Deep Inspection of the Noise in WiFi Time-of-Flight Echo Techniques

Domenico Giustiniano
IMDEA Networks Institute
Madrid, Spain

Theodoros Bourchas
ETH Zurich
Zurich, Switzerland

Maciej Bednarek
ETH Zurich
Zurich, Switzerland

Vincent Lenders
armasuisse
Thun, Switzerland

ABSTRACT

Time-of-flight (ToF) echo techniques have been proposed to estimate the distance between a local and a target station using regular WiFi radio devices. As of today, there is little understanding of the noise sources from ToF measurements. We conduct extensive experimental tests based on a customized WiFi echo technique implementation residing in the core of the 802.11 MAC processor and a high-resolution signal analysis of the WiFi traffic captured with a wideband oscilloscope. We discern the root of the error components in WiFi echo technique measurements and statistically characterize the offset noise added by the target station. Our measurements provide key insights to model the sources of noise and guidance for the design of robust distance estimators.

Categories and Subject Descriptors

C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design—*Wireless Communication*

Keywords

Indoor localization; 802.11; Time of Flight; Measurement noise; Analysis; Implementation; Evaluation

1. INTRODUCTION

Location-based services have recently experienced a huge growth of interest to support entertainment, social media, rescue, advertisement, sport and navigation applications. This unprecedented interest is witnessed by huge commercial and scientific efforts to provide a pin-point indoor positioning system. Following this trend, WiFi Positioning Systems (WPS) based on the signal strength have achieved wide commercial success to complement GPS in indoor and urban canyon environments, even despite being proved to be error-prone and offering limited performance figures [1].

MSWiM'15, November 02-06, 2015, Cancun, Mexico.

DOI: http://dx.doi.org/10.1145/2811587.2811621.

In order to alleviate this problem, solutions such as [2–12] have attempted to use the time-of-flight (ToF) principle to devise an echo technique for ranging measurements that leverages the existing 802.11 protocol. While WiFi echo techniques may offer a cost-effective alternative to estimate the distance, they suffer from severe noise, which may lead to low accuracy and precision of the estimate, particularly when the time to collect samples is limited by the mobility of the target user.

The goal of this work is to dissect the current limitations of WiFi echo techniques. We conduct an in-depth experimental inspection with a customized WiFi echo technique operating in the core of the 802.11 MAC state machine. With controlled tests in the laboratory, we then extract and characterize the offset originated by commercial 802.11 radio chipsets operating as target stations. We study the statistical error distribution of the ToF and characterize how it affects the ranging accuracy and precision. In particular, we aim to answer the following key questions related to timing information extracted for 802.11 ranging:

- How deterministic and predictable are the time offsets for ranging measurements using regular IEEE 802.11 chipsets?
- What are the dominating sources of unpredictiveness in the offset noise?
- How do different IEEE 802.11 chipsets and different physical modulations behave in terms of timing accuracy and precision?
- How stable are the device's offsets over time and according to the network traffic conditions?

Our statistical analysis concludes that the chipset-dependent offset noise at the target station and the additional quantization error of the WiFi clock for timing measurements at the local station are the dominant sources of noise. We provide guidance to discern which features must be taken into account in order to design a robust estimator of the distance.

The rest of this paper is organized as follows. Section 2 introduces the principles of WiFi echo techniques. Our experimental platform is described in Section 3. In Section 4, we investigate the accuracy and the precision of the WiFi echo technique using our platform. We then introduce the setup we use to measure the error of the target station in Section 5 and present our main insights in Section 6. In Section 7, we provide statistical evidence that the target station introduces significant error in ranging measurements, and we finally discuss the implications and conclusions in Section 8.

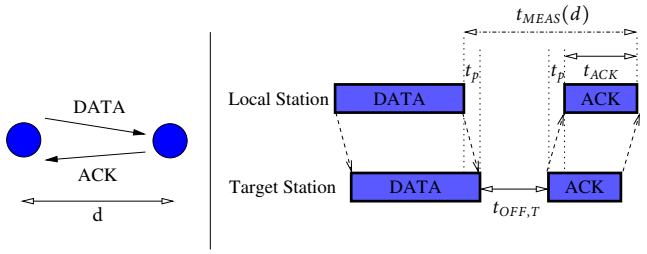

Figure 1: WiFi ToF echo technique. The local station measures $t_{MEAS}(d)$ and computes the distance to the target station. $t_{OFF,T}$ is originated at the target station, and it depends on the 802.11 SIFS time. Not shown in the figure, $t_{OFF,L}$ appears from quantization errors at the local station and errors in the operation of starting and stopping the timer.

