

Exploiting Social Navigation

Meital Ben Sinai
Technion, Israel

Nimrod Partush
Technion, Israel

Shir Yadid
Technion, Israel

Eran Yahav
Technion, Israel

Abstract

We present two attacks against social location based services: (a) An effective Sybil attack [14] aimed at creating fake congestion and influencing benign user routing (b) A privacy targeted attack that allows tracking fellow users movement over a desired area and time frame. Both our attacks are based on creating a large number of “bot drivers”, and controlling their reported locations using fake GPS reports. We applied our attacks to WAZE [10] - a prominent social navigation application used by over 50 million drivers. We show that our first attack can fake traffic jams and dramatically influence routing decisions, and that our second attack effectively tracks users. We present several techniques for preventing the attacks, and show that effective mitigation likely requires the use of additional carrier information for the first attack, and removing fellow user location reporting altogether for the second attack.

1 Introduction

Social navigation is rapidly becoming a prevalent approach for navigation. A social navigation system collects navigation data such as route duration, traffic congestion, road obstacles, and even map layout from users, and then uses this information to route other users.

WAZE is a prominent social navigation system used by over 50 million users worldwide (recently acquired by Google). At this scale, social navigation data is considered a reliable source of traffic information trusted by many. For example, WAZE traffic information is used to feed reports for various Radio and TV stations [11], as well as the widely used Google Maps navigation application [3]. This propagates the influence of social navigation to a much wider community.

In this paper, we present the first Sybil attack on social navigation systems, and demonstrate its effectiveness by applying it to WAZE. In a Sybil attack the attacker subverts the reputation system of a network by creating a large number of pseudonymous identities, using them to

gain a disproportionately large influence [14]. Our main idea is to create a large number of reputed “bot drivers” and use fake GPS information to drive them in desired traffic patterns. For example, the “bots” can be used to simulate patterns detected by WAZE as “heavy traffic”.

In contrast to previous attacks on WAZE [1], our approach does not require reverse engineering the client, or the protocol, and cannot be mitigated by hardening the protocol as it uses legitimate clients.

We demonstrate the effectiveness of our attack by applying it to WAZE, and showing that our technique can simulate traffic jams, and dramatically influence the routing decisions presented to users.

The attack has vast security and financial implications. For example, the attack could be targeted towards specific businesses or toll roads, falsely reporting the area as congested and discouraging drivers from nearing it. The attack also affects drivers safety, where drivers may be sent along a less suited road, or distracted by spurious obstacle reports.

Additionally, we discuss privacy issues due to disclosure of user information on the WAZE “live map” and present a second attack aimed at compromising user privacy. This attack is carried out by spreading a large number of bots over a desired region. The bots operate by querying their surrounding live map area, finding fellow WAZE users and capturing their data.

Again, this attack required no reverse engineering of the WAZE application or network data and therefore resistant to any hardening techniques one may apply such as code obfuscation or protocol encryption. Instead, we used OCR techniques to search the image data provided by the live map to our bots to extract user information.

Main Contributions The main contributions of this paper are:

- We present a Sybil attack on social navigation systems. The attack is cheap and can be facilitated using free off-the-shelf emulation software, a simple

fake GPS player application, controlled by the Android Debug Bridge, running on a 16 core machine.

- We present a privacy attack, aimed at tracking fellow users. The attack is proportional to the target area and requires a single device to track an area of 140000 square meters.
- We demonstrate our approach by simulating traffic jams, affecting the routing and tracking users on the WAZE social navigation system.
- We present two approaches for mitigating our first attack: using carrier information and user behavioral analysis. We evaluate them based on parameters of simplicity, user experience, effectiveness and cost.
- We discuss mitigation of our second attack and show why use anonymization is not sufficient, and only removal of all user data can effectively mitigate the attack.

2 Attacks

In this section we describe all successfully implemented attacks. We first describe how to create bot drivers and increase their reputation, and then describe how bot drivers are used for various attacks.

2.1 Creating Bot Drivers

Becoming an influential part of the WAZE community requires a single click. Just by installing and starting the application on her device, a user is able to view surrounding user and traffic information, influence nearby users by reporting various road obstacles, and affect routing and traffic information. This allowed us for automated creation of bot drivers simply by starting an emulator, running WAZE, and running a script that clicks the “start driving” button.

