Filtering Noisy 802.11 Time-of-Flight Ranging Measurements From Commoditized WiFi Radios

Maurizio Rea, Aymen Fakhreddine, Domenico Giustiniano, *Member, IEEE*, and Vincent Lenders, *Member, IEEE*

Abstract—Time-of-flight (ToF) echo techniques have been recently suggested for ranging mobile devices over WiFi radios. However, these techniques have yielded only moderate accuracy in indoor environments because WiFi ToF measurements suffer from extensive device-related noise which makes it challenging to differentiate between direct path from non-direct path signal components when estimating the ranges. Existing multipath mitigation techniques tend to fail at identifying the direct path when the device-related Gaussian noise is in the same order of magnitude, or larger than the multipath noise. In order to address this challenge, we propose a new method for filtering ranging measurements that is better suited for the inherent large noise as found in WiFi radios. Our technique combines statistical learning and robust statistics in a single filter. The filter is lightweight in the sense that it does not require specialized hardware, the intervention of the user, or cumbersome on-site manual calibration. This makes our method particularly suitable for indoor localization in large-scale deployments using existing legacy WiFi infrastructures. We evaluate our technique for indoor mobile tracking scenarios in multipath environments and, through extensive evaluations across four different testbeds covering areas up to $1000\,m^2$, the filter is able to achieve a median 2-D positioning error between 2 and 3.4 m.

Index Terms—Indoor localization, 802.11, time-of-flight, analysis, design, implementation, evaluation.

I. Introduction

LOCALIZATION using the Time-of-Flight (ToF) of RF signals is today the most popular technique to track moving objects in outdoor environments. The most prominent usage of ToF is the Global Positioning System (GPS) which exploits signal propagation times from different satellites to localize mobile devices on earth. Also, radar systems commonly rely on the propagation time of RF signals to localize aircrafts in the airspace. While the ToF technique has been widely adopted in outdoor environments, its application for indoor localization in WiFi environments has been relatively modest so far. When it comes to WiFi based localization, the research community has instead focused more intensively on different approaches such as the signal strength [1]–[5], the angle of arrival [6]–[9] or combining WiFi signals with inertial sensors as found in smartphones [10], [11]. The problem with the latter approaches is that they either require extensive manual on-site calibration, the need for intervention of the user, or specialized hardware. These factors limit the deployment of these approaches at larger scale.

The challenge with ToF is that very precise signal propagation time estimates are required. At the speed of light, a measurement error of $1\,\mu s$ already results in a distance estimation error of 300 meters which is intolerable for most indoor applications. In order to achieve meter-level localization accuracy, a precision in the order of a few nanoseconds is therefore needed. However, at this level the ToF is affected by various sources of noise [12]–[17]. For example, the relatively small bandwidth of WiFi signals and the limited clock rate of commercial off-the-shelf (COTS) WiFi radio receivers hinder the exact determination of the time of arrival of a signal. In echo techniques, the target may further add significant jitter before acknowledging the reception of a data. Specially indoors, the signal may additionally be obstructed and reflected over multiple paths which adds a positive bias to the estimated range.

In order to mitigate the impact of noise in ToF measurements, several techniques have been proposed in the literature. To combat the device-related Gaussian noise, unbiased estimators that rely on the mean or median of the collected ToF samples provide a good estimate of the true ToF [17]. To combat the multipath noise, biased estimators that aim at identifying the direct path such as the MUSIC algorithm [18], [19] can provide good results as well. Other approaches such as the expectation-maximization algorithm have been successfully applied to signal strength-based localization [4], and we could envision their application to combat the multipath noise of ToF measurements. However, we show in this work that, in WiFi ToF, the noise follows a Gaussian mixture model with each Gaussian component being in the same order of magnitude or larger than the multipath bias, and we demonstrate that these estimators fail to provide an accurate estimate of the true ToF as we need for ranging.

In this paper, we present a new filter that is specially designed for estimating the range in WiFi indoor environments based on noisy ToF measurements. Our filter relies on a combination of statistical learning techniques to train an environmental-specific linear regression model for the
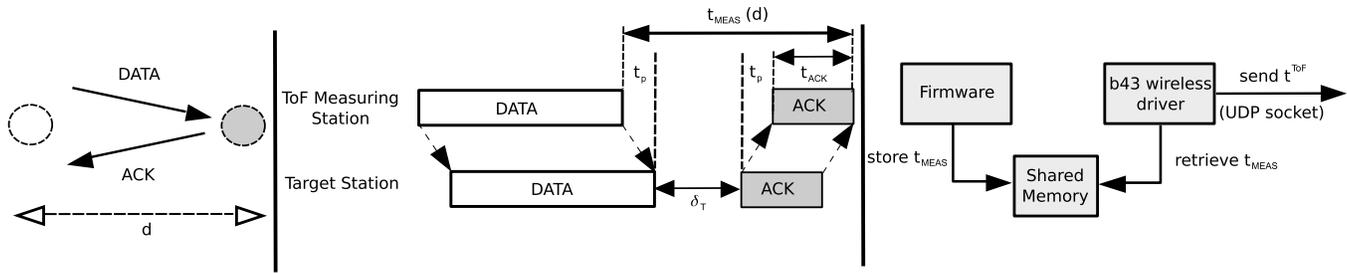
Fig. 1. WiFi ToF echo technique. Offset δ_T is originated at the target device, and it is a function of the hardware delay to schedule the ACK at the target device. The time offset δ_L (not shown in the figure) is due to the measuring device. On the right, the figure shows the interaction between firmware, driver and user space. Measured ToF data by the firmware is transferred to the driver through the shared memory (SHM). Buffered ToF measurements are then sent to user space.

multipath bias and robust statistics for range estimation. Our approach does not require any specialized hardware and no intervention from the user, and it relies exclusively on the ToF information that can be extracted from COTS WiFi radios.

Our solution can independently run on the WiFi infrastructure side and it has the advantage of disentangling the core of the positioning tracking from the mobile application where additional services and higher accuracy could be provided only if the user is interested to install the application. In addition, the need for intervention of the user may be unwanted in setups where the objective of positioning the user is to understand and better plan the physical space as input for data analytics. Finally, our environmental-specific linear regression model can be learned in-situ and thus does not require any manual calibration.

We have used our filter to estimate the location of WiFi mobile devices in four different experimental testbeds. Our results show that our technique is able to significantly reduce the median ranging error over classical estimators. Our main contributions in this paper are the following:

- We present a firmware-based ranging architecture for round-trip time measurements running in the MAC processor of WiFi chipsets, and we quantify and compare the amount of noise that comes from the WiFi devices and the noise from the WiFi radio signal propagation over the wireless channel.
- We propose a filter based on statistical learning and robust statistics to estimate the distance range from a series of noisy ToF measurements. Our filter does not require any manual calibration and manages to estimate the distance to a remote WiFi device with median error between 1.7 and 2.4 meters.
- Finally, we evaluate our system for indoor mobile device tracking. We demonstrate across four different experimental testbeds that mobile devices associated to the WiFi infrastructure can be tracked continuously using our filter with a median accuracy between 2.0 and 3.4 meters.

II. TOF ECHO TECHNIQUE

This section describes the principles of ToF ranging and the noise effects that arise when applying this technique to off-the-shelf WiFi devices.

While traditional ToF-based echo techniques as employed in radar systems rely on uncoded RF signals and their reflections,

WiFi echo techniques use regular frames of communication [12]–[17]. In WiFi communication, every DATA frame is acknowledged by the receiver with an ACK frame. In this work, we denote as local or measuring station the node sending DATA and receiving ACK, and target station the node receiving DATA and sending ACK. Since the interframe time between the DATA and ACK frames is fixed by the 802.11 standard (the Short InterFrame Symbol, SIFS, time), the delay between DATA and ACK frames can be used by the local station to infer the distance to the target.

In reality, sources of time offset exist at the local and target stations and they affect the accuracy of the ranging measurement. Let us refer to Fig. 1. If d is the true distance between a local station and a target, the local station measures the time $t_{MEAS}(d)$ between the end of a sent DATA frame and the end of a received ACK frame:

$$t_{MEAS}(d) = 2 \cdot t_P(d) + t_{ACK} + \delta, \qquad (1)$$

where $t_P(d)$ is the signal propagation time between the local station and the target (channel reciprocity is assumed), t_{ACK} is the time needed to transmit the ACK, and δ is the measurement offset. The offset δ can be expressed as:

$$\delta = \delta_L + \delta_T, \qquad (2)$$

where δ_L is the offset at the local station and δ_T the one at the target. δ_L arises at the measuring device due to the local imprecision in the timing information extraction, and δ_T is equal to the SIFS time plus any device-dependent deviation from this number. The presence of multipath effect causes a positive offset and increases δ_L and/or δ_T.

