RTT ranging to Wi-Fi APs beats GNSS localization

Berthold K. P. Horn[a]

***Abstract*—Wi-Fi round-trip-time (RTT) ranging has proven successful in indoor localization. Here it is shown to be useful outdoors as well — and more accurate than smartphone code-based GNSS when used near buildings with Wi-Fi access points (APs). A Bayesian grid with observation and transition models is used to update a probability distribution of the position of the user equipment (UE). The expected value (or the mode) of this probability distribution provides an estimate of the UE location.**

Localization of the UE using RTT ranging depends on knowing the locations of the Wi-Fi APs. Determining these positions from floor plans can be time-consuming, particularly when the APs may not be accessible (as is often the case in order to prevent unauthorized access to the network). An alternative is to invert the Bayesian grid method for locating the UE — which uses distance measurements from the UE to several APs with known position. In the inverted method we instead locate the AP using distance measurements from several known positions of the UE.

In localization using RTT, at any given time, a decision has to be made as to which APs to range to, given that there is a cost associated with each "range probe" and that some APs may not respond. This can be problematic when the APs are not uniformly distributed. Without a suitable ranging strategy one can get into a dead end state where there is no response from any of the APs currently being ranged to. This is a particular concern when there are local clusters of APs that may "capture" the attention of the RTT app. To avoid this, a strategy is developed here that takes into account distance, signal strength, time since last "seen," and the distribution of the directions to APs from the UE — plus a random contribution.

We demonstrate the method in a situation where there are no line-of-sight (LOS) connections and where the APs are inaccessible. The localization accuracy achieved exceeds that of the smartphone code-based GNSS.

***Index Terms*—outdoor location; outdoor position; round trip time; fine time measurement; RTT; FTM; IEEE 802.11mc; IEEE 802.11–2016; IEEE 802.11az; IEEE 802.11–2024; Bayesian grid; Bayesian grid update; observation model; transition model; relative permittivity; refractive retardation;**

I. Introduction

Indoor localization using fine-time-measurement (FTM) round-trip-time (RTT) with respect to Wi-Fi access points (APs) [1]–[31] has been shown to work well, achieving 1–2 meter accuracy [9], [19], [20] in houses and office spaces. All this without the impractical detailed surveying that some competing localization methods require. FTM RTT depends on Wi-Fi APs implementing the IEEE 802.11mc protocol ("two-sided" RTT) incorporated in IEEE 802.11-2016 [32]. More recently, it has also become possible to range to "uncooperative" APs using "one-sided" RTT [27]. This allows one to make use of more types of APs, albeit with somewhat reduced accuracy, typically 3–4 meters. There are applications available that implement localization methods using RTT ranging, including Google's `WiFiRttLocator` [33], which uses least-squares (LSQ) multilateration, and `FTMRTT` [34] which uses a Bayesian grid update (BGU) method (in 2-D and 3-D). The BGU method is described in papers on indoor localization [19], [20], and in the chapter "Updating Occupancy Grids Using Bayesian Estimation" in Murphy's "An Introduction to AI Robotics" [35].

The accuracy of RTT measurements is roughly inversely proportional to the bandwidth. However, in practice, the results are not nearly as good as suggested by the Cramér-Rao bound [27], [36] for a given bandwidth. There are several reasons for this. One is the so-called "position-dependent" error [19] which can induce meter-size changes in reported distance when the user equipment is moved just a few millimeters. This is prevalent indoors where reflections off various surfaces lead to interference patterns with texture elements about one-half wavelength in size [37]. A related problem is multi-path where the algorithm for detecting the first arrival may be misled by contributions from later arrivals (Note that there is no path diversity in IEEE 802.11mc since it does not exploit the multi-input multi-output (MIMO) method, introduced in IEEE 802.11ac Wave 2).

Another contribution to measurement error is the retardation of radio frequency signals by material of high refractive index along the path. The refractive index n equals $\sqrt{\varepsilon_{\mathrm{r}}\mu_{\mathrm{r}}}$, where ε_{r} is the relative permittivity and μ_{r} is the relative permeability. Many materials have high relative permittivity at Wi-Fi frequencies [38]–[41] so the refractive index can be quite large (e.g. $\varepsilon_{\mathrm{r}} \approx 79$ and $n \approx 8.9$ for water at 2.4 GHz). The resulting "refractive retardation" adds $(n-1)\ \delta$ to the measured distance, where δ is the thickness of the material. This is less of a problem for outdoor application of RTT ranging since the radiation typically has to go through just one or two walls near the AP but does not have to traverse additional obstacles.

So far there has been little effort to apply RTT methods outdoors. RTT potentially can work even better outdoors than indoors because typically most of the path between the AP and the UE is free of obstacles. It is also desirable to provide a smooth transition from navigating indoors to navigating outdoors. Methods based on RTT ranging will work both indoors and outdoors. Naturally, there are limitations on the distance one can stray from areas that have Wi-Fi APs.

II. Background

In the indoor situation, there are several alternative localization methods. Perhaps the most frequently reported upon is based on signal strength (RSSI). Since signal strength depends on many factors other than distance, it is not possible to unambiguously recover the distance of a smartphone from an access

[a] Department of Electrical Engineering and Computer Science and CSAIL, Massachusetts Institute of Technology, Cambridge, MA 02139, USA, e-mail: bkph@csail.mit.edu

* The corresponding author is Berthold K. P. Horn.

point. To overcome this problem, so-called finger-printing methods have been developed which require measurement of signal strenghts of multiple APs at each point on a grid. The accuracy achievable depends in part on how fine a grid one is willing to survey. Spatial fluctuations in the electromagnetic field due to interference of signals travelling along different paths have a "texture" on the order of half a wavelength [37] (which is 60 mm in the 5 GHz band). So, ideally, one should sample the field at that scale — something that is not practical, particularly if carried out for "all" locations on a 3-D grid rather than just on a single horizontal plane.

Further, modern enterprise Wi-Fi systems optimize their performance and minimize interference by periodically reassigning AP channel frequencies and by changing AP transmit power. Such adjustements make pre-determined fingerprinting data much less useful. Also, fingerprinting needs to be redone when there are changes in the radio frequency environment, such as caused by movement of absorbers or reflectors.

Another approach to indoor localization is based on scattering a large number of low power beacons (typically Bluetooth) over the area of interest. The location reported for the device is then the (known) location of the beacon with the strongest signal (perhaps also taking into account the locations of the beacons with the second and third largest signal). The accuracy of this approach depends on how many beacons one is willing to spread around. This approach requires additional investment in hardware, installation, and maintainence — unlike methods exploiting pre-existing Wi-Fi infra-structure.

Since the accuracy of time-of-flight methods is inversely proportional to the bandwidth of the signal, there is an advantage to very wide bandwidths. So-called ultra wide-band (UWB) systems (> 500 MHz bandwidth) can make accurate distance measurements. However, the Federal Communication Commission (FCC) limits the RF power of such systems (-41.3 dBm/MHz) to the point where they are only useful for short distances — the order of a few meters, as opposed to a few tens of meters for Wi-Fi signals.

None of these methods have found wide-spread use for indoor localization — and certainly not for outdoor localization.

Meantime, IEEE Wi-Fi standards are continuing to increase the available channel bandwidth, with IEEE 802.11ac (2013) supporting 80 MHz, IEEE 802.11ax (2021) supporting 160 MHz, and the next standard, IEEE 802.11be (2024), supporting 320 MHz bandwidth. Thus FTM RTT is approaching the ranging accuracy possible with UWB, but without the stringent power limitations, thus making it useful for distances of interest here.

