Open WiFi Networks: Lethal Weapons for Botnets?†

Matthew Knysz Xin Hu Yuanyuan Zeng Kang G. Shin
University of Michigan, Ann Arbor

***Abstract*—This paper assesses the potential for highly mobile botnets to communicate and perform nefarious actions using only open WiFi networks, which we term *mobile WiFi botnets*. We design and evaluate a proof-of-concept mobile WiFi botnet using real-world mobility traces and actual open WiFi network locations for the urban environment of San Francisco. Our extensive simulation results demonstrate that mobile WiFi botnets can support rapid command propagation, with commands typically reaching over 75% of the botnet only 2 hours after injection—sometimes, within as little as 30 minutes. Moreover, those bots able to receive commands usually have ≈40-50% probability of being able to do so within a minute of the command being issued. Our evaluation results also indicate that even a small mobile WiFi botnet of only 536 bots can launch an effective DDoS attack against poorly protected systems. Furthermore, mobile WiFi botnet traffic is sufficiently distributed across multiple open WiFi networks—with no single network being over-utilized at any given moment—to make detection difficult.**

I. Introduction

Botnets—large networks of compromised computers under the control of a single botmaster—are one of the gravest threats facing Internet security professionals today. Their vast size and use of a Command and Control (C&C) channel bestows botmasters an unprecedented amount of power and versatility to launch numerous sophisticated scams and attacks. Currently, much research has been done to understand and detect existing botnets, as well as predict their future capabilities and C&C channels. However, at this time, little research has been done examining mobile botnets, i.e., botnets composed entirely of compromised mobile devices. With their booming application markets, standardized OSes and rapid advances in processing power and memory, mobile devices are capable of increasingly sophisticated tasks. Combined with their multiple communication interfaces (i.e., WiFi, 3G/4G, Bluetooth, SMS and MMS messaging) and always-on connectivity, they are capable of sophisticated attacks not possible with traditional computers. As a result, mobile devices are quickly becoming an attractive target for botmasters, and it is only a matter of time before mobile botnets emerge on the Internet threatscape.

In this paper, we evaluate the potential for mobile botnets to communicate and perform nefarious actions solely over open WiFi Access Points (APs), which we term *mobile WiFi botnets*; to the best of our knowledge, we are the first to research this scenario. The use of open WiFi networks for mobile botnets can provide a higher level of stealthiness and has fewer barriers to entry than other communication mediums, which we discuss further in Section II-A. In assessing the feasibility of mobile WiFi botnets to support botnet C&C and DDoS attacks, we make the following contributions. First, we design a proof-of-concept mobile WiFi botnet, including its AP-selection algorithm and C&C and DDoS attack protocols. Second, we build a mobile WiFi botnet simulator, using accurate timing models for AP association, Internet communication and achievable wireless throughput based on the mobile bot's distance from the open AP. Third, we run this simulator for various attack scenarios using real-world cab mobility traces and actual open WiFi AP locations for the urban environment of San Francisco. Fourth, through in-depth simulations, we demonstrate that mobile WiFi botnets can support rapid command propagation, can mount DDoS attacks against poorly protected systems with as few as 536 reachable bots, and are sufficiently distributed across open WiFi networks—with no single network being over-utilized at any given moment—to make detection difficult. Moreover, bots able to receive commands usually have ≈40-50% probability of doing so within a minute of the command being issued.

† The work reported in this paper was supported in part by the ONR under Grant No. N000140911042

II. WiFi-Based Mobile Botnets

A. Why WiFi

WiFi has certain advantages over other communication mediums available to mobile botnets. Compared to cellular channels, such as 3G/4G and SMS/MMS, botnet activities over WiFi are more discrete and difficult to detect. Mobile devices must use a non-spoofable mobile ID to connect to cellular networks; since cellular providers posses an omniscient view of their devices' network activity, this makes it easier to detect and shutdown bot-infected devices. This is in contrast to a WiFi botnet, where IPs can be spoofed and malicious activity hidden behind many different open networks and NAT routers, where it is difficult to mitigate. While Bluetooth can provide the necessary stealth required for C&C, it cannot function as a medium for other botnet activities, such as spam or DDoS attacks. Furthermore, due to its limited range and transmission rates, a large number of identically infected, slowly moving, devices must continuously come in close proximity to one another to be effective. This is unsuitable for newly emerging botnets, where infected devices may be few and geographically dispersed. In contrast, WiFi botnets can achieve faster transfer speeds, support multiple botnet activities, and be incrementally deployed with less effort; a small WiFi botnet, even when spread across multiple cities, can immediately contribute to an existing botnet, unlike Bluetooth.

