SECURITY AND COMMUNICATION NETWORKS
Security Comm. Networks. 2009; **2**:1–12
Published online in Wiley InterScience
(www.interscience.wiley.com) DOI: 10.1002/sec.149

malnets: large-scale malicious networks *via* compromised wireless access points

Patrick Traynor[1], Kevin Butler[2*,†], William Enck[2], Patrick McDaniel[2] and Kevin Borders[3]

[1]*School of Computer Science, College of Computing, Georgia Institute of Technology, Atlanta, GA 30332, U.S.A.*
[2]*SIIS Laboratory, Pennsylvania State University, University Park, PA 16802, U.S.A.*
[3]*Department of Electrical Engineering and Computer Science, University of Michican, Ann Arbor, MI 48109, U.S.A.*

Summary

Densely populated areas are increasingly filled with vulnerable wireless routers set up by unsophisticated users. In isolation, such routers appear to represent only a minor threat, but in aggregate, the threat can be much greater. We introduce the notion of *malnets*: networks of adversary-controlled wireless routers targeted to a physical geography. Similar to Internet worms such as Slammer and Code-Red, malnets are created by the recursive compromise of targeted devices. However, unlike their traditionally wired counterparts, malnet worms exploit only other routers that are within their transmission range. The malnet thus creates a parallel wireless infrastructure that is (*a*) completely under control of the adversary, and (*b*) spans a targeted physical area, creating a valuable infrastructure for a variety of virtual and physical attacks. We initially study the propagation characteristics of commercial routers and model inter-router connectivity using publicly available war-driving data. The resulting characterization is applied to well-known epidemiological models to explore the success rates and speeds of malnet creation across cities such as New York, Atlanta, and Los Angles. Finally, we use a sampling of available exploits to demonstrate the construction of multi-vector, multi-platform worms capable of targeting wireless routers. Our analysis show that an adversary can potentially deploy a malnet of over 24 000 routers in Manhattan in less than 2 h. Through this work we show that malnets are not only feasible but can be efficiently deployed. Copyright © 2009 John Wiley & Sons, Ltd.

KEY WORDS: Q1 malware; routing; security

Q1

1. Introduction

Wireless routers provide an easy way to introduce small networks into homes, offices, and public spaces. These networks change the relationship between users and the Internet: they free us to move about our personal environments without loss of connectivity. Such utility has enormous positive and negative social impact. In Manhattan alone, there are more than 29 000 discoverable routers south of 59th street [1]. The dangers of such devices often lie in their simplicity. Because there is little technical or monetary barrier to introducing these networks, users do so with impunity. They often lack the training, knowledge, or motivation to properly secure those networks. In addition, devices and the protocols they run are in

*Correspondence to: Kevin Butler, 344 IST Building, University Park, PA 16802, U.S.A.
†E-mail: butler@cse.psu.edu

many cases poorly designed, buggy, and vulnerable [2,3]. Hence, the vast majority of deployed wireless routers are vulnerable to compromise.

This work explores the feasibility and use of *malnets*: networks consisting solely of compromised commodity wireless routers that route and manipulate traffic over targeted physical areas. Because all malicious traffic is only wirelessly routed between malnet nodes, it is outside the visibility and control of ISPs and law-enforcement. Such infrastructure, when isolated or (more dangerously) when combined with other malicious infrastructure such as botnets, represents a potentially catastrophic new piece of adversarial apparatus, allowing adversaries to invade privacy, launch attacks, and mask nefarious activities within the compromised region. Such attacks not only could support traditional cyber-crimes such as spam delivery and phishing, but also attacks on critical infrastructure and vehicles for command and control of coordinated physical attacks on real world targets.

A malnet is constructed as follows: the adversary initially locates and compromises a single wireless router through the wireless[‡] or wired interface. Such compromise is achieved *via* attacks ranging from exploiting default/poor passwords to a vast array of publicly documented vulnerabilities, allowing the adversary to gain administrative control of the router. New firmware is then loaded. The newly compromised router then attempts to compromise all routers within its transmission range using the same or similar attacks. Similarly to Internet worms, this infect and propagate cycle spreads across routers covering a potentially large physical area. Figure 1 summarizes this process. Once constructed, the malnet forms an *ad hoc* network and routes, filters, and manipulates traffic as is desired by the adversary. Note that once a malnet is installed, its removal would be exceedingly difficult. Maliciously flashed routers can mask the presence of malcode and prevent new images from being burned. Finding, checking, and physically replacing flash memory on all 29 000 routers in southern Manhattan would be intractable in any practical sense.

In this paper, we evaluate the degree to which malnets are feasible to construct. malnet construction is crucially dependent on the density and placement of routers in the targeted area and the adversary's ability to systematically compromise routers. We investigate these facets of current environments through three dependent efforts. First, we study the relative propagation characteristics of commodity routers and develop a model of router density and connectivity in metropolitan areas based on publicly available wardriving data. In so doing, we confirm our hypothesis that router deployments are of sufficient density to support malnets in major cities such as New York and Atlanta. Second, we apply epidemiological models to evaluate the probability that a single or small number of initial infections will lead to a malcode outbreak, and quantify the speed of the malnet growth. Our results show that a relatively simple piece of malcode could span the majority of Manhattan in approximately one day. More aggressive malcode can span the same physical area in less than 2 h. Finally, we use a sampling of available exploits to demonstrate the construction of multi-vector, multi-platform worms capable of targeting wireless routers. These attacks include a new low-bandwidth denial of service vulnerability discovered during the course of this work. The joint result of these analyses is that with very high probability, adversaries with only moderate sophistication can today deploy malnets covering large cities such as New York.

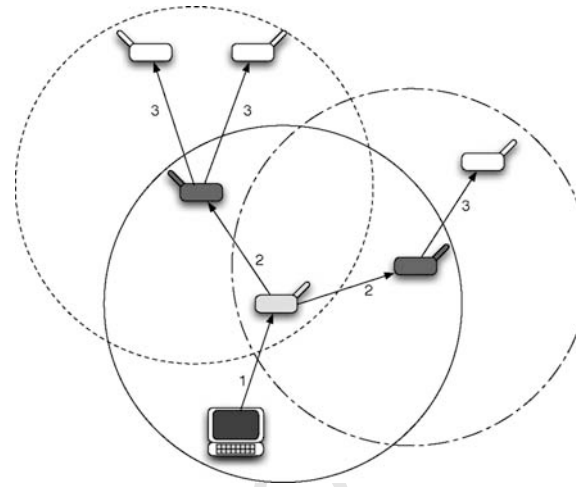

Fig. 1. A high-level overview of malnet creation. An adversary, represented as the laptop, finds and compromises a wireless router (light gray). The compromised router then finds nearby routers (dark gray) and compromises them. The process of penetration and propagation continues until all routers in an area are absorbed into the malnet.