2. TOF LOCALIZATION

While traditional echo techniques like radar systems are based on uncoded RF signals and their reflections, WiFi echo techniques use regular frames of communication. These approaches exploit the fact that frames propagate in the air as electromagnetic waves, and thus their propagation times linearly depend on the distance traveled. Echo techniques were first introduced by [3] to resolve the absence (or low precision) of clock synchronization in WiFi chipsets, and then extended by [4–6]. The most recent and promising approaches in the field rely on regular 802.11 DATA frames for the echo requests and on ACK control frames for the echo replies [7, 8, 10]. The distance estimation procedure of these techniques is visualized in Fig. 1. A local station measures the time $t_{MEAS}(d)$ elapsed from the instant that the DATA frame has been transmitted to the instant that the ACK has been received. Depending on the implementation, the local station may either wait until the ACK is completely received or it may stop its timer as soon as it detects the signal energy of the ACK. In this work, we make use of the first approach. Waiting until the ACK has been completely received has the advantage of reducing the time estimation uncertainty for very weak and very strong signal-to-noise ratios inherent in approaches which rely on energy detection [7, 13, 14]. The measurement time is therefore given by

$$t_{MEAS}(d) = 2 * t_P(d) + t_{ACK} + t_{OFF}, \quad (1)$$

where $t_P(d)$ is the signal propagation time between the local and target station, t_{ACK} indicates the duration of the ACK frame and t_{OFF} is an offset given by

$$t_{OFF} = t_{OFF,T} + t_{OFF,L}, \quad (2)$$

where $t_{OFF,T}$ and $t_{OFF,L}$ are the target and local offsets caused by the two stations, respectively. $t_{OFF,T}$ originates from hardware processing delays at the target station. It reflects the Short InterFrame Space (SIFS) time according to a tolerance level defined in the 802.11 standard [15]. $t_{OFF,L}$ originates from hardware and software processing delays at the local station. The distance between the local station and the target station is then inferred as

$$\hat{d} = \frac{c}{2} \cdot (t_{MEAS}(d) - t_{ACK} - t_{OFF}), \quad (3)$$

where c is the speed of light.

3. PLATFORM FOR RANGING MEASUREMENTS

In this section, we present the experimental platform for the ToF echo technique ranging measurements. For cost reasons, an important requirement in the design of the platform is that it must work with commercial off-the-shelf chipsets, such that it can be integrated in legacy networks with a software upgrade. Our implementation meets this requirement, since we rely on a customized software of a low-cost Wi-Fi chipset, with a cost per unit of less than six dollars.

In order to alleviate any source of instability, WiFi echo techniques have to work as close as possible to the radio hardware, and should ideally be integrated in the 802.11 MAC state machine such that the ranging computation is independent of the main CPU processes running in the local and target stations. Our experimental platform is based on the open source openFWWF firmware code for Broadcom chipsets and written in assembler [16]. $t_{MEAS}(d)$ is measured by monitoring two events, the end of the DATA transmission and the end of the ACK reception. The timing of the WiFi echo technique is defined by the General Purpose Timer (GPT), that operates in the wireless chipset at the resolution of the internal clock (88 MHz in our chipset). The GPT is launched in the MAC state machine of the firmware just after the 802.11 processor sets up the COND_TX_DONE condition register (a frame has been sent) which indicates that the timer starts to count clock cycles. The firmware does other operations as required by the 802.11 protocol. Once the ACK frame has been received (or the ACK timeout has elapsed), the COND_RX_COMPLETE register gets updated and the timer gets stopped. The value is saved as the number of clock cycles elapsed since the end of the DATA transmission.

The 802.11b implementation is affected by a quantization error of one clock cycle, equal to a resolution of approximately 11 ns. The 802.11g mode has instead a quantization error of 4 clock cycles, caused by the internal circuit design. Given the speed of light $c = 300\,\text{m}/\mu\text{s}$ and the factor 2 of eq. (3), an error of 20 ns corresponds to a distance error of 3 m. Hence, we have a resolution of approximately 1.7 m for 802.11b and 6.4 m for 802.11g[1]. In addition to this, the operation of starting and stopping the timer may introduce an error of one clock cycle in the ToF measurement.

4. RANGING ACCURACY AND PRECISION

Before digging into the detailed offset noise analysis, we first look in this section at the general distance estimation uncertainty that arises with a firmware-based ToF approach as described in the previous section. The uncertainty in distance estimation has two components, the accuracy (the degree of closeness to the true value) and the precision (the degree of certainty of the accuracy). In this section, we provide the results for these metrics. Unless specified otherwise, we use an estimator of the distance based on the median in order to provide estimates which are robust to non-Gaussian distributions of the offset noise components. In addition, outliers in the measurements are filtered out applying the Thompson Tau technique, a statistical method for deciding

[1]Other chipsets, such as the Atheros used by [7], have a distance resolution of 3.4 m for both 802.11b and 802.11g modulation schemes.

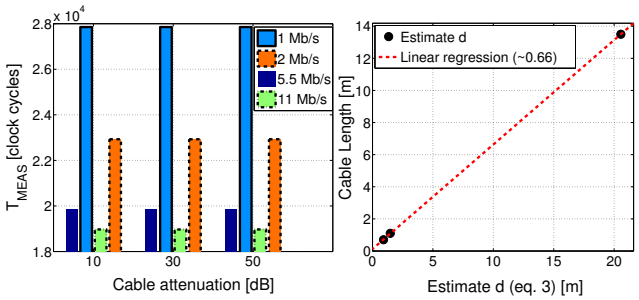

Figure 2: On the left, bar plot of the cable tests (short cable). The measurement at a given distance is independent of the signal strength between the two stations. On the right, plot of the cable length versus the estimate distance. It shows that the estimated distance d increases linearly with the cable length (scaling factor of 0.66 due to cable refractive index).