In this paper, we will examine mobile bots traveling in vehicles, as their rapid pace ensures bots will be within range of open WiFi AP for typically just a few seconds; if this rapid

pace is sufficient to perform botnet activities, then so will the more leisurely pace of bikers and pedestrians. With such a limited duration at each AP, it isn't practical for WiFi bots to attempt hacking into encrypted or closed networks. Rather, by only utilizing open and unencrypted WiFi networks, bots can quickly perform a small subset of their activities on multiple networks when they are in range. Despite the tightening of home WiFi network security, necessity dictates that many businesses (e.g., restaurants, cafes) must leave their WiFi networks open and unencrypted; inconveniencing customers with a network password could potentially drive them to competitors with open networks. Therefore, while decreasing, open WiFi networks will not entirely disappear, making them a useful medium for mobile botnets.

B. Mobile Botnet

Since most businesses with open WiFi networks limit online activity to web traffic over port 80, a mobile WiFi botnet should use an HTTP-based C&C channel; this provides the added benefit of hiding C&C traffic among other, benign HTTP traffic on the network. The command server is typically polled periodically for new commands using a deterministic yet changing domain name, making it difficult for defenders to predict and block the malicious domains using such techniques as a DNS sink hole. A mobile botnet can improve stealthiness further by limiting its C&C polling to open WiFi networks, purposely avoiding such activity when on home or office networks, where it might be more easily noticed. Likewise, DDoS attacks can be issued from multiple open networks as they are in range, issuing small portions of the total attack at each AP. In Sections III-C and III-D, we will discuss in detail how our prototype mobile WiFi botnet actually achieves HTTP-based C&C and performs DDoS attacks.

C. Threat Model

Our focus in this paper is to determine the effectiveness of open WiFi networks for botnet C&C and DDoS attacks using high-mobility devices (i.e., those traveling in vehicles). We assume all devices are infected with the same bot malware, allowing us to ignore the complications of infecting heterogeneous devices. It seems obvious that mobile devices will have sufficient online access when connected to a user's home or office WiFi network. However, at this point, little/no research has been done to determine if botnet activities can be supported using *only* open WiFi networks, which can significantly encumber detection. In our model, we assume the existence of a small mobile botnet operating solely over open WiFi networks in a single metropolitan city. The botmaster's adversarial goals are three-fold. The first goal is to obtain a sufficient amount of control over a modest portion of his total botnet during the bots' most transient and high-speed period, which is the most difficult to control. This period occurs when bots are traveling in a vehicle, and we will determine if this goal can be achieved by designing and simulating a C&C protocol for a mobile WiFi botnet. During weekday office-commute hours and weekends, we will explore what amount of the total potential botnet is actually reachable, how frequently bots can receive commands and how quickly new commands propagate throughout the botnet. The second adversarial goal is to use the small mobile WiFi botnet to issue a successful DDoS attack. Likewise, we will design and simulate such an attack for a mobile WiFi botnet to evaluate its feasibility. The final goal is that both botnet C&C and the DDoS attack can function in a stealthy manner, which we will evaluate by ensuring that the malicious botnet's traffic is adequately dispersed over multiple WiFi APs.

III. Experimental Setup

A. Description of Datasets

We now describe the datasets used in our experiments to determine open WiFi APs and simulate vehicular mobility patterns in an urban environment.

