

Why Allowing Profile Name Reuse Is A Bad Idea

Enrico Mariconti[‡], Jeremiah Onalapo[‡], Syed Sharique Ahmad[†], Nicolas Nikiforou[‡],
Manuel Egele^{*}, Nick Nikiforakis[†], Gianluca Stringhini[‡]

[‡]University College London

[†]Stony Brook University

^{*}Boston University

[‡]{e.mariconti,j.onalapo,n.nikiforou,g.stringhini}@cs.ucl.ac.uk

[†]{syahmad,nick}@cs.stonybrook.edu

^{*}megele@bu.edu

ABSTRACT

Twitter allows their users to change profile name at their discretion. Unfortunately, this design decision can be used by attackers to effortlessly hijack user names of popular accounts. We call this practice *profile name squatting*. In this paper, we investigate this name squatting phenomenon, and show how this can be used to mount impersonation attacks and attract a larger number of victims to potentially malicious content. We observe that malicious users are already performing this attack on Twitter and measure its prevalence. We provide insights into the characteristics of such malicious users, and argue that these problems could be solved if the social network never released old user names for others to use.

1. INTRODUCTION

When creating a Twitter profile, a user is asked to pick multiple elements to help other people recognize that the account is linked to that particular user. The user is required to input his name (*Justin Bieber* in the example in Figure 1), a profile picture, and a profile name, which will make the user memorable by assigning him a handle (*@justinbieber* in the example in Figure 1). Twitter incorporates the user's selected name in an easy to remember URL, such as <https://twitter.com/justinbieber>. At the same time, the account is assigned a numeric identifier by the social network. This identifier is used by the website to keep track of the account and is usually never seen by the users of the network, who instead identify accounts based on profile names, actual names, and profile pictures. In the case of the account in Figure 1, the user ID that is internally used by Twitter is 27260086. A user can also get his account verified. If the account verification process succeeds, a verification seal is displayed on the user's profile page, increasing the trustworthiness of that account (see Figure 1).

Previous research showed that malicious actors can fool social network users by setting up accounts that look similar to popular accounts on the same network [15]. In this case,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Copyright 2016 ACM 978-1-4503-4295-7/16/04 ...\$15.00.



Figure 1: Different identity-related elements on a Twitter profile: the profile name (①), the user's actual name (②), and the verified account seal (③).

attackers would set up profile names that are similar to the ones of popular accounts and use the same profile pictures of their targeted accounts. By looking very similar to popular accounts and using the same profile name, malicious users can attract a large number of victims and have a higher success rate in their illicit activities (e.g., having a higher fraction of users clicking on links pointing to malware).

Profile name squatting on Twitter. Twitter allows its users to change their profile names, freeing up the old, no longer used profile names. The ability to change their profile names can be useful to users, but it has security repercussions if the old profile name returns back to the “pool” of available profile names. Malicious users can wait for a popular account to change their profile name and immediately seize the old one, by changing the profile name of an account that they control to this recently-relinquished name. Such attacks provide multiple advantages to the perpetrator. First, controlling a popular profile name allows the attacker to perform more successful impersonation attacks than the ones described in previous work [15]. Second, profile names also determine the URL of the profile page on Twitter. When a user changes his profile name, multiple URLs on the web are likely to still point to the profile page associated with the old profile name. These links now au-

tomatically point to the attacker’s profile who can abuse incoming visitors in a wide range of ways.

In this paper, we show that malicious users are actively using this feature to increase their reputation on Twitter. Since these attacks involve the actual abandoned profile name of the victim, they can be much more effective than other types of impersonation, such as the ones shown in [15].

In summary, we make the following contributions:

- We show that profile name squatting on Twitter is possible, because the social network frees changed profile names after their users pick new ones.
- We show that profile name squatting is happening in the wild on Twitter, by identifying 39,453 accounts that took someone else’s old profile name over a period of one month.
- We argue that, from a security perspective, freeing profile names after they are abandoned is not a good design decision, and that it should be avoided to prevent the described attacks.