2. Applications of malnets

Because compromised routers control the flow of traffic between clients and the Internet, malnets can more

[‡]Some routers limited access to configuration operations to the wired interface. However, as detailed in Section 5, such mechanisms are very often easy to circumvent.

DOI: 10.1002/sec

efficiently implement many of the attacks common to traditional botnets. Denial of service attacks can be mounted by injecting volumes of traffic into the wired network through many connected clients, or directly by filtering client traffic. Simple phishing attacks using compromised routers, which have been previously suggested [4], could be extended to occur exclusively within the malnet so as to circumvent ISP blacklists. Adversaries attempting to steal sensitive user data could probe attached devices without concern for provider-level firewalls or intrusion detection systems. Current client-based countermeasures would also provide little defense or diagnostic ability against such attacks.

The geographic proximity of compromised nodes in a malnet creates the potential for many new attack vectors. By recording the MAC addresses of attached devices, malnets could be used to physically track large numbers of users throughout a city. An adversary attempting to physically harm a user could keep a near constant watch as their target is at home, work, or their local coffee shop. Such surveillance could be coordinated between multiple adversaries by also using the malnet as a city-wide "push-to-talk" voice over IP (VoIP) network. Crimes including kidnapping, targeted mugging, and burglary (informed by the fact that the user is currently across the city from their home) could all be assisted by such a network.

Should targets within the wired Internet become desirable, the owner of a malnet could obfuscate the source of their attacks. For example, traffic for a distributed denial of service attack on a website could be spread across thousands of backhaul connections. If a specific provider were to begin filtering traffic in the middle of an attack, the malnet could quickly reroute its efforts to other nodes. The sea of anonymity provided by such a network would make solutions short of dropping all connectivity to an entire city simply ineffective.

3. Modeling Connectivity

In this section we characterize the ability of existing wireless routers to support malnets. We begin by using wireless signal propagation models to establish realistic transmission ranges of deployed routers.

We use reports from the Wireless Geographic Logging Engine (WiGLE) [1] database throughout. As is true of all war-driving data, the WiGLE database provides a conservative sample of the actual wireless networks. Wardriving can, at best, provide an incomplete map of an area because of physical limitations. For example, wireless access points within tall buildings may be unobservable from street level. The devices actually listed in the data set are the result of numerous observations—the location of each access point has been reported an average of approximately 50 times [1]. Given this level of independent scrutiny, we believe that the location information is as accurate as collection techniques allow. Although the location of routers is generally static, some small portion of the devices in our snapshot are likely no longer located where previously denoted. As we will show, the loss of such nodes is statistically insignificant given the density of remaining devices. Moreover, because the number of deployed wireless routers continues to rise [1], it is credible to believe that as of yet unrecorded nodes exist in close proximity to the lost nodes. Accordingly, we assert that the data is a representative portrayal of reality.

3.1. Modeling Radio Propagation

The propagation of wireless data is governed by a number of factors including distance, attenuation through materials such as concrete and steel, diffraction and reflection, shadowing, and various fading behaviors. While router manufacturers advertise transmission ranges of up to 40 m, the influences of the above elements can substantially decrease these ranges [5,6].

Because of the presence of complex, varying and irregular construction materials and building designs, the only means of perfectly modeling any metropolitan area is through meticulous knowledge of the environment—an intractable process in all but the most constrained cases. Several highly accurate models have been suggested to approximate connectivity conditions in such environments. Sridhara *et al*. [7] developed estimations for indoor coverage *versus* distance from access points in urban 802.11 mesh networks based on a comprehensive measurement study. At a rate of 1 Mbps[§], the authors were able to demonstrate probabilities of approximately 100, 95, and 75 per cent for direct connectivity between an indoor and outdoor node separated by 50, 75, and 150 m in urban settings. Similar observations, with a median connectivity distance of 150 m, were recorded in a measurement study by Bychkovsky *et al*. [8]. Accordingly, we consider this range of values as realistic for the remainder of this study.

[§]Different modulations, each with an increased amount of error correcting mechanisms, are available as the transmission rate is reduced. Lower transmission rates therefore propagate better as a function of distance.

DOI: 10.1002/sec

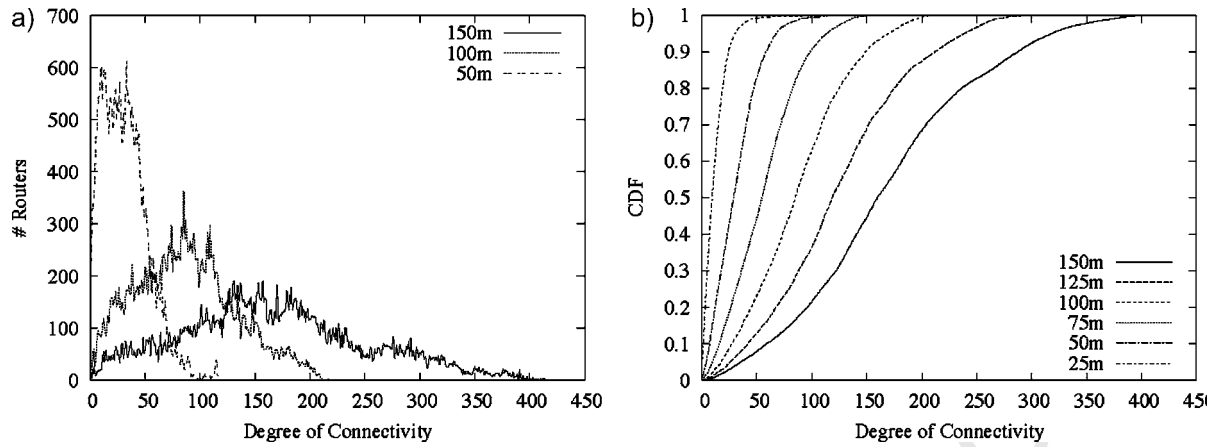

Fig. 2. The modeled number of reachable peer routers (degree of connectivity) by transmission range, and the cumulative distribution function (CDF) for node degree varying maximum transmission radii. (a) Reachable peer routers; (b) CDF for node degree.