1) Open WiFi APs: To obtain a comprehensive dataset of open WiFi APs, we use the Wireless Geographic Logging Engine (WiGLE) dataset for downtown San Francisco [3]. WiGLE contains an extensive database of wireless AP information, built from a collaborative effort of thousands of researchers and WiFi enthusiasts. For downtown San Francisco, 2,349 WiFi APs were unanimously identified as both "open" and unencrypted, ideal for the proposed mobile WiFi botnet.

2) Mobility Traces: To simulate vehicular mobility patterns in an urban environment, we use the mobility traces of taxi cabs in the San Francisco Bay area [1], first published in [10]. Unfortunately, the cabs' location granularity ranges from seconds to minutes, making it too inconsistent and course-grained for our purposes. To overcome this limitation, we converted the traces into the TIGER [2] map coordinate system. Combining this with the VanetMobiSim [7] vehicular simulation system, we simulated the path and speed between any two cab trace points based on San Francisco's actual road topology. The resulting dataset contains detailed location information at a one-second granularity for 536 cabs over 24 days in downtown San Francisco.

B. Prototype Mobile WiFi Botnet

We utilize 802.11's simple AP-selection algorithm for use in our mobile WiFi botnet experiments. When a mobile bot is within range of open WiFi networks, it chooses the network with the strongest signal strength, performing botnet activities until it is out of range and must connect to a new AP. Such a naive AP-selection approach is easily implemented by botmasters and works equally well in any urban environment, requiring no *a priori* knowledge. If a mobile WiFi botnet can operate under this simple AP-selection algorithm, more complicated mechanisms can potentially achieve better results.

With the requisite information unavailable, we simplify our simulation environment by assuming each AP has the same signal strength, ignoring attenuation due to competing wireless signals and environmental obstructions.[1] Then, by fitting exponential curves to the empirical data in [4]—which gives achievable throughput based on distance for 802.11b/g

[1]It is trivial for us to relax this assumption if and when the required information becomes available.

routers—we find Eq. 1, allowing us to calculate a mobile bot's average throughput (Mbit/s) based on its current distance in meters, d, from an open WiFi AP. Using our improved mobility dataset (Section III-A), our naive AP-selection algorithm and Eq. 1, our simulator calculates each mobile bot's throughput at a one-second granularity, which we use to determine if open WiFi networks are suitable for supporting botnet activities.

$$f(d) = \begin{cases} 37.5 & : d \leq 1.52 \\ 37.97 * e^{-8.15*10^{-3}*d} + 4.154 * 10^{-4} * e^{0.33*d} & : 1.52 < d \leq 22.86 \\ 3379 * e^{-0.22*d} + 22.48 * e^{-0.042*d} & : 22.86 < d \leq 53.34 \\ 0 & : d > 53.34 \end{cases} \quad (1)$$

C. C&C Protocol

Fig. 1 gives a high-level representation of our HTTP-based C&C protocol between a mobile bot and an open WiFi AP. Mobile bots connect to an open AP (STAGE 1), locate the C&C server using DNS (STAGE 2) and, finally, connect to the server and receive new commands (STAGES 3 & 4). Notice that mobile bots only successfully receive commands when connected to a given WiFi network long enough to complete all four stages. Interruption of the protocol due to loss of connectivity resets it to STAGE 1 with the nearest open AP. In our simulations, after successfully completing STAGE 4, our mobile bots immediately attempt to receive another fresh command by returning to STAGE 2; this process continues, so long as they remain connected to the same open AP, allowing us to determine the finest level of control available.

To determine the timing for STAGE 1, we turn to [6], which reports an average of 2.757 seconds for a mobile device (traveling in a vehicle) to scan for an open WiFi AP, associate with it and obtain an IP. For the later stages, we must account for 802.11b/g's overhead and transmission rates, as well as wired communications between the open WiFi AP and Internet servers (i.e., DNS and botnet command servers). In 802.11b/g, a 24-byte PLCP preamble and header must be transmitted at a constant 1 Mbit/s before any subsequent transmission; this constant, 0.192-second overhead is incurred for every wireless transmission in STAGES 2, 3 and 4. Messages are then transmitted at a rate determined by the bot's distance from the WiFi AP and Eq. (1), recalculating the rate every second. To simplify the complexity of Internet communications, we use the median Round Trip Time (RTT) of 0.086 second for cable connections in the USA [5] to estimate the RTT between the WiFi AP and Internet servers. The packet sizes of TCP's 3-way handshake in STAGE 3 and the FIN message in STAGE 4 are well defined. We chose a message size of 100 bytes for the DNS messages in STAGE 2 and 512 bytes for the HTML messages in STAGE 4, providing adequate space for complicated messages within a single packet.