2. DATA COLLECTION

Our dataset consists of a 10% random sample of all public tweets posted on Twitter over a period of one month, between February and March 2015. In total, the dataset contains 985,538,449 tweets posted by 66,679,721 distinct users. Each tweet was collected in a json format and included both information about the message and about the account that posted it. We call the dataset of the 66 million users \mathbf{U} . We analyzed \mathbf{U} to identify accounts potentially performing name squatting. In the following, we describe our methodology in detail.

To identify a set of accounts that adopted a profile name previously used by another account, we proceeded as follows. First, we created a bipartite graph to represent the mapping between profile names and user IDs on Twitter. From the accounts in \mathbf{U} we extract two set of vertices, a set \mathbf{S} of the user names used by the accounts in \mathbf{U} and a set \mathbf{I} of the user IDs of such accounts. Each vertex s in \mathbf{S} is connected to a vertex i in \mathbf{I} if s was used as a profile name by the account identified by i . We call this bipartite graph \mathbf{G} .

To identify which accounts have used the same profile name at different points in time, we perform the weighted one mode projection of the bipartite graph \mathbf{G} onto the account vertex set \mathbf{I} . This projection generates a graph in which each vertex identifies a user ID and two vertices are connected by an edge if they shared at least one profile name. The weight of the edge is the number of profile names that the two accounts have shared in the past [31]. This graph-based approach has been used to detect malicious accounts in past work and has two advantages. First, it allows us to quickly identify accounts of interest: any account that is connected to at least another one has either taken someone else’s old profile name or has given up a profile name that has been taken by someone else. The second advantage is that accounts that have taken multiple profile names over time form connected components in this graph, and can therefore be easily identified and further analyzed. We then extracted the set of accounts that used a certain profile name for the first time in our dataset. We did this by looking at the tweets sent by the users in \mathbf{U} . We call this set of accounts, \mathbf{F} . We consider all the other accounts that obtained a profile

name after someone else owned it in the past as potential name-squatters. We call this set of accounts, \mathbf{A} .

In our dataset, we found 39,453 accounts that shared 20,891 distinct profile names on Twitter. 5,207 accounts shared more than one profile name with each other, showing that the ecosystem of name squatters is quite complex. In Section 3 we analyze the identified accounts in detail.

Limitations. Our dataset allowed us to gain a good overview of the name squatting practices on Twitter. However, our data has some limitations. First of all, we can only identify two accounts as using the same profile name if they tweeted during the observation period. If an account took another user’s previous profile name and never posted a tweet, we would not detect it. If the original account stopped posting messages before our data collection started, or an account took a profile name after the data collection ended, we would not identify it. Similarly, many of our analyses (such as the ones performed in Section 3) rely on the fact that the first user observed using a profile name is its original holder. It could be, however, that this account has taken a profile name that was abandoned after our data collection started. We believe that although these cases are possible they have not affected the correctness of our results and overall takeaway message. In fact, most of the accounts in our dataset only appeared with a single profile name during the entire observation period.

Ethics. Dealing with online social network data raises ethical concerns. In this paper, we only used publicly-available Twitter data, and we successfully obtained ethics approval from UCL (Project ID 6521/004). To treat data ethically, we followed the guidelines by Rivers et al. [28]. In particular, we ensured not to link multiple datasets together with the goal of further deanonymizing the users contained in them, and we stored our data according to the UCL Data Protection Officer guidelines.

3. ANALYSIS OF THE DATA

In this section, we explain the differences on basic features of name squatting users. Previous research showed that malicious accounts that get suspended on Twitter, account for roughly 3% of all accounts on the social network [35]. Twitter routinely suspends accounts that engage in malicious activities such as spreading spam or links pointing to malware, as well as accounts that more broadly violate its terms of service [1]. Although taking an account’s old profile name does not violate Twitter’s terms of service by itself, if an account uses such a profile name to impersonate other users, this is a violation of the terms and can lead to the account getting suspended. Typically, the original holder of the profile name has to file a complaint to Twitter which then checks the claim and decides the validity of the impersonation claim.