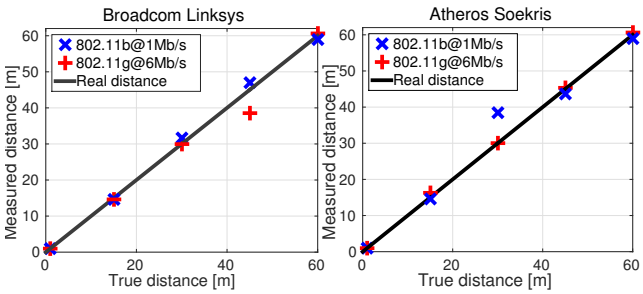

Figure 3: The indoor tests show that most of the ranging measurements have an accuracy below 2 m. Samples have been gathered for different 802.11 modes and different data rates. The distance is estimated calculating the median out of 10000 samples per test.

whether to keep or discard samples based on the expected value and the expected deviation of the sequence of samples.

4.1 Accuracy

We first measure the distance accuracy using an access point (AP) as local station and a laptop as remote station. Both stations are equipped with a Broadcom AirForce54G BCM4318 802.11 chipset and run the customized firmware for ToF ranging as presented in Section 3. In order to avoid any environmental effects, we connect the two stations using coaxial cables. The cables are of length 0.7 m (short), 1.1 m (medium) and 13.5 m (long), and are based on the standard RG-58. We produced different gain amplifier and signal strength values by introducing different attenuators before the cable of 10, 30 and 50 dB. For each test, we send a total of 10^4 frames and we fix the PHY rate to 1, 2, 5.5, or 11 Mb/s.

On the left of Fig. 2, we report the results for a fixed cable length of 0.7 m. From the bar plot, we draw the conclusion that the results are independent of the signal strength between the two stations, which confirms our expectations (cf. Section 2) that the timing triggered at the end of the ACK reduces the uncertainty for very weak and very strong signal-to-noise ratios with respect to approaches such as [7]. We then measure the estimated distance as a function of the cable length. An example is given on the right of Fig. 2 for the tests at 11 Mb/s. Given that the results are independent of the signal strength, we include at each distance all the measurements with different attenuators and compute the average. We find that the estimated distance increases linearly with the cable length. Note also that we measure a slope coefficient of ≈ 0.66 rather than one, as expected from the refractive index of our coaxial cables.

We then perform tests using a large number of samples sent over antennas in a long corridor with local and target stations in line-of-sight channel conditions where we expect that most of the ranging samples use the direct path between the two stations. This measurement is performed to get a rough feeling on the accuracy that our WiFi echo technique implementation can achieve. We place the local station measuring the ToF ranges and two types of target stations with different 802.11 chipsets at few selected distances (1, 15, 30, 45, 60 m), and collect measurements in order to compute $\hat{d}$ according to eq. (3). At each distance, we gather around 10^4 samples by flooding the channel. We then convert the clock difference between the reference value at 1 m and the median at each distance from clock cycles to meters (one clock cycle corresponds to ≈ 1.7 m) and compute median ranging values of gathered samples.

Figure 3 shows the accuracy by comparing the measured distance with the expected one for the 802.11b modulation at 1 Mb/s and 6 Mb/s. We do not report tests we conducted at other 802.11b and 802.11g PHY rates since they have similar results to the one illustrated here. The results in Figure 3 show that most of the measurements report an accuracy below 2 m. The measurement at 30 m with the Atheros Soekris is likely affected by multipath, which tends to cause an overestimation of the distance. Note also that, using an estimator based on the median, the accuracy is a step function of the distance, with steps equal to 1.7 m for 802.11b. The steps increase to 5.1 – 6.4 m for 802.11g, according to the 802.11g quantization noise and the noise caused by starting and stopping the timers. For instance, the measurement at 45 m with Broadcom Linksys 802.11g is affected by the lower clock resolution of 802.11g timing measurements.

4.2 Precision

We study the 95% confidence interval (CI) of the median as an indicator of the precision of the estimate. Fig. 4 depicts the 95% CI as a function of the number of collected samples for both 802.11 modes at 60 m. Since the measurements are affected by the clock resolution of the WiFi echo technique, we observe step functions. The figure shows that, in order to reach higher precision than 2 m, we need from 80 up to more than 600 samples. Besides the high uncertainty of the measured median for small number of samples, we further observe some fluctuation of the precision in some of the tests - even for large sets of samples. While some multipath in the measurement may be one of the reasons, 802.11g is further affected by the coarser grained clock resolution using OFDM modulation of our firmware. Similar results have been encountered at other distances.

As a result of this investigation, it is evident that the quantization error of $t_{OFF,L}$ plays an important role. However, it is less clear what are the other sources of offset noise. As

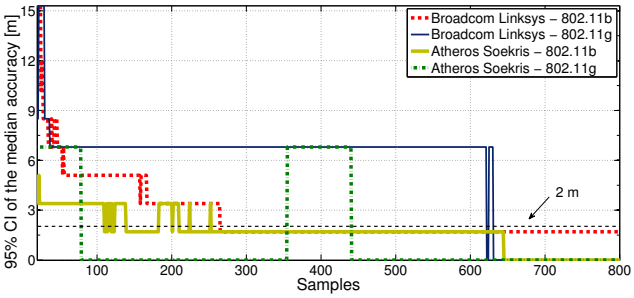

Figure 4: Precision of the estimate as a function of the number of samples for the Broadcom Linksys and Atheros Soekris, operating in 802.11b and 802.11g modes. The distance between the local and target station is 60 m.

such, in the next sections, we first study and quantify how much the target station contributes to the low precision of the distance measured for relatively small number of samples.