D. DDoS Protocol

Botmasters achieve the required coordination for DDoS attacks via the C&C channel, informing bots of future attack targets and times. Later, at the appropriate times, mobile bots with access to open WiFi networks perform the synchronized DDoS attacks. Unlike the C&C protocol in Fig. 1, mobile bots only require an online connection (STAGE 1) before they can begin a DDoS attack. As they move, bots continue to connect to different open WiFi APs, issuing a stream of SYN packets at the target for the duration of the attack. Hidden among benign traffic and dispersed across multiple networks, such a DDoS attack would be difficult to mitigate.

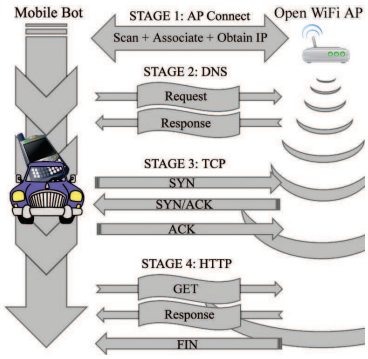

Fig. 1. Mobile bot's C&C protocol

IV. Experimental Results

In this section we will take a close look at our evaluation results. Our simulator makes use of fine-grained cab mobility traces, as described in Section III-A. We examine the feasibility of mobile WiFi botnets during weekends and the office-commute hours (i.e., rush hours) of 6:00-10:00am and 3:30-7:30pm on weekdays. Ignoring the issue of infection (Section II-C), we treat each cab as a mobile bot. Notice that our cab mobility traces are, in actuality, a collection of different people's trips around the city. Therefore, we are actually treating each passenger as an infected mobile device during the duration of his/her trip in the cab. Because we only have traces for 536 cabs, the maximum mobile WiFi botnet size at any point in time is limited to 536 bots in our simulations. This is a very small fraction of the total mobile devices in the city, but it will give us a good first look at the potential capabilities of future mobile WiFi botnets. Next, we will go over the HTTP-based C&C simulation results. Finally, we examine the DDoS attack simulation results.

A. Command and Control

In this section, we simulate the C&C protocol described in Section III-C to answer the following questions:

- What level of control does the botmaster have, both in the number of bots reachable and how frequently the bots can receive commands?
- How long does it take a command to propagate through the reachable botnet (i.e., maximum of 536 bots)?
- How is the botnet distributed across open WiFi APs?

First, for each day of the week, we calculated the average number of unique cabs/bots that could receive *at least one* command for the hours examined. For the weekday rush hours, the results were fairly consistent, with typically ≈70-75% of the total cabs receiving commands. For the weekend, the results demonstrate a clear diurnal trend, shown in Fig. 2.

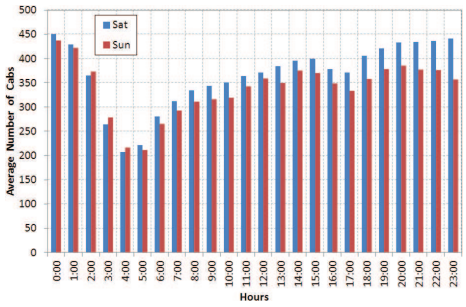

Fig. 2. Avg num cabs recv commands (weekends).

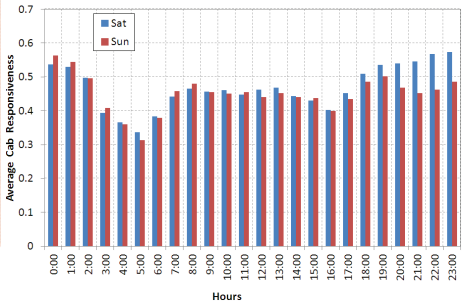

Fig. 3. Avg cab responsiveness (weekends).

