

Reducing FTM Ranging and Location Attack Exposure With Crowd-Wisdom

Jerome Henry¹, Yann Busnel², Romaric Ludinard³ and Nicolas Montavont⁴

¹Cisco Systems, Research Triangle Park, NC 27560, USA

^{2,3,4}IMT Atlantique, IRISA, Cesson-Sévigné, France

Abstract

802.11 Fine Timing Measurement is an indoor ranging technique. Because it is unauthenticated and unprotected, an adversary can implement ranging and location attacks, by inserting one or more rogue responders and causing an unsuspecting client to incorporate forged values into its location computation. We show in this paper that protection intended for attacks on comparable ranging techniques, like GPS, are ineffective in the case of FTM. However, we also show that a crowd-sourcing technique that confirms that one AP is known by the others can mitigate the attack exposure.

Keywords

802.11az, FTM, indoor location, ranging

1. Introduction

Outdoor location is commonly possible with methods leveraging GPS. Indoor however, GPS signal is often not available and other techniques have been sought for decades. Among several proposed methods, Fine Timing Measurement (FTM) specifies an indoor location procedure based on Time of Flight (ToF). Defined in 802.11-2016 [1] and augmented in the 802.11az amendment (planned for publication in 2022), FTM enables an initiating station (ISTA, typically a mobile Wi-Fi client) to perform ranging exchanges with a responding station (RSTA, typically a Wi-Fi system set at a fixed location, e.g., an access point) and also query the RSTA location. Performing such exchange with multiple RSTAs allows the ISTA to compute its location.

FTM does not include an AP validation mechanism, and an attacker can send invalid range or location information, driving the ISTA off course. Such attack could have dramatic consequences in some environments, for example self-driving shuttles in convention centers and airports that use FTM to assess their position. This paper shows that 1. an attacker can easily drive an FTM station to the destination of the attacker's choice 2. protection techniques intended for similar ranging technologies (GPS) are mostly ineffective, because of FTM fundamental properties and assumptions 3. 802.11 security techniques also provide only limited protection, but 4. a crowd-sourcing technique that augments the Pre-Association Security Negotiation (PASN) process (an unauthenticated 802.11 security procedure) to make it 802.11r, or Fast Transition

IPIN 2021 WiP Proceedings, November 29 – December 2, 2021, Lloret de Mar, Spain


✉ jerhenry@cisco.com (J. Henry); yann.busnel@imt-atlantique.fr (Y. Busnel);

romaric.ludinard@imt-atlantique.fr (R. Ludinard); nicolas.montavont@imt-atlantique.fr (N. Montavont)

🆔 0000-0001-8157-8530 (J. Henry); 0000-0001-6908-719X (Y. Busnel); 0000-0002-4997-4813 (R. Ludinard)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

equivalent, coupled with a better AP sorting algorithm, can dramatically reduce the attack surface.

The rest of this paper is organized as follows: Section 2 exposes how FTM computes range and location and why it is vulnerable to attacks. Section 3 details why current protection techniques, for GPS or 802.11, cannot protect FTM efficiently. Section 4 presents an improvement to 802.11 security that can limit FTM exposures to ranging and location attacks. Section 5 demonstrates how this proposed method reduces exposure in various experimental scenarios, while Section 6 concludes this paper.

2. FTM Vulnerabilities to Ranging and Location Attacks

FTM is particular in the 802.11 world. It is the first 802.11 protocol where a station (STA) exchanges unassociated user-targeted content with an access point (AP). Prior provisions of the Standard allowed an AP and a STA to exchange elements pre, or without association. However, these provisions were either intended to provide general information to the STA about the services expected to be found at or beyond the AP (e.g., clause 11.22.3. in [1]), or were focusing on radio parameters exchanges to allow the STA to associate to the best AP (e.g., clause 4.3.11.1 in [1]). But with FTM, the result of the unassociated exchange is directly leveraged for user-centric purposes (finding one's location indoors). This trend has since appeared with other provisions (e.g., 802.11bc allowing a STA to discover and receive video broadcast without association). 802.11 was not engineered to protect such exchanges, leaving the STA exposed to attacks.

2.1. FTM Principles

Understanding the principles of FTM is necessary to understand the attack vector. The ISTA starts by negotiating with the RSTA some ranging session parameters (intended number of exchanges over a given number of bursts). Then, in each burst, the RSTA first sends a frame at time t_1 , which is received by the ISTA at time t_2 . The ISTA replies at time t_3 with an acknowledgement frame, received by the RSTA at time t_4 . In the subsequent frame, the RSTA communicates the values (t_1, t_4) to the ISTA. The ISTA can then compute the messages times of flight (ToF) $[(t_4 - t_1) - (t_3 - t_2)]$ and thus its distance to the RSTA. In a related exchange, the ISTA can request from the RSTA its Location Configuration Information (LCI), a set of geographical coordinates which logic is similar to that defined in RFC 6225 [2]. The exchange repeats several times. In a typical implementation, the ISTA then retains the smallest ToF, with the reasoning that direct line of sight (LoS) path always produces the shortest ToF.

The 802.11-2016 or 802.11az Standards do not define a method for computing location from these elements. However, it should be clear that the position is computed solely from the ToF and LCI values extracted from the frame exchanges with the local RSTAs.