From eq. (1), we can then define the ToF measurement as:

$$t^{ToF}(d) = \frac{t_{MEAS}(d) - t_{ACK} - \delta}{2}. \qquad (3)$$

At distance $d = 0$, we then have:

$$t^{ToF}(d = 0) = \frac{t_{MEAS}(d = 0) - t_{ACK} - \delta}{2} = 0. \qquad (4)$$

And therefore eq. (3) can be rewritten as:

$$t^{ToF}(d) = \frac{t_{MEAS}(d) - t_{MEAS}(d = 0)}{2}. \qquad (5)$$

A sample is collected per 802.11 DATA/ACK handshake. For the generic sample m, the distance from the measuring station to the target device can be computed as:

$$d_m = c \cdot t_m^{ToF}(d), \qquad (6)$$

where c is the speed of signal propagation which is close to the speed of light in air. In the more general case, let $\mathcal{L}$ denote the set of links to a target station and $\mathcal{M}$ the set of samples. We can then express the generic sample as:

$$d_{i,m} i \in \mathcal{L}, \quad m \in \mathcal{M}.$$

Once M samples have been collected for a link $i \in \mathcal{L}$, the distance $\widehat{d}_i$ between the local station and the target is estimated as follows:

$$\widehat{d}_i = f(d_{i,1}, d_{i,2}, \ldots, d_{i,M}), \quad i \in \mathcal{L} \qquad (7)$$

where f is an estimator of the distance that aims at filtering out the measurement noises. The major contribution of this work is *the design, implementation and evaluation of a robust estimator f that estimates the distance to a target device using WiFi ToF, in the presence of rich multipath as found in common indoor environments, and severe noise as found in commodity WiFi chipsets.*

III. ToF on COTS WiFi Radios

The design of a robust estimator f requires first to characterize the noise sources of WiFi ToF. For this reason, we implement a ToF ranging system using commodity hardware. To alleviate any unnecessary source of noise or instability from the operating system, ToF measurements have to be performed as close as possible to the radio hardware. Rather than in the driver [17] or upper layers [14], the best method is therefore implementing the code for ToF measurements in the firmware of the WiFi radio chipset. To measure $t_{MEAS}(d)$ in the firmware, we have customized the open-source 802.11 openFWWF firmware.[1] This firmware is written in assembler and runs on off-the-shelf 802.11 Broadcom chipsets. Our customized firmware reports $t_{MEAS}(d)$ for each successful DATA-ACK frame pair. The timing is regulated by the general purpose timer, running on the wireless card's internal clock at a rate of 88 MHz. The timer starts to count clock cycles just after the 802.11 processor sets up a register to indicate that a frame has been sent. Once the ACK frame has been received (or the ACK timeout has elapsed), another register gets updated and the timer gets stopped. Every time a measurement is made, the firmware writes $t_{MEAS}(d)$ into a defined address of the shared memory (SHM). The architecture is shown on the right of Fig. 1. Since the driver has also access to the shared memory block, it can retrieve the measurement every time an ACK is received. In the driver, we gather additional data about the incoming ACK such as the data rate, the AP (Access Point) MAC address, etc, and store them all in a buffer. Once this buffer is full or a timeout has elapsed, the data is transferred to the user space with a UDP socket.

A. System Setup

We have built a prototype ToF system that consists of COTS APs operating as ToF measuring stations. In our deployments, the APs are static stations. The APs use Soekris

net5501 embedded machine with a 500 MHz AMD Geode LX single chip processor. The APs are equipped with Broadcom AirForce54G 4318 mini PCI type III cards and an omnidirectional antenna. The Broadcom chipset is operated with our customized firmware and the b43 driver presented above. The APs are connected over Ethernet to the location computing unit, that is responsible to process the raw data and compute the position.

As target station we use unmodified Dell Inspiron 5150 laptops equipped with Broadcom AirForce54G 4318 mini PCI type III cards and the integrated antenna of the laptop. In our deployments, the target station may be static or mobile according to the type of experiment (see Section III-B on deployment scenarios for the details). In order to perform ranging measurements, the APs use regular DATA frames that are acknowledged by ACK frames from the targets. The ranging measurements are performed in a round-robin fashion among the APs. The target is associated to a single AP, and the other APs send their DATA frames with the source MAC address of the AP to which the target is associated. In order to configure the DATA transmission, we rely on raw sockets and PCAP library.[2] This approach does not require any specific feature of the Broadcom chipsets used in our deployment, and thus it can be extended to other 802.11 firmwares.

B. Deployment Scenarios

We consider two types of scenarios for the deployment. The first scenario considers controlled experiments, where we avoid any environmental effects by connecting two stations using coaxial cables. The cables are based on the standard RG-58. In the second scenario, we perform experiments with transmissions over the air in four indoor testbeds, namely Testbed I, II, III and IV. The maps are shown in Fig. 2. Testbed IV uses the same environment as in Testbed I, but APs are partially placed at different locations. We deploy 9 APs in Testbed I, 9 APs in Testbed II, 10 APs in Testbed III, and 9 APs in Testbed IV. Testbed I, II and IV are office environments.

The environment of Testbed I is shown in Fig. 2(a). It covers a surface of almost $1000 \, \text{m}^2$. We use 25 randomly selected locations (marked as a cross) to test our algorithms. We conduct tests over two different days, with some positions repeated again with different locations of some furniture. We could achieve connectivity for a subset of the total set of links (207 links out of $25 \cdot 9$).

Testbed II is depicted in Fig. 2(b). It features 180 links and it covers a smaller space of around $200 \, \text{m}^2$. The target station is placed in 20 different positions.

Testbed III is shown in Fig. 2(c). It has been deployed at the facilities of the IEEE/ACM IPSN 2014 - Microsoft Indoor Localization Competition [20]. The testbed has 200 links and it covers $320 \, \text{m}^2$. The target device is placed at 20 positions in two rooms and a hallway.

Testbed IV is deployed to study the performance of mobile device tracking. In this testbed, we measure the position when a mobile device is moving from position 1 to 19 as marked

---

[1] http://www.ing.unibs.it/openfwwf/

[2] http://www.tcpdump.org/

Fig. 2. Testbeds to assess the ranging and positioning capabilities of our system in static conditions. Circles correspond to AP locations and crosses to target locations. (a) Testbed I. (b) Testbed II. (c) Testbed III. (d) Testbed IV (mobile tracking).

in Fig. 2(d). For these mobility tests, a user moves at a speed of approximately 0.4 m/s. At the time that the user passes at one of the marked positions, we estimate the position and record the value. We then repeat the tests in the following setups: when the user stops at the marked locations for 2 and 5 seconds, respectively.

In all the testbeds, a mixture of line-of-sight and non-line-of-sight wireless links are present. During all experiments, people are moving within the testbed areas. The testbeds also contain several obstructions, including concrete walls, tables and glasses. All experiments are conducted with other active WiFi networks in the neighborhood. We operate the testbeds on a fixed frequency channel of the 2.4 GHz ISM band. The physical (PHY) automatic selection rate is active, such that the measurements include DATA sent at different PHY rates.

IV. SOURCES OF NOISE

While the firmware-based approach introduced in Section III can allow us to measure the ToF with the best precision using commoditized WiFi radios, the ToF measurements are still affected by large noise coming from severe sources which we discuss in the following in more details. In order to collect samples $\{d_{i,m}\}$ for this analysis, we first determine the reference value $t_{MEAS}(d = 0)$ (cf. eq. (5) and eq. (6)). To this end, we directly connect the AP operating as ToF measuring station to the target device to be calibrated and perform reference measurements over three coaxial cables of different lengths. We then infer $t_{MEAS}(d = 0)$ for the distance at 0 m by means of linear regression. This process is required only once per chipset.
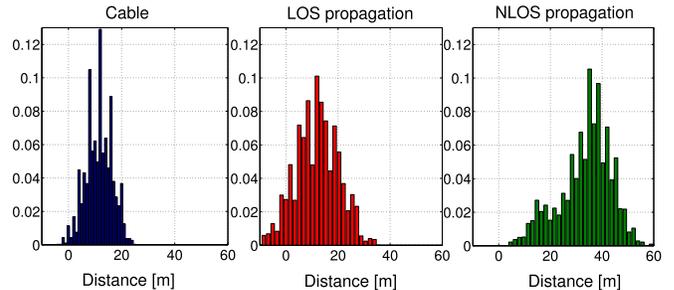

Fig. 3. Noise introduced by ToF ranging. Tests in a controlled environment (Cable) show that there are sources of large dispersion in the estimation. Tests over the air (LOS and NLOS propagation) show that the distribution greatly depends on the channel conditions. All the tests are in addition subjected to quantization noise and other spurious noise sources.

Target ACK Delay: The 802.11 standard specifies the SIFS time between the reception of a DATA and the transmission of an ACK at the receiver as a fixed interval. In 802.11b, for example, this time is specified as $10\,\mu s$ [21]. However a relatively high tolerance of 1 μs is tolerated which can result in significant noise and distance estimation errors up to 300 meters if the target would fully exploit this specified tolerance level. While most chipsets may not fully exploit this tolerance, the dispersion is still quite significant. To illustrate this, Fig. 3 on the left represents the resulting dispersion of a typical Broadcom WiFi chipset. The shown histogram was obtained by collecting sequences of samples $\{d_{i,m}\}$ according to Eq. (6) for $10,000$ packets. To avoid any dispersion from environmental effects, the measurements were performed over a coaxial cable of 13.5 meters. We observe that there are sources of noise that can lead to distance estimations that

range from 0 to 25 meters, even under ideal signal propagation conditions such as cables.

Multipath Reflections: It is well known that signal propagation in complex indoor environments is subject to multipath effects in which multiple copies of the transmitted signal arrive at the receiver over different reflected paths. It is even possible that the direct component is entirely attenuated and the signal is received only over indirect paths. Since signals that travel over indirect paths will take longer time to arrive at the receiver, they introduce a non-Gaussian error in the distance estimation when considering the time-of-flight. This situation is shown in Fig. 3 in the middle and on the right where the same experiment as on the left was repeated but for a line-of-sight (LOS) and non-line-of-sight (NLOS) signal propagation link over omnidirectional antennas. In the LOS experiment, there is a visible connection between the measuring station and the target, while in the NLOS experiment, they are obstructed. The dispersion spans a range of 40 and 60 meters for the LOS and NLOS links, respectively, and the large noise in the measurement can also result in negative $d_{i,m}$ samples (see Fig. 3 in the middle), especially when the true distance d_i is relatively small. In addition, the NLOS link shows a non-Gaussian distribution. Multipath effects must therefore be taken into consideration in order not to overestimate the distance when dealing with reflected signal propagation paths. Finally, multipath may also happen in LOS links, and thus a method robust to the propagation conditions must be designed.