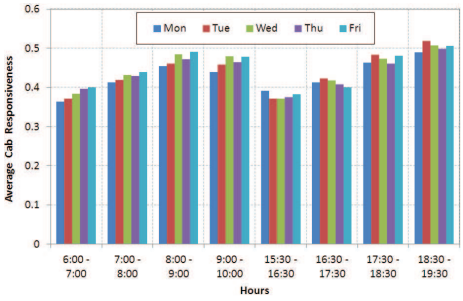

Fig. 4. Avg cab responsiveness (weekday rush hrs).

Next, for those cabs/bots able to receive commands within a given hour, we calculated the percentage of minutes when a new command could be received, which we term the cab/bot *responsiveness*. It can be interpreted as a bot's probability—for a given hour—of receiving commands within a minute of them being issued. Figures 3, 4 and 5, show the average cab responsiveness per day during both the weekday rush hours and weekend, as well as the CDF of the responsiveness per cab across all weekdays and weekends for particular hours. From the figures, we see that the evening hours afford the greatest responsiveness due to the combined traffic congestion of people returning from work and going out for the evening. From Figs. 2 and 3, we find that the peak evening and early morning weekend hours grant access to a large percentage (≈75-84%) of the botnet, with the average bot responsiveness exceeding 50%. This fine level of control is also observed during weekday evenings between 6:30-7:30pm. While evenings impart the greatest level of control, all the hours examined provide relatively quick control over a significant portion of the botnet. Even between 4:00-6:00am on weekends, when people are asleep, ≈39% of bots are reachable with an average responsiveness above 30%.

In Fig. 6, we demonstrate how quickly the WiFi botnet can receive commands by plotting the average number of cabs/bots a newly injected command propagates to over a 2-hour period for various command-injection times. Nearly all commands reach over 55% of the botnet after 30 minutes and over 75% after 2 hours. At 9:00am or 6:30pm on weekdays, commands reach over 65% within 30 minutes of injection. Commands issued on Saturdays at 10:00pm propagate even more quickly, reaching ≈76% after only 30 minutes. Even commands issued at 7:00am on weekends reach over 61% after 2 hours. This is a significant improvement over the Bluetooth-based C&C in [12], achieving ≈67% propagation within 24 hours of injection. These results demonstrate that a relatively fine level of control is possible over a significant portion of the total reachable botnet.

Lastly, we are interested in determining how the botnet is distributed over the open WiFi networks. If only a small set of networks are used, then they have a significant view of the overall botnet activity, making detection and mitigation easier. Figure 7 is a CDF plot of the average number of cabs using an open AP per minute during weekday mornings. It shows that 90% of the open APs are used by fewer than 6 bots during any given minute. Of the remaining 10%, less than 2% ever have more than 10 simultaneous mobile bots, and even they never have more than 23. From these results, we observe that a moderately sized mobile WiFi botnet can successfully be controlled using only open WiFi APs, spreading its traffic across many different networks to hinder detection.

B. DDoS Attack

In this section, we aim to answer the following questions concerning mobile WiFi botnet DDoS attacks:

- What is the botnet's capacity for DDoS attacks?
- How are DDoS attacks distributed across open WiFi APs?

We simulate the DDoS protocol described in Section III-D and plot the average number of SYN packets sent per hour by the mobile WiFi botnet for weekday rush hours (Fig. 9) and the weekend (Fig. 10). We can see that during the peak weekday commute hours of 8:00-10:00am and 5:30-7:30pm, the botnet is capable of issuing ≈1.4 million SYN packets per hour (≈389 per sec). During prime weekend hours, the capacity is even greater, achieving ≈1.4-1.7 million SYNs per hour (≈389-472 per sec). Usually, it can achieve between 0.8-1.2 million SYNs per hour (≈222-389 per sec). Even between 4:00-6:00am on weekends, it can issue ≈500,000 SYNs per hour (≈139 per sec). According to [9], an unprotected server—or one using a default firewall configuration—can survive DDoS attacks of only 100 SYNs per sec. However, a properly configured firewall can survive DDoS attacks of 500 SYNs per sec, effectively defeating the DDoS capacity of our small botnet. Nevertheless, even a small WiFi botnet could prove valuable when used to augment the DDoS attacks of larger, traditional botnets or mobile botnets in other cities.