5. METHODOLOGY TO EVALUATE THE TARGET OFFSET

To discern the individual sources of offset noise, we present in this section a methodology that we have developed to independently study the offset at the target station $t_{OFF,T}$. Measuring the target offset in practice is quite challenging because we have to remove the channel noise as much as possible and make sure that our measurement system does not add any significant source of noise itself. To this end, we have developed a sophisticated measurement system using a high performing oscilloscope.

5.1 Measurement setup

We measure $t_{OFF,T}$ by monitoring the actual distribution of the PHY SIFS time over the wireless channel. In the rest of this paper, we interchangeably use the term $t_{OFF,T}$ and PHY SIFS time. One constraint of the methodology to monitor PHY SIFS time is that we must assure that the captured DATA/ACK time series have a resolution of at least the main 802.11 clock, so that a direct comparison with $t_{MEAS}(d)$ is viable. One approach could be to access the pins of the baseband signal with an oscilloscope. However, these pins are chipset dependent, and mostly available in older WiFi chipsets. Besides, measurements of the baseband signal would exclude the noise that is generated in the radio front-end.

We therefore employ the setup shown in Fig. 5 that allows us to collect DATA/ACK traffic over the air by measuring the PHY idle time. It consists of:

(i) Traffic monitoring: An Infiniium 90000A Oscilloscope is used to capture signal samples at a rate of 10 GSamples/sec and store data of 20 MB/50 MB per trace, which results in several GBs of total traces collected. A horn antenna is plugged to the I/O port of the oscilloscope and operates as baseband filter in the 2.4 GHz band.

(ii) Local station: Atheros Soekris, configured as access point (AP).

(iii) Target station: One among a) Broadcom Linksys with the Broadcom AirForce54G BCM4318 802.11 chipset

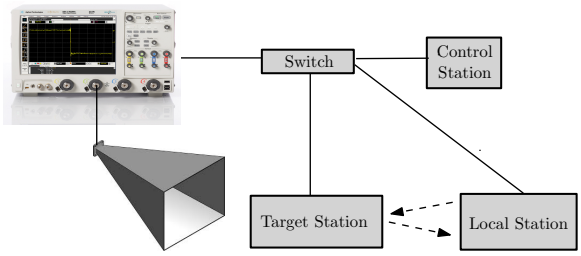

Figure 5: The experimental setup to investigate $t_{OFF,T}$ (PHY SIFS).

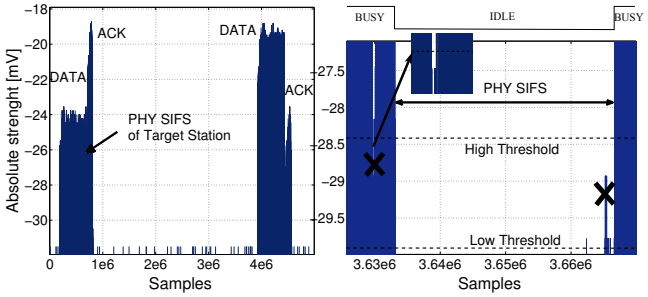

Figure 6: Two traces gathered with the Infiniium 90000A. Figure on the left: the local station sends a ping request and the target answers with a ping reply. We are interested in the distribution of the PHY SIFS ($t_{OFF,T}$) at the target station, acknowledging the reception of the DATA frame (ping request) with an ACK frame. Figure on the right: the Schmitt trigger avoids false detections of busy and idle state, indicated with crosses in the figure. We then compute $t_{OFF,T}$, how long the channel is idle at the target station between the DATA and the ACK.

b) Atheros Soekris with Atheros AR9220 802.11 chipset
c) Samsung GALAXY Nexus smartphone, with Broadcom BCM4335 chipset.

(iv) Control station: a laptop connected to both the oscilloscope and the local station (and the target station, in case of the Broadcom Linksys and the Atheros Soekris). It is used to issue commands to the oscilloscope and to copy the data traces collected by the oscilloscope. The laptop further configures the local and target station, setting the PHY data rate, the number of frames per second (fps) and the frame size. A commodity switch is used to connect through Ethernet cables the oscilloscope, the local station, the target station and the control station.

We place the local station and the target station at a distance of one meter and send ping echo requests from the local to the target station. Traces are collected using both the 802.11b 1 Mb/s and the 802.11g 6 Mb/s PHY rates. Simultaneously, the control station enforces the oscilloscope through a python script to capture traces at a fixed time interval and store them in the control station's hard disk. We collect from 134 to 400 samples per setup. In addition, the target station is in close proximity to the horn antenna, so that its radio signals are immediately captured by the oscilloscope, and the effect of reflections is minimized. Once the

Table 1: Statistical analysis of PHY SIFS time.

	Atheros Soe.	Broadcom Lin.	Broadcom Sam.	Atheros Soe.	Broadcom Lin.	Broadcom Samsung
Mode	802.11b	802.11b	802.11b	802.11g	802.11g	802.11g - display ON, display OFF
Number of samples	310	296	146	350	400	189 – 134
Median (μs)	10.209	10.326	11.581	16.173	16.816	16.190 – 16.200
CI of median (μs)	[10.206 – 10.211]	[10.319 – 10.335]	[11.464 – 11.801]	[16.170 – 16.176]	[16.805 – 16.823]	[16.183 – 16.194] – [16.193 – 16.207]
Standard Deviation (ns)	12.4	47.7	340.2	13.4	34.8	17.4 – 17.1
Skewness	−0.102	0.186	−0.188	0.017	0.078	0.568 – 0.161
Normality Hypothesis	rejected	not rejected	rejected	rejected	rejected	rejected (dis. ON) - not rejected (dis. OFF)

traces are collected, we downsample them by a factor of five to reduce the computational burden of the data processing. An example of trace is shown on the left of Fig. 6. We are interested in the PHY idle time between the DATA received by the target station and the ACK sent back. Note that we do not process the PHY SIFS of the local station, since the value measured by the oscilloscope may not reflect the one observed by the local station (given that it is not in close proximity).