3.1 Maliciousness of name squatting accounts

To assess whether accounts that engage in name-squatting are more likely to be suspended by Twitter, we proceeded as follows. For each of the 39,453 accounts that we identified, we checked whether the account was suspended or not. We performed this test in September 2015, six months after the accounts were observed to be active in our dataset. We allowed this time to pass to make sure that possible impersonation claims of profile names in our dataset would have been sorted out by Twitter during that period. In total,

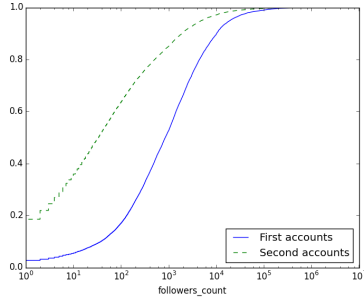


Figure 2: Cumulative distribution function of the number of followers of the first account using a certain username compared to the ones of the second account using the same username. Second accounts generally have less followers.




Figure 3: Cumulative distribution function of the number of statuses of the first account using a certain username compared to the ones of the second account using the same username. Second accounts have usually been less active on the social network than first accounts.




Figure 4: Cumulative distribution function of the time since the opening of the account of first and second accounts, compared to the last time the account was observed active in our dataset. Second accounts have generally been active for a shorter period of time than first accounts.

7,673 accounts (19%) in our dataset were suspended. This number shows that accounts reusing someone else’s old profile name are more likely to engage in malicious activity and get suspended by Twitter. On the other hand, it also shows that the majority of accounts engaging in such practices were still active on the platform six months after the profile name change. This might indicate that many accounts are used for legitimate purposes. In Section 3.3, however, we provide evidence that most of these accounts are indeed malicious, and that the lack of a clear threat model and the need for users to manually report impersonator accounts might be the reason that many of these accounts manage to avoid detection and stay active for a long period of time.

3.2 Characteristics of name squatting accounts

As we mentioned in Section 2, we separate accounts that used the same profile name into two categories: the accounts that were the first ones using a profile name in our dataset (referred to as **F**) and the accounts that used a profile name after the first accounts abandoned it (referred to as **A** or second accounts). Putting aside the possible biases due to the observation window that we discussed in Section 2, we consider the first accounts as the legitimate ones, while the second accounts are likely to be malicious.

In this section, we characterize the differences between the first accounts and the second ones that shared the same profile name. To show these differences, we produced the Cumulative Distribution Function (CDF) plots of some basic characteristics of these Twitter accounts. Figure 2 shows the CDF of the followers of such accounts. As we can see, the number of followers is usually lower in second accounts than in first ones. Considered that some second accounts are clearly malicious, people often do not follow them and this is why almost one out of five second accounts does not have followers at all. The beginning of the line is different, but it is easy to notice that both lines are smooth and quite similar, indicating similar distributions with different median values. This graph is similar to the ones related to favorites and friends’ counts, therefore we did not include these two graphs to avoid redundancy. Favorites and friends are other two classic profile features: the favorite number is the sum of all the favorites that each tweet of the profile received while

the friends are those Twitter profiles that are followed by the account we are studying.

Figure 3 shows the differences in the number of statuses of the two types of accounts: similar to the previous plot, the line for the second accounts starts from a higher y-intercept but, in this case, only a small number of second accounts did not publish any status updates. As for the followers figure, the two lines get close only when almost all the accounts are taken into consideration; there are differences in the shapes of the lines: while the line related to the first accounts is again smooth, the one related to the second accounts has a high slope in the first part, and at about 20 statuses the slope changes into a less strong one.

In Figure 4, we checked the life in days of the two types of accounts. The definition of second account using the same name implicitly considers a temporal question that probably affects the fact that this graph is the one where the difference between first and second accounts is still easily visible when all the accounts are taken in consideration. However, it is necessary to underline that, while the first accounts line maintains a smooth shape, the second accounts one changes again: as in Figure 2 the first part of the line is a smooth curve, while there is a discontinuity point from where the slope is more linear as in Figure 3.

All three plots show differences between the two types of accounts. This indicates that there are characteristics that identify the different nature of these accounts. The fact that accounts that take someone else’s abandoned profile name present a lower level of popularity and engagement from other Twitter accounts, however, casts a doubt on the effectiveness of name-squatting as a tool for malicious users. This is in line with what was discovered by previous work regarding other forms of reputation boost on social networks [32].