To determine if the DDoS attack is sufficiently distributed to encumber detection, we have plotted the average and maximum number of APs used per minute between 6:00-10:00am on weekdays in Fig. 11. The complementary CDF in Fig. 8 shows the average number of cabs using an open AP per minute during this same period. From Fig. 11, it is apparent that the DDoS attack is spread across multiple APs, averaging from 220 to over 340 per minute. Furthermore, Fig. 8 shows that over 50% of the open APs used in the attack only service a single mobile bot per minute, and over 86% service fewer than 6 bots per minute. Thus, we find that mobile WiFi botnet DDoS attacks are distributed across many different open networks, with each network participating in only a small portion of the attack; obviously, this makes detection increasingly difficult for defenders. Considering their

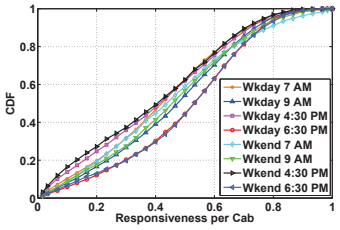

Fig. 5. CDF, responsiveness per cab.

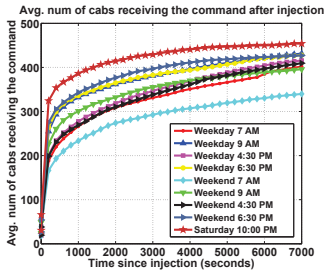

Fig. 6. Command propagation for diff injection times (weekdays/weekends).

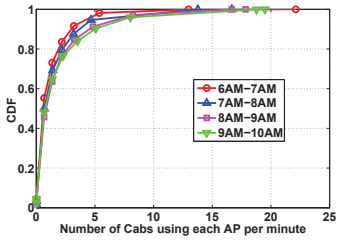

Fig. 7. CDF, avg num cabs using open APs per min for C&C (weekdays).

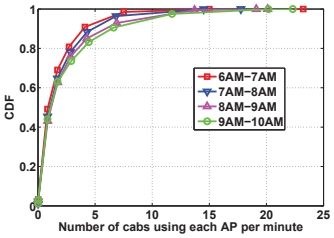

Fig. 8. CDF, avg num cabs using open APs per min for DDoS (weekdays).

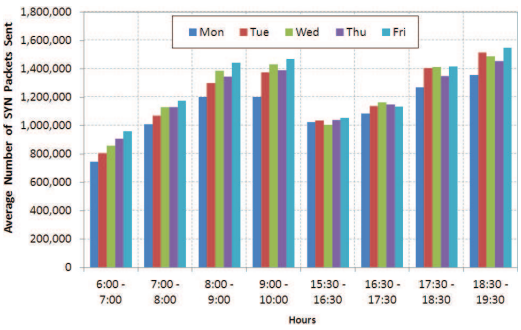

Fig. 9. Avg num SYNs sent (weekday rush hrs).

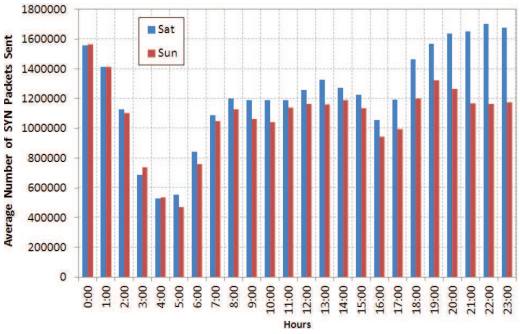

Fig. 10. Avg num SYNs sent (weekends).

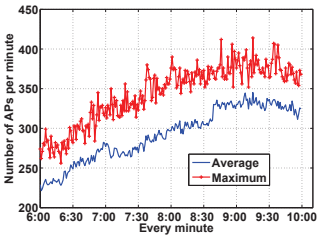

Fig. 11. Avg & max num of APs used per min for DDoS attacks (weekdays).

overall stealth and capacity, mobile WiFi botnets could prove a serious future threat as a DDoS-attack mechanism.