Registration does not require validation WAZE offers a registration process, encouraged by a rating system where users receive points for using the application and reporting obstacles. The registration process however, employs no validation mechanisms. We assume WAZE intentionally avoids user validation, in hopes of encouraging new users to try the application. Interestingly, deleting a WAZE account does require the user to pass a CAPTCHA. This further establishes our assumption that the WAZE registration process was designed to gather and maintain users, and not to validate them.

Training Reputed Bots Due to lack of validation, we were able to automate the process of *creating a registered WAZE user*. We used random bot usernames to avoid clashes with pre-existing names. Registration does require a well-formed email address, but it is not further verified. Throughout our experiments, we re-used the same user accounts and were able to gain rating, by simulating driving and reporting obstacles. The ability to

automatically register and gain rating, makes it harder to mitigate attacks by relying on user rating and tenure. In fact, by emulating benign user behavior, attempts to eliminate bots on behavioral basis are likely to fail.

2.2 Creating Traffic Jams and Influencing Routing

Next, we describe how we were able to simulate traffic jams and influence the WAZE routing algorithm. WAZE deduces traffic congestion and routing time information from location and movement data reported by its users [4]. This algorithm resides on the server side of the WAZE system and was never publicly disclosed. The main challenge of this work, was in fact experimentally deducing and exploiting this algorithm.

Our experiments consisted of explorative adjustment of the following parameters: (i) Data set size (number of bots), (ii) Drive duration, (iii) Speed and movement pattern. All parameters were aimed at matching the WAZE traffic jam scenario. We show that with a small amount of resources, we were able to simulate a traffic jam.

Gaining reputation The WAZE application states: “As you advance to higher levels your reports get greater influence”. We therefore trained our bot drivers by driving on campus in average speeds for several hours.

Simulating standstill traffic Our initial attempts at simulating congestion consisted of spawning botnets (groups of “bot drivers”) of increasing sizes (5, 10, 15 etc.) and statically placing them at the target location. We assumed WAZE would acknowledge these static drivers as stuck in traffic, and report a traffic jam. This approach did not work, even after scattering the bots along the target road in different patterns, attempting to mimic a more realistic standstill traffic scenario. We attribute this result to WAZE interpreting the scenario as users entering their vehicle and starting WAZE, but have yet to move - a common scenario among navigation users.

Simulating slowdowns Our next round of experiments consisted of simulating a gradual slowdown in traffic. As before, we sent increasingly larger groups of bots to the target location, but this time they moved through the area in gradually slower speeds of 15, 8 and 2 kph. At first, each bot passed through the area once for each speed in the set, then twice, etc., up to the point where each bot spends 20 minutes passing the route (again and again) at a certain speed before moving on to the next slower speed. Lastly, we incorporated short periods of standing still (10 seconds) into each pass, also in increasing durations. This scenario still failed to yield congestion, as it required one final addition, described next.

Speed is relative. We conducted our experiments on the closed campus area (see Sec. 3.4). Initially, we saw that even when consistently simulating gradually slowing traffic, with large botnets of reputed drivers, no con-




Figure 1: Speed pattern of bots participating in the attack

gestion was reported. One way this result could be explained is that WAZE takes other, non user-related parameters into effect. These parameters could include: time of day, date, weather conditions and *road information*. We in fact learned, that since the campus area is a very low speed zone, WAZE regarded all slow driving traffic through it as normal rather than congested. In fact, it seems that the WAZE congestion reporting algorithm is a relative one, i.e., a route is congested if its current average speed is considerably lower *relative to former known speeds*. Thus, we added another phase to the bots driving pattern, where driving was initiated at 70kph. After a period of 20 minutes, we reverted to the previous pattern of simulating slowdown in traffic, finally leading to successful simulation of a traffic jam. Fig. 1 shows the speed over time of each bot in the successful implementation of the attack. The successful attack consisted of only 15 bots, driving at the pattern shown in Fig. 1, where we reset a bot location to the route origin whenever it has reached the end of the route (and thus enable it to drive the route again according to the desired speed pattern).

Simulated traffic jams on the WAZE map. Fig. 2 shows a traffic jam created by our attack. The congestion is marked by a bold red line, annotated with arrows showing the congested direction, along with a notification marking the average speed as 8kph. The congested route also features the botnet used for the attack, showing as group of WAZE users along the start of congested route (we further discuss users shown on the live map in Sec. 2.4). We observed that traffic jams created by the attack persists for roughly 20 minutes past the point where the bots stop driving.