III. Motivation

The method that comes to mind first for outdoor navigation is, of course, GNSS. The typical accuracy of smartphone's code-based GNSS under an open sky is 4.9 meters [42]. There are methods such as carrier phase enhanced GPS (CPGPS), Real-Time Kinematic Positioning (RTK), Continuously Operating Reference Stations (CORS), GDGPS, TDRSS, PPP-AR, etc. that can increase the accuracy of GNSS localization, but these typically add to cost and complexity, require subscription to auxiliary services, or do not operate in real-time [9]. There has also recently been interest in alternate methods for localization given the susceptibility of GNSS to jamming and spoofing [43].

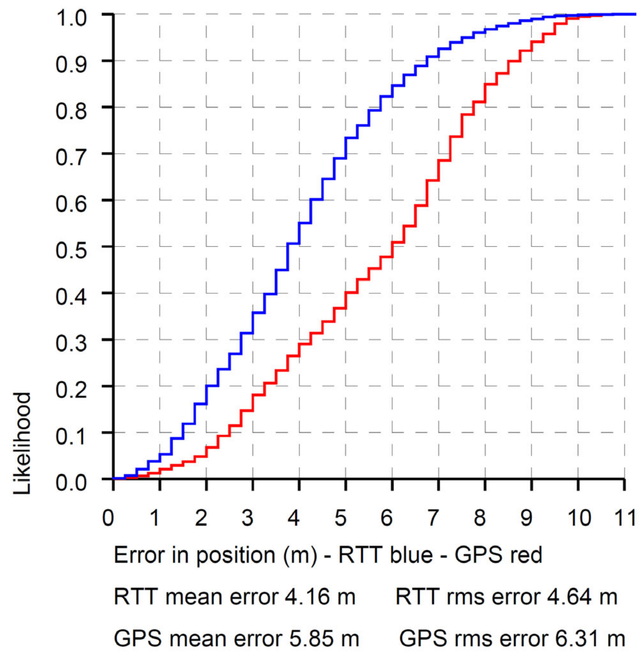

Fig. 1. Cumulative Distribution Function (CDF) for position errors (in meters) using (i) RTT and the expected value of the Bayesian Grid Update method (BGU) (blue) and (ii) smartphone GNSS (red). The mean and RMS error are smaller for the BGU results than for the smartphone's GNSS. The CDFs are based on measurements in a hotel/condo complex described in section XI.

In the experiments reported here, the positions determined using the BGU method with RTT ranging have smaller errors than those of the smartphone's code-based GNSS. In Fig. 1, the blue curve is the Cumulative Distribution Function (CDF) for the error in expected value of the probability distribution of the RTT measurements, while the red curve is the CDF for the error in the positions reported by the smartphone's GNSS. These CDFs are based on 6140 measurements in a hotel/condo complex described in section XI.

IV. Where are the APs?

One issue that may have dampened enthusiasm for the deployment of localization using Wi-Fi RTT is that the positions of the APs need to be known. Installation of APs may be based on a marked-up floor plan with the desired coordinates. But such plans may not be available, may not be accurate or may not be up to date, given that APs are replaced when faulty, and are moved when "holes" in Wi-Fi coverage are noted.

Wi-Fi APs supporting IEEE 802.11mc do have the opportunity to encode their own location and provide it to the UE using the Location Configuration Information (LCI) or the Location Civic Report (LCR) fields in the "Measurement Report Element" [32], [44], but setting these fields has not yet been widely supported by AP makers. This is about to change as major vendors such as HPE/Aruba [45], Cisco/Meraki [46], Juniper [47], Extreme Networks [48], and Zebra [49] start to support AP location information and automated determination of AP locations. These systems can use RTT ranging between APs, in addition to known coordinates of a small number of "anchor" APs (whose position may, for example, be obtained via some form of high-accuracy GNSS) [50].

For APs in areas where GNSS is available (e.g. near the outside walls of buildings or under non-metallic roofs), it is also possible to get estimated positions using crowd-sourced data such as those made available by Google's geolocation API [51]. There are some limitations to this, since what is reported is based in part on the weighted centroid of the GNSS-determined locations of smartphones that were near enough to the AP to receive a signal, not the location of the AP itself. This can be problematic when directional antennas (such as panel antennas) are used or when most of the pedestrian traffic is off to one side of the AP. Also, this service can only be provided for "universally-administered" MAC addresses. Typically APs support several Basic Service Set Identifiers (BSSIDs) many of which are "locally administered." These are not recorded by the geolocation API because they are not globally unique.

In any case, what is needed is an automated way of determining where the APs are.

V. What is New?

The following are contributions of this paper:

- The Bayesian grid update (BGU) method of localization using FTM RTT ranging is extended for use outdoors;
- An "inverse" Bayesian grid method is used to determine the locations of the APs;
- A simplified observation model for the outdoor application of the Bayesian grid method is presented;
- A strategy is introduced for selecting the set of APs to range to;
- A way of avoiding Wi-Fi scans, using instead range probes to determine what frequencies APs operate on;
- Localization results are shown for a large complex (210 $\times$ 120 m) containing multiple buildings with over a hundred APs;
- Results are compared with coordinates provided by the smartphone's code-based GNSS.

VI. Inverse Bayesian grid for finding AP locations.

In localizing the UE, RTT ranges to several APs with known positions are measured. These ranges are then used either in least-squares multilateration or in a Bayesian grid method for estimating the position of the UE. Generally speaking, the accuracy of the estimated location is inversely proportional to the square root of the number of APs that respond. Naturally, the distribution of APs also affects accuracy. For example, there won't be sufficient constraint if most APs are located in more or less the same direction as seen from the UE. More generally, if the APs lie more or less along a straight line there will be a two-way ambiguity (related by reflection about that line). The dependence of accuracy on geometry is similar to the "dilution of precision" in GNSS parlance [52].

Indoors, distributions for APs that are good for satisfactory Wi-Fi coverage tend to be useful for localization also. Such distributions tend to have roughly the same density throughout, avoid dead zones and avoid clustering of APs. (Although things tend not to work quite as well near corners of the convex hull of the AP locations, since constraints there are not available from all directions). By contrast, in outdoor localization, many APs may be arranged in clusters inside buildings and may lie more or less to one side of the UE.

The localization process can be "inverted" and the positions of APs obtained from ranges measured from known positions of the UE. RTT ranging data for this purpose can be collected while walking around with the UE in hand. For such ranging data to be useful there must be a way of independently locating the UE during this calibration step. This location information can be based on prominent landmarks, such as intersections of paths, corners of buildings, and so on (with a way of marking the times at which the UE passes them).

The method of locating an AP is very similar to that used for localization of the UE in that a Bayesian grid of probabilities of possible AP position is updated using Bayes' rule based on a suitable observation model.

VII. Simplified Observation Model for Outdoor Use.

Ideally, the RTT-reported distance would be equal to the true distance, perhaps with a small amount of measurement noise added — and a small offset that may depend on the equipment used [53]. Inside buildings, multi-path and refractive retardation complicate matters in that measured RTT ranges are typically larger than the actual distances. This can be taken into account in the BGU method if statistical information is available relating observed distances to actual distances — i.e. if we have an observation model [20].

In outdoor localization, while paths between APs and UEs may not be line of sight (LOS), typically *most* of the path is unobstructed. That is, there may very well be one or two walls between the AP and the outside, but there are typically no additional obstacles along the path to the UE — as there would be inside a building.

Experiments confirm that the observation model can be much simpler in the outdoor application. In the results reported below (see e.g. Fig. 8), a good model is one where the measured distance is the actual distance plus a measurement noise with standard deviation of about 3 meters. This does not take into account outliers contributing to "long tails" on the probability distribution. If we ignore these outliers for the moment, the conditional probability can reasonably be modeled as

$$p(r_{\mathrm{rtt}}|r_{\mathrm{true}}) \approx \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{1}{2}\left(\frac{r_{\mathrm{rtt}}-r_{\mathrm{true}}}{\sigma}\right)^2\right) \tag{1}$$

where r_{rtt} is the RTT-measured distance while r_{true} is the true distance and $\sigma \approx 3$ m.