3.3 Characteristics of account clusters

We then analyzed the relationships between the accounts that reuse the same user names. We identified a number of accounts that engage in interesting behaviors by sharing the same set of profile names and adopting it at different points in time. We found 1,587 groups of such accounts, composed of 5,208 accounts. Most of the groups were not larger than five accounts each but the largest one was composed of 37 ac-

counts. Figure 5 reports a CDF summarizing these groups. This practice of sharing a set of profile names among different accounts is particularly suspicious, and we were not able to discern a legitimate reason for why regular accounts should engage in such practices. Interestingly, only 10% of the accounts belonging to these groups were suspended by Twitter. We therefore decided to further investigate these groups, by analyzing their activity.

We first analyzed the largest of the groups, accounting for 37 profile names. Interestingly, we found that the accounts in this group are not performing any clearly malicious activity: the reused profile names are used by players interested in sniper video games. Players set up Twitter accounts and use them to link videos of their games and interact with the other players; sometimes, when an expert player quits playing the game and deletes their Twitter account, another user would take her name, therefore generating this group. Although they are not posting any malicious content, the accounts in this group take advantage of the reputation of the previous owners of profile names. This behavior violates the 10th chapter of Twitter’s terms of service [1] and creates opportunities for attackers to post videos in place of the player who abandoned. These videos may contain malicious contents that would be downloaded at least by the other gamers of the group.

The other groups analyzed were either composed of malicious accounts or by accounts that may have been created with a benign intent but that violated the Twitter terms of service [1]. It is interesting to observe common aspects in those clusters: all the accounts portrayed young people (especially female) who are part of the East Asia show-business: models, dancers, singers and so on. In some cases, profile names in groups were similar to famous video games (such as World of Warcraft). In general, these accounts performed various questionable activities: from sharing pornography to posting spam, passing through advertising services that promise to increase the number of followers [32]. In some cases, the accounts in groups would not advertise a Twitter follower market, but would post messages advertising some sort of re-follow ring (e.g., “retweet my message, and follow me, I will follow you”). These messages are in clear violation with Twitter’s terms of service [1]. Another interesting feature of the communications is that all the people tagged in the tweets are not part of the group of name squatting accounts, but always external accounts.

4. DISCUSSION

We showed that name squatting is a real, albeit under-exploited, problem on Twitter. Detecting name squatting accounts is challenging due to a vague threat model and the variety of malicious activities (more or less evident) that a name squatting account can engage in on the social network. In the following, we discuss possible mitigations to the name squatting problem, as well as some potential future work.

Mitigating profile name squatting. The best solution to solve the problem of profile name squatting on Twitter is not to allow profile name reuse. Although allowing users to change their profile name can have legitimate applications, we do not see any particular reason why it should be possible to take someone else’s abandoned profile name on a social network. We therefore consider the current policy adopted by Facebook and LinkedIn — to allow users to change their

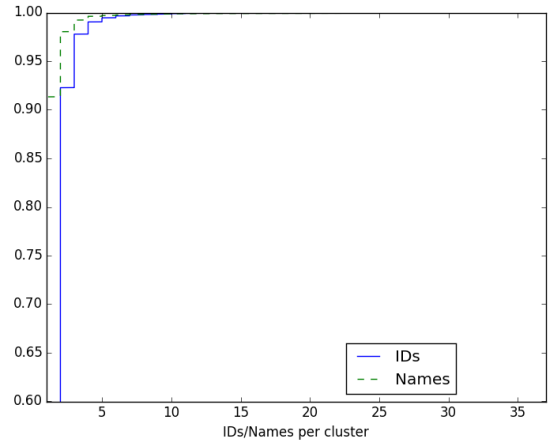


Figure 5: Relative Cumulative Distribution Function of the number of IDs and the number of names in each analyzed cluster.

profile name but not releasing it afterwards — as the optimal one, and we encourage Twitter to adopt it.

In case Twitter considers freeing abandoned profile names a possibility that they do not want to renounce, there are other partial mitigations that could be taken. For example, accounts adopting someone else’s old profile name could be automatically vetted similarly to what happens for accounts that are flagged for impersonating another user [15]. In particular, accounts that change their user names multiple times to different abandoned profile names are extremely likely to be malicious and should be suspended.