Influencing routing We were able to influence the WAZE routing algorithm and potentially send benign users on different routes. This has vast security and financial implications. The attack can be used to ease actual congestion in roads of interest to the attacker, by falsely reporting them as congested, and sending away benign users. Alternately, an attacker may congest areas causing po-

¹WAZE map images included Hebrew and thus were edited for readability.




Figure 2: Reported traffic jam, created by the attack¹

tential financial damage, due to time and fuel wasted by drivers taking longer routes, or drivers avoiding using tool roads or visiting businesses residing in congested areas. This also a safety issue, where drivers may be sent on a less suited road. Furthermore, using the attack, one is able to completely determine the route planned by WAZE for a user trying to get to a destination. Consequently, users will be directed along a desired path controlled by a malicious attacker. Fig. 3a shows the recommended route, prior to an attack. The route is marked by a bold purple line, annotated with street names, and reported to be 1.5km long with 3 minutes traversal time. Fig. 3b shows the change in route due to the fake traffic jam we created along the previous route. The new route is 1.6km long and takes 4 minutes to traverse.

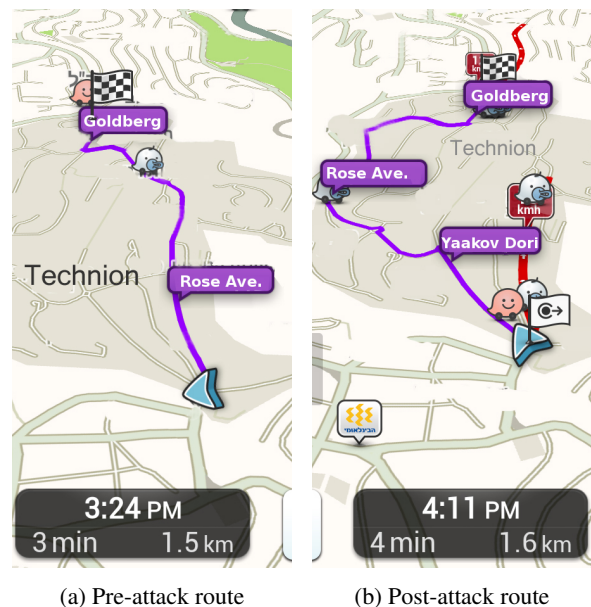


Figure 3: WAZE routing, influenced by the attack

Thwarting machine-like behaviour Throughout our experiments, we explored how WAZE reacts to machine-like behavior, such as spawning all bots at the exact same

time or driving at the same speed, along the same route, all at once. We spent a significant amount of time planning our experiments to simulate benign human behavior. We added a mechanism for gradual creation of the driver botnets, with random time spacing between each spawn. We also implemented our command and control interface such that each bot could receive a different movement pattern, along with randomization. After successfully implementing the attack, we proceeded to verify the necessity of these mechanisms. We retried the traffic jam scenario, while disabling each of these mechanisms, and learned that a congestion was reported regardless. This means that WAZE makes no observable effort to differentiate human from machine behavior. Ultimately, this may have been a wise design choice, as applying such mechanisms is destined to fail against a highly skilled attacker. However, in Section 4 we describe two techniques that can mitigate such attacks.

2.3 Reporting Obstacles




Figure 4: WAZE obstacle reporting menu

Reporting persistent spurious obstacles The WAZE application allows notifying nearby users of the presence of obstacles on the road. Fig. 4 depicts the WAZE obstacle reporting interface. Reporting obstacles can be done by any user, whether registered or not. The reporting process consists of a few clicks, and users are also able to add a note or take a picture of the reported obstacles. Once the report is complete, the obstacle immediately appears on the map of all nearby users, and stays there for roughly 20 minutes. The duration can be extended, however, if reported as credible by other users. Thus, fake sustainable reports can be easily produced by a small group of bot drivers, one reporting an obstacle with a designated note and the rest confirming the report. We note that reporting a “Traffic Jam” obstacle does not effectively create a traffic jam as in Fig. 2, or affect routing. Instead, a notification is shown on the map, left to the discretion of the user.

Indirectly influencing traffic Although obstacle reports do not influence routing (WAZE will not prefer a non-obstructed road over an obstructed one), they have a strong mental effect on human users. A user observing multiple obstacles along a certain route, some containing notes supposedly written by other users, may prefer tak-

ing a different route, or drive much slower through the obstructed route. This could have an indirect effect on traffic congestion and can be performed with a minimal amount of resources.