When an RTT range measurement r_{rtt} is reported for an AP, the BGU method steps through the cells of the grid and updates the probabilities stored there using Bayes' rule [19], [20], [35]. At each cell it uses the corresponding position of the UE to compute the distance r_{true} from the AP (this can be precomputed if desired). Then, it applies Bayes' rule, which, in the 2-D case, comes to

$$p_{i,j}^{(n+1)} = p_{i,j}^{(n)} \; p(r_{\text{true}}|r_{\text{rtt}})/N \tag{2}$$

where

$$N = \sum_{i,j} p_{i,j}^{(n)} \; p(r_{\text{true}}|r_{\text{rtt}}) \tag{3}$$

Division by N can be taken care of simply by renormalising after the multiplicative part of the update step. Here $p_{i,j}^{(n)}$ is the probability of the UE being in cell (i, j), at step n (In the 3-D case we have instead $p_{i,j,k}^{(n)}$).

The BGU method needs an estimate of $p(r_{\text{true}}|r_{\text{rtt}})$, which, using Bayes' rule, can be written

$$p(r_{\text{true}}|r_{\text{rtt}}) = p(r_{\text{rtt}}|r_{\text{true}}) \; p(r_{\text{true}})/p(r_{\text{rtt}}) \tag{4}$$

Now, if we assume that the distribution of APs is uniform in the plane out to some radius R_0 say, then

$$p(r_{\text{true}}) = 2 \; r_{\text{true}}/R_0^2 \tag{5}$$

Further, if the RTT distance is more or less equal to the true distance (modulo some measurement error) then we have

$$p(r_{\text{rtt}}) \approx 2 \; r_{\text{rtt}}/R_0^2 \tag{6}$$

and so

$$p(r_{\text{true}}|r_{\text{rtt}}) \approx p(r_{\text{rtt}}|r_{\text{true}}) \; (r_{\text{true}}/r_{\text{rtt}}) \tag{7}$$

Since the Gaussian (eq. 1) is small except near $r_{\text{rtt}} = r_{\text{true}}$ (because the distances are mostly large compared to σ) one may be justified in further simplifying this to

$$p(r_{\text{true}}|r_{\text{rtt}}) \approx p(r_{\text{rtt}}|r_{\text{true}}) \tag{8}$$

and just use $p(r_{\text{rtt}}|r_{\text{true}})$ from eq. 1 for $p(r_{\text{true}}|r_{\text{rtt}})$.

In any case, given a measurement r_{rtt}, values of $p(r_{\text{true}}|r_{\text{rtt}})$ can be precomputed for a range of values of r_{true} in order to speed up the grid update (the saved result of this precomputation has been called a "rate vector" in [20]). Note that the simple observation model can be used both for initially locating the APs from RTT measurements, and later for locating the UE from RTT measurements.

The above shows what to do with positive evidence that the AP may be in certain grid locations based on a response to a ranging request from the UE. There may also be negative evidence: An AP is unlikely to be nearby if it does not respond to a ranging request (i.e. the absence of an RTT response is also an "observation"). In this situation one may want to *reduce* the probabilities of cells near the current UE positions — by an amount that decreases with distance. Experiments with a difference-of-Gaussian probability distribution for this did not prove especially beneficial in the cases described below.

VIII. Wi-Fi installations used for AP localization

We determined the positions of APs in three installations using the "inverse" Bayesian method described above. In each case, the APs where inside buildings in condo developments or hotels (see Fig. 3). We report here specifically on one large complex (210 × 120 m) with 108 APs concentrated in three main buildings, (with another 5 APs in two auxiliary buildings). The APs in this case were from Ruckus Wireless according to their Organizational Unique Identifier (OUI) (94:B3:4F). These APs do not advertise support for IEEE 802.11mc in the Wi-Fi beacon; but do respond to two-sided RTT ranging requests. The RTT measurements in this case were determined to have an offset or bias of 2.8 m (to be subtracted from the RTT measurement) [53] — and a standard deviation of about 3 m. (Note that the standard deviation reported in the RTT measurement itself is typically smaller because it is based on multiple measurements in rapid succession under near identical conditions and from basically the same position).

We experimented with several smartphones including Pixel 5, Pixel 6 Pro and Pixel 7a. The data reported here was obtained using a Google Pixel 5.

The ground truth is based on coordinates from Google Maps of landmarks — such as intersections of paths — along with interpolation over short distances based on step counting and headings obtained using the smartphone's acceleration and magnetic sensors. Coordinates determined from Google Maps are claimed to be accurate to about 1 m in some regions of the developed world (with a world-wide average of about 4.38 m) [54].

Because many RTT measurements were available for locating each AP, the accuracy of the derived AP position is better than that of the raw RTT measurement. If the probability distribution was Gaussian, and the directional geometry ideal, we would expect

$$\sigma_{\text{AP}} = \frac{1}{\sqrt{N/D}} \; \sigma_{\text{RTT}} \tag{9}$$

where σ_{RTT} is the standard deviation of the RTT measurements, σ_{AP} is the standard deviation of the computed position of the AP, N is the number of different UE positions, and $D = 2$ in the 2-D case, and $D = 3$ in the 3-D case. Based on measurements on three different days, the error in AP position was estimated to be around 1–2 meters. (This uncertainty in positions of the APs naturally induces some small error later when ranging to these APs in order to determine the position of the UE).

Fig. 2 shows four sample Bayesian grids for locating APs in the first building. In each subfigure, the red cross marks the mode of the probability distribution, while the outlines of the buildings are shown in blue. The grid cells here are 1 × 1 m. To accommodate the large dynamic range of the probabilities, the grey level in the figure is made proportional to the square root of the logarithm of the probability.

RTT measurements often have outliers far from the correct distance (an example may be seen in the lower left of the bottom right image). Some of these are easy to reject (e.g. very large RTT standard deviations or negative RTT distances) but

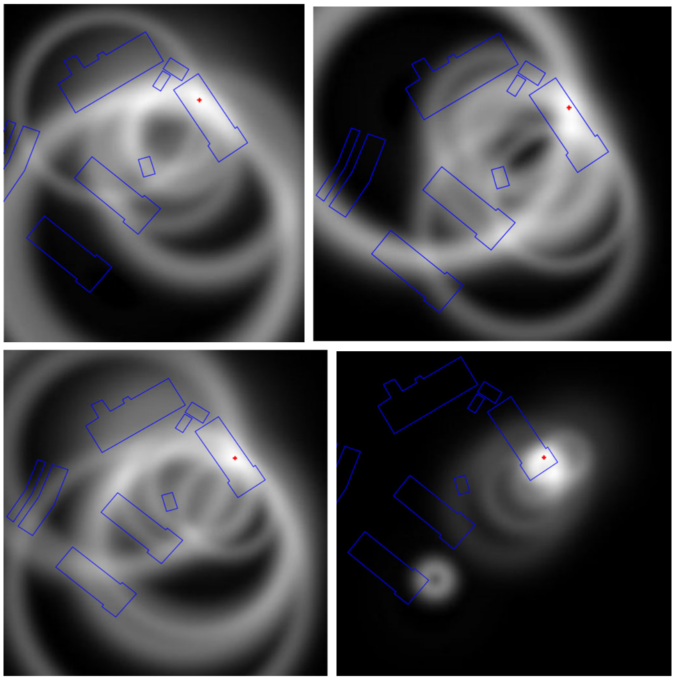

Fig. 2. Inverse Bayesian grid distributions for four different Wi-Fi access points. The red crosses mark the peak values in each probability distribution. Outlines of the buildings containing the APs are shown in blue.

many masquerade as legitimate measurements (with reasonable looking standard deviations) even when they are dozens of meters off the correct value. These outliers do not present much of problem for the BGU method, but can potentially be problematic for LSQ-based multilateration because of the large contributions to the least-squares loss function (An effect that can be ameliorated to some extent by weighting the contributions with the inverse of the standard deviation of the RTT measurement, and by dropping the worst of a group of related range measurements.).