Figure 2(a) and (b) illustrate the application of the Sridhara *et al*. model to routers in Manhattan south of Central Park (below 59th Street). This histogram in Figure 2(a) indicates the degree (i.e., number of reachable peer routers) of nodes with maximum transmission radii of 50, 100, and 150 m. A density of 2321.35 nodes/km^2 and an average degree of 31.63, 87.37, and 167.91 suggests that a highly connected wireless network is possible. The high average degree shown in both Figures is the result of dense populations (e.g., 150 m radii $\approx$ 70 000+ meters^2).

Note that, Figure 2(b) shows that for the 50, 100, and 150-meter cases respectively, 98.5, 99.85, and 99.9 per cent of all nodes in the network have *at least* one neighbor within 50, 100, and 150 m, respectively. Even with our conservative radio ranges, the probability that nodes are isolated is very small.

3.2. Approximating Network Topology

We now extend the model from a single router to model the potential connectivity of an arbitrary metropolitan area. In spite of the regular grid pattern common to most cities, the placement of wireless routers is not predictable. To account for this complexity, we considered numerous models that characterize complex networks [9]. We began with the power-law topology, given that it has been used to approximate larger networks such as the Internet [10]. However, examination of Figure 2(a) shows no correspondence to this topology as node degree is not exponentially distributed.

We turn to random-graph theory—a technique used to successfully *approximate* connectivity in other wireless networks with very similar characteristics [11]. Such a technique allows us to make estimates about the required density of nodes for malnets to successfully propagate. A random graph $G(n, p)$ is composed of n nodes with probability p that a link exists between any two vertices. Graphs in which $p = 1$ are fully connected, whereas nodes in graphs with $p = 0$ are entirely isolated. There is a threshold value for p in which the connectivity of $G(n, p)$ transitions from "nonexistent" to "certainly true" [12]. Given a very large random graph with monotone properties, the probability that it is connected can be defined as

$$P_{\text{conn}} = \lim_{n\to\infty} P[G(n, p) \text{ connected}] = e^{-e^{-c}} \text{ and}$$

$$p = \frac{\ln(n) + c}{n}$$

where c is any real constant. For any $G(n, p)$, it is therefore possible to determine an average node degree $d = p * (n - 1)$ such that the graph is connected with probability P_{conn}. Figure 3 illustrates the increasing connectivity for a network of 30 000 nodes as the average number of edges per node increases.[‖] For the network to be fully connected with a probability of 0.99999, nodes would require approximately 22 neighbors. As the nodes in the Manhattan data set have an average degree of 18.23, 72.93, and 164.08 for 50, 100, and 150 m transmission radii, such a network is possible.

[‖]There are approximately 30 000 nodes in Manhattan between the southern edge of Central Park and the southern tip of the island [1].

DOI: 10.1002/sec

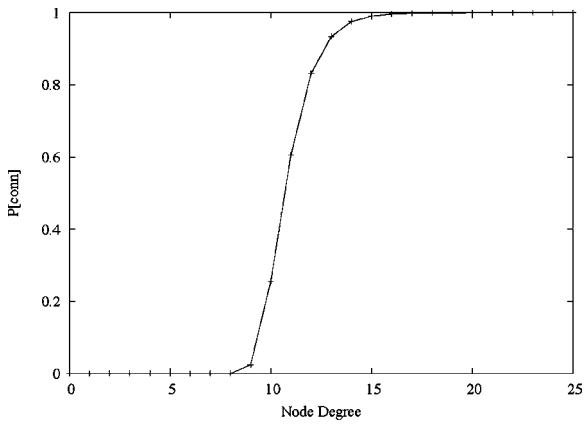

Fig. 3. The probability that a random graph is fully connected given a varying average node degree.

3.3. Characterizing Network Connectivity

The WiGLE data is now used to validate our connectivity model. The analysis above demonstrated that there exist nodes that are isolated. We therefore characterize the largest observable subgroup as a means of validating our general model of network connectivity. Figure 4 shows the proportion of the nodes that are connected to the largest partition as a function of the maximum transmission radius. For the 50, 100, and 150 m cases, approximately 93.18, 98.55, and 99.68 per cent of the nodes in Manhattan are reachable through the largest subgraph. With an average observed degree of 4.56, the largest subgraph in the 25 m case only connects 33 per cent of the wireless routers and therefore fails to span the city. Thus, the measured connectivity of Manhattan is accurately approximated by the random graph model.

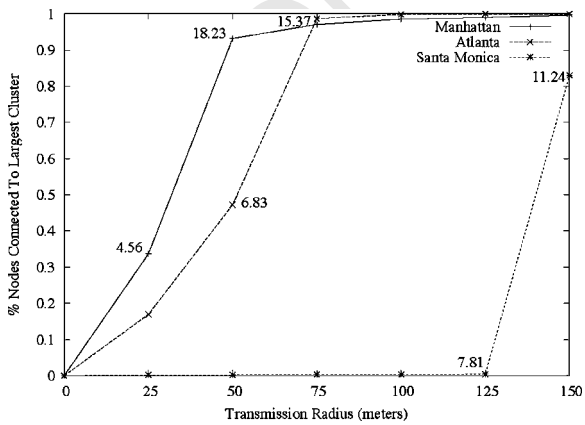

Fig. 4. The percentage of nodes in the largest cluster in the network *versus* transmission radius. When the density falls below 22 neighbors, the percentage of nodes comprising the largest subgraph falls quickly.

Samples from two additional metropolitan locations, the Buckhead/Downtown region of Atlanta and the Santa Monica area of Los Angeles [1], were then selected for analysis. With a density of 870.09 and 159.04 nodes/km^2, these samples were arbitrarily chosen to represent urban areas with medium and low node density, respectively. For a transmission radius of 150 m, an average node degree of 61.49 and 11.24, respectively, are calculated. Connectivity experiments in Figure 4 again corroborate the theoretical results from Figure 3 for both cases by illustrating the rapid loss of connectivity as a function of decreased node degree. Note that for all cases, an average node degree below 10 (or approximately 140 nodes/km^2) always results in the inability to create a significantly sized malnet. Local anomalies (e.g., a isolated district on the Lower East Side of Manhattan, seen in Figure 4) aside, the random graph model tells us that any city in which wireless routers have a degree of greater than 20 are almost certain to support a malnet.

As a final test of connectivity, we examined the possibility of malnet creation with only a single brand of wireless router. For these experiments, we selected Linksys because of its position as the market leader. Using MAC addresses to filter for this brand, our results demonstrate that a subgraph connecting 49.94, 96.08, and 97.84 per cent of the Linksys routers across Manhattan is possible using transmission radii of 50, 100, and 150 m. While targeting all routers clearly creates a network more able to deal with node failure, attacks on a single brand of router can be used to form an effective, city-wide malnet.