Several techniques exist for such location computation. Geometric methods, like the simple three spheres method, consider the distances as the radii of matching spheres and find their intersection (when it exists). Another technique organizes the measured distances in a matrix, then attempts to minimize the error between the positions computed from the distance to each RSTA. A third common technique is the use of an extended Kalman filter (or an alpha-beta filter). Shareef and Zhu [3] provide a good introduction to this technique. Kalman filters are

popular in technologies where the subject is expected to move, which is the case with FTM. For this technique like the previous ones, because the measured distances are noisy [4], using more than 5 to 6 responders offers diminishing incentive, as ranging to each additional RSTA increases the energy and ranging (airtime) cost with a decreasing accuracy gain.

2.2. Ranging and Location Attacks on FTM

Thus, these techniques use the distances to 3 to 6 APs. However, 802.11 specifies that a STA can only associate with one AP at a given time. Therefore, even if the STA associates to a first AP, its exchanges with the other APs will be unassociated, leaving the ranging exchanges unprotected.

A simple attack vector is thus to add an AP that returns manipulated (t_1, t_4) or LCI values. The victim's machine then unwittingly integrates these values into its location computation. Because the exchanges are unprotected, the attacker can also spoof the MAC address of a valid AP and add its answers to that of the real AP. Client stations prefer the shorter range as explained above, and thus will retain the attacker's returned (t_1, t_4) values if they are smaller than the real AP's. In an environment where the AP positions are known, and the victim path predictable (e.g., along hallways), ensuring such conditions is trivial.

We studied such effect in [5] with a setting represented in Figure 1, where 5 valid APs are deployed along the victim path. One to three temporal attacker APs are then introduced, sending invalid (t_1, t_4) values or invalid LCI. As the STA starts detecting the invalid AP (labelled AP6 in the figure) and integrating its values into the computation mix, the victim's STA calculated location deviates from its true position.

When knowing the type of algorithm that the victim would use (e.g., because the victim location app is known), the attacker can insert targeted false values selectively and progressively, pushing the victim away from its real trajectory and toward a destination of the attacker's choice. This possibility is illustrated in the upper part of Figure 1, where the attacker drives the victim toward a target 13 meters away from the intended path (after a 38-meter walk) and represented by a cross (envisioning location determination through the three-sphere method [top left], a least square technique [top center] and a Kalman filter [top right]).

Such attack may not be of consequence if the victim is a human in a casual setting, who also uses other input for trajectory determination (e.g., visual inspection of the environment in search of a particular object, like a store in a shopping mall), although a human victim may not carefully validate the trajectory provided by a navigation device and still be driven off course. However, the consequences can be more severe if the victim is a robot guided through FTM to operate a particular task (e.g., self-driving delivery robots in hotels or convention centers, self-driving shuttles in indoor amusement parks, airports and other public venues, etc.), without a multiplicity of other validation techniques. In all these cases, an attack on FTM used as the primary navigation source can cause the robots to bring about damage, hit equipment, hurt people, or lose certainty about its location and paralyze operations. In a world with aggressive business competition and hostility between state actors, such an attack is an easy and cheap vector to conduct.

3. Insufficiency of Existing Protection Solutions

FTM is not the only location solution relying on distances, and not the only 802.11 technology where data validation is required, so one could think that the problem is well-known and possibly solved. However, FTM is a young protocol and the literature has not yet studied the issues that this paper exposes.

3.1. Related Work

Although there are multiple methods to determine the location of an RF object indoor [6], the most scrutinized location technique based on ToF is GPS. Just like FTM, GPS uses the estimated distance to multiple satellites, which positions are known, to determine station location. Similar to FTM, the GPS signal available for civilian usage is neither encrypted nor authenticated. The same limitation is currently true for the other systems, Galileo, GLONASS, Michibiki, BeiDou and NavIC, although protection schemes have been proposed for the BeiDou protocol[7].

There are many documented attacks techniques against GPS that are similar in concept to the attacks against FTM that this article studies. These attacks spoof the identity of one or more valid satellites [8], and provide signals which timing and timestamps cause the station to miscalculate its distance to the source [9] and compute an incorrect location [10].

In some cases, these attacks can be detected. A simple technique is to use Receiver Autonomous Integrity Monitoring (RAIM) [11]. With this technique, the client uses a rolling subset of satellites, and alerts if any particular satellite contribution provides results inconsistent with the others.

Additional techniques look at each individual signal. For example, Foruhandeh *et al.* [12], fingerprint each satellite signal, and recognize the impersonation by matching the received signal against the expected signature. Other techniques use fusion to compare the GPS-computed location with estimations from other sources, like cellular towers [13] or accelerometers and other sensors internal to the station [14]. More complex, multi-factor detection methods naturally augment the reliability of the detection [15], for example when neural networks combine all contributing elements together to determine if one of them is surfaced as an outlier [16].

All these techniques mitigate the spoofing attack, but do not completely remove it. A large trend is to apply to the civil GPS the same authentication and encryption scheme as military grade GPS [17], thus disabling the spoofing risk, and complement security with resilience features to disable the jamming risk [18].

3.2. Applicability Limitation of GPS-attack Solutions

Unfortunately, GPS-attack protections that apply fusion techniques [5] prove often insufficient for FTM. These techniques use a primary source, and one or more secondary (less accurate or reliable) sources. The values coming from the secondary sources are used to spot outliers in the primary source, but then the primary source data is used to compute location. In other words, the secondary sources augment, but do not entirely replace, the primary source. GPS is commonly not available inside buildings, and cell tower triangulation is also usually not possible.