Report system DDos We also considered attacks where the report mechanism is rendered useless. This can be performed by blasting the application with a huge amount of spurious reports, in the form of a DDos attack. Fig. 5 shows the effect of reporting multiple police units and hazards over a small region. WAZE limits the number of reports one user can produce in a short time frame, but this limitation can be easily overcome by closing the application and logging as a different user (reports remain after the reporting user closed the application).




Figure 5: Fake reports of police units on the WAZE map

Invalidating benign reports Another similar attack can be performed by invalidating benign user reports. Whenever a WAZE user approaches an obstacle, they are given the option to update the report as outdated. An attack where the WAZE map is constantly combed for reports and those are, in turn, invalidated by a bot driver, will incapacitate the report system and in fact give the attacker full control of said system.

We note that the WAZE report system also allows directly influencing map layout, by reporting road closures and map issues. However, we refrained from exploring such scenarios as the risk of damaging benign WAZE users was too great. Should this attack be made possible, financial and safety implications are dire, where users may completely avoid a closed road or be sent along a non-existing path.

2.4 Tracking Users

Statistically tracking users via the live map Fig. 6 shows the WAZE “live map”, which functions as the main view of the application. The live map features details of fellow WAZE users, over all locations, randomly appearing in and out of the map. Further, one can receive detailed information on a fellow user by clicking on their icon, as seen in Fig. 6 (blurred for privacy reasons). It is also possible to chat with or message the selected user, which

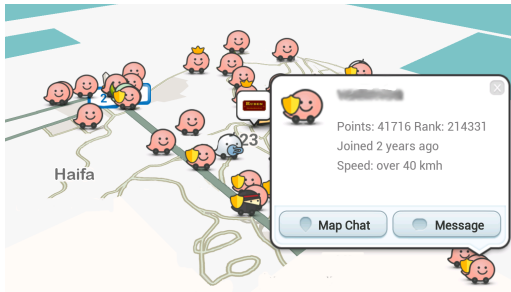


Figure 6: WAZE live map

adds social network-like features to WAZE. Revealed user information includes username, points and ranking, seniority and speed. We leverage this information towards compromising user privacy.

Capturing user data In order to capture private user data we started by defining the following bot scheme: When a bot is deployed to a location, it starts by capturing the image of its surroundings and applying image processing using the `openCV` [5] library to search for the fellow user marker in the entire image. Once a user’s coordinates were found, the bot proceeded to click on it in order to display the user information dialogue specifying the aforementioned user data, as can be seen in Fig. 6. This operation needed to be fast enough to overcome the live map refresh rate, which changes the reported surrounding every few seconds. Next, the bot captures the image of the displayed dialogue and proceeds to apply OCR using the `Tesseract` [6] library in order to extract username, ranking, seniority and speed. We experimented with a training phase for the `Tesseract` library, in order to specialize it for the `Waze` font. Once these were extracted, the data was sent to a centralized database, along with the map location and time of day, for further correlation, and the bot was ready to be redeployed in a different location in the desired range. The screen capturing, clicking operations and mock GPS locations were all provided through the `ADB` interface. A centralized command script handled the deployment of the bots to cover the targeted area. Each bot iteration elapsed around 10 seconds to capture a user in an area of 140000 square meters, resulting in a data retrieval rate of 400 records per hour, per device. We note that we could not carry out this attack on the Android emulator, as it wasn’t responsive enough and presented poor images which could not be processed. Thus the attack cost is proportional to the area being monitored, in number of devices needed.

Correlating user data and extrapolating routes So far we have collected roughly 200000 records from about 30000 users over areas of roughly 5000 square kilometers and a period of two months. We proceeded in building a framework for extracting and presenting user

routes, searchable according to username. This framework will be publicly available, after anonymization of user data, at www.trackmywaze.com. Our current algorithm operates by searching the database for records of users that appear within a time frame of 5 hours in an area of 750 kilometers (maximum distance allowed between two points). According to this criteria, we detected about 20000 routes overall. The points are then ordered according to time of appearance and a route is built and presented on the map, as shown in Fig. 7.

Tracking specific users Currently, tracking a specific user requires the knowledge of their handle. Therefore, one could claim that changing one’s handle frequently enough is sufficient to avoid tracking. That is not the case in the WAZE system. Applying such name change strategies will not be sufficient as the live map supplied further data, including seniority and rank. An attacker can use these as alternate identifiers for singling out a user. Even if these are approximated, the attacker can cross-reference this data with expected location for the target user (say their home or work address) and search the area until the user appears.