4. Propagation

Equally critical to knowing that malnets *can* be built is understanding how they will be built. A malnet requires the network reach a point in which the vast majority of routers in the metropolitan area are compromised. We assume a model in which the routers are compromised by a multi-vector worm that infects and propagates across routers in a manner similar to Internet worms (e.g., Blaster [13], Slammer [14], and Code Red [15]). In this section, we characterize the propagation of those worms. Like previous work (e.g., [16,17]), we begin by applying epidemiological models to demonstrate the potential for widespread infection.

4.1. Epidemic Models

A first order approximation analysis of the malnet spread can be made by studying the network's *epidemic*

DOI: 10.1002/sec

threshold. Represented as the birth/infection rate (β) over the death/cure rate (δ), an epidemic threshold represents the point over which a pathogen will become persistent within a network. Any network that does not meet the threshold cannot probabilistically sustain the epidemic—the malcode will die out and the malnet will fail to form. In practical terms, the birth rate is the rate at which routers are infected by malcode, and the death rate is the rate at which the compromised routers are discovered and either repaired or shut off.

Epidemic thresholds can be estimated using well-known techniques. Kephart and White [16] investigated the propagation of viruses in Erdös-Rényi random graphs. As demonstrated by Wang *et al*. [18], the epidemic threshold τ in these topologies is equal to the inverse of the average node degree $< k >$, such that $\tau = \frac{1}{<k>}$. In Section 3 we showed that the average node degree for the connectivity graph was 167.91 given a transmission radius of 150 m, meaning that for $\beta = \frac{\delta}{167.91}$ the epidemic threshold will be reached; if β exceeds this value then the network will experience an infection epidemic.

Wang *et al*. [18] generalize a model of determining epidemic thresholds to be topology-independent, and assert that by considering the adjacency matrix of the graph, the epidemic threshold is the inverse of the largest eigenvalue in the adjacency matrix; that is,

$$\tau = \frac{1}{\lambda_{1,A}}$$

given an adjacency matrix **A**. From our graph of wireless connectivity in midtown and downtown Manhattan, we derived an adjacency matrix based on the same 150 m transmission radius. The largest eigenvalue in the matrix was found to be 321.13, thus making the corresponding epidemic threshold $\frac{1}{321.13}$.

Intuitively, the death rate must be greater than the birth rate times the inverse threshold for the network **not** to saturate. Table I summarizes he epidemic thresholds for these models, shown as their inverse (τ^{-1}). These exceedingly low thresholds coupled with a natural low death rate for router compromise (see Section 5) indicates that the malcode will spread throughout the network with very high probability.

4.2. Modeling Infection

We now consider the *rate* at which malcode will propagate across Manhattan. We use the *parasim* simulator [19], which takes as input a graph topology.

Table I. Inverse epidemic thresholds using the Erdös-Rényi random graph model and the topology-independent eigenvalue model from Wang *et al*. [18].

Tx Range (m)	Random graph	Topology-independent
50	31.63	110.65
100	87.34	178.36
150	167.91	321.13

We assume an initial number of randomly placed infected nodes I and that in each one minute epoch t an infected node can infect one of its neighbors with probability P_i. For simplicity, we assume the malcode does not possess the state to be cognizant of which hosts it has previously infected; an infected router will randomly pick one of its neighbors to infect. Simulations were run on the connectivity graphs for Manhattan using wireless transmission ranges of 150, 75, 50, and 25 m. For each set of parameters, we performed 100 simulation runs and plotted the mean of the results. We do not show errorbars in these plots for clarity, but a 95 per cent confidence interval shows that deltas are on the order of two magnitudes less than the mean, and as time progresses, these converge to roughly zero as saturation is reached. Note that the highly conservative range estimates developed in Section 3 are lower bounds on infection rates; actual propagation speeds could be substantially higher.

4.2.1. Simple attack model

We begin our analysis by examining a single attack vector using traditional epidemiological models. We chose an attack exploiting poor router administrative password selection, given the ubiquity of this in the wild [20]. To demonstrate the attack, we installed a custom build of OpenWRT firmware on a Linksys WRT54GS router and communicated wirelessly with a Linksys WRT54GL router running Linksys firmware version v4.30.02. We ran one million attempts over 802.11b and recorded a mean of 5200.64 attempts per minute, with standard deviation of 515.73.

Given the ability to launch password cracking attacks between routers, we conservatively model the impact of such attacks using an infection rate of $P_i = 0.02$. We seed infection in 10 and 100 random nodes in Manhattan, and show the resulting propagation rates in Figure 5(a) and (b). The 75 m transmission radius test shows that within 3 h over 2000 routers have been infected. By minute 520 (8 1/2 h), more than half of all routers in midtown and lower Manhattan (≈15 000)

DOI: 10.1002/sec

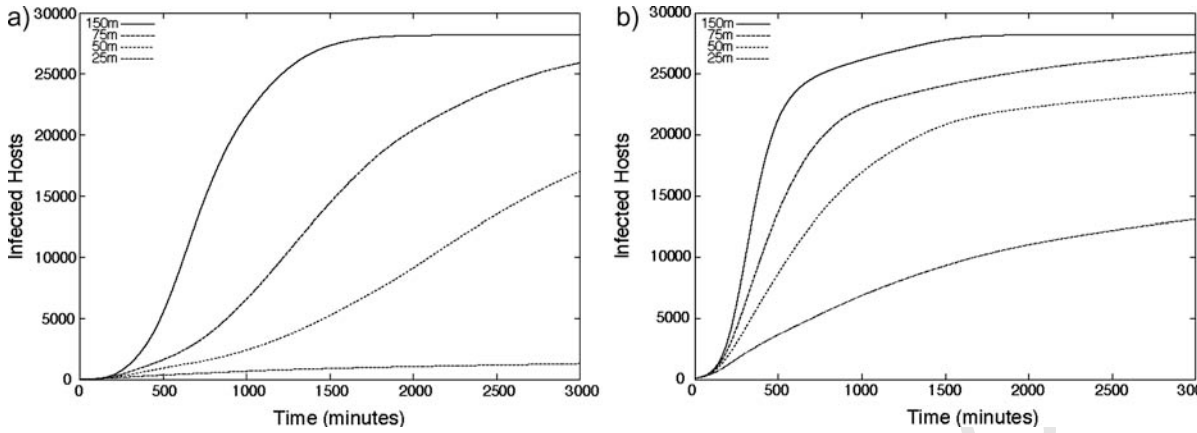

Fig. 5. Propagation of malcode in a wireless network given 2 per cent infection probability and 10/100 initially infected nodes. (a) 10 initally infected nodes; (b) 100 initially infected nodes.

have been infected. 12 h into the attack, over 24 000 routers, or more than 80 per cent of all routers, have been compromised. Increasing the transmission radius increases both the infection rate and the resulting equilibrium of infected nodes, due to the increased graph connectivity.