The above figures are based on data collected during a single "walk-through" of the property. For improved accuracy, the final AP locations were based on data collected during several such data collecting sessions.

IX. Which of the APs should be used for ranging?

When using FTM RTT for localization, one needs a strategy for selecting APs. It takes time to perform RTT measurement; so one typically can't afford to request RTT responses from all of the responders in an area of interest.

In addition, the Android API limits the number of BSSIDs in a given ranging request (`maxPeers=10`). To address more APs, multiple ranging requests have to be used, with considerable associated overhead. A tempting strategy is to include in the current set of APs just the ones that have a high probability of responding again. One could, for example, range to the APs that have responded recently, are nearby, and have good signal strength. But this means that, as the ranging process is repeated, the number of APs in the current "ranging set" can only go down. This creates a problem when one moves away from those APs, since once far enough away, no responses will be received at all. There must be a mechanism for adding "new" APs into the "current set" (and for retiring "old" ones).

One way to attack this problem is to perform frequent Wi-Fi scans to see which APs are "within range." But, as discussed below, Wi-Fi scans are slow and may interfere with ranging.

Another approach is to sort APs on their distance and select the nearest ones. The distance can be calculated for all APs — including ones not in the current set — using an estimate of the current position of the UE. However, if the APs come in clusters rather than being uniformly distributed, then the APs nearest to the old estimate of the UE position may form a group, none of which responds once the UE is far enough away. When moving away from such a cluster, all ranging requests may be for "old" APs that eventually are too far away to respond.

An alternative is to "probe" APs that, while not in the current set, are likely to respond. There is a cost associated with ranging to APs. As a result, there is a trade-off between addressing mostly APs that are expected to respond — in order to get the highest localization accuracy — versus adding in some "range probes" — to avoid losing contact and becoming stuck ranging to a cluster of non-responsive APs.

What criterion to use? We already discussed selection based on distance (see also [27]). We can also use signal strength and the time since a response was last obtained. Some of these criteria only apply to APs that have been "seen" recently. One alternative is to sort on a "quality" factor that incorporates all available information. A simple version of a quality factor is a linear weighted combination of distance, time since last seen, and the signal-strength when last seen (The stored signal-strength of an AP not seen recently is not trustworthy, so should be forced to algorithmically "fade away" with time.). Other factors may be included, some with positive weights and some negative. Examples include whether the frequency of the AP is known or whether it is yet to be determined (see below). Higher weight can be given to an AP whose frequency has already been found.

The directions from the UE to the APs are also important, in that a set of APs lying in more or less the same direction do not provide a strong constraint on the position of the UE. The best is a set of directions spread all around the compass. A criterion based on this notion is harder to incorporate since it is not dependent on just the direction to an individual AP. One approach is to first determine the set of directions of all APs currently responding (since these are likely to respond again to the next ranging request) and then compute some criterion favoring directions that are *different* from those in the current set of directions. That is, favor APs that lie in large gaps in the set of current directions to APs.

There may also be prior information about how reliable the information from a particular AP is (e.g. a radio with different accuracy may be used in waterproof APs installed outdoors than those installed indoors e.g. [27]). If BSSIDs with different frequencies and bandwidths are used, then one may favor those with higher bandwidth since the RTT range is more accurate with higher bandwidths.

Finally, to avoid getting stuck in a situation where a set of APs yields no response (perhaps because the currently estimated position of the UE is outdated) it is useful to add a random component to the overall quality criterion. This ensures that successive ranging requests will have different AP composition even when there has been no response (and hence no change in the quality factors). The random component is also useful when there are just a few responses. The weight of the random component can be made to decline with the number of APs responding, so that it is quite large when there are no responses (or when there are just one or two), but does not have much of a disturbing effect when several APs are providing good data.

In summary, the selection criterion may be based on the following: (i) distance; (ii) time since last seen; (iii) signal-strength (faded since last seen); (iv) whether the frequency is known; (v) directional factor; (vi) prior quality factor (weight); and (vii) a random perturbation;

Some other factors that could be taken into account in the quality factor calculation include the frequency band, the Wi-Fi standard supported by the AP, the selection of preamble, and whether the Service Set Identifier (SSID) is hidden. Further, ranging with APs operating in the Dynamic Frequency Selection (DFS) part of the 5GHz band may be slower on some platforms, so one may want to assign a lower weight to those APs. Also, since switching the frequency that the radio operates on takes time, there may be some advantage to picking a set of APs that operate on the same small set of different frequencies. Overall, while some strategy is needed to avoid getting stuck in a situation where no APs respond to successive ranging requests, the exact details appear to be less important.

X. Determining frequency of APs in enterprise Wi-Fi

The frequency of each AP needs to be known in order to construct an RTT ranging request. While APs in some home Wi-Fi systems operate on one (or two) fixed frequencies, in enterprise-level installations, some central optimization periodically (often overnight) reassigns frequencies to improve throughput and reduce conflicts between APs.

One way to obtain the necessary information is to scan the Wi-Fi bands and determine the frequencies of APs within range. Wi-Fi scans take the order of 100 msec per channel This means that it takes about 4 to 5 seconds to perform a Wi-Fi scan covering the 2.4 GHz band (11 possible channels) and the 5 GHz band (25 possible channels) — longer if the 6 GHz band (15 preferred scanning channels) is also supported. At an average pace of 1.4 m/sec, a person walks 6 to 7 m during a Wi-Fi scan. So one can't afford to simply stop RTT ranging to wait for Wi-Fi scan results, and, at this point in time there doesn't appear to be a way to reliably interlace RTT ranging and Wi-Fi scanning on certain platforms (nor a way to restrict Wi-Fi scanning to specific channels).

An alternative is to probe the different possible frequencies for each AP using range requests. A trial frequency is verified if there is a valid RTT response (while lack of a response may be due to any number of reasons, not just that the wrong frequency was selected). This may be a reasonable alternative since RTT ranging responses are much faster (50–400 msec — depending on burst size, number of responders, radio model, and frequency) than Wi-Fi scans (4–5 sec). It helps that often enterprise Wi-Fi systems limit the number of channels used (e.g. using only channels 1, 6, and 11 in the 2.4 GHz band, and only 8 of the 25 channels in the 5 GHz band). This speeds up the search for the correct frequencies.