5.2 Inferring the PHY SIFS time

Since the analysis is conducted based on the traces of the oscilloscope, we have to discriminate the PHY SIFS of the target station (first DATA/ACK observed in the left of Fig. 6) from the one of the local station (second DATA/ACK observed in the left of Fig. 6). The target station is in close proximity to the horn antenna and its measured power is higher than the one of the traffic sent by the local station. Hence, we can infer which station transmitted the frame by monitoring the power level. We further double check this fact by slightly moving one between the local and the target station, and observing the power level's changes in the oscilloscope.

The second problem is to convert the traces from the raw data into busy and idle channel states as needed to sense the presence (busy channel) or absence (idle channel) of a frame on the medium. In addition, the conversion should be robust against any false detection of state change, due to spikes of noise, which would affect the accuracy of the observed PHY SIFS. We therefore apply a Schmitt trigger filter over smoothed samples with exponential weighted moving average and forgetting factor α equal to 0.9. In order to have a common baseline, we apply the same high threshold for the Schmitt trigger, 1.5 dB over the noise (low) threshold, for all the tests. An example of trace is illustrated on the right of Fig. 6. Once the idle and busy state have been determined as above, we finally derive the idle time between the DATA and the ACK based on the delimiters of the busy-to-idle and idle-to-busy transitions, and finally remove outliers in the measurements using again the Thompson Tau technique.

6. EVALUATION OF THE TARGET OFFSET

In this section, we explore the results we obtain with the setup presented in the previous section. Our evaluation shows that:

- The statistical distribution of the PHY SIFS time is chipset (and partially modulation) dependent, and the distribution may or may not pass the normality test according to the specific chipset/ modulation/ traffic rate configuration.
- The median of the PHY SIFS is affected by the number of samples considered for computation (error of up to 10 – 20 ns) and by the selection of the frame rate at the local station (the dispersion increases with the frame rate).
- Tests using a smartphone with display on and off show that the median PHY SIFS varies by 10 ns as a result of the power saving state of the device.
- We find that some chipset/modulation introduces an unacceptable level of degradation of the PHY SIFS.

While the errors above may seem small, their impact on the WiFi ToF echo technique is not negligible. Hence, we remark (cf. Section 3) that an error of 20 ns is mapped to a distance error of 3 m.

6.1 Analysis of PHY SIFS and impact on WiFi echo technique

Table 1 reports a summary of the median and the CI of the median (confidence level 95%) of each test. A first general conclusion we can draw is that the 802.11b (802.11g) PHY SIFS time does not have the nominal value of 10 μs (16 μs) [15] but we observe a systematic error[2]. This result is expected, since the 802.11 standard tolerates up to 1 μs of error of the SIFS time with respect to the nominal value [15]. While in the 802.11g Atheros Soekris and the 802.11g Broadcom Samsung, this bias is approximately 0.17 – 0.2 μs, the 802.11g Broadcom Linksys has a higher bias of approximately 0.8 μs, still in the range tolerated by the standard. On the other hand, the 802.11b Broadcom Samsung does not respect the 802.11 standard (bias of 1.58 μs). Concluding, we can state that the median PHY SIFS time is chipset and modulation dependent, which indicates that this value must be measured/calibrated per chipset and modulation in order to perform ToF-based distance estimation.

Another result can be observed comparing statistically the 802.11g PHY SIFS for Broadcom Samsung in the two cases when the display is switched on and off. When the display is switched off, the smartphone transmits frames periodically in a burst. The reported median varies by 10 ns, which entails a variation of the estimated ToF distance of 1.5 m. We hope that the next generation of WLAN chipsets will reduce this error, allowing to provide an estimate of the distance regardless of the smartphone's power state.

6.2 Evolution of the median of PHY SIFS

Figure 7 displays the evolution of the median of the PHY SIFS over the collected samples for the Atheros Soekris and

[2]Note that the 802.11g SIFS observed at MAC layer is 10 μs, for backward compatibility with the 802.11b standard [15]. This is achieved with 6-μs virtual SIFS extension.

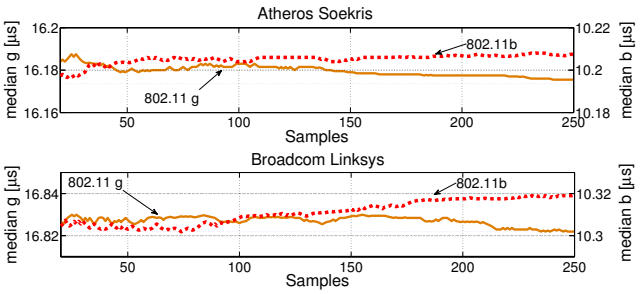

Figure 7: Evolution of the median of PHY SIFS. Median g (median b) indicates the median for 802.11g (802.11b). The median deviates by 10 – 20 ns with respect to the values reported after the very few samples, which would tend to reduce the accuracy of the WiFi echo technique for small number of samples.