V. Related Work

Recently, there has been a surge of mobile malware which have already started to demonstrate botnet-like traits. For instance, SymOS.Yxes, discovered in early 2009, reports user-sensitive information back to a centralized server through an HTTP-based C&C protocol. `Ikee.B`, targeting jailbroken iPhones, uses a similar HTTP-based mechanism to connect to a control server, download additional components and steal user information. There have also been several research efforts to design advanced C&C protocols for mobile botnets. Singh *et al.* [12] evaluated the feasibility of using Bluetooth as a medium for botnet C&C. Muliner [8] proposed SMS and SMS-HTTP hybrid C&C protocols to facilitate the communication between compromised smartphones. Taking a different angle, Traynor *et al.* [11] demonstrated the impact of DDoS attacks against the core of cellular networks utilizing compromised mobile phones. The focus of these work is on the characterization of large-scale attacks, whereas our work investigates the effectiveness of using open WiFi networks as a stealthy channel to coordinate a large number of moving bots and launch DDoS attacks.

VI. Concluding Remarks

In this paper, we leverage real-life cab mobility traces and actual open WiFi AP locations to successfully simulate the C&C and DDoS attack of a mobile botnet using only open WiFi networks. We have shown that such a mobile botnet, traveling quickly through an urban environment in vehicles, can successfully achieve an HTTP-based C&C channel with a fine level of control and responsiveness. We have shown that even a small mobile botnet can successfully mount a DDoS attack against unprotected (or default firewall) systems. Finally, our simulations have demonstrated that botnet traffic is adequately distributed across open WiFi networks, with no single AP over-utilized at any given moment. Together, these results affirm the stealthy nature of mobile WiFi botnets, making them especially alluring to botmasters.

References

[1] Crawdad: A community resource for archiving wireless data at dartmouth. http://crawdad.cs.dartmouth.edu/.

[2] Tiger: Topologically integrated geographic encoding and referencing system. http://www.census.gov/geo/www/tiger/.

[3] Wigle: Wireless geographic logging engine. http://wigle.net/.

[4] Netgear/atheros: 108mbps/802.11g wireless router performance testing. Technical report, VeriTest, 2004.

[5] M. Arlitt, B. Krishnamurthy, and J. C. Mogul. Predicting short-transfer latency from tcp arcana: A trace-based validation. In *In Proceedings of Internet Measurement Conference (IMC 2005)*.

[6] V. Bychkovsky, B. Hull, A. K. Miu, H. Balakrishnan, and S. Madden. A measurement study of vehicular internet access using in situ wi-fi networks. In *In Proceedings of 12-th ACM MobiCom*, 2006.

[7] J. Härri, F. Filali, C. Bonnet, and M. Fiore. Vanetmobisim: generating realistic mobility patterns for vanets. In *Proceedings of VANET*, 2006.

[8] C. Mulliner and J. Seifert. Rise of the iBots: 0wning a telco network. In *Proceedings of the 5th IEEE International Conference on Malicious and Unwanted Software (Malware)*, Nancy, France, October 2010.

[9] R. Oliver. Countering syn flood denial-of-service attacks. Technical report, Tech Mavens Inc., 2001.

[10] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser. CRAWDAD data set epfl/mobility (v. 2009-02-24). Downloaded from http://crawdad.cs.dartmouth.edu/epfl/mobility, 2009.

[11] P.Traynor, M.Lin, M.Ongtang, V.Rao, T.Jaeger, P.McDaniel, and T.L.Porta. On cellular botnets: Measuring the impact of malicious devices on a cellular network core. In *Proceedings of the 12th ACM Conference on Computer and Communications Security (CCS'09)*.

[12] K. Singh, S. Sangal, N. Jain, P. Traynor, and W. Lee. Evaluating bluetooth as a medium for botnet command and control. In *Proceedings of DIMVA 2010*.