XI. Sample outdoor localization results

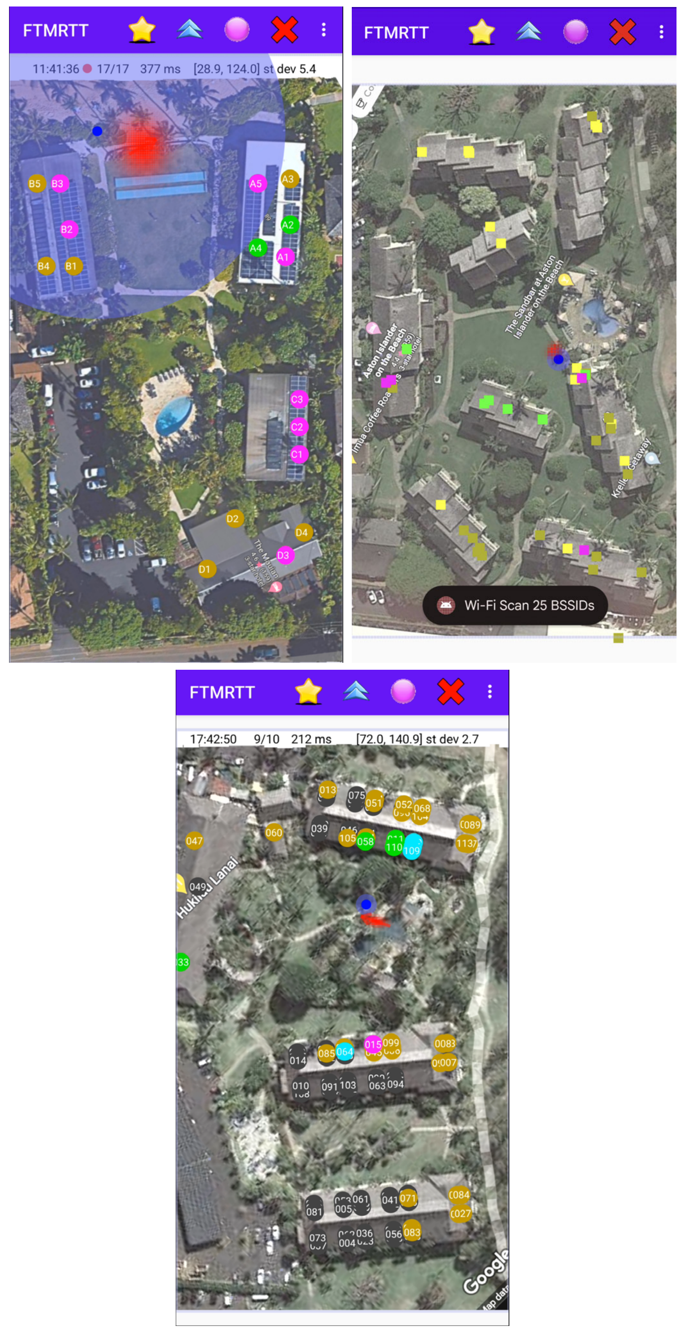

Fig. 3. Sample FTMRTT screen shots at the three locations: (i) with 17 APs; (ii) with 29 APs; (iii) with 113 APs. The red blob is the probability distribution of the UE position. The GNSS position is shown as a small dark blue dot. APs currently responding are shown in green (or cyan); APs not currently responding to range requests are shown in magneta (or red);

We experimented with the method in three locations (see Fig. 3):

- A small hotel with three main buildings (133 × 72 m) and 17 APs supporting only one-sided RTT (video at [55]).
- A condominium development (182 × 108 m) with 29 active APs, supporting only one-sided RTT.
- A hotel/condominium complex (210 × 120 m) with 113 APs supporting two-sided RTT (videos at [56], [57]).

In each of the above referenced screen recordings, the red blob is the probability distribution of where the app thinks the UE may be. APs currently responding are shown in green (or cyan), while APs not responding are shown in magenta (or red). APs not in the current selection, but whose frequency is known, are shown in a brownish yellow. When available, the GNSS location is shown as a dark blue dot. We focus here on the third, most challenging case.

The locations of the 113 APs were determined using the inverse BGU method described above in section VI. In this particular location, additional prior information was available to constrain the results for the 108 (= 36 × 3) APs in the three main buildings: the layout of rooms and the position of APs within the rooms was the same in each of the buildings. This made it possible to use a least-squares method to find that regular pattern using data from all three buildings. The positions of the APs could then be fine-tuned to fit this common pattern. The results could not be checked directly since the APs were not visible from outside, but comparing AP positions obtained using data on three different days indicated that the AP positions were accurate to within 1–2 m.

For comparison, the Google geolocation API applied to the "universally-administered" MAC addresses of these APs yielded results mostly within 3–5 m of the AP positions determined as above (albeit with a few large outliers). So, actually, one could just use the locations supplied by the geolocation API if a slightly larger error in the estimated location of the UE was acceptable.

The known positions of the APs were then used in the `FTMRTT` app [34] to determine the position of the UE as it was carried around the property. In this installation APs use channels in both the 2.4 GHz and 5 GHz bands. The 5 GHz band potentially provides for higher ranging accuracy because it can support wider bandwidths (up to 80 MHz), but in this installation, channels in *both* bands were restricted to 20 MHz bandwidth. BSSIDs operating in the 2.4 GHz band were chosen for ranging because, for the same transmitted power, lower frequencies provide a stronger signal [58]. and so allow for ranging at larger distances (on the other hand, using both 2.4 GHz and 5 GHz channels on an AP can provide improved accuracy because of frequency diversity [19]).

GNSS coordinates were recorded, as well as the times at which easy-to-identify landmarks on the path were passed (using the repurposed "volume down" key). The geographical coordinates of these were determined using Google Maps.

In Fig. 4, the ground truth is indicated in blue, the GNSS coordinates in red, and the results of the BGU method in green. Note that some parts of the path are traversed twice, and that there are two gaps in the recording apparently resulting from firmware freezes leading to 5 second watch-dog timer events.

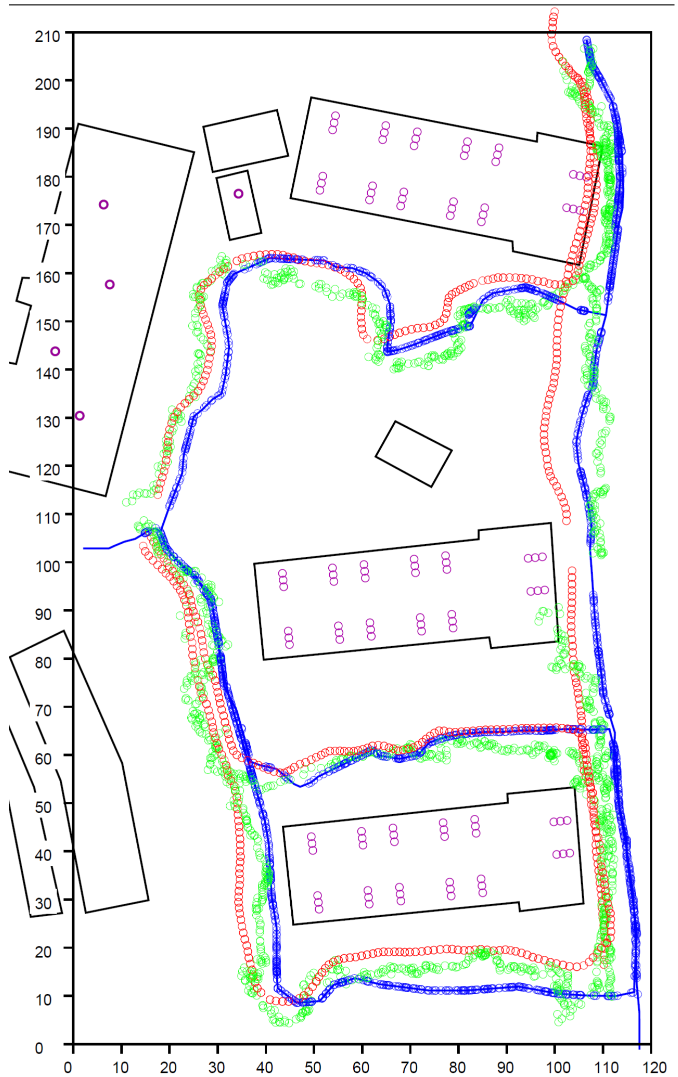

Fig. 4. Ground truth (blue), GNSS (red) and RTT (green) positions with outlines of buildings in black. The path folds back on itself and traverses some sections twice. The large outer rectangle (210×120 m) delineates the area covered by the Bayesian grid.

The locations of the five APs (top left) not in the three main buildings are not as well known because they are not part of a repeating pattern that could be used for fine-tuning.

Fig. 5 compares the errors in the estimated position of the UE based on RTT measurements and the smartphone's GNSS. The mean error for determining the position of the UE using the BGU method here was 4.16 m (RMS error 4.64 m). For comparison, the mean error in the GNSS track was 5.85 m (RMS error 6.31 m). The error in the GNSS positions is about what is expected of smartphone's GNSS (4.9 m with open sky [42]). The corresponding CDF plots are shown in Fig. 1 above.

Note that any remaining uncertainty in the position of the APs, although small, contributes a bit to the errors in UE positions.