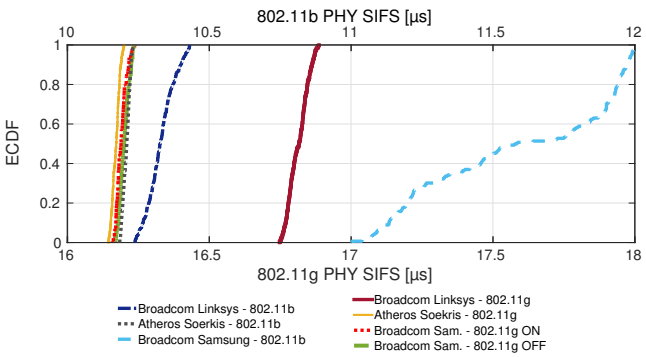

Figure 8: The experiments show that the ECDF of the PHY SIFS is chipset and (partially) modulation dependent. The dispersion is an undesired effect for the WiFi ToF echo technique.

the Broadcom Linksys. The plot clearly shows that the value reported after a few samples for the four scenarios differs from the one measured over several samples by less than 10 ns. A slightly higher variation is encountered for 802.11b Broadcom Linksys. These variations would affect the median accuracy of the WiFi ToF echo technique by up to 1.5 m (higher for 802.11b Broadcom Linksys) when a small number of samples are collected for the estimation.

6.3 PHY SIFS distribution

We then study the distribution of the PHY SIFS time varying the chipset and the modulation. The empirical cumulative distribution function (ECDF) for the three target stations is plotted in Fig. 8. We observe different types of distributions but none of them provides a constant bias over multiple samples: While for the Atheros chipset and the 802.11g Broadcom Samsung we observe a high density of values around the median value, the Broadcom Linksys exhibits a wider deviation of the values. Since the Broadcom chipset in the Linksys is older than the others, this may indicate the tendency towards smaller dispersion of the target noise in newer chipsets, with a standard deviation that is reduced by a factor of at least 3, as shown in Table 1. However, this is not always the case: in 802.11b Broadcom Samsung, the very large dispersion of the PHY SIFS samples ($\approx 1\,\mu$s) is unacceptable to reliably estimate the distance. We further observe that in most of the experiments there is little skewness, with up to 0.18 in most of the cases. Only for the 802.11g Broadcom Samsung with display on, we measure a higher (right) skewness of 0.56.

We then evaluate whether any of the configurations of chipset and modulation is normal by means of both the Lilliefors and Anderson-Darling tests. We find that some of the distributions do not reject the normality hypothesis for a significant level of 0.05 (802.11b Broadcom Linksys with a p-value of 0.43, and 802.11g Broadcom Samsung with display off with a p-value of 0.31). However, all the other configurations of chipset and modulation reject the normality hypothesis.

6.4 Impact of number of frames per second

We evaluate how the number of frames per second affects $t_{OFF,T}$ at the target station. To this end, we compare the

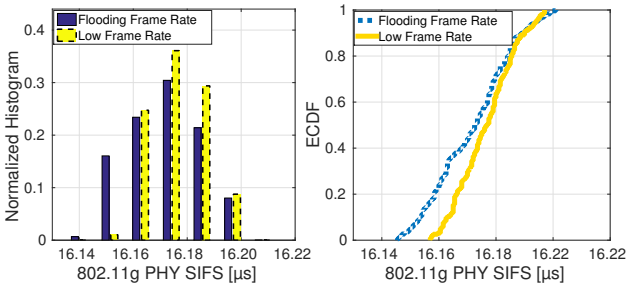

Figure 9: Impact of frame rate. Histogram (on the left) and ECDF (on the right). The dispersion increases with the frame rate.

results of 802.11g Atheros Soekris of previous section with a test where we send 10 fps as opposed to flooding. For ease of representation, we group PHY SIFS samples in bins, with a bin size equal to the 802.11 clock rate, and plot the normalized histogram on the left of Fig 9. On the right of Fig 9, we further show the ECDF for the same tests.

The resulting median for the low frame rate is 16.177 μs (95% CI equal to [16.174 – 16.179]), which is close to the results observed in Table 1 for the flooding frame rate. We further observe a reduction of the standard deviation from 13.4 to 9.9 ns, which would tend to increase the precision of the WiFi echo technique, at the cost of requiring longer time to collect the same number of samples as for the flooding rate.

With regard to the normality test, it is remarkable that the distribution for the low frame rate does not reject the hypothesis test with a significant level of 0.05 and a p-value of 0.09. This is in contrast to the distribution for the flooding rate. This suggests that the high load of the chipset at the target station for the flooding rate does impact the normality of the distribution. However, normality of the PHY SIFS distribution is not guaranteed at the low rate. For instance, tests with the 802.11b Atheros Soekris (not shown in the figure) do reject the normality hypothesis also at the low transmission rate.