Deducing user identity We claim that the collected data can be used to find a correlation between a WAZE handle and real person identities. If an attacker wants to deduce the handle for a specific target, and assuming they have knowledge of some regular route taken by the target (say their home and work address), they can place bots surrounding that area and search for re-occurring users. After a sufficient period of time, the correct handle will appear along the expected path enough times to establish it as the target user.

3 Experiment System Description

In this section we will shortly describe the experiment parameters and tools developed to carry out attacks.

3.1 Emulating WAZE Clients

For the experiment purposes, we required a non-costly, easily deployed system that will allow us to create and control a large amount of WAZE clients. This suggested a software emulation solution and after some research we arrived at using `Android Developer Tools` software emulator. This solution complied with our requirements as it allowed us to run multiple instances over our `Linux` servers. We used the `Android Developer Tools` version 22.3.0 with `Android API 17` and ran the `WAZE` Android client version 3.5.1.4.

3.2 Command and Control Interface

To allow for rapid development, we developed our system in `Python`. We controlled the android clients through the `Android Debug Bridge (adb)`, a versatile command line tool that allows communicating with an emulator in-

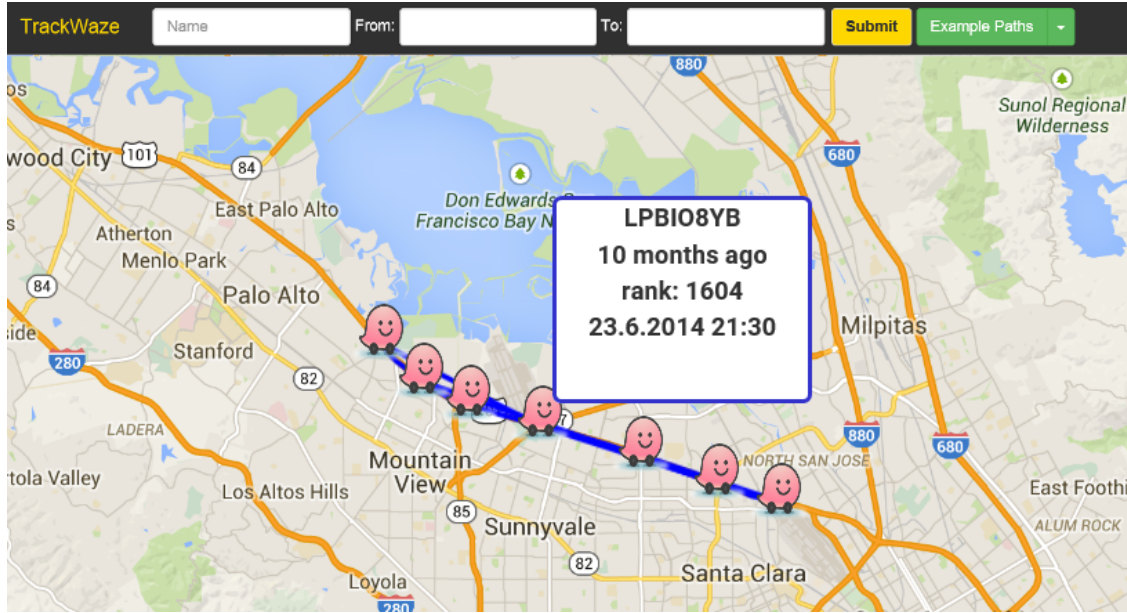


Figure 7: WAZE user tracking framework, with displayed (anonymized) user route

stance. After starting each emulator instance, WAZE was installed on it through a series of ADB commands, simulating text and touch input. In order to successfully automate controlling the WAZE client on the emulator instance, we used the WAZE application activity log, as it was updated with every completed operation (we assume for debugging). We suggest removing said logging from future versions of WAZE.

3.3 GPS Data

In order to emulate driving along a route, we constructed a small Android application which generates mock GPS locations adhering to driving on the target route, according to a given movement pattern. Our application, appropriately named *TrafficJam*, receives a start point, end point and direction and then proceeds to generate mock GPS locations along that route. WAZE considers these mock GPS locations to be valid and moves the client accordingly on the map. The application also receives a movement pattern, composed of an average speed, duration of movement and standstill intervals, and generates the mock GPS locations accordingly, emulating moving along the desired route, at the given speed, for the given time duration, while standing still at certain intervals.

