Towards Secure Localization Using Wireless 'Congruity'

Arunesh Mishra, Shravan Rayanchu, Ashutosh Shukla, Suman Banerjee
University of Wisconsin, Madison.
Email: {arunesh, shravan, shukla, suman}@cs.wisc.edu.

Abstract

Traditional methods for localization in wireless networks rely on the correlation of the received signal strength with physical distance. It is also well known, that these mechanisms fail in an adversarial setting due to the lack of robustness of the signal strength property to malicious intent. In this paper, we present a property of the wireless medium, which we call 'wireless congruity', that captures the *relative similarities* in wireless media characteristics (such as packet receptions, idle channel time, etc.) as observed by two receivers that are in physical proximity of each other. We show that wireless congruity holds promise for secure localization by presenting an initial yet encouraging set of results obtained through extensive experimentation in a rich indoor wireless environment.

I. Introduction

With the growth of m-Commerce applications, systems that provide location information for a mobile user are becoming popular. Such localization systems can be classified into two: the ones that use dedicated hardware for localization, such as Cricket [1] which uses ultrasound, and the ones that operate through off-the-shelf 802.11 hardware. The systems in the latter category are of particular commercial interest due to their ease of deployment over the widely available 802.11 WLANs and the resultant cost-savings. Thus, a lot of prior research has focused on building accurate indoor localization systems [2], [3], [4] which use the signal strength property of wireless transmissions to assist in location inference. While signal strength is a good indicator of physical distance, its predictability has become an Achilles heel when faced with the issue of validating it in the presence of malicious intent – it is possible for an attacker to 'guess' the signal strength at a location without being there physically.

The signal strength property of wireless transmissions suffers from *temporal* and *spatial* predictability. Temporal predictability refers to the possibility of inferring behavior at time T by observing behavior at an earlier (or later) time $T - t$, for example, by building a radio-map of the environment. Next, spatial predictability refers to the ability to predict signal strength properties at location L by observing it at a different location L'. Both properties are, in fact, exploited by localization systems to aid location prediction [5], [2]. Such predictabilities allow an attacker to visit a WLAN installation and collect sufficient information (such as a radio map) to launch two types of attacks: (i) *Against authentication:* An attacker can forge his location by 'guessing' the signal strengths at the location being spoofed and reporting this to the localization system, and (ii) *Against privacy:* An attacker could monitor another user's communication with the localization system and use that information to his advantage in determining that user's location.

Looking beyond signal strength: Philosophically, it is not good security practice to build systems over a property that is known to have weaknesses. While it might be possible to address security through dedicated hardware such as infrared or ultrasound, the space of building accurate yet secure localization systems over commodity 802.11 hardware is challenging and interesting from a research perspective apart from being commercially attractive. Thus, the challenge is – *can we build such a system ?* In this paper, we identify a unique property of the wireless environment, called *Wireless Congruity* which we believe

All authors were partially supported by the following NSF grants: CNS-0520152, CNS-0627102, CNS-0639434, and CNS-0627589.

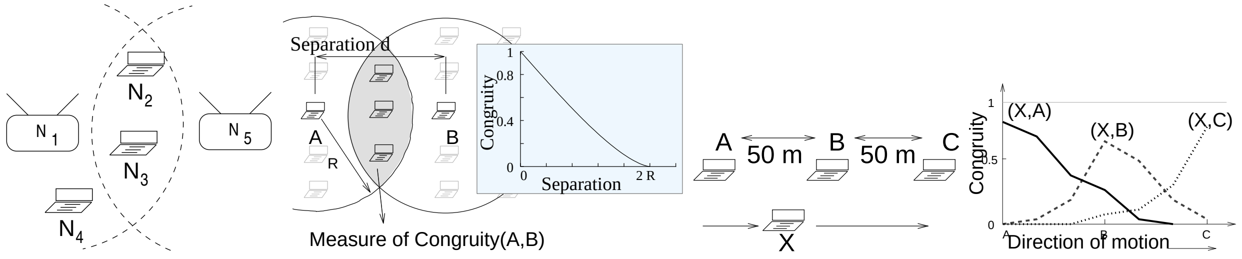

(a) Illustrating congruity. (b) Localization prospects. (c) Lab Experiment. (d) Lab Results.