These simulations do not model a probability of inoculation, as users will be very unlikely to detect and recover from compromises during the duration of the malnet setup. As discussed in Section 5, there is a low chance that users will actually attempt to change their firmware, particularly since router manufacturers themselves often actively discourage users from doing so. The inoculation rate in this case would therefore be the probability that a user buys a new router during the course of the attack. The infection rate and time period may be affected by other factors. If the user's password on a given router is not in the dictionary, the probability of infection could be lower, as the search space could be larger to determine the correct password. However, informal studies show that between 25 per cent [21] and 50 per cent [22] of routers use default passwords, allowing compromise within seconds.

4.2.2. *Multi-vector worms*

In reality, an adversary would most likely not rely upon a single attack vector, rather using a series of publicly available exploits (e.g., [23,24]) to create the malnet. We therefore model a more realistic set of parameters. Figure 6(a) and (b) show the results of propagation given $P_i = 0.05$ for $I = 10$ and 100, respectively. For a transmission radius of at least 50 m, there is sufficient connectivity such that over 85 per cent of nodes

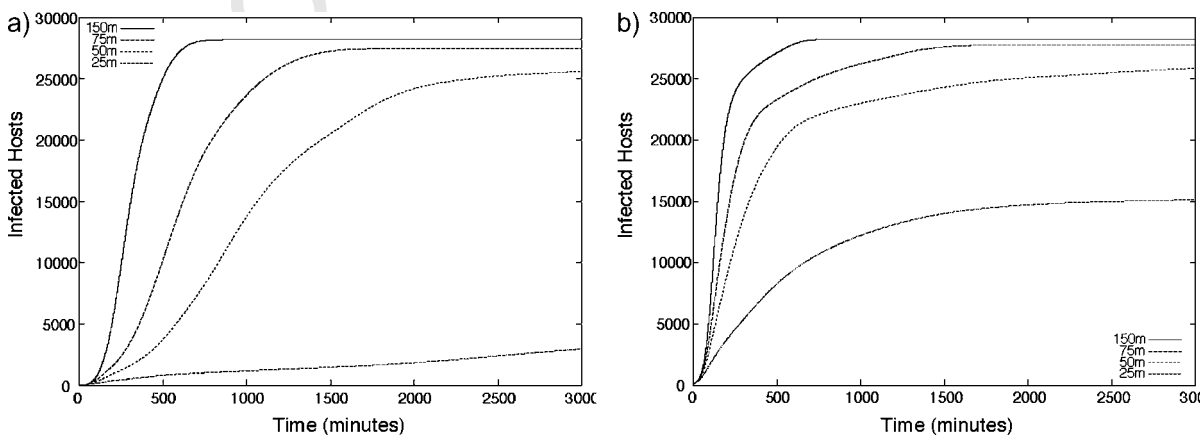

Fig. 6. Propagation of malcode in a wireless network given 5 per cent infection probability and 10/100 initially infected nodes. (a) 10 initally infected nodes; (b) 100 initially infected nodes.

DOI: 10.1002/sec

within Manhattan will be infected, and the distribution of routers is sufficiently dense that infections will happen quickly.

The main differences between the single (2 per cent) and multi-vector (5 per cent) infection cases are the decreased times to saturation and the flattening of the infection curves. Note that in either case the malnet will not saturate or exhibit the characteristic epidemic curve for 25 m transmission range; however, a range of at least 50 m leads to saturation.

4.2.3. Worst-case scenario

Consider an aggressive but plausible worst-case scenario. Suppose an attacker has specific exploits for all wireless routers. This is not far-fetched; in our examination of the WiGLE data for Manhattan, the top 10 wireless router manufacturers accounted for 18 844 of the 29 452 routers identified, or 64.0 per cent. This number could be higher given MAC-layer filtering by ISPs causing users to clone device addresses on their router. A motivated attacker could compile vulnerabilities for at least this many routers and create or discover an assortment of undocumented exploits. In conjunction with the attacks outlined in Section 5, the adversary could have the tools for a rapid mass-compromise. Consider that an infection rate of 20 per cent exists. Figure 7(a) shows that if the effective transmission rate is at least 50 m, the malnet can be very quickly established. Assuming a 75 m transmission radius, over 15 000 routers have been infected within 52 min. By minute 117, over 24 000 routers, or 80 per cent of all routers, will be infected—the adversary will have created a malnet that effectively covers lower Manhattan within 2 h.

4.2.4. WPA protection

The above simulations assume equal likelihood of compromising any wireless router. In reality, routers protected by WPA are considerably more difficult to attack, compared to the almost trivial ability to exploit WEP routers. We model WEP-protected routers to be as vulnerable as those with no encryption and only consider WPA protected-routers. However, until WPA is fully deployed, routers will still be susceptible to relatively simple attacks. There are also barriers to adopting WPA; certain devices in home networks (e.g., some video game consoles, DVRs) do not support WPA, potentially causing owners of these devices to degrade their wireless networks to WEP or to turn off encryption altogether. While studies have shown about 15 per cent of users adopting WPA [3], the WiGLE data only shows whether a router is protected or not. We created 25 new connectivity topologies where 15 per cent of protected nodes were removed.