Quantization and Measurement Uncertainty: Off-the-shelf WiFi chipsets have not been designed to provide accurate ToF measurements. A main source of noise comes from the coarse clock resolution of the radios. For example, the Broadcom chipset operates with a reference clock of $88\,\mathrm{MHz}$, corresponding to a maximal distance resolution of 1.7 meters. In addition to this quantization noise, off-the-shelf chipsets introduce all sorts of considerable additional noise. As we could see in the histograms of Fig. 3, the shape of the distribution is far from being smooth despite using $10,000$ samples to create the histograms, suggesting that the radios must have some bias when measuring the time. Other sources of noise must therefore also be factored in order to estimate the distance. On the other hand, the effects of the clock drift are negligible. See details in [17].

Congestion and Interference: Wireless congestion and interference have no impact on the ACK delay (that is, δ_T) since the 802.11 channel is reserved during the SIFS period to the receiving node (i.e. the target station) as dictated by the 802.11 carrier sense multiple access with collision avoidance (CSMA/CA) protocol [21]. However, collisions occur on the wireless medium, causing data retransmissions. Only acknowledged data frames are considered as valid ToF samples, while unacknowledged frames do not generate any ToF measurement. Since the wireless resources are shared, the presence of congestion can increase the time required to obtain a sufficient number of samples for ranging. We study the problem of wireless congestion in an open space room (to reduce as much as possible the multipath, and focus on the impact of interference), with one link of 19 meters in LOS. In this setup, the AP transmits DATA at PHY rate of $1\,\mathrm{Mb/s}$.
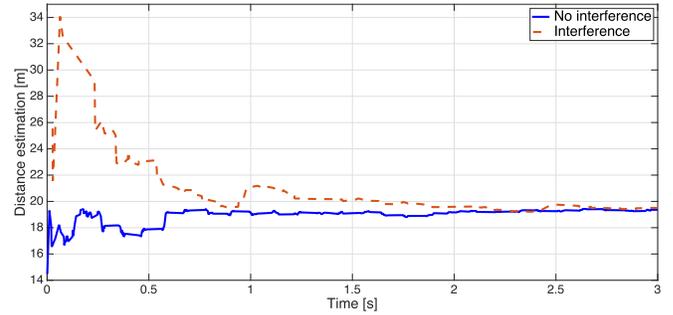

Fig. 4.   ToF distance estimation in a LOS link of 19 meters in absence and presence of interfering traffic. In the experiment, the ToF measuring station sends DATA at PHY rate of $1\,\mathrm{Mb/s}$.

The AP computes the distance averaging over the collected ToF samples. We then repeat the ranging measurement for the same link, adding 802.11 traffic from another wireless station that saturates the channel with interfering traffic (UDP traffic of $4\,\mathrm{Mb/s}$ sent at PHY rate of $1\,\mathrm{Mb/s}$). The results are presented in Fig. 4. The figure shows the average of the distance estimation, both in absence and presence of interference. While the long-term ranging accuracy is the same in absence of interference and in presence of interference saturating the channel, the latter causes a longer time to converge to the true distance d_i. We can conclude that, in order to efficiently handle interference, it is important to design a system that requires only a few samples M for the ranging estimation.

In Section V we provide an analysis about the target and measuring noise. We then describe how to deal with the Non-Gaussian noise resulting from the multipath propagation in Section VI to finally describe in details the robust shortest-path estimator we propose as a solution to mitigate the effects of the main sources of noise in Section VII.

V. DEVICE-RELATED NOISE ANALYSIS

In this section, we analyze in details the noise originated at the measuring station and at the target station. The goal of this analysis is to quantify the device-related noise and determine whether the non-Gaussian noise we observe in ToF measurements is only due to environmental effects such as multipath or the local and target devices also add non-Gaussian noise to the measured values.

In order to address this question, we use a low-noise oscilloscope to measure the offset distribution of $\{\delta_T\}$. This high-end wideband oscilloscope is an Infiniium 90000A model with a fast sampling rate of $10\,\mathrm{GS/s}$ [22]. The internal noise of the oscilloscope can be regarded as negligible compared to the noise introduced by the WiFi radios and therefore provides us a mean to analyze the target offset δ_T in isolation.

On the left of Fig. 5, we show the histogram of the target offset δ_T as measured with the oscilloscope using approximately 300 samples, and the resulting Gaussian fitting function. We run the Lilliefors test and find that the hypothesis that the distribution of δ_T is Gaussian *cannot be rejected*, with a high p-value equal to 0.43 while setting the significant level to 0.05. Therefore it is safe to assume that δ_T can be
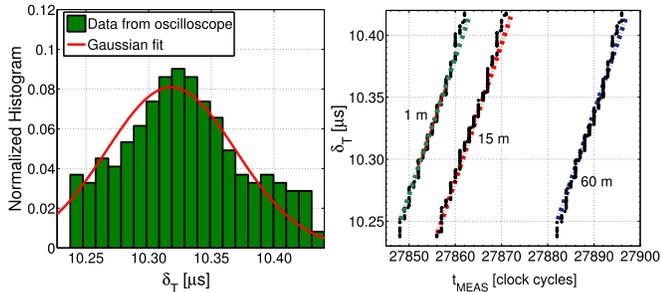
Fig. 5. On the left: distribution of $\{\delta_T\}$, as measured with our oscilloscope, and its Gaussian fitting function. On the right: QQ plots of δ_T versus $t_{MEAS}(d)$.

approximated with a Gaussian distribution $\mathcal{N}$ with standard deviation equal to σ_{δ_T}.

We then statistically compare the distribution of $\{\delta_T\}$ measured with the oscilloscope with the distribution $\{t_{MEAS}(d)\}$ measured at the local station with our firmware-based approach presented in Section III. For the comparison, we collect samples of $t_{MEAS}(d)$ in LOS links with limited multipath at $d = \{1, 15, 60\}$ m and draw the quantile-quantile (QQ) plots, a graphical method for comparing two probability distributions. If the distributions are linearly related, the points in the QQ plot will approximately follow a line.

The resulting QQ plots are shown on the right of Fig. 5. We observe a linear pattern, which indicates that $\{\delta_T\}$ *and* $\{t_{MEAS}(d)\}$ *have very similar noise distributions in presence of limited multipath*. We also observe that the dispersion of the two distributions is very similar, with a standard deviation of $4.0 - 4.1$ WiFi clock cycles in both cases. As a result of these investigations, we conclude that:

- The noise of $t_{MEAS}(d)$ in LOS links with limited multipath is largely dominated by the Gaussian noise of the target offset δ_T.
- The dispersion of the local offset δ_L in LOS links has a negligible impact on the distribution of $t_{MEAS}(d)$, and our approach to implement the ToF measurement in the firmware does not add a significant dispersion.
- The non-Gaussian ToF noise that we observe in many indoor links of our testbed is not related to the local or target noise of the WiFi devices but rather to environmental effects such as multipath.

VI. DEALING WITH WiFi MULTIPATH NOISE

Existing multipath-resistant estimators assume that the non-Gaussian noise from the reflected signal path components are the dominant sources of noise. However, as we have shown in the previous section, the device-related Gaussian noise is extensive in WiFi radios. In order to deal with these combined sources of noise, the focus of this section is to analyze different robust estimators f to estimate the true distance d for multipath-rich indoor environments.

A. Noise Model

While the device-related noise can be approximated as an environmental-independent Gaussian distribution, this is not the case for the noise that arises from the effect of multipath

reflections. The amount of multipath is very much dependent on the environment and the exact position of the devices. Under multipath propagation between two WiFi nodes, the distribution of ToF measurements will be non-Gaussian, as a result of those data frames where the 802.11 transceiver synchronizes to a delayed copy of the signal.

Estimating the true distance using this noisy data is not trivial with WiFi radios. In fact, we cannot directly measure the individual multipath components at the signal-level on off-the-shelf WiFi radios since the WiFi chipset only synchronizes to the strongest path, resulting in a positive or null bias $b_{i,m} \geq 0$.

In order to design a reliable estimator f of the distance, we model the ToF noise as follows. Let us consider a link $i \in \mathcal{L}$. The true distance d_i of this link is affected by an additive Gaussian noise $\mathcal{N}$, generated by the target station, with the standard deviation equal to σ_{δ_T} (cf. Section V). d_i is further affected by a distance bias $b_{i,m} \geq 0$ caused by the absence ($b_{i,m} = 0$) or presence ($b_{i,m} > 0$) of multipath. Therefore, for each sample, we can model the distance $d_{i,m}$ as follows:

$$d_{i,m} = d_i + b_{i,m} + \mathcal{N}(0, \sigma_{\delta_T}), \qquad (8)$$

where d_i is the true distance.

Grouping together the samples with the same bias after clock quantization (that limits the distance resolution per sample to 1.7 m, cf. Section III) results in a finite Gaussian Mixture Model (GMM) with a small number of modes. One of these modes corresponds to the samples synchronized to the direct path (the samples with $b_{i,m} = 0$), while the others correspond to the samples synchronized to any of the reflected paths. In general, no component is dominant and we can conclude that the Gaussian mixture is non-Gaussian distributed.

B. Reliability of the Estimation of the Gaussian Components

A fundamental difference with respect to classical GMM applications is that the multipath noise superimposes to the large noise σ_{δ_T} added by the target device. The deviation due to σ_{δ_T} can be much stronger than the bias error caused by synchronizing to a delayed copy in a multipath propagation.[3] Current techniques for learning the modes of the GMM model need instead that the modes are either sufficiently apart or consider that only a small number of samples is received as reflected paths representing a few outliers [23].