Fig. 1. An analysis of the congruity property.

might provide the right answer towards our question of building secure localization algorithms. In the next Section, we derive the concept of congruity. Later in Section III we validate its properties through experiments in the lab and in a rich wireless environment. We conclude with key research challenges that lie ahead in Section IV.

II. Wireless Congruity

Motivation and threat model: Suppose a group of people wish to conduct a conference or a meeting at a public place such as hotel or a community center equipped with wireless access. The attendees wish to grant access to confidential material (or just implement network access control and/or encryption) to wireless clients who are within a certain physical premise such as a conference hall. Here, validation of a client's location and maintainig his privacy in this process are both important. A successful attack in this scenario, for example, could be that an attacker convinces the localization system that he is present within the conference hall while being physically present at a different location.

We assume a reasonable and practical threat model for our location validation problem. We restrict ourselves to the case where the attacker has a single wireless interface to send/receive traffic. This is justified for the following reason: consider a case where the attacker uses a network of monitoring nodes who communicate among themselves. Firstly, it would be considered a breach of physical security if an attacker is able to place such nodes around and let them communicate. Also a reasonably good wireless intrusion detection system would be able to monitor extraneous traffic in the system. Finally, if an attacker is indeed able to place such nodes such as monitors and proxies and is able to communicate with them fast enough, its clear that a localization system that does not use dedicated hardware would not be able to thwart such an attack. In fact, its is nearly impossible for a localization system to distinguish between a legitimate user and a proxy device that possesses all the necessary credentials to act as the user.

Although our threat model and thus the proposed problem domain does not tackle sophisticated attacks, we argue that in fact attacks where its hard to detect abnormal behaviour either through survelliance or wireless monitoring are in fact the toughest to defend against. Thus, in this paper, we restrict ourselves to attacks where an adversary uses his/her laptop to passively capture some information over time in an *inconspicuous manner*. And he uses this to either (i) spoof his location to the system, or (ii) determing another legitimate client's location. We realize that augmenting signal strength based systems with strong cryptographic primitives or timing constraints [6] will not provide the answer. This is because even if a wireless transmission from a client that is part of a localization protocol is fully protected (through strong cryptographic mechanisms), the signal strength of that transmission still provides sufficient information to the system (and the attacker): source, destination and strength of transmission[1]. This is sufficient for both the system and the attacker to realize their respective goals.

[1] Note that, conceptually, the source and destination of all wireless packets have to be sent in the clear.

Concept of congruity: We illustrate the concept of congruity through a thought-experiment shown in Figure 1(a). There are five nodes in this wireless environment, possibly belonging to different networks. Nodes N_1 and N_5 are 802.11 access points (APs); the rest are laptop users. Also shown is the transmission range of the AP-nodes N_1 and N_5. Now suppose only one AP-node was present, say N_1. Since the nodes N_{1-4} are all in close vicinity of each other, they would experience similar 'behavior' of the wireless channel. By 'behavior', we mean the following: suppose alongside each node was a passive observer, who made a log of all possible events or observations reported by the wireless card used by that node. What events could such observers log ? To be precise, this would depend on the amount of information that the wireless interfaces export back to the host operating system today, but the following events would commonly available: packet receptions (with or without bit-level errors) for both data, management (beacon messages) and control frames (such as RTS, CTS, ACK) and observations of the medium being idle due to contention related backoffs (start, end and duration of the idle-times) – channel idle-time.

Now if two observers (at different nodes, say N_2 and N_3) were to compare their logs, they would find a large number of similar entries. This would happen because both observers more-or-less have the same *local* wireless environment; that is, they have very similar set of neighbors, contending stations, or sources of interference. Thus, they experience very similar events or behavior of the wireless medium. Now if AP-node N_5 were to join this experiment (with its range as shown in Figure 1(a)), this would increase the entropy of the network and change the local wireless environment for some of the nodes. In particular, nodes N_2 and N_3 would find greater similarities than nodes N_3 and N_4. This is because the transmissions made by N_5 would not reach node N_4 and such events would increase the differences in their respective logs. Through similar reasoning, we can find that out of all node-pairs, nodes N_1 and N_5 will have the least number of similarities because of the differences in their local wireless environment.