Future work. In this paper we shed light on the phenomenon of profile name squatting on Twitter. In future work, we will study the phenomenon in more detail by, for example, analyzing a longer period than the one month studied in this paper. We will also investigate whether other social networks have similar name squatting problems.

5. RELATED WORK

A number of researchers analyzed Twitter and other social networks and their properties, to understand its general characteristics [21, 38], understand reputation on the network [7], investigate how different Twitter is, compared to traditional social networks [22], or looking at patterns in the unfollow dynamics of users [18].

A line of research that has been followed particularly is understanding the security issues linked to social networks and the extent of the abuse happening on them. Bilge et al. showed that identity theft attacks on social networks are a concrete possibility [6]. Grier et al. performed the first large-scale analysis of abuse on Twitter [17], while Gao et al. performed a similar study on Facebook [14]. Thomas et al. studied the ecosystem of marketplaces where one can buy fake Twitter accounts [36]. Following these studies, a number of systems have been proposed to automatically detect fake accounts and malicious activity on social networks [5, 11, 13, 23, 24, 30, 34, 39].

Recently, cybercriminals abusing social networks developed techniques to look more trustworthy to potential victims and therefore be more successful in their illicit activities. Stringhini et al. studied services that sell compromised

accounts as followers to customers that are willing to pay for them [29, 32]. De Cristofaro et al. studied the ecosystem of services that deliver likes to the Facebook pages of their customers [8]. Oana et al. have recently studied the problem of impersonation on Twitter [15]. In their work, Oana et al. studied accounts that purposely look similar to popular ones to confuse Twitter users, and developed a classifier to automatically detect such impostor accounts. In this paper we analyzed a different type of Twitter impersonators, those that reuse the old username of another account. We showed that these accounts are particularly dangerous, because they have the capability of leveraging links on the web that are still pointing to the old username. To the best of our knowledge we are the first ones studying the security implications of allowing users to change their usernames on social networks.

Even though we are the first to report on how name squatting applies to social networks, name squatting has attracted a significant amount of research in the area of domain names. *Cybersquatting* originally referred to the early days of the internet (early nineties) when long-existing brick-and-mortar companies did not yet operate websites. Various opportunists registered their trademarks as domain names before them, so that they would sell the domains back to their rightful owners for profit [16].

Cybersquatting evolved into *typosquatting*, i.e., the act of registering domains that are mistypes of popular authoritative domains, with the intention of capturing the traffic of users who mistype URLs in their browser address bar. This practice can be traced back to over 15 years, since the 1999 Anticybersquatting Consumer Protection Act (ACPA) already mentioned URLs that are “sufficiently similar to a trademark of a person or entity” [2]. Researchers have extensively studied typosquatting in domain names [3, 4, 10, 20, 25, 33, 37], and have measured, among others, its prevalence, the way that attackers monetize typosquatting domains, whether typosquatting domains exchange hands and whether users would be victimized by malicious pages resulting from unintended typographical mistakes.

Finally, apart from typosquatting, there also exist other, less popular, types of domain squatting, such as domains that abuse the visual similarity of characters in different character sets [12, 19], domains that sound like authoritative domains [27], and domains that capture the traffic originating from erroneous bit-flips in user devices [9, 26].

6. CONCLUSION

In this paper we studied the security implications of allowing username reuse on online social networks, and showed that Twitter’s profile names are vulnerable to name squatting practices. We showed that the threat posed by name squatting is not only theoretical, but that malicious users are actively reusing abandoned profile names on Twitter for illicit purposes. To solve this problem, we advise Twitter to disallow profile name reuse.

7. ACKNOWLEDGMENTS

We wish to thank the anonymous reviewers for their comments. This work was funded by the H2020 RISE Marie Skłodowska Curie action (MSCA) grant number 691925, by the EPSRC grant number EP/N008448/1, and by the NSF

under grant CNS-1527086. Enrico Mariconti was funded by the EPSRC under grant 1490017, while Jeremiah Onaolapo was supported by the Petroleum Technology Development Fund (PTDF), Nigeria.