In order to verify how well current learning mechanisms for a GMM model could be applied to WiFi ToF, we generate in simulations $1,500$ samples using the GMM model and suppose that there exist only two modes, one mode representing the direct path and one mode representing the reflected path. We then apply the expectation–maximization (EM) algorithm [24] to cluster the samples (the details of the algorithm are presented in Appendix XII-A of the supplementary

---

[3]For instance, in the example in Fig. 5 (left), we have a standard deviation of 4 clock cycles for δ_T in LOS conditions. This results in a 99-percentile dispersion of the error of approximately 35.1 m, higher than a typical error caused by synchronizing to a delayed copy of the 802.11 signal in a multipath propagation in indoor scenario.
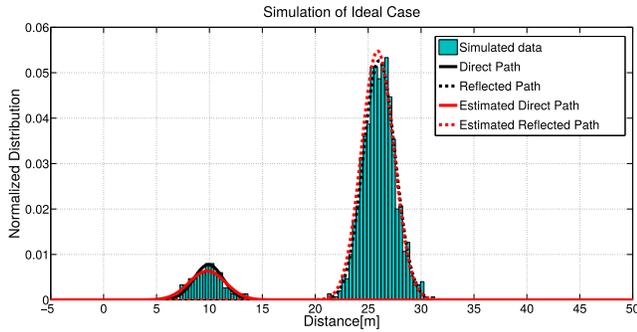
Fig. 6.　Generation of a direct path and a reflected path in simulation using the GMM model, with small Gaussian noise of the target station. As expected, the EM algorithm can reliably estimate the two modes.

Fig. 8.　Same simulation data set as in Fig. 7: We conjecture that, by taking a percentile below 50%, we can counterbalance the biased values $b_{i,m} > 0$ in the estimation process and estimate the mean of the direct path's Gaussian distribution.
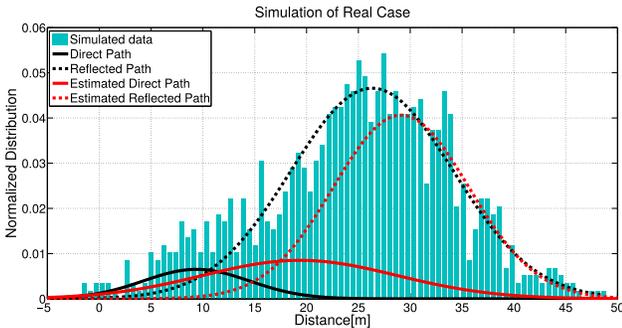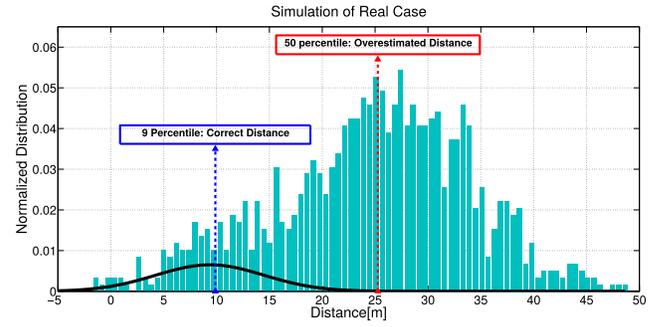

Fig. 7.　Generation of a direct path and a reflected path in simulation using the GMM model, with Gaussian noise of the target station as found in our experiments. The EM algorithm fails to reliably estimate the two modes.

material). In this study, we consider that the direct path has a distance of $10\,$m and that the reflected path has a traveled distance of $26\,$m. The mixture weights of the Gaussian components are equal to 0.12 and 0.88, respectively. This indicates that the 802.11 WiFi chipset synchronizes more often to the reflected path, which is in general the stronger signal component.

We first make the (ideal) assumption that δ_T is much smaller than in real measurements. We call it "Simulation of Ideal Case". In the second scenario, we consider that δ_T is equal to the one empirically measured in Section V. We call it "Simulation of Real Case". We plot the results of "Simulation of Ideal Case" in Fig. 6. We observe that the EM algorithm estimates the parameters of the components of the Gaussian mixture almost perfectly: the estimated mean of the direct path's distribution is equal to $9.8\,$m and the mean of the reflected path's distribution is $25.9\,$m. It follows that we are able to estimate the distance of the direct path with an error of only $0.2\,$m. We then plot the results of "Simulation of Real Case" in Fig. 7. Clearly visible from the figure, we do not have anymore a clear separation of the two Gaussian components. The parameters are here overestimated: the estimated mean of the direct path's distribution is $19.3\,$m and the mean of the reflected path's distribution is $29.1\,$m. Concluding, we have a high error in the estimation of the direct (and shortest) path, used for ranging, of $9.3\,$m.

VII. Robust Shortest Path Estimator

The results of the previous section have shown that classical algorithms fail to estimate the distance in WiFi echo techniques. In this section, we explore a different methodology for the estimation of the distance. First of all, in WiFi ToF, we are only interested in the direct path's distance or, in general, the shortest path's distance. Indeed, we do not need an algorithm (as the EM algorithm) that estimates the distance of all the paths, including the longer ones. In GMM links with only a direct path ($\forall m \in \mathcal{M}, b_{i,m} = 0$), the median or the mean would be a reliable estimator of the direct path's distance. However, as a result of the GMM model in WiFi ToF, the sequence is in general mixed with samples where the local station synchronizes to a delayed copy of the WiFi signal ($\exists\ m \in \mathcal{M} : b_{i,m} > 0$). In the latter case, using the median of the sequence, e.g. the 50% percentile of the distribution as estimator would result in an over-estimation of the distance. *We then conjecture that, by taking a percentile below 50%, we can counterbalance the biased values $b_{i,m} > 0$ in the estimation process.* In the example shown in Fig. 8, a percentile equal to 9 would be effective to counterbalance the biased value and report the true distance.

Formally, let us consider a link $i \in \mathcal{L}$ and collect a sequence of M measurements $\{d_{i,1}, d_{i,2}, \ldots, d_{i,M}\}$. We define the optimal percentile p_i^{opt} as the percentile that provides, for each link i, the minimum absolute distance estimation error with respect to the true distance d_i:

$$p_i^{opt} = \operatorname*{argmin}_{0 \le p \le 0.5} |d_i - \widehat{d_i(p)}|, \qquad (9)$$

where $\widehat{d_i(p)}$ is the estimated distance using the p-percentile of the distribution $\{d_{i,m}\}$. The goal is to design a statistical estimator f that estimates the optimal percentile p_i^{opt} based on some observables and gives as output the estimated distance $\widehat{d_i(p_i^{opt})}$.

We train various options for the observables in the estimator f in an extensive evaluation in one of our testbeds (Testbed I, Fig. 2(a)) with all links. As candidate observables, we consider the first three moments (median, standard deviation, and skewness) of the ToF as well as of the received signal strength indicator (RSSI). All three moments could in principle be indicators of multipath. For example, when the median of the ToF is low (or the median of the RSSI is high), the two ranging devices are close to each other and hence likely to have a short line-of-sight connection between each other with little multipath delay. An increased standard

TABLE I

THE ABSOLUTE VALUE OF PEARSON CORRELATION COEFFICIENT ρ BETWEEN DIFFERENT MOMENTS OF THE RSSI AND ToF VERSUS THE OPTIMAL PERCENTILE p_{opt} (0=NO CORRELATION, 1=MAXIMAL CORRELATION). IT IS WORTH NOTING THAT THE RANGING IS ALWAYS PERFORMED USING ToF MEASUREMENTS. FOR EXAMPLE, THE USE OF MOMENTS OF RSSI IS LIMITED TO INFER THE OPTIMAL PERCENTILE, IT SHOULD NOT BE CONFUSED WITH RSSI-BASED RANGING

	unfiltered	pre-filtered
median of RSSI	0.62	0.63
standard deviation of RSSI	0.04	0.05
skewness of RSSI	0.23	0.24
median of ToF	0.76	0.76
standard deviation of ToF	0.19	0.21
skewness of ToF	0.20	0.51

deviation of the ToF or the RSSI may indicate that signals are received over several paths with different propagation delays and/or attenuations. The skewness of the ToF and RSSI distributions will also be intuitively larger over links with multiple propagation paths. For each of these six observables, we evaluate two variants, leading to a total of twelve candidate estimators. In the first variant, we determine the moments directly on the raw samples $\{d_{i,1}, d_{i,2}, \ldots, d_{i,M}\}$. In the second variant, we attempt to pre-filter obvious outliers that arise from the device-related measurement noise (cf. Section IV) prior to determining the moments. These outliers are filtered out applying the Thompson Tau technique [25], a statistical method for deciding whether to keep or discard samples based on the expected value and the expected deviation of the sequence of samples.

We evaluate the precision of these estimators by determining their correlation to the a priori known optimal percentile $p_i{}^{opt}$. For this analysis, we compute $p_i{}^{opt}$ for a set of links with known distance, i.e. placing the mobile device at known positions. To quantify the correlation between the different moments and $p_i{}^{opt}$, we rely on the Pearson correlation coefficient [26]. The Pearson correlation coefficient ρ is an indicator of the linear correlation of the variables, where absolute values close to zero indicate a low correlation and absolute values close to one represent a high linear dependence of two variables. A value close to one thus indicates that a moment is a good estimator to predict the percentile that will filter out the multipath noise effectively.