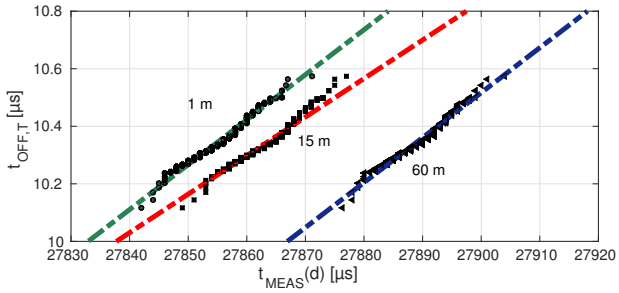

(a) Broadcom Linksys.

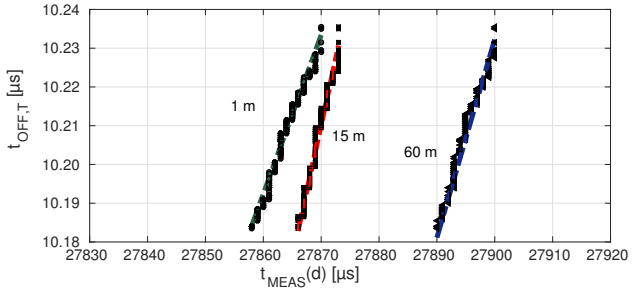

(b) Atheros Soekris.

Figure 10: Quantile-quantile plots. The tests show that $t_{MEAS}(d)$ and $t_{OFF,T}$ come from very similar distributions. It follows that the distribution of the error observed at the local station is mostly caused by the noise of the target offset.

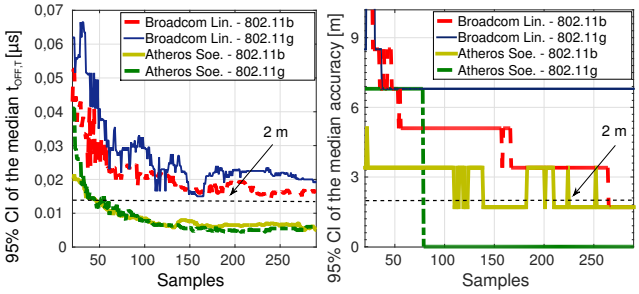

Figure 11: On the left: 95% confidence interval of the median $t_{OFF,T}$ (PHY SIFS time) as a function of the number of samples for the Broadcom Linksys and the Atheros Soekris. The results are in agreement with the precision of the distance measured at the local station (Fig. 4), and reported for ease of comparison on the right, which is further affected by quantization error and timing error in the operation of start and stop of the timer.

7. ANALYSIS OF TARGET VS LOCAL OFFSET NOISE

In this section, we aim to understand the relative importance of the error of the target offset with respect to the noise in the $t_{MEAS}(d)$ samples, locally reported by the WiFi echo technique. Our goal is to quantify how much of the low precision we have for small number of samples (Section 4.2) is originated by $t_{OFF,T}$ or may come from the local offset $t_{OFF,L}$.

7.1 Quantile-quantile plots for $t_{OFF,T}$ and $t_{MEAS}(d)$

We compare the distribution of $t_{OFF,T}$ with the one of $t_{MEAS}(d)$ using the method of the quantile-quantile plot. For $t_{MEAS}(d)$, we show the samples collected at $d = 1$ m, $d = 15$ m, and $d = 60$ m. Regarding $t_{OFF,T}$, we use the PHY SIFS samples collected with the experimental setup presented in Section 5.1. Results for the Broadcom Linksys and the Atheros Soekris for 802.11b modulation at 1 Mb/s are summarized in Fig. 10. (Similar results can be found in other configurations.) We observe a linear pattern in all the tests, which indicates that $t_{OFF,T}$ and $t_{MEAS}(d)$ have a very similar distribution.

7.2 Confidence interval of the target offset and precision of the WiFi echo technique

On the left of Fig. 11, we display the 95% CI of the median $t_{OFF,T}$. We observe that the CI for the Broadcom Linksys is larger than the one from the Atheros Soekris. For a target error of less than 2 m, we require around $45 - 50$ samples for the Atheros Soekris, while the Broadcom Linksys does not reach a precision higher than 2 m in the measurement set (despite being close to it). Besides, comparing these results with the the 95% CI of the median distance accuracy reported on the right of Fig. 11, we find an agreement between the CI of the median $t_{OFF,T}$ and of the median distance accuracy (measured with eq. (3)). However, in general, we observe further degradation of the CI for the median distance accuracy. As already stressed in Section 4, aside from some multipath in the experimental tests, the origin of this additional noise are the quantization error and the noise of starting and stopping the timer at the local station.

From the results on the left of Fig. 11, we also note that, for a given chipset, the CI follows a similar trend for the 802.11b and 802.11g modulation. Not shown in the figure, the 802.11g Samsung Broadcom has a performance in the middle between the other two platforms, while the Samsung Broadcom 802.11b has, by far, the worst 95% CI of the median.

Comparing the results on the left and on the right of Fig. 11, when higher WiFi clock rates for timing measurements are available at the local station, as for our 802.11b experiments, we see that the convergence rate of the CI of the median distance accuracy is still relatively slow, and hence, even higher WiFi clocks are desired. Concluding, the interplay of the noise of the $t_{OFF,T}$ and the quantization error of $t_{OFF,L}$ strongly limit the precision of the distance measurement. It is remarkable that both sources of noise are originated by intrinsic limitations of the underlying hardware, rather than by any software component of the WiFi ToF echo technique.