Thus, based on the above thought-experiment, we conclude that two nodes that are closer to each other in terms of their position in the overall distribution of the wireless nodes in a given environment (that possibly belong to different networks), will experience increasingly similar behavior as quantified by our passive observer thought-experiment on these nodes. As the nodes get farther apart, they will each have an increasingly different local wireless environment; that is, they will interfere, contend and communicate with a different set of wireless nodes, and thus, they will increasingly differ in their observations. We call this concept of two nodes experiencing similarities in the behavior of the wireless medium as *wireless congruity*. Although the concept of congruity is not previously explored, we note that prior work in [7] uses similar concepts in a very different problem domain.

From our discussion, it follows that two nodes that are in *sufficient* vicinity will have a very similar local wireless environment, and will thus experience good congruity. We further support this conjecture through experiments in a rich indoor library environment consisting of over 150 nodes, in Section III. The next question is in order to achieve good congruity, how close should two nodes be? To answer this, we consider the reverse question: If we somehow 'measure' and determine that two nodes have congruity, to what degree are they in the physical vicinity of each other? The answer depends on the density and the entropy of the wireless network. Going back to our thought-experiment of Figure 1, before AP-node N_5 was added to the network, nodes N_2 and N_4 were observing good congruity. The addition of node N_5 increased the entropy and the density of the network and thus, a higher degree of vicinity was needed for good congruity. Only nodes N_2 and N_3 could experience good congruity, while the congruity between N_2 and N_4 decreased due to insufficient vicinity between them. It can be seen that a dense network will be able to distinguish between small physical separations between wireless nodes when compared to sparse one. We support this further through experiments in Section III.

In order to better understand how congruity relates to spatial vicinity, we study this property in an idealized setting. Assume a large area populated with wireless nodes uniformly spread around, with a per unit area density of ρ. Assuming uniform transmit power and receiver radio characteristics, any node in this region will receive transmissions from another node located at most at a distance R (say). This environment is shown in Figure 1(b). Assume that nodes get roughly equal chances to transmit using the 802.11 distributed coordination function. Consider two nodes A and B. Say we measure the congruity between A and B as the number of packets they receive in common over a certain interval of time. Asymptotically, this value will be proportional to the number of nodes that are in the common region between A and B. This is equal to the area of intersection between the circles multiplied by the node density. The inset plot in Figure 1(b) shows the congruity as a function of the separation between A and B through elementary geometric analysis. The degradation in congruity with distance might appear to be close to linear (inset in Figure 1); this is precisely given by the expression (obtained through simple mathematics) $2R^2 sin^{-1}(d/2R) - d(R^2 - d^2)^{1/2} = O(R)$ which computes the desired area of intersection in Figure 1(b) (d is the separation between A and B). From this analysis, two key properties of wireless congruity follow:

Robustness: If congruity between two wireless nodes A and B is zero, they are then separated by a certain minimum distance, called the *Congruity Distance* D_c. This property can be proved easily through contradiction – if such a minimum distance did not exist then it would be possible for two nodes in maximum possible vicinity of each other to receive a totally different set of transmissions each. Practically this is hard to happen for the reasoning given above and thus the property follows. For the example environment of Figure 1(b), D_c is $O(R)$. Likewise if the congruity is non-zero, A and B are separated by a certain maximum distance D_c. The actual value of D_c is a function of the network topology, radio propagation and such wireless characteristics. We call this the security property for the reason that there exists a *practical and finite* value of D_c for every network environment. The value of D_c places a strict limit on the chances for an adversary to successfully guess the wireless events at a specific location – the adversary cannot predict the events for a wireless network at a location that is distance D_c apart from the adversary's current location.

Accuracy : Accuracy refers to how well can congruity predict physical distances. From the analysis above and as shown in the inset plot of Figure 1(b), we see a near-linear degradation of the metric with distance. In general, the exact nature of this relation would depend on the network topology and other characteristics. Building localization algorithms that take full advantage of congruity and its relation to physical proximity requires further thought and research; we refer to this later in Section IV. The goal of this paper is to discuss the security properties of wireless congruity within the research community.