We consider the entire set of links in one of our testbeds (Testbed I). Table I shows the resulting correlation coefficient ρ for all twelve variants. *The best correlation is provided by the median of the ToF*, followed by the median of the RSSI and the skewness of the ToF. All other moments have a correlation coefficient below 0.5 which indicates a low correlation. We note that not all moments profit from pre-filtering to remove the outliers. For example, the median of the RSSI and ToF perform worse when outliers are pre-filtered. On the other hand, the skewness of the ToF increases from 0.20 to 0.51 and is therefore considerably better when pre-filtering the outliers.
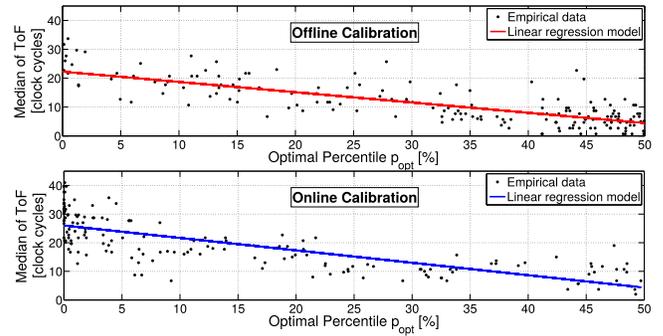

Fig. 9. Median ToF versus optimal percentile p_{opt} for all links of Testbed I - Linear regression using offline and online calibration

One may ask why the skewness of the ToF has a worse correlation than the median ToF. Intuitively, the skewness of the distribution should be a good indicator of the multipath, given that links with strong reflected (delayed) components are left-skewed, with p_{opt} smaller than for right-skewed link. Our results suggest that the combined device-related noise of the receiver and the measuring station have a strong negative effect on the correlation on the skewness. This is reflected in the pre-filtered version of the skewness which has a considerably better correlation than the unfiltered version. In contrast, the median of the ToF is much more robust to any device-related noise and therefore outperforms the skewness. The reason for the median ToF working best can be associated to the tendency of having longer reflected paths (and thus longer delays) for links with longer distances and vice versa. In other words, when the ToF is small, the devices are close to each other and the multipath noise will likely not affect much the ToF measurements. On the other hand, when the median ToF is large, the devices are further apart and the links may be affected more severely by the noise from reflected paths.

A. Proposed Filter Design

Provided the good correlation ρ between the median ToF and the optimal percentile $p_i{}^{opt}$, we design a linear model for the estimator f that relies on this correlation to estimate the percentile from ToF measurements. The model is specific to the environment and we therefore train a model for each testbed. The training works as follows. For each environment, we compute the median ToF and the corresponding $p_i{}^{opt}$ as in the tests in Section VII. Note that this requires extensive training at many locations, but we will show how to train a model without manual efforts in the next section by training using data collected only between the APs. The empirical distribution of the median ToF versus $p_i{}^{opt}$ for all the 207 links of Testbed I is shown in the top of Fig. 9.

We note that $p_i{}^{opt}$ is widely distributed between 0 and 50%. Therefore, it does not exist one value of percentile p that it is optimal for all the links, but it rather changes from link to link. Nevertheless, there is a clear linear trend which we will exploit in our estimator f. To get the model, we perform a linear regression on these data points and obtain an environment-dependent linear model, as represented by the continuous line in the figure. Formally, let $\widehat{p_i{}^{opt}}$ be the estimated optimal
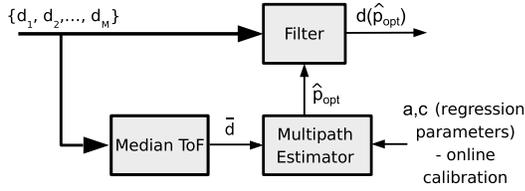
Fig. 10. A sequence of ToF measurements $\{d_1, d_2, \ldots, d_M\}$ are filtered by selecting a percentile p of the distribution according to the median of the measurements $\bar{d}$ which serves as an estimate of the amount of multipath on the link. The multipath estimator is trained using a linear regression model of the median ToF versus the optimal percentile that minimizes the error from ToF measurements between the APs.

percentile p_i^{opt} for a link $i \in \mathcal{L}$. $\widehat{p_i^{opt}}$ is computed with the following linear regression equation:

$$\bar{d}_i = a \cdot \widehat{p_i^{opt}} + c \qquad a \leq 0, c \geq 0 \qquad (10)$$

where $\bar{d}_i = d_i(\widehat{p = 0.5})$ is the median ToF (estimated distance using the median), and the regression parameters a and c are trained by environment calibration on a given set of known distances, minimizing the sum of squared residuals. Applying classical linear regression theory [27], we can then state that a ρ^2 ratio of the total variation of p_i^{opt} can be explained by looking at the median ToF $\bar{d}_i$. Small values of $\widehat{p_i^{opt}}$ indicate that stronger multipath is statistically expected using the median ToF for the distance estimation, while values closer to $p = 0.5$ indicate smaller multipath delay.

B. Automatic Model Calibration

In real deployments, it is desirable to avoid any manual offline calibration to estimate the regression parameters a and c of the linear regression model. We therefore propose to run the calibration methodology based on ToF measurements between pairs of access points from which we know the exact positions. (The assumption of known AP positions generally applies to many localization scenarios such as in airports, malls, museums or companies which have a single wireless network operator.) This can be therefore performed online without human intervention. The results of this online calibration between APs of Testbed I are shown in the bottom of Fig. 9. We find that the calibration results in a very similar linear regression. This suggests that *in order to calibrate the model for a particular environment, the ToF measurements between each pair of APs are sufficient*. We illustrate our filter in Fig. 10. The individual steps are as follows:

1) We train a linear regression model from online measurements between the APs and determine the parameters a and c of eq. (10). The model characterizes the relationship between the median ToF and the optimal percentile p_i^{opt} (cf. eq. (9)).
2) For each link $i \in \mathcal{L}$, we take a sequence of M measurements $\{d_{i,1}, d_{i,2}, \ldots, d_{i,M}\}$ and determine $\bar{d}_i$ (median ToF).
3) The median ToF $\bar{d}_i$ is used to estimate the multipath delay using the linear regression model from step 1. The output of the estimator is a percentile value $\widehat{p_i^{opt}}$.

4) For each link $i \in \mathcal{L}$, we apply a linear interpolation to the sequence of ToF measurements $\{d_{i,1}, d_{i,2}, \ldots, d_{i,M}\}$, select the $\widehat{p_i^{opt}}$-percentile of that sequence, and estimate the distance as $d_i(\widehat{p_i^{opt}})$.

VIII. EVALUATION METHODOLOGY

We first present the algorithms we use to evaluate the performance of our distance estimator and then present how we compute the position in our system.

A. Algorithms for the Distance Estimation Evaluation

For the evaluation, we study the performance of the following filters:

- our filter that uses the median ToF to estimate the optimal percentile (un-filtered),[4]
- two other versions of the filter that rely on the median of the RSSI (pre-filtered) and the skewness of the ToF (pre-filtered), as they provided the second- and third-best correlation coefficients in Section VII, Table I.

Recall that these different metrics all estimate the optimal percentile p_i^{opt} of the multipath filter and that all three variants rely on the ToF for the final distance estimation. For each of the testbeds in Fig. 2, we compute the linear regression coefficients of each moment above with the p_i^{opt} (cf. Section VII-A).

In order to compare our filter against other approaches proposed in the literature, we further implement the MUSIC algorithm [18], the EM algorithm for GMM model and CAESAR [17]. CAESAR has been specifically designed for WiFi ranging while the MUSIC and the EM algorithms represent popular approaches to mitigate the impact of multipath in wireless localization. We briefly describe these three algorithms in Appendix XII-A of the supplementary material.

B. Localization and Tracking

In addition to the ranging error, we also study the localization error. We define a coordinate system on a two-dimensional map. To ease the notation, we consider only one target and a number of links equal to the number of APs in range, denoted as L. Let $\mathbf{s_i} = (x_i, y_i)$ be the position of the AP i, $\mathbf{p} = (x, y)$ the position of the target device, and $r_i = \|\mathbf{s_i} - \mathbf{p}\| = \sqrt{(x - x_i)^2 + (y - y_i)^2}$ the distance between the AP i and the target device. Let $\hat{d}_i$ further denote the estimate of the distance from AP i to the target, as computed using eq. (7). Our objective is to find the coordinates that solve the following weighted least square optimization problem:

$$\hat{\mathbf{p}} = (\hat{x}, \hat{y}) = \underset{x,y}{\operatorname{argmin}} \sum_{i=1}^{L} \frac{(\|\mathbf{s_i} - \mathbf{p}\| - \hat{d}_i)^2}{w_i} \qquad (11)$$

where w_i is a weighting function of the samples $\{\hat{d}_i\}$. We explain in Appendix XII-B.1 of the supplementary material how the weights w_i are determined. For solving eq. (11),

---

[4]Since the correlation coefficient of the median ToF versus the optimal percentile is not improved by pre-filtering outliers, we apply this model to the raw, unfiltered samples.

we apply the Newton-Gauss method with line-search for the step size, a well-known method for this problem. As for the initial position for the computation, we use the one given by a linear least square algorithm and having the advantage of being computationally efficient.

To capture the localization error intrinsic to our filter and avoid including measurement data with high error due to a bad deployment of the AP positions, we calculate the horizontal dilution of precision (HDOP) values. The HDOP relates to the multiplicative effect of AP geometry on positional measurement precision error. We only consider the positions with an HDOP value smaller or equal to five.

For mobile device tracking, we apply an Exponentially Weighted Moving Average (EWMA) filter of the position:

$$\bar{\mathbf{p}}_{\mathbf{k}} = (1 - \alpha)\bar{\mathbf{p}}_{\mathbf{k-1}} + \alpha\hat{\mathbf{p}}_{\mathbf{k}},$$
$$s.t. \quad \bar{\mathbf{p}}_{\mathbf{1}} = \hat{\mathbf{p}}_{\mathbf{1}} \tag{12}$$

where $\bar{\mathbf{p}}_{\mathbf{k}}$ is the current position estimate, $\hat{\mathbf{p}}_{\mathbf{k}}$ is the last measured position and α is the filter weight. An EWMA filter puts more or less weight on historical values according to the filter weight α. A good choice of α depends on the mobility of the target devices. The filter weight α is computed by means of experiments. Details are provided in Appendix XII-B.2 of the supplementary material.