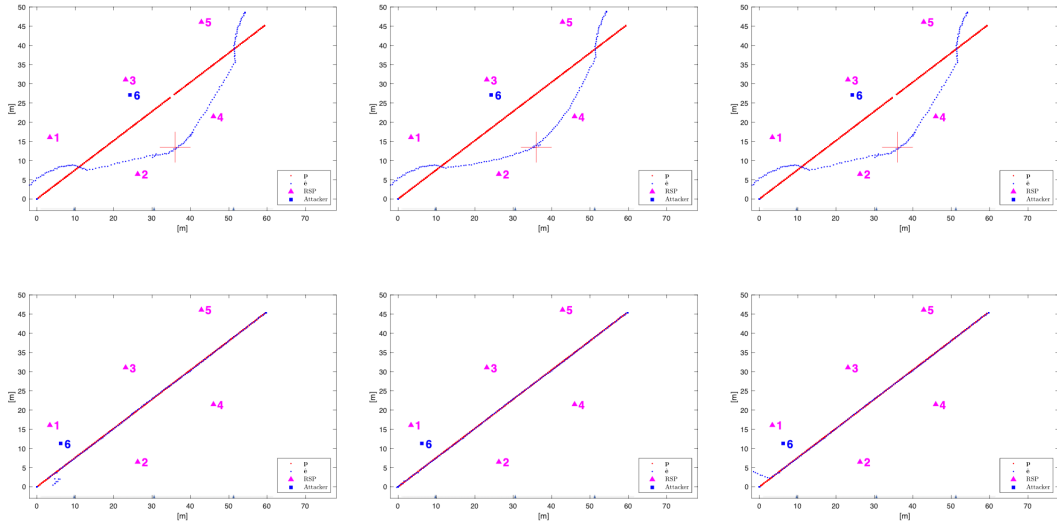


Figure 1: Top: Injection of misleading LCI with the three-sphere method (left), the matrix method (center) and a progressive injection for the Kalman filter method (right), leading target toward the location marked with a cross. Bottom: the same attack attempts against PASN-FT and the improved AP selection method proposed in this paper. The victim real position is marked p , the forged path e , real AP responders with a triangle and the attacker APs with a square.

The STA can collect information from movement sensors (gyroscopes etc.), and the combined information from these sources (Pedestrian Dead Reckoning, PDR) to compute the movement of the object containing the sensor. PRD is a very promising direction for indoor positioning, and provides good results in constrained scenarios *e.g.*, for example when the sensors are attached to the user’s foot [19]. In the cases addressed by FTM however, the trajectory can change rapidly, causing the accuracy of the values returned by the sensors to dissolve with distance [20]. For example, a user may have parked in the underground garage of a shopping mall, have a smartphone in their hand, then their pocket, then hands again, while the user undergoes a complex trajectory (stopping, moving aside to navigate through a crowd with constantly changing speed, slowing down and turning near store windows etc.), each change introducing errors into the positioning algorithm [21]. Without an additional reliable source of reference, the STA cannot compensate for the continuous injection of invalid parameters. Wi-Fi commonly cannot be the primary or the secondary source reliably. This difficulty is caused by the fact that Wi-Fi ToF values are usually noisy. With obstacles and multipath, there is no good correlation between the AP signal strength and the calculated distance, and the precision of measurement is low. Therefore, fusion is not a sustainable solution for FTM.

Other GPS attack protection techniques that rely on signature and fingerprinting are also not transferable. With billions of Wi-Fi APs on the planet, and the assumption that a core use case for FTM is a human with a handheld device (*e.g.*, a smartphone) in an unfamiliar venue, the designers of FTM have not implemented any validation technique, and there is no global database that the ISTA could use to verify the identity of each AP in range.

RAIM-like mechanisms are not widely implemented yet in FTM-based app. However, even when they are, we have shown [5] that the progressive insertion of invalid parameters could defeat the protection.

Therefore, although the attack exposure of FTM is similar to that of GPS attacks, the partic-

ularities of FTM makes that the solutions for GPS-attack mitigation cannot be transparently transposed to the FTM case.

3.3. Limitations of IEEE 802.11 Solutions

Another avenue to explore could be to leverage 802.11 existing solutions. For example, 802.11 association can be followed with a mutual authentication mechanism, where both the STA and the infrastructure prove that they know a shared secret (which is a proof that the AP is legitimate). It could be suggested to use this mechanism to protect the FTM exchanges. However, besides the difficulty that each ISTA would need to be configured with credentials for the local SSID, the entire procedure commonly takes 300 milliseconds (and can take more than one second when the 802.1X/EAP part relies on an external RADIUS server). Once association is complete, the STA can leave the AP cell and re-join with a re-association process, which delay is much shorter (80 to 100 ms).

This overhead can cause several problems. A first immediate issue is time consumption. In the FTM negotiation, the ISTA and the RSTA agree to meet on the channel at pre-determined points in time to exchange FTM frames in a burst. The designers understood that either side may be delayed. For this reason, the burst duration can last up to 128 ms. During that interval, each side will attempt to be on the channel and perform ranging. The ISTA will listen on the channel, waiting for the RSTA first FTM frame of the burst. The RSTA will send the first FTM frame, and wait for the first acknowledgement response from the ISTA (then retries repeatedly until the end of the burst when no acknowledgment is received). The ISTA and the RSTA also agree on an inter-burst interval, but if during that interval the ISTA goes to another channel and does not know if it will spend there 80 ms, 300 ms or more than a second before coming back, failed bursts are guaranteed to be a common occurrence. This would not be a reasonable design.

Another issue is the processing cost on the RSTA. The ISTA and RSTA negotiate a number of bursts (between 1 and 16384). If the ISTA reassociates for each new burst, then the AP potentially has to manage 16384 reassociation procedures per client and FTM session, each with a computing cost related to key verification (not to mention 80 ms inserted delay each time). A few stations performing rapid-fire ranging with multiple APs are susceptible to exhaust the APs computing resources. For these reasons, relying on 802.11 authentication has never been considered as a viable solution by the 802.11 designers of FTM.