Fig. 8 contains pseudo-code of the congestion creation script, implementing the attack described in Section 2.2.

3.4 Experiment Location

As experiments for our first attack affect real WAZE users, we avoided using major roads and highways. We ran our experiments on Technion campus in the city of Haifa, Israel. We chose this region as it is small enough and

```

CreateTrafficJam(num_bots, route, reputed_users) {
  emulators = StartEmulators(num_bots);
  InstallAndStartApp(emulators, "Waze");
  InstallAndStartApp(emulators, "TrafficJam");
  LoginWaze(emulators, reputed_users);
  stops="none";
  TrafficJam(emulators, route, 70kph, 20min, stops);
  stops="stop for 10s every 10s";
  TrafficJam(emulators, route, 8kph, 20min, stops);
  TrafficJam(emulators, route, 2kph, 20min, stops);
}

```

Figure 8: Congestion creation script pseudo-code

relatively closed off to have a minor affect on the general user population. However, the campus road system is large enough to allow several navigation options. More importantly, since the campus is frequented by a large and technology-friendly crowd, we had ample amount of legitimate WAZE users to establish our results as credible.

As the second attack did not affect or harm users in any way, we deployed it to several central locations including: New York City, Paris, London, California, Singapore and Northern and Central Israel.

4 Mitigating Attacks

In this section, we will discuss several means for mitigation attacks, while comparing them by parameters of simplicity, user experience, security and cost.

4.1 Verifying Drivers via Carrier Data

Relying on carrier data in social networks and mobile applications An established trend for verifying users of mobile applications and social networks is the use of mobile carrier data. Facebook [7], Google, [2] and many prominent mobile applications such as WhatsApp [8] and Viber [9], encourage and even require their users to verify their account using a legitimate mobile phone number. While users of traditional social networks like Facebook and Google+ may not always have a mobile number, for WAZE this is not the case, as it was designed to be run on a mobile device.

Linking users to a mobile carrier identity Retrieving the user's phone number upon registration and validating it would make the attacks of Section 2.1 considerably more complex and costly. WAZE could still allow for non-verified use of the application, but should give much smaller weight to these type of users. Note that this incurs a one time carrier cost per user.

Verifying user location via carrier Mitigating attacks in Section 2.2 and Section 2.3 requires validating the actual location of the reporting entity. This can be done by verifying user reported GPS locations with the location of the cellular antenna carrying it. To lower costs and net traffic, this validation can occur periodically, at random times selected by the application. This solution can also be used to mitigate user tracking (Section 2.4), by disclosing fellow user information only to verified drivers. Note that all attacks described in Section 2 would still be possible, but would require purchasing a large amount of mobile carrier equipment (SIM cards and equipment for operating them), that would then need to be located in the vicinity of the target area (to be considered as verified).

Ultimately, we consider using carrier data to be the most secure method (as it relies on external verified data), fairly simple to implement, and user friendly. The disadvantage is the possible incurred cost by the carrier.

4.2 Verifying Drivers via Behavioral Analysis

Distinguishing human and robot behavior in Sybil attacks has been the subject of much study [17, 16, 19, 18]. However, in the context of verifying driver behavior, the subject has remained fairly unstudied. We briefly discuss several means for mitigating attacks based on user behavior analysis.

Relying on existing validation mechanisms Like many other mobile application, WAZE offers the option to login via Facebook account. This can be further extended towards using Google and Apple accounts, or simply asking users to solve a CAPTCHA. Granting better standing to validated users could help mitigate attacks in Section 2, but is dependant on the validity of the underlying mechanisms [12, 15].

Network traffic analysis Our attack consisted of numerous client reports originating from a small number of static IP addresses. A simple network analysis could have thwarted these reports and treated them as unreliable. This approach could be further extended so that reports originating from IP addresses affiliated with 3G carriers [13] receive better standing. Although this method is somewhat related to carrier data, we did not include it in Sec. 4.1, as it could be performed relatively unbeknownst to the carrier. A skilled attacker could subvert this defence mechanism by using multiple 3G network access, incurring extra attack costs.

Analyzing user creation, movement and report patterns In Sec. 2.2 we discussed possible mechanisms of disguising bot behavior, such as gradual creation of botnets and different movement patterns for each bot. These could be mitigated by performing a more in-depth analysis of driver behavior and comparison with benign user behavior. Unfortunately, these break against a skilled opponent, successfully mimicking human drivers. The same technique could be applied to obstacle reports, preventing spurious reports and invalidation of benign reports.