A. Observed RTT distance measurement errors

We next consider the errors in the RTT measurements by comparing the reported RTT distance with the actual distance

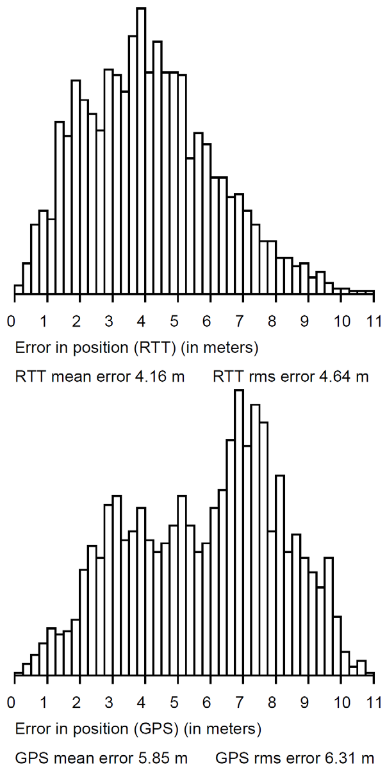

Fig. 5. Histogram of errors in the position of the UE (in meters) based on (i) RTT measurements and known positions of APs; and (ii) the smartphone's GNSS measurements. The mean and RMS error are smaller for the RTT based results than for the smartphone's GNSS.

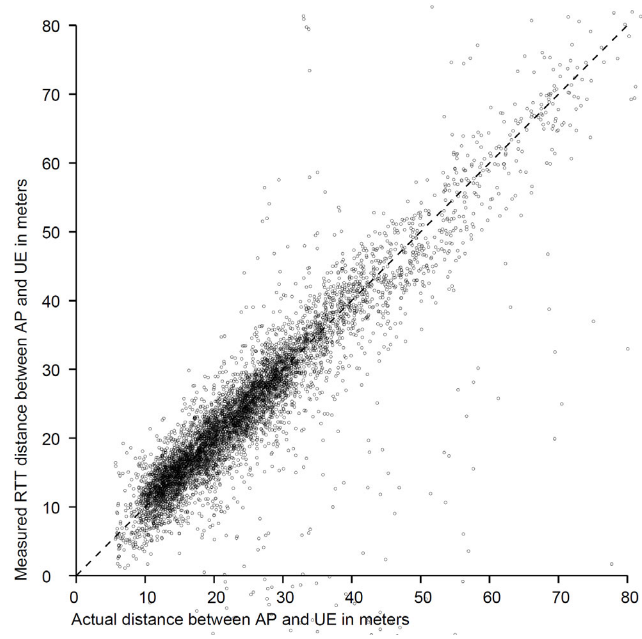

Fig. 6. Scattergram of RTT distance (vertical axis - meters) versus actual distance (horizontal axis - meters). The scatter is relatively large, in part because of the small bandwidth (20 MHz). The dashed line has slope one. There are many outliers far above and below the dashed line.

between the UE and the AP.

In Fig. 6, the vertical axis is the RTT measured distance while the horizontal axis is the actual distance. This shows that the RTT distance can be considered to be the true distance plus some form of measurement error (and a small offset). There are, however, also quite a number of large outliers (including some negative values).

We can present this information in another way by plotting the difference between the RTT measurement and the true distance (i.e. the measurement error) versus distance as shown in Fig. 7. Note that the vertical spread is fairly independent of distance. Overall, a histogram of RTT error, as in Fig. 8, looks approximately Gaussian with standard deviation $\sigma \approx 3$ m (aside from some long tails). This probability distribution was used as the observation model in the BGU method both for finding the APs from RTT measurements and then also for finding the UE from RTT measurements (section VII).

We can use the formula $\sigma \approx 2.45 + 0.035\ d$ if we wish to account for a slight increase of measurement error with distance d (in meters). Note, however, that there remain some large outliers that do not fit the Gaussian distribution model.

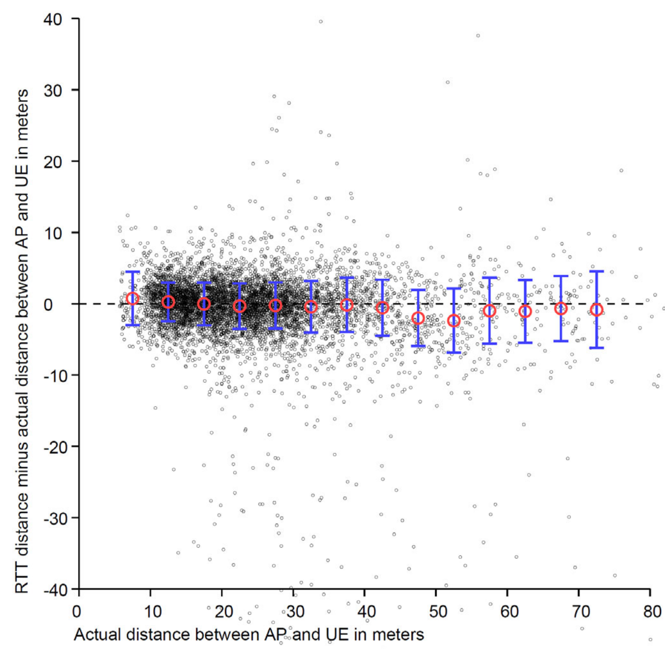

Fig. 7. Scattergram of error in RTT distance (vertical axis) versus actual distance (horizontal axis). The mean (μ) of each 5-meter wide histogram bin is shown in red, while the blue lines go from $\mu - \sigma$ to $\mu + \sigma$, where σ is the standard deviation.

XII. Boot-strapping

Obtaining measurement data for constructing the observation model — and for locating the APs — takes some effort. There is a limit to how many distinctive landmarks one may want to make measurements from. Once an approximate observation

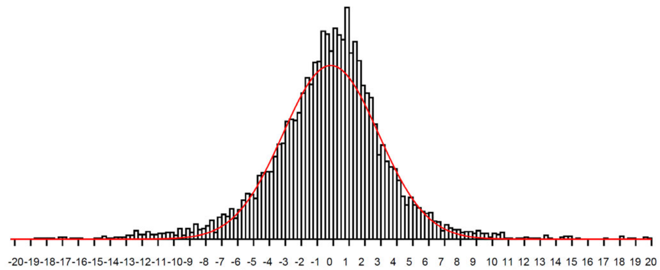

Fig. 8. The histogram of error in RTT measurement has a bell shape somewhat like a Gaussian, although also sporting long tails. (The error here is in the RTT measurement, given the AP positions). Superimposed in red is a Gaussian with standard deviation 3 meter.

model — and approximate positions for the APs — have been determined in this fashion, one can generate a lot of synthetic data: The estimate of the location of the UE can then be used as a proxy for the true position of the UE. This way, instead of a few dozen measurements obtained from known landmarks, one has available thousands of RTT measurements that can be used to improve the model — and refine the positions of the APs. This boot-strapping process for fine-tuning the model can be repeated; although the return-on-effort tends to be lower the second time around.

XIII. Nature of the measurement error

One striking feature of Fig. 4 is the different "fine structure" of the BGU tracks (green) and the GNSS tracks (red). Part of the difference is the result of GNSS positions being updated regularly at 1 second intervals, while the BGU estimates of position are computed irregularly, several times per second, as RTT responses come in.

The GNSS tracks are smoother but appear to have a relatively large offset or bias that persists over time. The offset changes when different satellites come into view, as is apparent in some sections of the path that were traversed more than once a few minutes apart. The smoothness of the tracks is most likely the result of some form of smoothing of the raw measurements, such as Kalman filtering. This filtering also manifests itself as a lag that causes the GNSS along-track error to be greater than the GNSS across-track error.

The BGU tracks, which appear more "jittery," could be made smoother using some form of filtering as well. But filtering introduces lag, which is undesirable when navigating in real-time. Unlike the GNSS tracks, the BGU tracks tend to pass through nearly the same positions when the path is traversed a second time. This is because the "position-dependent error" [19] does not change much with time.