These limitations do not mean that the ISTA cannot associate and perform FTM. The ISTA could associate with one AP, perform FTM with that AP, then go to other channels and perform unassociated and unprotected exchanges with the other APs. Quite obviously, this mode would offer very limited protection.

In an attempt to provide better exchange protection while limiting the overhead, 802.11az, the 802.11 Amendment for Enhancements for Positioning that expands FTM, defines a Pre-Association Security Negotiation (PASN), allowing an ISTA to establish a secure session with an AP/RSTA without association. This mode is efficient because it only relies on a two-frame exchange (plus one acknowledgement) with each AP. PASN exists in two modes. One mode supposes the existence of pre-existing shared keying material between the ISTA and the RSTA. In that mode, the ISTA and the RSTA then form what 802.11 calls a Robust Security Network

Association (RSNA) authentication, in that they create a secure link and authenticate each other. This mode is efficient to protect from spoofing, as an attacker AP would not have the valid keying material, and would therefore not be able to insert frames in the exchange between an ISTA and a valid RSTA.

One clear limitation of this mode is the requirement to pre-populate the ISTA with the keying material. This is easily possible in a private setting (*e.g.*, a factory with a single WLAN system). However, in a public venue, this constraint becomes difficult to overcome, even for a device that is expected to be familiar with the venue (*e.g.*, a self-driving shuttle or delivery robot in a shopping mall). It is likely that there will be many APs and SSIDs (for example, an American shopping mall commonly includes 4 or 5 anchor large stores, and several dozens of small shops in between). APs may communicate over the back-end, but it is unlikely that the ISTA would have credentials for all stores and SSIDs. Even in a private setting, it is common to find zones with different SSIDs (for example department names in a hospital), forcing the implementers to configure multiple profiles on the device, even if the APs communicate with the same backend infrastructure (and thus even if the device credentials would be the same for all SSIDs).

Thus RSNA PASN provides an interesting, but incomplete direction for FTM protection. PASN also exists in a non-RSNA mode, where no initial keying material exists between the STA and the APs. In that case, a protected, but unauthenticated link is established ad-hoc between the STA and each AP, creating what some call "trust at first sight". The frames that the STA and the AP exchange (including FTM) are then encrypted. In this scenario, (t_1, t_4) and possibly LCI values can be protected from eavesdroppers. A similar unauthenticated but encrypted link is formed between the ISTA and each and any other AP with which the ISTA needs to range. This mode simplifies the deployment (no pre-existing keying material required), but does not include any validation structure. Thus, an attacker can pretend to be an additional AP in the system, or can impersonate (spoof) a valid AP. The ISTA would then establish an unauthenticated but encrypted connection to that rogue AP, with no means to know that the AP is invalid and send forged (t_1, t_4) and LCI values. The only protections that this mode offers are exchange obfuscation (an observer cannot see the values exchanged between the ISTA and RSTA) and session hijacking protection (an attacker can spoof the identity of a valid AP at the time of a new FTM session establishment, but cannot insert rogue values once a session has started with a valid AP).

4. A Crowd-Wisdom FTM Attack Exposure Mitigation Solution

Clearly, RSNA PASN is a good option for environments with a single SSID and a device that is familiar with the venue (and thus can be programmed with keying material). However, in environments with a single WLAN infrastructure but multiple SSIDs, or multiple neighboring WLAN infrastructures, or environments with which the device is not familiar (thus does not possess any keying material), RSNA PASN shows its limitation, and non-RSNA PASN offers no real protection.

For these environments, it is difficult to establish trust when no APs are known. However, we propose a mechanism to reduce the likelihood that a given contributing AP would be an attacker's. In a setting where there is no authentication, and no protected exchange, all

exchanges are open to all abuses, and we do not believe that an easy remediation is possible. However, if an initial protection can be built between an ISTA and a first AP, we believe that partial further protection becomes possible. We write “partial” because without authentication, the protection stays limited, but it can be implemented in a way that makes the attack less trivial, and thus less attractive to an amateur attacker. Therefore, we next examine partial remediation options for cases where the ISTA is not associated to a network but can establish a protected link to APs.

4.1. Modified FTM AP Sorting Algorithm

One key assumption for this method is that the venue presents several APs announcing the same SSID. These APs communicate with each other (either in a mesh fashion, or because they are connected to the same management system, *e.g.*, a WLAN controller). Other APs may be present in the same RF space, forming individual islands of Wi-Fi coverage around individual or shared SSIDs. In this environment, it is unlikely that the attacker can become the dominant system, *i.e.*, deploy with impunity more APs than the valid network, without causing multiple rogue alarms on the main system and being detected. Most managed Wi-Fi systems can identify AP MAC address spoofing or SSID impersonation. An attacker can deploy one or a few temporal (physical or virtual) RSTAs for the duration of the attack, but is unlikely to have deployed a system larger than the venue legitimate WLAN infrastructure.