Figure 7(b) shows the results of the password cracking attack given these more realistic topologies. The confidence intervals were again two orders of magnitude less than the mean, showing that deviations were minor. As these results show, if a 75 m transmission radius is assumed, we will have infected 2000 routers in approximately 3 h. Over half of all routers

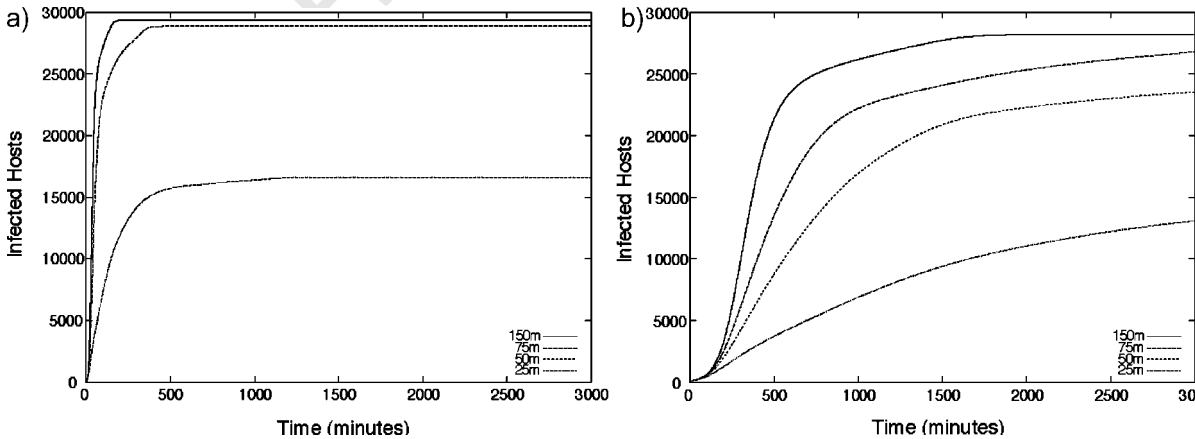

Fig. 7. Simulated results, first assuming a worst case scenario where the probability of infection is 20 per cent, then assuming 15 per cent of protected nodes in the Manhattan graph are using WPA and cannot be cracked, with a 2 per cent probability of infection. 100 initial node infections are used in both cases. (a) Worst-case: $P_i = 0.20$; (b) 15 per cent WPA-protected nodes.

DOI: 10.1002/sec

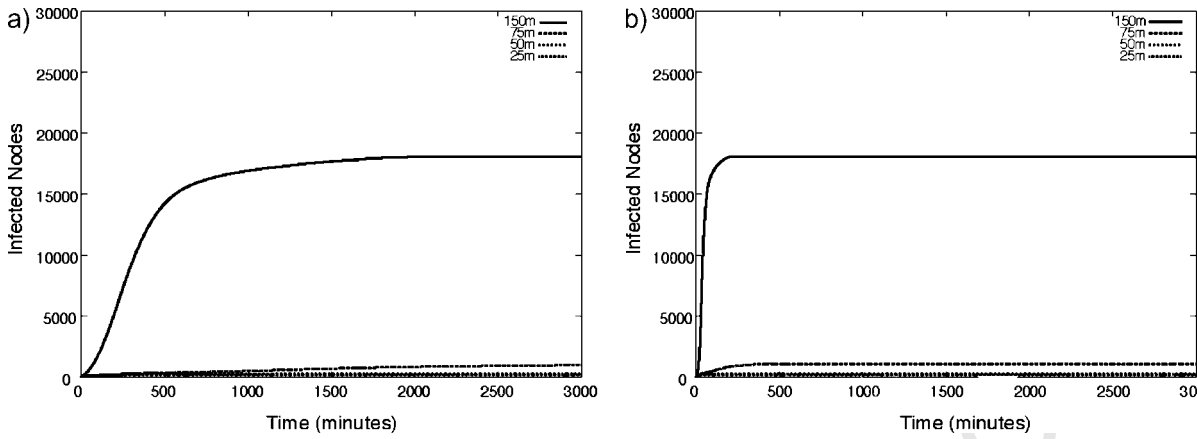

Fig. 8. Simulated propagation if all routers employ good security measures, given a 2 per cent infection probability and the previous worst-case scenario of a 20 per cent probability of infection. 100 initially infected nodes in both cases. (a) 2 per cent infection probability; (b) 20 per cent infection probability.

will be infected 9 h after the initial infection, while 24 000 routers, or over 80 per cent of the network, will be infected in just over 24 h.

Our final analysis considers a more optimistic scenario, where the public is cognizant of the vulnerabilities of their wireless routers. We assume that all wireless router users who currently have encryption enabled on their routers have securely configured their devices, have moved to using WPA and have chosen sufficiently good passwords to make them impervious to dictionary attacks. To model this, we consider any router that uses encryption from the WiGLE data to be impervious and remove it from the topology. Such protection effectively removes approximately 30 per cent of all routers from those exposed to our malnet worm. Figure 8(a) shows that given this scenario, the attack levels off after saturating a much smaller portion of the city, and only under the 150 m transmission radius scenario. Otherwise, very limited compromises of approximately 1000 routers will result. Even in the worst-case scenario, while the spread of the infection is faster, still only a small portion of the overall topology will be compromised, shown in Figure 8(b). The results seem encouraging; by employing better security techniques, the ability to form a large malnet is curtailed. However, as we describe in the next section, even solutions such as WPA are not immune to vulnerability.

5. Compromising Wireless Routers

Constructing a malnet depends on the ability to penetrate and propagate across wireless routers. Like traditional Internet worms, malnet malcode infects and spreads from device to device. To account for router diversity, such worms would contain multiple exploits. This is not a novel idea; even the original Morris worm infected hosts by attacking *rsh*, *fingerd* and *sendmail* [25]. Worms exploiting both Windows and Solaris machines have been known since 2001 [26]. Functionality such as JavaScript running on the router [27], can also be exploited due to a significant number of implementation weaknesses. Given that tools and tutorials [28] for building such worms are widely available, constructing a malnet worm is simply a matter of accumulating exploits.

As its first attack vector, a malnet worm would attempt to attack the administrative interface, allowing control of all aspects of the router, from DNS settings to firmware updates. Through the use of publicly available default administrative usernames and passwords for these devices [29], infected firmware could be loaded onto the targeted router with little difficulty. In cases where upgrading firmware over wireless links is not permitted by the device, remote administration settings can be modified from this interface to allow the adversary to load firmware through an Internet connection. Because an estimated 25– 50 per cent of users fail or are unable to change their default passwords [21,22], this initial attack would compromise a significant proportion of devices within seconds. Furthermore, because users that do change default passwords often replace them with weak substitutes, standard online dictionary attacks [30] would provide a similarly rapid compromise vector.

In the face of strong passwords, a malnet worm would then use one of a multitude of device-specific

DOI: 10.1002/sec

vulnerabilities. If the make of a targeted device is not immediately available from the ESSID, the manufacturer can quickly be identified *via* the MAC address [31] or through wireless fingerprinting techniques [32]. Vulnerabilities including buffer overflows (e.g., [23,24]) and authentication circumvention weaknesses (e.g., [33,34]) could then be launched. While manufacturers provide firmware upgrades to fix many such vulnerabilities, few upgrades are installed in practice,[¶] leaving many routers at risk.