Design of a congruity metric: A 'metric' function to measure congruity on a fine-grained basis is an essential ingredient for building localization algorithms around it. While designing such a function would require careful analysis and remains as a research challenge for future work (Section IV), for the purposes of this paper we use a simple yet efficient function: we compute congruity between two nodes A and B as$\zeta(A, B) = \frac{N_{AB}}{N_A + N_B - N_{AB}}$. Here, N_A (N_B) is the number of packets received by A(B) during a fixed duration of time. N_{AB} is the number of packets that were commonly received by both A and B during this fixed observation time. For the homogeneous setting shown in Figure 1, we note that the congruity function $\zeta(A, B)$ closely approximates the theoretical estimate shown as the inset plot in Figure 1. We shall use this later in Section III to study congruity in a rich wireless environment.

Design of a secure localization system: A practical system based on congruity can be built in the following manner. The target wireless installation is equipped with a certain set of monitors, or receivers, which constantly receive packets and create a *fingerprint* of the sequence of transmissions received. These monitors are placed at locations which need to be authenticated. One way of creating such a fingerprint would be use a suitable locality-preserving hash function [8]. The fingerprint preserves information on the set of packets received along with their sequence or ordering in a concise form. These fingerprints are sent to a central server periodically using a secure method. Similarly, a wireless nodes or clients compute this fingerprint and communicate them to the central server which computes its congruity with each of the monitors. The monitors that exhibit non-zero congruity with the client give a strong sense of its location. The exact value of the congruity could also be used to further triangulate the client's location coordinates. This is a straightforward and simple design of a system based on congruity, and it might be possible to combine this effectively with existing approaches [9], [2], [3], [4], [10], [11], [12] to get additional benefits (Section IV).

III. An Initial Empirical Evaluation

We present results from two sets of experiments – first in a Lab building (30 nodes) which studies the localization prospects, and second in a rich library environment (150 nodes) which study the robustness prospects for congruity.

Lab experiments: First, we study how physical separation impacts congruity. Figure 1(c) shows our experiment setup where a mobile node X moves through three passive monitors A, B and C. All four nodes make observations on the wireless medium. The experiment was performed in an office building which hosts a 802.11b/g wireless LAN built off 30 APs. No artificial traffic was injected; all measurements were thus based off the everyday network usage traffic on the wireless LAN. Figure 1(d) shows the results. From the plot, it is evident that as the separation to a monitor (such as A) increases the corresponding congruity (denoted by (X, A)) decreases and reaches zero after a certain distance. This distance can act as an estimate for the network's congruity distance metric. This trend is also reflected in the congruity measurements of (X, B) and (X, C). Also by comparing the values of $(X, A), (X, B), (A, B)$ we observed that triangle inequality was satisfied which allows practical systems to use congruity as a distance metric. The gradual degradation in the congruity values also illustrate how well physical separation affects congruity. Thus, by designing better congruity metrics combined with localization algorithms, it might be possible to get good localization accuracy with congruity.

Library environment: We conducted experiments in a rich library environment containing about 150 wireless nodes. About 10% of these nodes were access points (APs), the rest were laptops. During the experiment about 11% of the nodes were mobile, the rest were stationary (including the APs). Figure 2(a) shows the floor-plan and the landmarks where data was collected (shaded circle). As shown, the library consists of three large rooms. We noticed that due to good line of sight, two nodes could communicate with each other as long as they were in the same room. The circles shown acted as the landmarks where we captured wireless traffic. The experiment was performed using the following methodology: Using two laptops at any point of time, we selected two of these landmarks and collected wireless traces. These traces were used to compute the congruity distance between these selected landmarks. By repeating this setup for a comprehensive set of tuples, we covered the entire set of landmarks shown in Figure 2(a).

Figure 2(b) plots the results. The X-axis shows three sets of data – tuples in the same, adjacent and non-adjacent rooms. The Y axis (shown in log-scale) is the congruity metric computed between two landmark points. The results show that the congruity metric exhibits good security properties. The congruity between adjacent rooms is an order of magnitude lower than that of tuples within the same room. Also in the case of non-adjacent rooms the congruity is four orders of magnitude lower. The inset plot of Figure 2(b) magnifies on the values for adjacent versus non-adjacent (a separation of one-room) rooms. This essentially implies that one-in-ten-thousand packets was common between two points separated by a room