IX. EVALUATION

We analyze the performance of our statistical filter for ranging and mobile device position tracking.

A. Distance Ranging

We first evaluate the distance estimation error of our filter.

1) Distance Ranging Accuracy: We first evaluate the ranging accuracy using all links of Testbed I and offline calibration for our filters (the impact of online calibration is evaluated later). Similar experiments are performed in Testbed II and III. For each link, we compute the distance estimation error with 20 samples, and then calculate the average error using 500 sequences. We consider our filter that uses the median ToF, the median of the RSSI and the skewness of the ToF for estimating the optimal percentile. We also provide the error for CAESAR, the MUSIC algorithm, the GMM model using the EM algorithm, as well as classical estimators such as the mean and median of the ToF.

From Fig. 11 we see that our three new estimators outperform the mean and the median metrics by comparing the Empirical Cumulative Distribution Function (ECDF) of the ranging errors. We then compare our best performing filter, based on the median of ToF, against the EM algorithm, CAESAR and the MUSIC algorithm. We show in Figure 12 that our filter outperforms the other evaluated approaches. The best performance is achieved by our filter using the median ToF. We obtain a median error of $2.4\,\mathrm{m}$ and a 80-percentile error of $5.3\,\mathrm{m}$. The filter that uses the median RSSI slightly outperforms the skewness of the ToF. This is not surprising since Table I shows higher correlation coefficients of the filter with median RSSI. The mean and median have roughly equal estimation error. Their median error is approximately
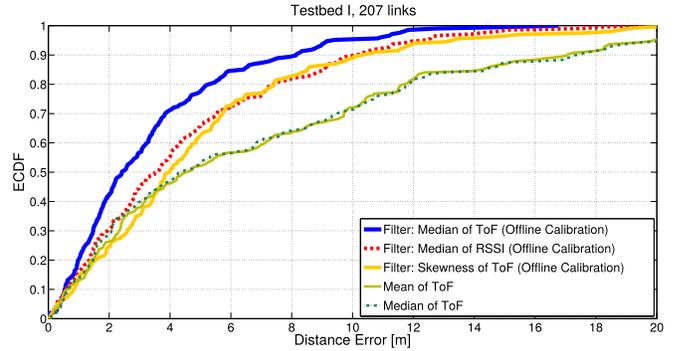

Fig. 11. ECDF of the distance estimation error for our adaptive filter compared to adaptive filters based on other metrics and classical estimators based on the mean and the median of ToF (Testbed I).
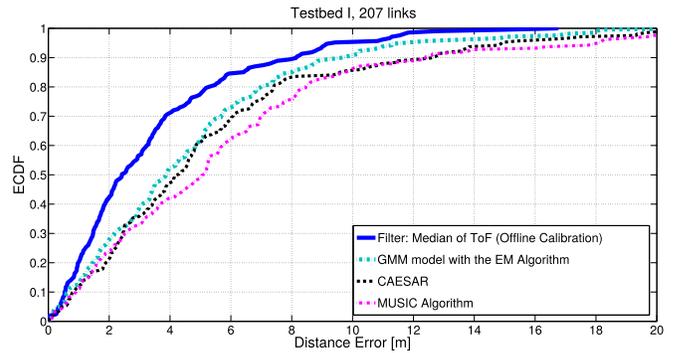

Fig. 12. ECDF of the distance estimation error for our adaptive filter compared to other algorithms (Testbed I).
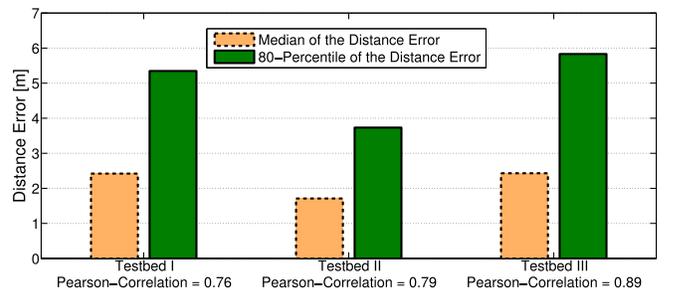

Fig. 13. Distance accuracy (50- and 80-percentile distance error) in Testbed I, II and III.

$4.5\,\mathrm{m}$ and the 80-percentile error is approximately $11.5\,\mathrm{m}$. CAESAR gets a median error of $4.1\,\mathrm{m}$ and a 80-percentile error of $7.3\,\mathrm{m}$. Our evaluation of CAESAR is also very consistent to the one recently presented in the indoor evaluation of [11]. With regard to the MUSIC algorithm, this approach is ineffective as a result of the large noise introduced by the target station (cf. Section V), which is not taken into account in the model used by the MUSIC approach. The 80-percentile of the MUSIC algorithm shows better performance than the mean and the median metrics, but still worse than our new estimators. Finally, the EM algorithm outperforms CAESAR and MUSIC. However, the median error of $3.9\,\mathrm{m}$ and an 80-percentile error of $7.1\,\mathrm{m}$ shows its inefficiency with respect to our filter using the median ToF.

2) Robustness to Different Environments/Testbeds: Fig. 13 shows the median and 80-percentile of the distance error for the three different testbeds using sequences of $M = 20$
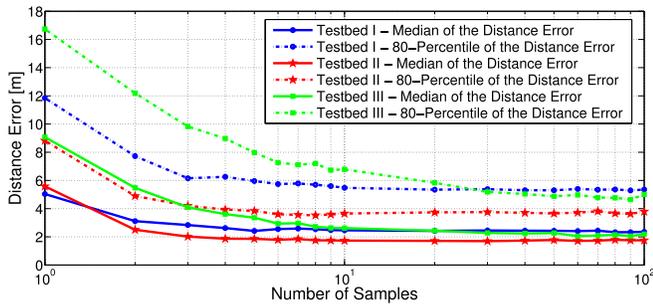
Fig. 14. Median and 80-percentile error of our estimator that filters the noise based on the median of ToF. Results are provided for different number of samples for Testbed I, II and III.

Fig. 15. Capacity analysis for ranging, considering N target stations and $M = 20$ samples for the ranging estimation.

Fig. 16. Comparison of offline and online calibration for Testbed I, and comparison with different placements of the APs (i.e., Testbed IV).

samples and offline calibration. As shown in the x-label of the figure, we measure a high Pearson correlation coefficient between the median of ToF and the optimal percentile, in the range of $0.76 - 0.89$. The median distance error is in the range $1.7 - 2.4$ m, and the 80-percentile error is in the range $3.7 - 5.8$ m. For comparison, the median (80-percentile) real distances for Testbed I, II and III are equal to 12.3 m (19.2), 8 (11.9) m, 10.2 (15.7) m, respectively. Concluding, our filter is robust across different environments.

3) Impact of Number of ToF Samples: We evaluate the number of samples M necessary in our filter. Figure 14 shows the error for our filter that relies on the median of the ToF as a function of the number of samples. The error is stable with ten or more samples for the median of the distance error. Only the 80-percentile of Testbed III can benefit from a higher number of samples.

4) Ranging Capacity: We study the capacity of the WiFi ranging technique, defined as the time C required to collect M samples in a WiFi network of N users. In order to find the capacity of the ToF ranging method, we apply Little's formula [28]:

$$C = \frac{M \cdot N}{S/P}, \tag{13}$$

where S is the throughput and P indicates the payload bits of the single frame. We conduct the analysis in saturation conditions and we apply the Bianchi's formula to compute S [29]. We also consider that the traffic for ranging does not have data content and it only consists of MAC overhead. Considering $M = 20$ for the ranging estimation, the results versus the number of users in the system are shown in Fig. 15. As expected, increasing the PHY rate helps reduce the time C required to compute ranging estimates for N users. For instance, $C = 0.25$ s is needed to compute the ranges to 30 users with PHY rate of 11 Mb/s, while $C = 0.1$ s is needed with PHY rate of 18 Mb/s.

Additional results are provided in Appendix XII-C of the supplementary material.

B. Online vs Offline Calibration

Our proposal for online calibration is to train our model for the adaptive filter using the links between the fixed APs. Since the APs are at fixed known locations, this type of calibration does not require any additional manual effort and can therefore be performed online.
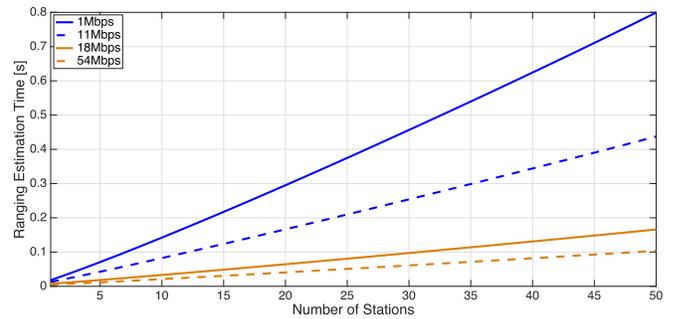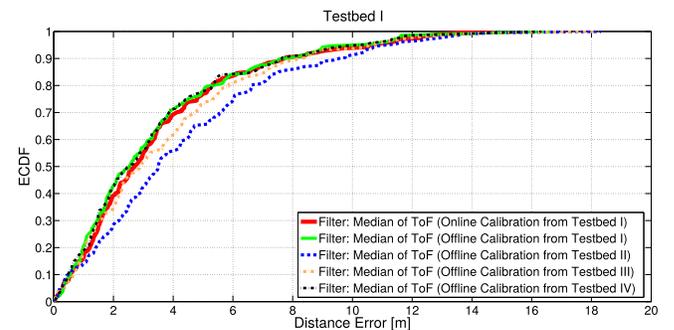
We compare online versus offline model calibration for Testbed I in Fig. 16. We have three findings. First, the online calibration achieves similar results with respect to the offline tests, and we can then apply the calibration online without significant performance loss. Second, online calibration outperforms the results we could achieve applying the linear regression achieved in Testbed II and III to Testbed I, which implies that it is better to calibrate online than using the calibration coefficients of other environments. Third, we find that the regression parameters $\{a, c\}$ of Testbed I and Testbed IV are very similar. This is mapped to very similar performance observed in Fig. 16 using the offline calibration of Testbed IV. These testbeds use the same environment but different placements of the APs. This result suggests that the calibration is largely AP-position independent, but rather a feature of the environment.