Absent the above attack vectors, a malnet worm could use a sustained denial of service (DoS) attack to compromise wireless routers. When packets cease to be delivered, owners may resort to using the device reset button, which restores device configuration to factory defaults. Upon reset, default passwords can once again be applied toward device compromise. Our testing of a Linksys WRT54GL revealed a previously undocumented low-bandwidth DoS vulnerability (detailed in our technical report [36]).

Standard layers of defense, including the Wired Equivalent Privacy (WEP) protocol, are easily circumvented. Attacks on WEP, which was initially broken by Walker [37] in 2000, have exploited weaknesses in key use, the RC4 cipher [2,38] and the 802.11 protocol itself [3] to reduce practical key recovery time to less than 60 s [39]. In spite of glaring vulnerabilities and continued improvements in key recovery times, WEP remains the only widespread security protocol for such networks.

WEP's replacements, the Wi-Fi Protected Access (WPA) protocol and IEEE 802.11i, provide notably improved security. Unfortunately, both protocols are only minimally deployed due to issues including complex configuration and a lack of universal hardware support. As the setup requires users to configure an external authentication server (e.g., RADIUS), such solutions are well beyond the technical skills of most wireless router owners. In an attempt to meet the needs of most users, WPA provides a "personal" mode using pre-shared keys; however, this configuration remains susceptible to both online and offline dictionary attacks [40].

[¶]One major manufacturer states: "If the Router's Internet connection is working well, there is no need to download a newer firmware version, unless that version contains new features that you would like to use. Downloading a more current version of Router firmware will not enhance the quality or speed of your Internet connection, and may disrupt your current connection stability." [35]

6. Discussion

In the short-term, there are a number of modifications to the firmware that would curb the creation of malnets. Most importantly, current practices, especially those discouraging users from updating firmware, must be abandoned. Second, exponential backoff on unsuccessful login attempts can throttle online password attacks. While the selection of weak passwords would still remain a problem in the long-term, an attacker would have to devote a significantly longer period of time to compromising even poorly protected devices. Third, new versions of firmware should include a password strength checker and prevent users from setting passwords below some entropy threshold. Finally, the administrative interface could be made unavailable to wireless users. Only those users physically able to access a LAN port could then access a device's privileged management interfaces. Such a solution is still vulnerable to attack from compromised hosts within the network. While a significant number of attacks would be prevented through such simple mechanisms, exploits such as buffer overflows and the circumvention of authentication mechanisms would still successfully compromise these devices.

The longer term evolution of security for wireless routers must account for poor default settings and insecure code. Currently, most routers are shipped with standard passwords and no security enabled. This is appropriate only if users can be expected to enable safe settings; however, most do not. Manufacturers could physically label devices with secure, randomly generated default administrative passwords and WPA passphrases. This mechanism would work, assuming that the wireless routers are physically protected from adversaries; however, such a mechanism may cause difficulties if passwords are lost (e.g., stickers removed from the device). By default, manufacturers could also enable WPA "personal" mode, if not "enterprise" mode. Given that third party router software for running a RADIUS server on a router is already available [41], such an option is entirely realistic if configured by the manufacturer. Note that older client wireless cards may not be able to support WPA, and may therefore require replacement. In addition to more secure defaults, commodity router manufacturers must develop firmware according to higher security standards. Similar to software supporting routers in the core of the Internet, more resources must be dedicated to security analysis of the codebase before each release. Finally, developing a mechanism for authenticated automatic updates would make routers less susceptible to attack.

DOI: 10.1002/sec

7. Related Work

Router vulnerability research has traditionally focused on nodes in the backbone of the Internet. For example, a recent investigation uncovered the susceptibility of systems running Cisco IOS to a buffer-overflow attack [42]. Other researchers have demonstrated vulnerabilities in protocols such as BGP [43,44]. Such exploits are attractive to adversaries because they hold the potential to affect large numbers of users per compromised node. In contrast, the exploits for wireless routers are typically designed to affect one or a small number of users. For example, Tsow examines a number of methods compromised routers can use to perform phishing attacks [4]. Others instead target weak passwords and or protocols as a means of simply gaining access to network resources such as bandwidth [2,3,45]. Independently but in parallel to this research, Akritidis *et al*. [46] use similar data to show that infected laptops could use open access points to spread throughout a metropolitan area; however, their work does not consider the potential for malicious routers. After this work was completed, Hu *et al*. [47] investigated a similar problem and focused on the speed of infection propagation. At this time, however, no researchers have explored connectivity or the extent of the attacks made possible only by the widespread compromise of such devices.

Emergent vulnerabilities are those weaknesses made possible only by the amalgamation of seemingly unrelated conditions. The resulting vulnerability is typically unexpected and often unrepresentative of any single component vector. The Slammer worm [14] displayed this quality; while the target of the exploit was machines running SQL Server, resources ranging from 911 services in Bellevue, WA to Bank of America ATMs were rendered inoperable. Traynor *et al*. [48,49] demonstrated another such attack on telecommunications networks, in which the transmission of a small volume of targeted text messages proved capable of denying voice service to large areas. Such vulnerabilities are often difficult to discover because of the subtle interplay between multiple weaknesses necessary for their formation.

8. Conclusion

This paper introduced and explored *malnets*, *ad hoc* netowrks Q4composed of malicious wireless routers. We used statistical methods and analyzed wardriving data to show the potential for city-wide malicious connectivity, demonstrating that cities such as New York and Atlanta were sufficiently dense to support malnets. More generally, we presented techniques to calculate the susceptibility of malnet creation based on router density, and used models to characterize the malnet growth. We showed that worst-case self-propagating malcode can establish far-reaching malnets in less than 2 h.

Traditional threat models for embedded devices such as commodity wireless routers no longer apply. As such devices are increasingly tasked with critical or sensitive operations, they must be protected with the same urgency as personal computers.