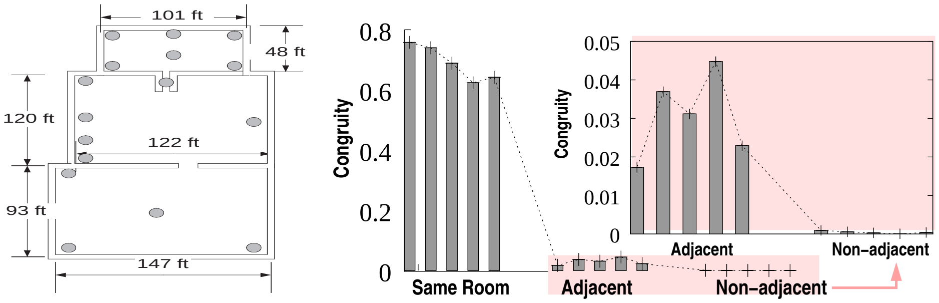

(a) The Library floorplan.

(b) Results – Library.

Fig. 2. Experiments in two operational wireless environments.

in between. Also about 3.6 packets were common in every 1000 packets for two tuples in adjacent rooms while 80 % of the packets were common for tuples in the same room. This gives us a sense of the difficulty that an attacker will have in making an educated guess about the wireless transmissions received at a different location. Using a sensitive receiver here will aggregate significant additional interference from nearby sources and will actually hamper an attackers chances of predicting the wireless transmissions.

The first set of results show promise for localization while the second results show its robustness. These observations provide us with a positive indication that we might be able to converge onto a suitable metric or an algorithm for robust localization. This could be either through congruity, signal strength or maybe an appropriate combination of these metrics.

IV. Research Challenges Ahead

In this paper, we have proposed congruity as a property of the wireless medium that could act as the basis for designing localization systems which exhibit robustness in the face of malicious intent. There are a number of interesting and challenging issues that remain before such a design can be accomplished. Firstly, we need to understand the various factors that affect congruity and to what extent. Factors such as hidden terminals, artificial reduction in node density due to usage of non-overlapping channels, mobility of users, and lack of sufficient traffic could have a negative affect on congruity. Second, we need to build fast methods (with practical convergence times) to estimate the congruity in a dynamic wireless environment. Finally, we need to also evaluate how congruity could be extended to non-802.11 networks and whether we could indeed take advantage of diversity in the underlying wireless technologies.

References

[1] Nissanka Priyantha, Anit Chakraborty, and Hari Balakrishnan. The cricket location-support system. In *Mobicom*, 2000.

[2] Paramvir Bahl and Venkata N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *Infocom*, 2000.

[3] A. Smailagic. Location sensing and privacy in a context aware computing environment, 2001.

[4] Y. Cheng, Y. Chawathe, A. LaMarca, and J. Krumm. Accuracy characterization for metropolitan-scale wi-fi localization. In *MobiSys*, 2005.

[5] Moustafa Youssef and Ashok Agrawala. The horus wlan location determination system. In *ACM Mobisys*, 2005.

[6] Yih-Chun Hu, Adrian Perrig, and David Johnson. Ariadne: A secure on-demand routing protocol for ad hoc networks. *Wireless Networks Journal*, 2005.

[7] Toshihiro Takada, Satoshi Kurihara, Toshio Hirotsu, and Toshiharu Sugawara. Proximity mining: Finding proximity using sensor data history. In *IEEE Workshop on WMCSA*, 2003.

[8] Rajeev Motwani Piotr Indyk and Prabhakar Raghavan. Locality-preserving hashing in multidimensional spaces. In *ACM Symposium on the Theory of Computing*, 1997.

[9] S. Capkun, Mario Cagalj, and M Srivastava. Secure localization with hidden and mobile base stations. In *Infocom*, 2006.

[10] S. Capkun and J.-P. Hubaux. Secure positioning of wireless devices with application to sensor networks. In *Infocom*, 2005.

[11] Donggang Liu, Peng Ning, and Wenliang Du. Attack-resistant location estimation in sensor networks. In *IPSN*, pages 99–106, 2005.

[12] A. Haeberlen, E. Flannery, A. M. Ladd, A. Rudys, D. Wallach, and L. Kavraki. Practical robust localization over large-scale 802.11 wireless networks. In *MobiCom*, 2004.