The main advantage of this approach is that it could be performed entirely in-house, without involving third party components which incur extra costs and net traffic. User experience is affected by registration validation. However, as mentioned, analysing user behavior is fairly complicated and can require constant maintenance and upgrade, in order to stay ahead of attackers.

4.3 Avoiding User Tracking

Although it is not the default option, WAZE allows users to opt out of appearing on the live map. However, it fails to notify them that unless they do so, their location can be tracked and eventually correlated to their identity. WAZE can offer effective mitigation for the attack by removing the option to see fellow user data, and maintain a part of the experience by only showing fellow user location, stripped of any identifying properties. We note that in this case, areas with small amounts of WAZE users are still exposed to a tracking attack, as they will inform the attacker that *some user* is in the target area, and in sparse areas this will point to the target user.

5 Conclusion

We presented two new attacks for social navigation. The first is a Sybil attack on a location based service is possible, and could affect large communities of users, causing financial damage and compromising security and privacy. The second tracks user location and can be used to correlate user handle to actual person identities. Our attacks require a low amount of resources, and can be performed by many attackers. We offer ways for mitigating attacks, balancing simplicity, effectiveness and cost.

Acknowledgement

This work was (partially) supported by the Israel Ministry of Science and Technology grant no. 3-9779. We would like to thank Omer Katz and Yaniv David for their feedback and suggestions for this work.

References

- [1] Floating car data from smartphones: What google and waze know about you and how hackers can control traffic. <https://media.blackhat.com/eu-13/briefings/Jeske/bh-eu-13-floating-car-data-jeske-wp.pdf>.
- [2] Google 2-step verification. <http://www.google.com/landing/2step/>.
- [3] Google maps and waze, outsmarting traffic together. <http://googleblog.blogspot.co.il/2013/06/google-maps-and-waze-outsmarting.html>.
- [4] How waze calculates routes. https://www.waze.com/wiki/How_Waze_calculates_routes.
- [5] Open source computer vision. <https://www.opencv.org>.
- [6] tesseract-ocr. <https://code.google.com/p/tesseract-ocr/>.
- [7] Verify facebook account. <https://www.facebook.com/help/398085743567023/>.
- [8] Verifying your phone number in whatsapp. <http://www.whatsapp.com/faq/android/20970873>.
- [9] Viber activation. http://activate.viber.com/viber/activate_user.html.
- [10] Waze official website. <http://www.waze.com>.
- [11] Waze on tv. <http://blog.waze.com/search/label/OnTV>.
- [12] Facebook has more than 83 million illegitimate accounts, August 2012. <http://www.bbc.co.uk/news/technology-19093078>.
- [13] BALAKRISHNAN, M., MOHOMED, I., AND RAMASUBRAMANIAN, V. Where's that phone?: geolocating ip addresses on 3g networks. In *Internet Measurement Conference* (2009).
- [14] DOUCEUR, J. R. The sybil attack. In *IPTPS* (2002), pp. 251–260.
- [15] MOTOYAMA, M., LEVCHENKO, K., KANICH, C., MCCOY, D., VOELKER, G. M., AND SAVAGE, S. Re: Captchas-understanding captcha-solving services in an economic context. In *USENIX Security Symposium* (2010), pp. 435–462.
- [16] TRAN, D. N., MIN, B., LI, J., AND SUBRAMANIAN, L. Sybil-resilient online content voting. In *NSDI* (2009), pp. 15–28.
- [17] WANG, G., KONOLIGE, T., WILSON, C., WANG, X., ZHENG, H., AND ZHAO, B. Y. You are how you click: Clickstream analysis for sybil detection. In *Proceedings of the 22Nd USENIX Conference on Security* (2013), SEC'13, USENIX, pp. 241–256.
- [18] YU, H., GIBBONS, P. B., KAMINSKY, M., AND XIAO, F. Sybil-limit: A near-optimal social network defense against sybil attacks. *IEEE/ACM Trans. Netw.* 18, 3 (2010), 885–898.
- [19] YU, H., KAMINSKY, M., GIBBONS, P. B., AND FLAXMAN, A. D. Sybilguard: defending against sybil attacks via social networks. *IEEE/ACM Trans. Netw.* 16, 3 (2008), 576–589.