When initiating an FTM session, the ISTA needs to first establish a list of possible RSTAs. The ISTA starts by scanning all channels (standard 802.11 discovery) and establishes the list of channels and AP MAC addresses (Basic Service Set identifiers, BSSIDs) offering RSTA services. In traditional FTM, the ISTA would then directly start ranging against the first AP (*i.e.*, the AP with the strongest signal, or the AP on the lowest channel in the band). We now propose instead that the ISTA sorts the BSSIDs by signal level, then considers BSSIDs within the same signal level (RSSI) range (*e.g.*, within 6 dB of each other) as being a single system (for a reason that will become apparent later). We then propose that the ISTA operates a second sort, grouping the BSSIDs by their announced SSID. Thus, we propose that the ISTA starts by ranging against the RSTAs representing the largest system (max BSSID count for the given SSID). This first step is precautionary and does not offer any strong protection, but decreases the likelihood of ranging against a lone attacker, as will be seen.

We then suggest that the ISTA selects one BSSID at random within the largest SSID group, and first establishes a protected link to the RSTA, using PASN (likely, the non-RSNA PASN flavor). The goal of such link is to protect the exchange from eavesdropping. Although the largest group is likely to be a valid system, the attacker may impersonate one AP of the valid system, and thus it is possible that the first AP may belong to the attacker.

4.2. PASN and 802.11r

We then propose an augmentation of PASN, that we call PASN-FT, to allow secure link pre-establishment to other APs. FT, or 802.11 Basic Service Set (BSS) Fast Transition (FT), was defined in the 802.11r-2008 amendment (integrated in the 2012 version of the 802.11 Standard), and is intended for RSNA, fast roaming key exchanges for associated STAs. In this mode, a

STA first establishes a secure association (RSNA) with an AP. During that process, the 802.11 choreography allows for AP and STA mutual validation (as they both have to prove to the other that they have the right temporal keying material). The 802.11-FT process incorporates two major changes to the previous 802.11 association process:

- The AP advertises a Mobility Domain Element (MDE), which is a string representing the domain, *i.e.*, the set of APs between which fast transition will be possible. The string is commonly an arbitrary set of characters (it does not need to have a meaning, and just needs to be common between APs participating to the same group). When roaming, the STA selects APs that advertise the same MDE.
- 802.11r establishes a new key hierarchy. Upon a STA first association, the WLAN infrastructure establishes a first Pairwise Master Key (PMK-R0). This key is derived from the Master Session Key (MSK), which is formed on the client side and the infrastructure side through the regular authentication process defined for 802.11. In a non-FT mode, the PSK is directly derived from the MSK (the PSK is the first 256 bits of the MSK). With FT, the PMK-R0 is derived by also integrating other elements, such as the value of the domain in the MDE, the SSID name, the STA MAC address and the identifier of the first entity with which the client establishes this first keying material (this can be the first AP MAC address, or a value for a centralized WLAN controller; this entity is later identified as the holder of the PMK-R0, or R0KH-ID). Then, for each AP, a PMK-R1 is established, built from the PMK-R0 value, the MAC address of the client and the MAC address of the target AP. When a STA needs to establish a communication with a first AP, it is provided the elements it needs to compute PMK-R0 from the MSK (that the client should be able to derive during the authentication phase, from its credentials or a pre-shared key). The client can then compute the PMK-R1 to associate with any AP in the domain, if the AP MAC address is known. From the PMK-R1, other keys are derived (temporal unicast keys).

This key hierarchy allows BSS-FT to enable a fast transition mode. When a STA needs to roam to a neighboring AP, a non-FT mode would mean that the STA should deassociate from the current AP, associate to the next AP, then undergo the full authentication exchange in order to derive a new PMK. This process can take a long time, as detailed in Section 3. With FT, the keying material required for association to the next AP can be derived while the STA is still associated to the first AP. BSS-FT allows two modes for that process:

- With the Over-the-air mode (OTA), the STA sends to the next AP an FT authentication request, that includes the PMK-R0 Name, the MDE, and a Fast Transition Information Element that includes the R0KH-ID. These elements allow the next AP to determine if it can build a PMK-R1. If the answer is positive, the next AP responds with an FT authentication response, that includes the elements the STA needs to derive the PMK-R1 value for the next AP.
- With the Over the Distribution System method (Over-the-DS), the STA first identifies the target next AP, then sends to its current AP an action frame requesting the establishment

of keying material with the target AP. The request also includes the PMK-R0 Name, the MDE value and the R0KH-ID. The current AP should relay this request over the wire to the target AP. Similar to above, the target AP should determine if it can build the PMK-R1, and reply through the current AP (over the wire) if the answer is positive, with a frame containing the elements the STA needs to derive the PMK-R1 value for the next AP.

In both cases, the STA is then ready to communicate securely with the next AP. At any time, the STA can deassociate from the current AP, and send a reassociation request to the next AP, mentioning the PMK-R1 Name, the MDE, the R0KH-ID and the MAC of the target AP. The STA also mentions a message integrity check (MIC) that proves that the STA has the right keying material. The next AP replies in kind, and data communication can resume immediately.

4.3. PASN FT

802.11 BSS-FT was intended for associated STAs. However, we propose to adapt FT principles to the PASN case by adding to PASN exchanges the FT elements that do not strictly force the STA to undergo an association.

Thus, in this method, all APs part of the same infrastructure include the MDE in their probe responses and beacons. This informs the STA about which APs are claiming to be part of the same domain.

In PASN, the STA seeking to establish a secure link sends a first PASN 802.11 authentication frame to the first AP. The frame may include base Authentication and Key Management (AKM) parameters (if there is a pre-existing keying material that the STA can use), but also includes an ephemeral public key that the STA wishes to use (the STA also generates internally the matching private key). We add the MDE value to this frame.

In PASN, the AP responds with a second PASN authentication frame that includes the AP temporal public key (the AP also generates internally the matching private key), and optionally base AKM parameters (if the STA included them). We also add the MDE value to this frame.