8. REFERENCES

- [1] Twitter Terms of Service. <https://twitter.com/tos>.
- [2] Anticybersquatting Consumer Protection Act (ACPA). <http://www.patents.com/acpa.htm>, November 1999.
- [3] AGTEN, P., JOOSEN, W., PIESSENS, F., AND NIKIFORAKIS, N. Seven months’ worth of mistakes: A longitudinal study of typosquatting abuse. In *Proceedings of the 22nd Network and Distributed System Security Symposium (NDSS 2015)* (2015).
- [4] BANERJEE, A., BARMAN, D., FALOUTSOS, M., AND BHUYAN, L. N. Cyber-fraud is one typo away. In *Proceedings of the 27th Conference on Computer Communications, IEEE INFOCOM* (2008).
- [5] BENEVENUTO, F., MAGNO, G., RODRIGUES, T., AND ALMEIDA, V. Detecting Spammers on Twitter. In *Conference on Email and Anti-Spam (CEAS)* (2010).
- [6] BILGE, L., STRUFE, T., BALZAROTTI, D., AND KIRDA, E. All Your Contacts Are Belong to Us: Automated Identity Theft Attacks on Social Networks. In *World Wide Web Conference (WWW)* (2009).
- [7] CHA, M., HADDADI, H., BENEVENUTO, F., AND GUMMADI, K. Measuring User Influence in Twitter: The Million Follower Fallacy. In *International AAAI Conference on Weblogs and Social Media* (2010).
- [8] DE CRISTOFARO, E., FRIEDMAN, A., JOURJON, G., KAAFAR, M. A., AND SHAFIQ, M. Z. Paying for likes?: Understanding facebook like fraud using honeypots. In *Internet Measurement Conference (IMC)* (2014).
- [9] DINABURG, A. Bitsquatting: DNS Hijacking without Exploitation. In *Proceedings of BlackHat Security* (July 2011).
- [10] EDELMAN, B. Large-scale registration of domains with typographical errors, September 2003.
- [11] EGELE, M., STRINGHINI, G., KRUEGEL, C., AND VIGNA, G. Compa: Detecting compromised accounts on social networks. In *Proceedings of the Network and Distributed System Security Symposium* (San Diego, CA, February 2013).
- [12] GABRILOVICH, E., AND GONTMAKHER, A. The homograph attack. *Communications of the ACM* 45, 2 (Feb. 2002), 128.
- [13] GAO, H., CHEN, Y., LEE, K., PALSETIA, D., AND CHOUDHARY, A. Towards Online Spam Filtering in Social Networks. In *Symposium on Network and Distributed System Security (NDSS)* (2012).
- [14] GAO, H., HU, J., WILSON, C., LI, Z., CHEN, Y., AND ZHAO, B. Detecting and Characterizing Social Spam Campaigns. In *Internet Measurement Conference (IMC)* (2010).
- [15] GOGA, O., VENKATADRI, G., AND GUMMADI, K. P. The doppelgänger bot attack: Exploring identity impersonation in online social networks. In *Proceedings of the 15th Internet Measurement Conference (IMC)* (2015).
- [16] GOLINVEAUX, J. What’s in a domain name: Is cybersquatting trademark dilution? In *University of*

San Francisco Law Review 33 *U.S.F. L. Rev.* (1998-1999) (1998).