We then study the variability of the coefficients of the automatic model calibration over time. These tests rely on experiments with 5 APs across an area of 250 m^2 in an office space. We perform the online calibration each hour for a total of 48 consecutive hours. The results of this investigation, in terms of the evolution of the regression parameters $\{a, c\}$ versus the time are shown in Fig. 17. From the figure, we observe a dynamic trend during the working hours of the day and a static one during out-of-office hours. As a result of this experiment, we can conclude that frequent updates, in order of every hour, may be required during the day due to the presence and mobility of people as well as objects moved around, both increasing the dynamics of the environment.
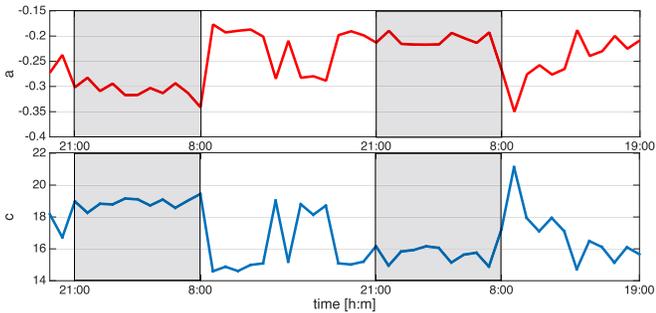
Fig. 17. Evolution of the coefficients of the linear regression model used for online calibration.
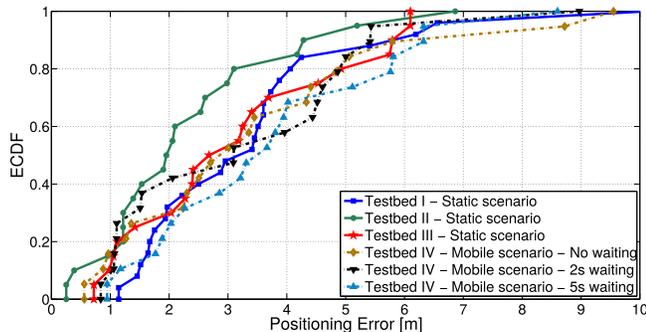

Fig. 18. Positioning error in static (Testbed I, II and III) and mobile scenarios (Testbed IV).

C. Localization and Tracking

Following the ranging accuracy, we investigate next the positioning and tracking performance.

1) Positioning Accuracy: We study the position accuracy in static setting and summarize the results in Fig. 18. The figure shows that we have a median positioning error between 2.0 and 3.1 m, and 80-percentile positioning error between 3.2 and 4.8 m, depending on the testbed. Compared to ranging, the median positioning accuracy is slightly worse ($1.7 - 2.4$ m for ranging). However, the 80-percentile error is slightly better ($3.7 - 5.8$ m for ranging).

2) Mobile Tracking: Fig. 18 shows the ECDFs for the different waiting times in Testbed IV at each location, no waiting, two seconds and five seconds waiting. We observe that the tests under full mobility (label 'no waiting') give comparable positioning accuracy to the cases where the user waits for a few seconds to collect more data at a given position (label '2s' and '5s'), with a median error in the range $2.6 - 3.4$ of meters. The reason is that the moving device experiences different multipath constellations which cancel out more effectively than when the device remains fixed. This effect compensates for the increased tracking error which arises when measurements are slightly lagging behind due to the measurement delays when the device constantly moves ahead.

X. RELATED WORK

The indoor localization literature is vast, including techniques using signal strength [1]–[5], the angle of arrival [6]–[8] or combining WiFi signals with inertial sensors as found in smartphones [10]. In this section, we focus on time-based localization techniques which are most related

to ours in addition to some general NLOS mitigation based solutions.

ToF Echo techniques based on packet exchanges in WiFi networks were first proposed in [12] and [13] and refined in [14]–[17]. However, unlike our work, none of these approaches address the effect of non-Gaussian noise such as in multipath-rich environments. A direct comparison to [17] as presented in this work shows that the error with our statistical filter can be reduced in indoor environments by a factor of more than two compared to classical estimators that do not compensate for the bias of the multipath. Gallo *et al.* [30] introduced directional Yagi antennas to eliminate the effect of multipath and other noise sources from WiFi echo techniques. They achieved a positioning accuracy of less than 5 m in 8 from 10 positions. In contrast, our filter works with omnidirectional antennas in environments with multipath.

There have been many attempts to harness the time-of-flight of wireless signal in indoor propagation environments despite multipath. Common proposals to combat the multipath problem are for example the use of ultrawideband signals [31]–[33] or frequency diversity [34], [35]. However, these approaches require specialized hardware or software-defined radios which increase costs and hinders localization at larger scales. SAIL [11] is a ToF system using WiFi that has been designed for localization in multipath environments. However, SAIL requires inputs from the inertial sensors in the smartphone. SAIL achieved median error of ≈ 1 m and 80-percentile error of ≈ 5 m, which is comparable to our filter, at the drawback of requesting the collaboration from the mobile user through the installation of a dedicated application on smartphones. ToneTrack [19] tries to overcome the problem of limited bandwidth, inherent to WiFi time-based localizations. It combines channels to form virtual larger bandwidths without increasing the radio's sampling rate, taking benefit of frequency hopping, to increase the resolution of ToF profiles. The work SpotFi [9] uses APs equipped with 3 antennas and commodity WiFi chipsets. It jointly estimates the Angle of Arrival (AoA) and ToF pairs of each path using the channel state information, and estimates the likelihood that these pairs correspond to the direct path between the AP and the target. In contrast to our work, SpotFi does not rely on ToF for ranging. Other solutions try to identify and mitigate the NLOS effects of WiFi RSS measurements [36]. The proposed solution combines a machine learning technique to first extract typical features from the training data collected during extensive indoor measurement campaigns and estimate the ranges using a regression model, and an identification approach based on hypothesis testing. The approach still needs a training phase that requires an offline classification or calibration. Still in RSS based ranging, [37] proposes to use a GMM model to filter corrupted range estimations caused by NLOS radio propagation by modeling distributions of LOS and NLOS sets of estimates. In this work, we have demonstrated that the GMM model does not work well in WiFi ToF and we have proposed a combination of statistical learning and robust statistics that outperforms the classical expectation-maximization algorithm used by the GMM model.

XI. Conclusion

We have developed a new filter based on statistical learning and robust statistics to improve the ranging accuracy in the presence of noisy ToF measurements. Our filter does not require additional information form the devices or the user besides the ToF values which makes it applicable to a wide range of applications. We have shown how to apply our filter for indoor localization and tracking of COTS WiFi devices with legacy 802.11 Access Point infrastructures. In indoor deployments with multipath, our filter outperforms conventional ToF based range estimators by a factor of more than two. We have demonstrated the accuracy of the filter to estimate the position in static and mobile settings. In static conditions, the filter achieved a median error of $2.0 - 3.1$ meters. In mobile settings, the error was only slightly higher with a median error of $2.6 - 3.4$ meters. We have shown that the performance of our filter can be achieved with online model calibration, and hence does not require any cumbersome onsite pre-calibration efforts. The system has also participated in the Microsoft Indoor Localization Competition 2016, achieving an average error of 3.17 meters in a challenging and uncontrollable environment with 5 APs covering $500 \, \text{m}^2$. This resulted in ranking 5th out of ten teams in the final. Our system was the only solution presented at the competition that did not require any customized software in the mobile device, neither inputs from inertial sensors in the mobile device. Since ToF is becoming readily available in WiFi chipsets, we foresee that our approach can be applied by the provider of positioning systems in airports, malls, museums, homes for context-aware networking as well as data analytics that require positioning data. The advent of new WiFi chipsets operating on wideband channels (such as the 160 MHz clock of 802.11ac) will greatly help to increase the accuracy and the integration with phase information will also be needed to improve the system performance.

References

[1] P. Bahl and V. N. Padmanabhan, "RADAR: An in-building RF-based user location and tracking system," in *Proc. IEEE INFOCOM*, Mar. 2000, pp. 775–784.

[2] M. Youssef and A. Agrawala, "The horus WLAN location determination system," in *Proc. MobiSys*, New York, NY, USA, 2005, pp. 205–218.

[3] H. Lim, L.-C. Kung, J. C. Hou, and H. Luo, "Zero-configuration, robust indoor localization: Theory and experimentation," in *Proc. 25th IEEE Int. Conf. Comput. Commun. (INFOCOM)*, Apr. 2006, pp. 1–12.

[4] A. Goswami, L. E. Ortiz, and S. R. Das, "Wigem: A learning-based approach for indoor localization," in *Proc. CoNEXT*, New York, NY, USA, 2011, pp. 3:1–3:12.

[5] K. Chintalapudi, A. P. Iyer, and V. N. Padmanabhan, "Indoor localization without the pain," in *Proc. ACM MobiCom*, 2010, pp. 173–184. [Online]. Available: http://doi.acm.org/10.1145/1859995.1860016

[6] J. Xiong and K. Jamieson, "ArrayTrack: A fine-grained indoor location system," in *Proc. 10th USENIX Symp. Netw. Syst. Design Implement. (NSDI)*, Lombard, IL, USA, 2013, pp. 71–84.