XIV. Conclusions

The Bayesian grid method of localization using FTM RTT has been extended for use outdoors. An "inverse" Bayesian grid method has been used to determine the locations of the APs. A simplified observation model for outdoor use has been presented. A strategy has been introduced for selecting the set of APs to range to in order to avoid dead ends where no AP being ranged to respond. Experimental results are presented and analyzed for a large complex ($210\times$ 120 m) containing multiple buildings with over a hundred APs. The methods were demonstrated in a situation where there are no line-of-sight (LOS) connections, where the APs are inaccessible, and where the APs are concentrated in a few clusters. The RTT-based localization results are compared with the smartphone's code-based GNSS.

Of course, many factors influence the performance of GNSS and FTM RTT. For real-time use, smartphones provide code-based GNSS. Much greater accuracy can be achieved with geodetic grade GNSS devices with survey-grade circularly polarized antennas if locations are *not* required in real-time and data can be collected over a long time. Such setups typically require equipment that costs several thousand dollars, plus a second device in a known fixed location, a subscription to a service providing correction data, as well as post-processing (i.e. positions are not available in real-time).

Some Wi-Fi APs do not support the FTM RTT protocol and so can be used only with "one-sided" RTT, which has lower accuracy [27]. Going in the other direction, the IEEE 802.11az standard will enable ranging with even higher accuracy. The results in this paper apply to current smartphone implementations of code-based GNSS positioning and current "two-sided" IEEE 802.11mc FTM RTT distance measurements with respect to Wi-Fi access points.

The adoption of IEEE 802.11mc may not have been as rapid as one might have hoped. Meantime much progress has been made in finalizing IEEE 802.11az, the next generation positioning (NGP) protocol [59], [60]. IEEE 802.11az promises to further increase accuracy, likely to sub-meter level, by exploiting MIMO path diversity and channel sounding repetition. In addition 802.11az will provide reduced measurement time and added privacy and security. This will further improve localization using Wi-Fi APs, both indoors and outdoors.

References

[1] J. P. Tejedor, A. B. Martinez, S. M. Franco, R. M. L. Toledo, P. F. Reguero, and E. J. Abril, "Characterization and mitigation of range estimation errors for an RTT-based IEEE 802.11 indoor location system," *Progress In Electromagnetics Research B*, vol. 15, pp. 217–244, January 2009.

[2] A. B. Martinez, S. M. Franco, J. P. Tejedor, R. M. L. Toledo, P. F. Reguero, and E. J. Abril, "Indoor location based on IEEE 802.11 round-trip time measurements with two-step NLOS mitigation," *Progress In Electromagnetics Research B*, vol. 15, pp. 285–306, January 2009.

[3] T. Wang, Y. Shen, S. Mazuelas, H. Shin, and M. Z. Win, "On OFDM ranging accuracy in multipath channels," *IEEE Systems Journal*, vol. 8, no. 1, March 2014.

[4] L. Banin, U. Schatzberg, and Y. Amizur, "WiFi FTM and map information fusion for accurate positioning," in *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Alcale de Henares, Spain, October 2016.

[5] Wi-Fi Alliance. (2017, February) Wi-Fi certified location brings Wi-Fi indoor positioning capabilities. Wi-Fi Alliance. [Online]. Available: https://wi-fi.org/news-events/newsroom/wi-fi-certified-location-brings-wi-fi-indoor-positioning-capabilities

[6] A. K. Paul and T. Sato, "Localization in wireless sensor networks: A survey on algorithms, measurement techniques, applications and challenges," *Sensor and Actuator Networks*, vol. 6, no. 24, September 2017.

[7] W. Sakpere, M. Adeyeye-Oshin, and N. B. Mitwa, "A state-of-the-art survey of indoor positioning and navigation systems and technologies," *South African Computer Journal*, vol. 29, no. 3, December 2017. [Online]. Available: https://doi.org/10.18489/sacj.v29i3.452

[8] M. Ibrahim, H. Liu, M. Jawahar, V. Nguyen, M. Gruteser, R. Howard, B. Yu, and F. Bai, "Verification: Accuracy evaluation of WiFi fine time measurements on an open platform," in *MobiCom-18*, New Delhi, India, October 2018.

[9] F. V. Diggelen, R. Want, and W. Wang. (2018, July) How to achieve 1-meter accuracy in Android. [Online]. Available: https://www.gpsworld.com/how-to-achieve-1-meter-accuracy-in-android/

[10] J. Cobb. (2018, September) Testing WiFi RTT on Android P for indoor positioning. [Online]. Available: https://www.crowdconnected.com/blog/testing-wifi-rtt-on-android-p-for-indoor-positioning/

[11] K. Han, S. M. Yu, and S. Kim, "Smartphone-based indoor localization using Wi-Fi fine timing measurement," in *2019 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, September 2019.

[12] D. Campbell. (2019, October) Indoor positioning with WiFi RTT and Google WiFi. [Online]. Available: https://medium.com/@darryncampbell_83863/indoor-positioning-with-wifi-rtt-and-google-wifi-a638f1147b84

[13] G. Guo, R. Chen, F. Ye, and Y. Pan, "Indoor smartphone localization: A hybrid WiFi RTT-RSS ranging approach," *IEEE Access*, December 2019.

[14] S. Xu, R. Chen, Y. Yu, G. Guo, and L. Huang, "Locating smartphones indoors using built-in sensors and Wi-Fi ranging with an enhanced particle filter," *IEEE Access*, vol. 7, pp. 95 140–95 153, July 2019.

[15] H. Luo, W. Shao, F. Zhao, and A. Crivello, "Wi-Fi RTT based indoor positioning with dynamic weighted multidimensional scaling," in *International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, Pisa, Italy, Sept–Oct 2019.

[16] K. Jiokeng, G. JakHari, A. Tchana, and A.-L. Bylot, "When FTM discovered MUSIC: Accuracte WiFi-based ranging in the presence of multipath," *IEEE International Conference on Computer Communications*, August 2020.

[17] C. Gentner, M. Ulmschneider, I. Kuehner, and A. Dammann, "WiFi-RTT indoor positioning," in *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, April 2020. [Online]. Available: https://doi.org/10.1109/PLANS46316.2020.9110232

[18] M. Bullmann, T. Fetzer, F. Ebner, M. Ebner, F. Deinzer, and M. Grzegorzek, "Comparison of 2.4 GHz WiFi FTM- and RSSI-based indoor positioning methods in realistic scenarios," *Sensors*, vol. 20, p. 4515, August 2020.

[19] B. K. P. Horn, "Doubling the accuracy of indoor positioning: Frequency diversity," *Sensors*, vol. 20, no. 5, March 2020.

[20] ——, "Observation model for indoor positioning," *Sensors*, vol. 20, no. 14, July 2020.

[21] I. Martin-Escalona and E. Zola, "Ranging estimation error in WiFi devices running ieee 802.11mc," in *GLOBECOM 2020*, December 2020.

[22] M. Garcia-Fernandez, I. Hoyas-Ester, A. Lopez-Cruces, M. Siutkowska, and X. Banqué-Casanovas, "Accuracy in WiFi access point position estimation using round trip time," *Sensors*, vol. 21, p. 3828, June 2021.

[23] K. Han, S. M. Yu, S.-L. Kim, and S.-W. Ko, "Exploiting user mobility for WiFi RTT positioning: A geometric approach," *IEEE Internet of Things*, vol. 8, no. 19, October 2021.

[24] X. Liu, B. Zhou, P. Huang, W. Xue, and Q. Li, "Kalman filter-based data fusion of Wi-Fi RTT and PDR for indoor localization," *IEEE Sensors Journal*, January 2021.