In PASN, the STA then responds with a third PASN 802.11 authentication frame, that serves as an acknowledgment to the exchange. We reuse this frame unchanged.

At this point, the STA has established keying material with the AP and can undergo protected exchanges. With PASN, the secure link is established with a single AP at a time (the AP with which the STA wishes to communicate). We augment this procedure by allowing an Over-the-DS key pre-establishment with other APs. An OTA mode would also be possible, but presents limited added value. From an airtime consumption standpoint, the amount of frames to exchange would be equivalent to a direct PASN exchange with the next AP (thus bringing no airtime consumption, and no process time consumption, advantage). From a security standpoint, both the Over-the-DS and OTA modes prove that the first AP has a trusted backend relationship with the second AP, thus that they belong to the same infrastructure. The OTA mode thus does not surface specific advantages.

The Over-the-DS PASN FT is illustrated in Figure 2, and works as follows:

- When in need to communicate securely with a second AP, the STA sends to the current AP a PASN FT authentication frame wrapped in a protected (robust) action frame. The frame

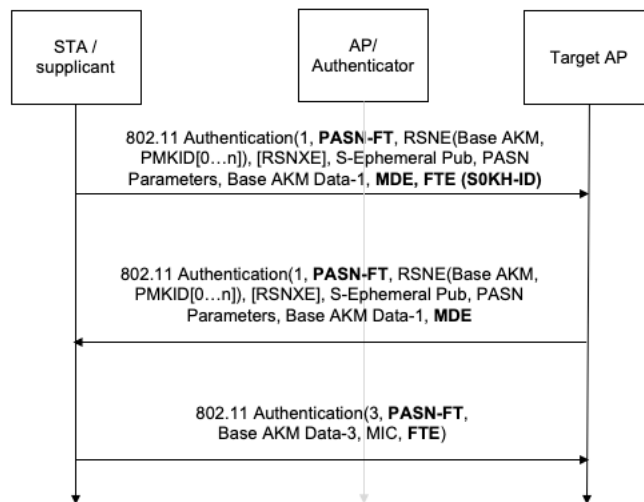


Figure 2: Proposed PASN-FT choreography. The elements that diverge from strict PASN are labelled in boldface.

resembles the PASN first frame, but also includes the MDE, and a Fast Transition Element (FTE) that includes the target AP BSSID (MAC address), and also the client intended MAC address (called S0KH-ID) for exchanging with the next AP. This last element offers an interesting additional protection. The STA can decide to use a different MAC address for its dialog with the next AP (using locally administered MAC addresses). As the frame is encrypted with the current AP public key, this value is obfuscated from an eavesdropper's view. This way, the infrastructure can keep track of the client queries while observers do not see a single STA.

- The current AP forwards this frame over the backend to the next AP. The next AP responds over the DS with a frame that resembles the PASN second frame, but also includes the MDE, the second AP MAC address, the target STA current MAC address (used by the STA to send the first PASN FT frame to the current AP) and the S0KH-ID. The frame also includes the FTE that includes a timeout value. This value tells the STA the interval for which the current keying material will be valid.
- The current AP relays the frame to the STA. The STA validates the frame components, and returns an acknowledgement frame. The frame resembles the PASN third authentication frame, but also includes the FTE that mentions the next AP MAC address, and the client intended MAC address (S0KH-ID).

At this point, the STA is ready to communicate securely with the next AP. Within the timeout interval specified by the next AP, the STA can switch to the next AP channel and directly send protected data (using as a source MAC the S0KH-ID value). The STA can this way pre-establish secure links with a multiplicity of APs, then switch to their respective channel in turn to proceed to protected FTM exchanges.

This process does not guarantee that no attacker will insert in the exchange. However, this methods greatly limits the risks, as four permutations and scenarios are possible:

Current AP \ Next AP	Legitimate	Attacker
	Legitimate	(1)
Attacker	(3)	(4)

1. The APs can communicate over the backend and the STA can successfully establish a secure connection to the next AP.
2. The legitimate AP does not have a trusted relationship to the attacker AP, and rejects the PASN-FT request from the STA.
3. The attacker AP may attempt to pass the PASN FT request to the legitimate AP, but they do not have a backend trusted relationship. The attempt fails and the STA does not receive a response from the next AP.
4. The APs may communicate over the backend, and the STA may be able to successfully establish a secure connection with the next AP.

It should be noted that, because the STA merges RSTAs with the same RSSI as indicated above, scenario (4) only succeeds if the attacker has positioned 2 different physical APs. Scenarios (2) and (3) do not directly allow the STA to determine that one AP is illegitimate. From the STA viewpoint, the PASN FT process failed, possibly because one of the APs is illegitimate, or because both APs are legitimate but in disjoint systems (a less common, but possible case).

However, as the process repeats with more APs, the STA surfaces groups of APs that have a backend trusted relationship, and outliers APs (AP_O) that are not trusted by others (because requests made to an AP_O to PASN FT toward other APs will usually fail, and requests made to other APs to PASN FT toward an AP_O will also usually fail, unless the other AP is also an AP_O). The STA can then use these groups of largest APs having a trusted relationship as the set of RSTA from which location is computed. These are likely to be legitimate, unless the attacker is the dominant system in the venue.

5. Experimental Validation

We tested this method in two different environments. The first setting is similar to the attack test setup described earlier (5 legitimate APs, and the attacker emulating one to three APs), and leads to the following observations:

1. The attack fails in 100% of cases where the attacker presents a single AP. This is likely because that AP cannot form a group large enough to be usable (and of course the AP also fails to establish PASN-FT with the other APs).

