[47] G. Stringhini, M. Egele, C. Kruegel, and G. Vigna. Poultry markets: On the underground economy of twitter followers. In *Proc. of WOSN*, 2012.

[48] G. Stringhini, C. Kruegel, and G. Vigna. Detecting spammers on social networks. In *Proc. of ACSAC*, 2010.

[49] Thebot.net. Botting world. http://thebot.net/forums/botting-world.8/.

[50] K. Thomas et al. Suspended accounts in retrospect: An analysis of twitter spam. In *Proc. of IMC*, 2011.

[51] K. Thomas, C. Grier, and D. Nicol. unFriendly: Multi-Party Privacy Risks in Social Networks. In *Proceedings of the 10th Privacy Enhancing Technologies Symposium*, 2010.

[52] K. Thomas, C. Grier, and V. Paxson. Adapting social spam infrastructure for political censorship. In *Proceedings of the 5th USENIX Workshop on Large-Scale Exploits and Emergent Threats*, 2012.

[53] K. Thomas, D. Iatskiv, E. Bursztein, T. Pietraszek, C. Grier, and D. McCoy. Dialing back abuse on phone verified accounts. In *Proceedings of the 21st Annual Conference on Computer and Communications Security*, 2014.

[54] C. Velazco. Google now lets you prove your humanity with a single click. `http://www.engadget.com/2014/12/03/google-nocaptcha-recaptcha/`.

[55] Vpnbook.com. Free vpn 100% free pptp and openvpn service. `http://www.vpnbook.com/`.

[56] K. Wagstaff. 1 in 10 twitter accounts is fake, say researchers. `http://www.nbcnews.com/tech/internet/1-10-twitter-accounts-fake-say-researchers-f2D11655362`.

[57] G. Wang, C. Wilson, X. Zhao, Y. Zhu, M. Mohanlal, H. Zheng, and B. Y. Zhao. Serf and turf: crowdturfing for fun and profit. In *Proc. of WWW*, 2012.

[58] Z. Yang, C. Wilson, X. Wang, T. Gao, B. Y. Zhao, and Y. Dai. Uncovering social network sybils in the wild. In *Proc. of IMC*, 2011.

[59] S. Yardi, D. Romero, G. Schoenebeck, and D. Boyd. Detecting spam in a twitter network. *First Monday*, 15(1), 2010.