[7] S. Sen, J. Lee, K.-H. Kim, and P. Congdon, "Avoiding multipath to revive inbuilding WiFi localization," in *Proc. 11th Annu. Int. Conf. Mobile Syst., Appl., Ser. (MobiSys)*, New York, NY, USA, 2013, pp. 249–262.

[8] J. Gjengset, J. Xiong, G. McPhillips, and K. Jamieson, "Enabling phased array signal processing on commodity WiFi access points," in *Proc. Mobicom*, New York, NY, USA, 2014, pp. 153–164 .

[9] M. Kotaru, K. Joshi, D. Bharadia, and S. Katti, "SpotFi: Decimeter level localization using WiFi," in *Proc. ACM Conf. Special Interest Group Data Commun. (SIGCOMM)*, New York, NY, USA, 2015, pp. 269–282.

[10] A. Rai, K. K. Chintalapudi, V. N. Padmanabhan, and R. Sen, "Zee: Zero-effort crowdsourcing for indoor localization," in *Proc. Mobicom*, New York, NY, USA, 2012, pp. 293–304.

[11] A. T. Mariakakis, S. Sen, J. Lee, and K.-H. Kim, "SAIL: Single access point-based indoor localization," in *Proc. 12th Annu. Int. Conf. Mobile Syst., Appl., Ser. (MobiSys)*, New York, NY, USA, 2014, pp. 315–328.

[12] X. Li, K. Pahlavan, M. Latva-Aho, and M. Ylianttila, "Comparison of indoor geolocation methods in DSSS and OFDM wireless LAN systems," in *Proc. 52nd IEEE Veh. Technol. Conf.*, vol. 6. Sep. 2000, pp. 3015–3020.

[13] D. D. McCrady, L. Doyle, H. Forstrom, T. Dempsey, and M. Martorana, "Mobile ranging using low-accuracy clocks," *IEEE Trans. Microw. Theory Techn.*, vol. 48, no. 6, pp. 951–958, Jun. 2000.

[14] A. Günther and C. Hoene, "Measuring round trip times to determine the distance between WLAN nodes," in *Proc. NETWORKING*, Berlin, Germany, 2005, pp. 768–779.

[15] M. Ciurana, F. Barcelo-Arroyo, and F. Izquierdo, "A ranging system with IEEE 802.11 data frames," in *Proc. IEEE Radio Wireless Symp.*, Jan. 2007, pp. 133–136.

[16] S. A. Golden and S. S. Bateman, "Sensor measurements for Wi-Fi location with emphasis on time-of-arrival ranging," *IEEE Trans. Mobile Comput.*, vol. 6, no. 10, pp. 1185–1198, Oct. 2007.

[17] D. Giustiniano and S. Mangold, "CAESAR: Carrier sense-based ranging in off-the-shelf 802.11 wireless LAN," in *Proc. 7th Conf. Emerg. Netw. EXperim. Technol. (CoNEXT)*, New York, NY, USA, 2011, pp. 10:1–10:12.

[18] X. Li and K. Pahlavan, "Super-resolution TOA estimation with diversity for indoor geolocation," *IEEE Trans. Wireless Commun.*, vol. 3, no. 1, pp. 224–234, Jan. 2004.

[19] J. Xiong, K. Sundaresan, and K. Jamieson, "ToneTrack: Leveraging frequency-agile radios for time-based indoor wireless localization," in *Proc. 21st Annu. Int. Conf. Mobile Comput. Netw. (MobiCom)*, New York, NY, USA, 2015, pp. 537–549.

[20] IPSN. (2014). *Microsoft Indoor Localization Competition*. [Online]. Available: http://research.microsoft.com/en-us/events/ipsn2014indoorlocalizatinocompetition/

[21] *IEEE Standard for Information Technology—Telecommunications and Information Exchange Between Systems Local and Metropolitan Area Networks—Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Standard 802.11-2012 (Revision of IEEE Std 802.11-2007), 2012, pp. 1–2793.

[22] T. Bourchas, M. Bednarek, D. Giustiniano, and V. Lenders, "Poster abstract: Practical limits of WiFi time-of-flight echo techniques," in *Proc. 13th Int. Symp. Inf. Process. Sensor Netw. (IPSN)*, Piscataway, NJ, USA, Apr. 2014, pp. 273–274.

[23] I. Guvenc and C.-C. Chong, "A survey on TOA based wireless localization and NLOS mitigation techniques," *IEEE Commun. Surveys Tuts.*, vol. 11, no. 3, pp. 107–124, Aug. 2009.

[24] J. A. Bilmes, "A gentle tutorial of the EM algorithm and its application to parameter estimation for Gaussian mixture and hidden Markov models," Int. Comput. Sci. Inst., Berkeley, CA, USA, Tech. Rep. TR-97-021, 1998.

[25] A. J. Wheeler, A. R. Ganji, V. V. Krishnan, and B. S. Thurow, *Introduction to Engineering Experimentation*. Englewood Cliffs, NJ, USA: Prentice-Hall,1996.

[26] J. Benesty, J. Chen, Y. Huang, and I. Cohen, "Pearson correlation coefficient," in *Noise Reduction in Speech Processing* (Springer Topics in Signal Processing), vol. 2. Berlin, Germany: Springer, 2009, pp. 1–4.

[27] N. R. Draper, H. Smith, and E. Pownell, *Applied Regression Analysis*, vol. 3. New York, NY, USA: Wiley, 1966.

[28] I. Tinnirello, D. Giustiniano, L. Scalia, and G. Bianchi, "On the side-effects of proprietary solutions for fading and interference mitigation in IEEE 802.11b/g outdoor links," *Comput. Netw.*, vol. 53, no. 2, pp. 141–152, Feb. 2009.

[29] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 3, pp. 535–547, Mar. 2000.

[30] P. Gallo, D. Garlisi, F. Giuliano, F. Gringoli, and I. Tinnirello, "WMPS: A positioning system for localizing legacy 802.11 devices," *Trans. Smart Process. Comput.*, vol. 1, pp. 106–115, Oct. 2012.

[31] J.-Y. Lee and R. A. Scholtz, "Ranging in a dense multipath environment using an UWB radio link," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 9, pp. 1677–1683, Dec. 2002.

[32] S. Gezici *et al.*, "Localization via ultra-wideband radios: A look at positioning aspects for future sensor networks," *IEEE Signal Process. Mag.*, vol. 22, no. 4, pp. 70–84, Jul. 2005.

[33] D. Dardari, A. Conti, U. Ferner, A. Giorgetti, and M. Z. Win, "Ranging with ultrawide bandwidth signals in multipath environments," *Proc. IEEE*, vol. 97, no. 2, pp. 404–426, Feb. 2009.

[34] S. Lanzisera, D. Zats, and K. S. J. Pister, "Radio frequency time-of-flight distance measurement for low-cost wireless sensor localization," *IEEE Sensors J.*, vol. 11, no. 3, pp. 837–845, Mar. 2011.

[35] P. Pettinato, N. Wirström, J. Eriksson, and T. Voigt, "Multi-channel two-way time of flight sensor network ranging," in *Proc. EWSN*, 2012, pp. 163–178.

[36] Z. Xiao *et al.*, "Identification and mitigation of non-line-of-sight conditions using received signal strength," in *Proc. 9th IEEE Int. Conf. Wireless Mobile Comput., Netw. Commun. (WiMob)*, Lyon, France, Oct. 2013, pp. 667–674.

[37] Q. Wang, I. Balasingham, M. Zhang, and X. Huang, "Improving RSS-based ranging in LOS-NLOS scenario using GMMs," *IEEE Commun. Lett.*, vol. 15, no. 10, pp. 1065–1067, Oct. 2011.

Domenico Giustiniano (S'06–M'13) received the Ph.D. degree in telecommunication engineering from the University of Rome Tor Vergata in 2008. He was formerly a Senior Researcher and Lecturer with ETH Zurich and a Post-Doctoral Researcher with Disney Research Zurich and with Telefonica Research Barcelona. He is currently a Research Associate Professor with IMDEA Networks Institute, Spain, and the Leader of the Pervasive Wireless Systems Group. He has authored over 70 international papers and is an inventor of five patents. His current research interests are in the areas of visible light communication systems, distributed spectrum monitoring systems, and mobile indoor localization systems.

Maurizio Rea received the M.Sc. degree in telecommunications engineering from the University of Palermo, Italy, in 2015, and the M.Sc. degree from the University Carlos III of Madrid in 2016. He is currently pursuing the Ph.D. degree with IMDEA Networks Institute and the University Carlos III of Madrid. Before joining IMDEA, he was a Researcher with ETH Zurich, where he focused his research on indoor localization systems. His interests include data analysis, mmWave networks, and context-aware mechanisms.

Aymen Fakhreddine received the M.Sc. degree in advanced wireless communications systems from CentraleSupélec, Paris, France, in 2012, the M.Eng. degree in telecommunications from INPT Rabat in 2012, and the M.Sc. degree in telematics engineering from the University Carlos III of Madrid in 2014. He is currently pursuing the Ph.D. degree with IMDEA Networks Institute and the University Carlos III of Madrid. He is currently involved in pervasive localization systems. Prior to joining IMDEA Networks Institute, he held a Research Intern position with Bell Labs France, where he focused his research on LTE mobility optimization.

Vincent Lenders received the M.Sc. and Ph.D. degrees in electrical engineering from ETH Zurich. He was Post-Doctoral Research Faculty with Princeton University. He is currently a Research Director with Armasuisse, where he leads the cyberspace and information science research activities for the Swiss Federal Department of Defense. His current research interests are in the fields of cyber security, information management, big data, and crowdsourcing. He is the co-founder and in the board of the OpenSky Network and Electrosense associations.