2. When the attacker emulates three APs, the attack fails if the APs are all emulated from the same physical system (e.g. virtual APs on the same laptop, or physical APs at the same location). Despite RF signal stochasticity, all APs then present an RSSI in the same range and get merged by the filtering procedure (thus leading the STA to the same conclusion as above).
3. When the attacker deploys 3 non-co-located APs, the attack fails for any location method using 4 APs or more (e.g., least square or Kalman filter with 4 or more RSTAs, Figure 1 bottom center). This outcome is expected. For location method using 3 APs, the attacker system becomes self-sufficient and may partially succeed. The STA temporarily follows the attacker's data, then suddenly jumps back to the correct trajectory as soon as contributors from the valid system are introduced (Figure 1 bottom left). When using Kalman filtering, the slide toward the correct path is progressive, as the system arbitrates between the observed and the computed values (Figure 1 bottom right).
4. The above attack succeeds only if the attacker system is within the first 3 APs to be attempted by the STA, and if the attacker system is large enough to be entirely sufficient for the STA calculations (i.e., 3 APs for the 3-sphere methods, and up to 6 APs for other methods). The STA then forms 2 groups of non-compatible sets (the legitimate APs, and the attacker's APs) and can then randomly consider the location by the valid AP set, or the attacker's AP set.

The second setting is a shopping mall with 4 major store anchors and a multiplicity of smaller stores, each with an individual SSID. The mall is set on two floors, and Figure 3 represents the ground floor. A user walks along the main corridor path, from bottom to top (represented by a dashed line).

In this environment, without changes, the attack fails in all attempts. One likely reason is that, at any point of the path, 9 to 11 APs can be heard that form groups larger than 4 APs, from a combination of the main mall Wi-Fi and one of the anchors'. In such environment, the attacker is unable to establish a system large (and distributed) enough to compete with one of the valid groups.

Forcing the ISTA to ignore the main mall Wi-Fi (supposing a targeted denial of service) does not allow for an attack vector either, because at least one of the anchors' Wi-Fi can be detected from any location along the path (commonly with 4 APs or more). In some areas (mid-point on the walking path), only one anchor's SSID is detectable, and its APs are all in the same direction (to the right). Location precision dilution occurs in this zone if only the anchor's SSID is used. Adding the smaller stores RSTAs restores the precision. Even when an individual store does not allow for PASN-FT with others, the anchor's RSTAs serve as a reference and the attacker is also identified as an outlier (when injecting forged LCI or (t_1, t_4) values causing an error larger than that resulting from computing location with valid RSTAs).

Forcing the ISTA to also ignore the anchors' Wi-Fi systems provides only a mild attack vector. Although each store display a specific SSID, in effect, many of the smaller stores use APs managed by the mall (and therefore would be displaying the same MDE value), even as most stores complement the system with home or small-business grade APs. These other APs may not communicate with one another over the wire. It is only when removing all small store APs

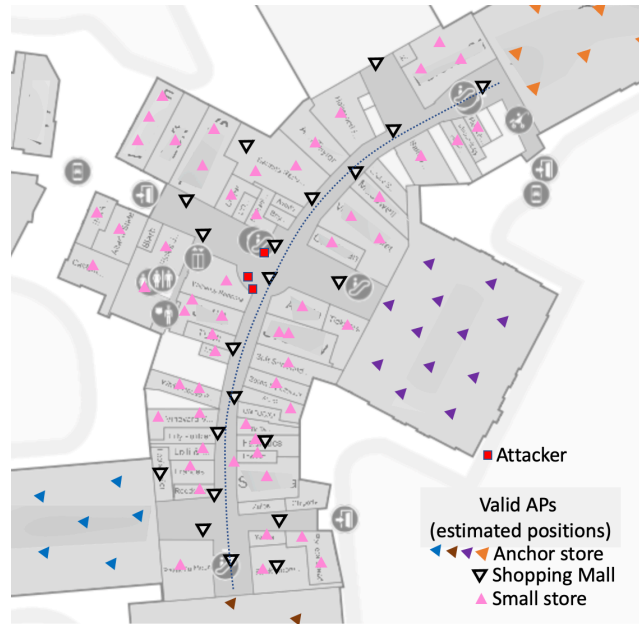


Figure 3: Attack attempts with PASN-FT.

that communicate over the wire that the attack succeeds. However, such a dramatic scenario (all major stores, and the entire shopping mall main WLANs are disabled, only isolated single-AP SSIDs remain) is unlikely in a real environment.

6. Conclusion and Future Work

In this paper, we have shown that 802.11 FTM is vulnerable to ranging and location attacks. As the client does not know in advance the APs, and as the exchanges are neither authenticated nor protected, an attacker can easily insert an additional AP that provides invalid ranging or position (LCI) information. The client has limited ability to distinguish the attacker's from valid APs, and tends to integrate the data provided by the attacker, if it is not excessively implausible, into its location computation.

The existing IEEE 802.11 Standard, along with a modification of the PASN protocol to allow for fast transition between APs, can be used to mitigate the attack exposure. The effect of these simple changes is to turn the attack from “trivial” to “challenging”. There is still a weak vector left, where the attacker is one of the first APs picked by the ISTA and has deployed a system as large as the main WLAN. Future work will examine how deeper protection can be achieved, for example by moving the LCI information into a trusted database that the ISTA could interrogate (*e.g.*, while retrieving the venue map), instead of relying of over-the-air interrogation of APs that the ISTA cannot completely trust.