[25] X. Guo, H. Wu, H. Yuan, and J. Song, "Impact of clock drift on WiFi round-trip-time ranging and positioning," *GPS solutions*, vol. 26, August 2022.

[26] S. Aggarwal, R. K. Sheshadri, K. Sundaresan, and D. Koutsonikolas, "Is WiFi 802.11mc fine time measurement ready for prime-time localization?" *ACM*, October 2022.

[27] B. K. P. Horn, "Indoor localization using uncooperative Wi-Fi access points," *Sensors*, vol. 22, no. 8, July 2022.

[28] D. Schepers and A. Ranganathan, "Privacy-preserving positioning in Wi-Fi fine timing measurement (FTM)," in *Proceedings on Privacy Enhancing Technologies*. Berlin, Germany: De Gruyter, April 2022, pp. 325–343.

[29] M. Orfanos, H. Perakis, V. Gikas, G. Retscher, T. M. andf Ioanna Syropoulou, and V. Papathanasopoulou, "Testing and evaluation of Wi-Fi ranging technology for personal mobility applications," *Sensors*, vol. 23, 2023.

[30] R. Jurdi, H. Chen, Y. Zhu, B. L. Ng, N. Dawar, C. Zhang, and J. K.-H. Han2, "WhereArtThou: A WiFi-RTT-based indoor positioning system," *IEEE Access*, January 2024.

[31] K.-M. Park, B. ho Lee, E. Lee, S.-C. Kim, and J. Choi, "GPS-Aided automatic site survey method for Wi-Fi RTT-Based positioning," *IEEE Transactions on Vehicular Technology*, vol. 72, no. 10, October 2023.

[32] *IEEE 802.11 - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Std., December 2016. [Online]. Available: https://standards.IEEE.org/standard/802_11-2016.html

[33] Google. (2022) WifiRttLocator app. [Online]. Available: https://play.google.com/store/apps/details?id=com.google.android.apps.location.rtt.wifirttlocator

[34] BKPH. (2022) FTMRTT app. [Online]. Available: https://play.google.com/store/apps/details?id=com.welwitschia.ftmrtt

[35] R. R. Murphy, *An Introduction to AI Robotics*. Cambridge, MA, USA: The MIT Press, 2000.

[36] E. G. Larsson, "Cramér–Rao bound analysis of distributed positioning in sensor networks," *IEEE Signal Processing Letters*, vol. 11, no. 3, 2004.

[37] J. Cole. (2014) Helmhurts. [Online]. Available: https://jasmcole.com/2014/08/25/

[38] R. Wilson. (2002, August) Propagation losses through common building materials—2.4 GHz vs. 5 GHz. [Online]. Available: http://www.am1.us/Papers/E10589PropagationLosses2and5GHz.pdf

[39] R. Rudd *et al.*, "Electrical properties of materials," in *Building Materials and Propagation Final Report - Ofcom*, R. Rudd, K. Craig, M. Ganley, and R. Hartless, Eds. Office of Communications, London, UK., September 2014, ch. 4, p. 31. [Online]. Available: http://www.qostic.org/Qostic/wp-content/uploads/Qostic6/AHQ-78-05-Building_Materials_and_Propagation.pdf

[40] ITU-R, "Effects of building materials and structures on radiowave propagation above about 100 MHz," July 2015. [Online]. Available: https://www.itu.int/dms_pubrec/itu-r/rec/p/R-REC-P.2040-1-201507-S!!PDF-E.pdf

[41] T. Koppel *et al.*, "Reflection and transmission properties of common construction materials at 2.4 GHz frequency," *Energy Procedia*, vol. 113, pp. 158–165, 2017.

[42] F. van Diggelen and P. Enge, "The World's first GPS MOOC and worldwide laboratory using smartphones,," in *Proceedings of the 28th International Technical Meeting of the Satellite Division of The Institute of Navigation*, Miami, Florida, September 2015. [Online]. Available: https://www.ion.org/publications/abstract.cfm?articleID=13079

[43] M. Thurber. (2024, March) GNSS jamming and spoofing events present a growing danger. [Online]. Available: https://www.ainonline.com/aviation-news/air-transport/2024-03-04/gnss-jamming-and-spoofing-events-present-growing-danger

[44] Google. (2024) RangingResult.getUnverifiedResponderLocation(). [Online]. Available: https://developer.android.com/reference/android/net/wifi/rtt/RangingResult#getUnverifiedResponderLocation()

[45] HPE/Aruba. (2024) Solution overview: Self-locating wireless access points. [Online]. Available: https://www.arubanetworks.com/assets/so/SO_Self-Locating-Wireless-Access-Points.pdf

[46] Cisco/Meraki. (2024) Set up access point auto location. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/wireless/spaces/config-guide/ciscospaces-configuration-guide/m-device-placement.html

[47] Juniper. (2024) Autoplacement: Position new access points. [Online]. Available: https://www.juniper.net/documentation/us/en/software/mist/mist-wireless/topics/task/floorplan-auto-placement.html

[48] Extreme Networks. (2024) AP location settings for fine timing measurement. [Online]. Available: https://documentation.extremenetworks.com/XIQC/10.10.01/UG/GUID-B01B9710-D7EC-4978-A0DE-FE5269CDC9A2.shtml

[49] ZEBRA. Zebra FTM Wi-Fi location: Integrator guide. [Online]. Available: https://www.zebra.com/content/dam/support-dam/en/documentation/unrestricted/guide/software/integrator-guide-ftm-wifi-location-en.pdf

[50] L. Banin, O. Bar-Shalom, N. Dvorecki, and Y. Amizur, "Scalable Wi-Fi client self-positioning using cooperative FTM-sensors," *IEEE Transactions on Instrumentation and Measurement*, December 2018.

[51] Google. (2022) Geolocation API overview. [Online]. Available: https://developers.google.com/maps/documentation/geolocation/overview

[52] P. D. Groves, *Principles of GNSS, Inertial and Multisensor Integrated Navigation Systems*. Norwood, MA, USA: Artech House, 2013.

[53] BKPH. (2021) WifiRttScanX app. [Online]. Available: https://play.google.com/store/apps/details?id=com.welwitschia.wifirttscanX

[54] X. W. Wang and F. Wang, "The precision of Google Earth map analysis with the coordinates of IGS stations," *ISPRS International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. XLII-3/W10, pp. 1053–1056, February 2020.

[55] B. K. P. Horn. (2023) Walking around location A. MP4. [Online]. Available: https://people.csail.mit.edu/bkph/movies/FTMRTT-Mauian-20240130-114220~3.mp4

[56] ——. (2024) Walking around location C-1 (pixel 5). MP4. [Online]. Available: https://people.csail.mit.edu/bkph/movies/FTMRTT_Kauian-20240225-092455~3.mp4
[57] ——. (2024) Walking around location C-2 (pixel 5). MP4. [Online]. Available: https://people.csail.mit.edu/bkph/movies/KSR_Pixel_5_screen-20240324-174934.mp4
[58] Circuit Design Inc. (2020) RF design guide. [Online]. Available: https://www.cdt21.com/design_guide/friis-equation-and-antenna-effective-area/
[59] R. Want. (2023, July) Next generation Wi-Fi positioning: An overview of the IEEE 802.11az protocol. Paul G. Allen School of Computer Science & Engineering. [Online]. Available: https://www.youtube.com/watch?v=sdpOYeaHNpw
[60] J. Segev and R. Want. (2023) Newly released IEEE 802.11az standard improving Wi-Fi location accuracy is set to unleash a new wave of innovation. [Online]. Available: https://standards.ieee.org/beyond-standards/newly-released-ieee-802-11az-standard-improving-wi-fi-location-accuracy-is-set-to-unleash-a-new-wave-of-innovation

Acknowledgment

I would like to thank Roy Want of Google for helpful feedback on a draft of this paper and discussions relating to FTM RTT and the IEEE 802.11 standards. I also wish to express my gratitude for the helpful suggestions of the reviewers.