References

1. WiGLE.net. Wireless Geographic Logging Engine. http://www.wigle.net/, 2006.
2. Stubblefield A, Ioannidis J, Rubin A. Using the Fluhrer, Mantin, and Shamir Attack to Break WEP. In *NDSS*, 2002.
3. Bittau A, Handley M, Lackey J. The Final Nail in WEP's Coffin. In *Proceedings of IEEE Symposium on Security and Privacy*, 2006.
4. Tsow A. Phishing with consumer electronics—malicious home routers. In *WWW*, 2006.
5. Molkdar D. Review on radio propagation into and within buildings. *Microwaves, Ant. and* Q3*Prop.* 1991; **138**: 61–73. Q3
6. Hashemi H. The indoor radio propagation model. *Proceedings of IEEE* 1993; **81**(7): 941–968.
7. Sridhara V, Kim J, Bohacek S. Performance of urban mesh networks. In *MSWiM*, 2005.
8. Bychkovsky V, *et al*. A measurement study of vehicular internet access using in situ Wi-Fi networks. In *MobiCom*, 2006.
9. Newman MEJ. The structure and function of complex networks. *SIAM Review* 2003; **45**(2): 167–256.
10. Faloutsos M, Faloutsos P, Faloutsos C. On power-law relationships in the internet topology. In *Proceedings of ACM SIGCOMM*, Boston, MA, USA, September 1999.
11. Eschenauer L, Gligor V. A key management scheme for distributed sensor networks. In *CCS*, 2002.
12. Erdös P, Rényi A. On the evolution of random graphs. *Publ. Math. Inst. Hungarian Acad.* Q3*Sci.* 1960; **5**: 17–61. Q3
13. Bailey M, Cooke E, Jahanian F, Watson D, Nazario J. The blaster worm: then and now. *IEEE Security & Privacy* 2005; **3**: 26–31.
14. Moore D, *et al*. Inside the slammer worm. *IEEE Security and Privacy* 2003; **1**(4).
15. Moore D, Shannon C, Brown J. Code-Red: a case study on the spread and victims of an Internet worm. In *IMW*, 2002.
16. Kephart JO, White SR. Directed-graph epidemiological models of computer viruses. In *Proceedings of IEEE Symposium on Security and Privacy*, Oakland, CA, USA, 1991.
17. Staniford S, Moore D, Paxson V, Weaver N. The top speed of flash worms. In *WORM*, 2004.
18. Wang Y, Chakrabarti D, Wang C, Faloutsos C. Epidemic spreading in real networks: an eigenvalue viewpoint. In *SRDS*, 2003.
19. Butler K, McDaniel P. Understanding mutable internet pathogens, or how I learned to stop worrying and love parasitic behavior. In *ICISS*, 2005.
20. Password Cracker's Inc. http://www.pwcrack.com/. Q4
21. Carlini J. Wardrivers now exploiting your wireless service with ease. http://wistechnology.com/article.php?id=889, June 2004.

DOI: 10.1002/sec

22. Haskins W. Router Hack Attack Could Expose Home Network Users. http://www.technewsworld.com/story/55820.html, 2007.
23. MacManus G. Linksys WRT54G Router Remote Administration apply.cgi Buffer Overflow Vulnerability. http://www.idefense.com/intelligence/vulnerabilities/display.php?id=305, 2005.
24. McLaughlin K. D-Link Hardens Firmware For Wireless Bug. http://www.crn.com/showArticle.jhtml?articleID=192201446, 2006.
25. Spafford EH. The internet worm program: an analysis. *Technical Report CSD-TR-823*, November 1988.
26. CERT. CERT Advisor CA-2001-11 sadmind/IIS Worm. http://www.cert.org/advisories/CA-2001-11.html, May 2001.
27. Mullikin G. Test drive: D-Link DWL-922 Wireless G Network Starter Kit. http://mobile.newsforge.com/article.pl?sid=06/01/27/1544241&from=rss, February 2006.
28. The Metasploit Project. http://www.metasploit.com. Accessed [31 January 2007].
29. Default password list. http://www.phenoelit.de/dpl/dpl.html, 2006.
30. Klein DV. "Foiling the cracker"—A survey of, and improvements to, password security. In *Proceedings of 2nd USENIX Workshop on Security*, pp. 5–14, Summer 1990.
31. IEEE Standards Association. IEEE registration authority—IEEE OUI and company_id assignments.http://standards.ieee.org.
32. Franklin J, *et al*. Passive Data Link Layer 802.11 Wireless Device Driver Fingerprinting. In *USENIX Security Symposium*, 2006.
33. Strine J. D-Link DWL-1000AP Wireless LAN Access Point Plaintext Password Vulnerability. http://www.securityfocus.com/bid/3735, 2001.
34. Knienieder T. Netgear WG602 Wireless Access Point Default Backdoor Account Vulnerability. http://www.securityfocus.com/bid/10459, 2004.
35. Linksys. Wireless-G Broadband Router—User Guide. http://www.linksysbycisco.com, 2005.
36. Traynor P, Butler K, Enck W, McDaniel P, Borders K. *malnets*: large-scale malicious networks via compromised wireless access points. *Technical Report* NAS-TR-0048-2006, Penn State University, September 2006.
37. Walker J. Unsafe at any key size; an analysis of the WEP encapsulation, October 2000.
38. Fluhrer S, Mantin I, Shamir A. Weaknesses in the key scheduling algorithm of RC4. In *SAC*, 2001.
39. Tews E, Weinmann R-P, Pyshkin A. Breaking 104 bit WEP in less than 60 seconds. *Technical Report*, Technische Universitiät Darmstadt, 2007.
40. Moskowitz R. Weakness in Passphrase Choice in WPA Interface. http://wifinetnews.com/archives/002452.html, 2003.
41. OpenWRT. http://www.openwrt.org.
42. Lynn M. The Holy Grail: Cisco IOS Shellcode And Exploitation Techniques. http://www.jwdt.com/~paysan/lynn-cisco.pdf, 2005.
43. Kent S, Lynn C, Seo K. Secure Border Gateway Protocol (S-BGP). *J. Sel. Areas Comm.* Q3 2000; **18**(4): 582–592.
44. McDaniel P, Aiello W, Butler K, Ioannidis J. Origin authentication in interdomain routing. *Computer Networks* 2006; **50**(16): 2953–2980.
45. Cam-Winget N, Housley R, Wagner D, Walker J. Security Flaws in 802.11 Data Link Protocols. *Communications of the ACM* 2003; **46**(5).
46. Akritidis P, *et al*. Proximity breeds danger: emerging threats in metro-area wireless networks. In *USENIX Sec.*, 2007.
47. Hu H, Myers S, Colliza V, Vespignani A. WiFi epidemiology: can your neigbor's router make yours sick? *Technical Report*, 2007.
48. Traynor P, Enck W, McDaniel P, La Porta T. Exploiting open functionality in SMS-capable cellular networks. *Journal of Computer Security* 2008; **16**(6): 713–742.
49. Traynor P, Enck W, McDaniel P, La Porta T. Mitigating attacks on open functionality in SMS-capable cellular networks. In *MobiCom*, 2006.

DOI: 10.1002/sec