References

- [1] Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification, Technical Report 802.11, Piscataway, NJ, 2016.
- [2] Y. Polk, M. Linser, M. Thomson, B. Aboba, Dynamic Host Configuration Protocol Options for Coordinate-Based Location Configuration Information, RFC 6225, RFC Editor, 2011. URL: <https://www.rfc-editor.org/rfc/rfc6225.txt>.
- [3] A. Shareef, Y. Zhu, Localization using extended kalman filters in wireless sensor networks, 2009. URL: https://digitalcommons.library.umaine.edu/gradstudent_pub/5.
- [4] B. K. Horn, Doubling the accuracy of indoor positioning: Frequency diversity., *Sensors (Basel)* (2020).
- [5] J. Henry, Y. Busnel, R. Ludinard, N. Montavont, Ranging and Location attacks on 802.11 FTM, 2021. URL: <https://hal-imt-atlantique.archives-ouvertes.fr/hal-03241630>, research report.
- [6] A. Yassin, Y. Nasser, M. Awad, A. Al-Dubai, R. Liu, C. Yuen, R. Raulefs, E. Aboutanios, Recent advances in indoor localization: A survey on theoretical approaches and applications, *IEEE Communications Surveys Tutorials* 19 (2017) 1327–1346.
- [7] R. Hu, L. Ju, P. Chen, A security transmission system for beidou short message based on SM9, *Journal of Physics: Conference Series* 1345 (2019) 022014. URL: <https://doi.org/10.1088/1742-6596/1345/2/022014>. doi:10.1088/1742-6596/1345/2/022014.
- [8] J. Su, J. He, P. Cheng, J. Chen, A Stealthy GPS Spoofing Strategy for Manipulating the Trajectory of an Unmanned Aerial Vehicle, *IFAC-PapersOnLine* 49 (2016) 291 – 296. URL: <https://doi.org/10.1016/j.ifacol.2016.10.412>.
- [9] ACM (Ed.), A Practical GPS Location Spoofing Attack in Road Navigation Scenario, 2017.
- [10] E. Horton, P. Ranganathan, Development of a GPS spoofing apparatus to attack a DJI Matrice 100 Quadcopter, *The Journal of Global Positioning Systems* 16 (2018) 1446–1464. URL: <https://doi.org/10.1186/s41445-018-0018-3>.
- [11] S. Hewiston, J. Wang, Gnss receiver autonomous integrity monitoring (raim) performance analysis, *GPS Solutions* 10 (2006) 155–170.
- [12] M. Foruhandeh, A. Z. Mohammed, G. Kildow, R. Gerdes, Spotr: GPS Spoofing Detection via Device Fingerprinting, in: *Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks (WiSec'20)*, Linz (Virtual Event), Austria, 2020. arXiv:2005.087875.
- [13] G. Oligeri, S. Sciancalepore, O. A. Ibrahim, R. Di Pietro, Drive me not: Gps spoofing detection via cellular network: (architectures, models, and experiments), in: *Proceedings of the 12th Conference on Security and Privacy in Wireless and Mobile Networks, WiSec '19*, Association for Computing Machinery, New York, NY, USA, 2019, p. 12–22. URL: <https://doi.org/10.1145/3317549.3319719>. doi:10.1145/3317549.3319719.
- [14] J.-H. Lee, K.-C. Kwon, D.-S. An, D.-S. Shim, GPS spoofing detection using accelerometers and performance analysis with probability of detection, *International Journal of Control, Automation and Systems* 13 (2015) 951 – 959. URL: <https://doi.org/10.1007/s12555-014-0347-2>.
- [15] K. Hu, Y. Huang, A composite detection method for direct GPS deception attack, *IOP Conference Series: Materials Science and Engineering* 790 (2020) 012028. URL: <https://doi.org/10.1088/1757-899X/790/1/012028>.

//doi.org/10.1088/1757-899x/790/1/012028. doi:10.1088/1757-899x/790/1/012028.

- [16] E. Shafiee, M. R. Mosavi, M. Moazedi, Detection of spoofing attack using machine learning based on multi-layer neural network in single-frequency gps receivers, *Journal of Navigation* 71 (2018) 169–188. doi:10.1017/S0373463317000558.
- [17] K. Wesson, M. Rothlisberger, T. Humphreys, Practical cryptographic civil gps signal authentication, *NAVIGATION* 59 (2012) 177–193. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/navi.14>. doi:<https://doi.org/10.1002/navi.14>. arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/navi.14>.
- [18] A. D. Molina-Markham, Probabilistic models for assured position, navigation, and timing, in: M. C. Dudzik, J. C. Ricklin (Eds.), *Autonomous Systems: Sensors, Vehicles, Security, and the Internet of Everything*, volume 10643, International Society for Optics and Photonics, SPIE, 2018, pp. 148 – 167. URL: <https://doi.org/10.1117/12.2301254>. doi:10.1117/12.2301254.
- [19] Y. WU, H. Zhu, Q. Du, S. Tang, A survey of the research status of pedestrian dead reckoning systems based on inertial sensors, *International Journal of Automation and Computing* 16 (2019).
- [20] W. Simoes, G. Machado, A. Sales, M. de Lucena, N. Jazdi, V. de Lucena, A review of technologies and techniques for indoor navigation systems for the visually impaired, *Sensors* 20 (2015) 3935. doi:10.3390/s20143935.
- [21] L. Ma, Y. Fan, Y. Xu, Y. Cui, Pedestrian dead reckoning trajectory matching method for radio map crowdsourcing building in wifi indoor positioning system, in: *2017 IEEE International Conference on Communications (ICC), 2017*, pp. 1–6. doi:10.1109/ICC.2017.7996457.