- [17] GRIER, C., THOMAS, K., PAXSON, V., AND ZHANG, M. @spam: the underground on 140 characters or less. In *ACM Conference on Computer and Communications Security (CCS)* (2010).
- [18] H.KWAK, LEE, C., PARK, H., AND MOON, S. What is Twitter, a social network or a news media? In *World Wide Web Conference (WWW)* (2010).
- [19] HOLGERS, T., WATSON, D. E., AND GRIBBLE, S. D. Cutting through the confusion: a measurement study of homograph attacks. In *Proceedings of the 2006 USENIX Annual Technical Conference* (2006).
- [20] KHAN, M. T., HUO, X., LI, Z., AND KANICH, C. Every second counts: Quantifying the negative externalities of cybercrime via typosquatting.
- [21] KRISHNAMURTHY, B., GILL, P., , AND ARITT, M. A Few Chirps About Twitter. In *USENIX Workshop on Online Social Networks* (2008).
- [22] KWAK, H., CHUN, H., AND MOON, S. Fragile online relationship: a first look at unfollow dynamics in twitter. In *SIGCHI Conference on Human Factors in Computing Systems (CHI)* (2011).
- [23] LEE, K., CAVERLEE, J., AND WEBB, S. Uncovering social spammers: social honeypots + machine learning. In *International ACM SIGIR Conference on Research and Development in Information Retrieval* (2010).
- [24] LEE, S., AND KIM, J. WarningBird: Detecting Suspicious URLs in Twitter Stream. In *Symposium on Network and Distributed System Security (NDSS)* (2012).
- [25] MOORE, T., AND EDELMAN, B. Measuring the perpetrators and funders of typosquatting. In *Financial Cryptography and Data Security* (2010), vol. 6052, pp. 175–191.
- [26] NIKIFORAKIS, N., ACKER, S. V., MEERT, W., DESMET, L., PIESSENS, F., AND JOOSEN, W. Bitsquatting: Exploiting bit-flips for fun, or profit? In *Proceedings of the 22nd International World Wide Web Conference (WWW)* (2013).
- [27] NIKIFORAKIS, N., BALDUZZI, M., DESMET, L., PIESSENS, F., AND JOOSEN, W. Soundsquatting: Uncovering the use of homophones in domain squatting. In *Proceedings of the 17th Information Security Conference (ISC)* (2014).
- [28] RIVERS, C. M., AND LEWIS, B. L. Ethical research standards in a world of big data. *F1000Research* (2014).
- [29] STRINGHINI, G., EGELE, M., KRUEGEL, C., AND VIGNA, G. Poultry Markets: On the Underground Economy of Twitter Followers. In *SIGCOMM Workshop on Online Social Networks* (2012).
- [30] STRINGHINI, G., KRUEGEL, C., AND VIGNA, G. Detecting Spammers on Social Networks. In *Annual Computer Security Applications Conference (ACSAC)* (2010).
- [31] STRINGHINI, G., MOURLANNE, P., JACOB, G., EGELE, M., KRUEGEL, C., AND VIGNA, G. EvilCohort: Detecting Communities of Malicious Accounts on Online Services. In *USENIX Security Symposium* (2015).
- [32] STRINGHINI, G., WANG, G., EGELE, M., KRUEGEL, C., VIGNA, G., ZHENG, H., AND ZHAO, B. Y. Follow the Green: Growth and Dynamics in Twitter Follower Markets. In *ACM SIGCOMM Conference on Internet Measurement* (2013).
- [33] SZURDI, J., KOCSO, B., CSEH, G., SPRING, J., FELEGYHAZI, M., AND KANICH, C. The long “taile” of typosquatting domain names. In *23rd USENIX Security Symposium (USENIX Security 14)* (San Diego, CA, 2014), USENIX Association, pp. 191–206.
- [34] THOMAS, K., GRIER, C., MA, J., PAXSON, V., AND SONG, D. Design and Evaluation of a Real-Time URL Spam Filtering Service. In *IEEE Symposium on Security and Privacy* (2011).
- [35] THOMAS, K., GRIER, C., SONG, D., AND PAXSON, V. Suspended accounts in retrospect: an analysis of twitter spam. In *Internet Measurement Conference (IMC)* (2011).
- [36] THOMAS, K., MCCOY, D., GRIER, C., KOLCZ, A., AND PAXSON, V. Trafficking fraudulent accounts: The role of the underground market in twitter spam and abuse. In *USENIX Security Symposium* (2013).
- [37] WANG, Y.-M., BECK, D., WANG, J., VERBOWSKI, C., AND DANIELS, B. Strider typo-patrol: discovery and analysis of systematic typo-squatting. In *Proceedings of the 2nd conference on Steps to Reducing Unwanted Traffic on the Internet - Volume 2* (Berkeley, CA, USA, 2006), SRUTI’06, USENIX Association, pp. 5–5.
- [38] WILSON, C., BOE, B., SALA, A., PUTTASWAMY, K., AND ZHAO, B. User Interactions in Social Networks and Their Implications. In *Annual Computer Security Applications Conference (ACSAC)* (2010).
- [39] YANG, C., HARKREADER, R., AND GU, G. Die Free or Live Hard? Empirical Evaluation and New Design for Fighting Evolving Twitter Spammers. In *Symposium on Recent Advances in Intrusion Detection (RAID)* (2011).