8. DISCUSSION AND CONCLUSION

We have demonstrated that, with the recent evolution of WiFi echo techniques, the estimation of the distance between WiFi devices is not bounded by any software implementation, but rather by the intrinsic limitations of the underlying hardware. This problem adversely affects the application of WiFi echo techniques for tracking the distance and position of mobile devices. Our in-depth statistical analysis and evaluation with a customized experimental platform for ToF ranging measurements and independent tools to inves-

tigate the origins of noise has dismantled the different components that are limiting the performance figures of current ranging solutions. We have identified and statistically characterized the offset noise of the target station. We have further shown that, together with the quantization error of the main WiFi clock, the target offset is a main source of noise in environments with modest multipath. Given that Fig. 8 shows a target offset's dispersion in the order of $50-200$ ns in most of the chipsets, this dispersion will strongly impact the distance estimation even in harsh environments with large multipath (e.g. 3 m of multipath errors would add a delay of 20 ns to the WiFi echo technique measurement). This would prevent from using standard methods to estimate the distance in harsh multipath environments.

While the quantization error of the local station may be alleviated by newer WiFi chipsets operating on wideband channels (such as the 160 MHz clock of 802.11ac), the variability of the target offset implies that some level of intervention of the target station would be needed to increase the precision of the measurement, such as the 802.11v amendment which would allow location-related timing information to be shared between local and target stations. However, as of today, there is very limited support of 802.11v features in chipsets and drivers. As a result of our investigation, the estimation of the distance must currently be based on a statistical analysis of the noise. Considering that estimators relying on the Gaussian noise assumption of the target device would be suboptimal, robust statistics would be necessary here. The chipset-dependent statistical distribution of the target offset implies that different approaches are used by vendors to implement the SIFS time specification of the standard. Luckily, vendors do not usually exploit all the available dispersion allowed by the standard itself. However, we are far from a very precise ACK generation of the target station to incoming DATA frames.

Acknowledgments

We are grateful to Francesco Gringoli for sharing his openFWWF code and supporting the project with his enlightening comments. We further thank Srdjan Capkun and Aanjhan Ranganathan from the Information Security Laboratory at ETH Zurich, who let us have access to the Infiniium 90000A oscilloscope. This work has been funded in part by the European Commission in the framework of the H2020-ICT-2014-2 project Flex5Gware (Grant agreement no. 671563) and in part by Ministerio de Economía y Competitividad grant TEC2014-55713-R.

9. REFERENCES

[1] D. Lymberopoulos, J. Liu, X. Yang, R. R. Choudhury, V. Handziski, and S. Sen, "A realistic evaluation and comparison of indoor location technologies: Experiences and lessons learned," in *Proceedings of the 14th International Conference on Information Processing in Sensor Networks*, ser. IPSN '15. New York, NY, USA: ACM, 2015, pp. 178–189.

[2] X. Li, K. Pahlavan, M. Latva-aho, and M. Ylianttila, "Comparison of indoor geolocation methods in dsss and ofdm wireless lan systems," in *IEEE Fall VTC 2000*, vol. 6, 2000, pp. 3015–3020 vol.6.

[3] D. McCrady, L. Doyle, H. Forstrom, T. Dempsey, and M. Martorana, "Mobile ranging using low-accuracy clocks," *Microwave Theory and Techniques, IEEE Transactions on*, vol. 48, no. 6, pp. 951–958, 2000.

[4] A. Günther and C. Hoene, "Measuring round trip times to determine the distance between wlan nodes," in *Proceedings of the 4th IFIP-TC6 international conference on Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; Mobile and Wireless Communication Systems*, ser. NETWORKING'05. Berlin, Heidelberg: Springer-Verlag, 2005, pp. 768–779.

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

[6] 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.

[7] D. Giustiniano and S. Mangold, "Caesar: Carrier sense-based ranging in off-the-shelf 802.11 wireless lan," in *Proceedings of the Seventh COnference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT '11. New York, NY, USA: ACM, 2011, pp. 10:1–10:12.

[8] A. Marcaletti, M. Rea, D. Giustiniano, V. Lenders, and A. Fakhreddine, "Filtering noisy 802.11 time-of-flight ranging measurements," in *Proceedings of the 10th ACM International on Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '14. New York, NY, USA: ACM, 2014, pp. 13–20.

[9] P. Gallo, D. Garlisi, F. Giuliano, F. Gringoli, and I. Tinnirello, "WMPS: A positioning system for localizing legacy 802.11 devices," in *Transactions on Smart Processing and Computing,*, October 2012.

[10] A. T. Mariakakis, S. Sen, J. Lee, and K.-H. Kim, "Sail: Single access point-based indoor localization," in *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services*, ser. MobiSys '14. New York, NY, USA: ACM, 2014, pp. 315–328.

[11] L. Sun, S. Sen, and D. Koutsonikolas, "Bringing mobility-awareness to wlans using phy layer information," in *Proceedings of the 10th ACM International on Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '14. New York, NY, USA: ACM, 2014, pp. 53–66.

[12] D. Vasisht, S. Kumar, and D. Katabi, "Sub-nanosecond time of flight on commercial wi-fi cards," *arXiv preprint arXiv:1505.03446*, 2015.

[13] P. Gallo, S. Mangione, and G. Tarantino, "Widar: Bistatic wi-fi detection and ranging for off-the-shelf devices," in *2013 IEEE World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, June 2013, pp. 1–6.

[14] S. Sen, D. Kim, S. Laroche, K.-H. Kim, and J. Lee, "Bringing cupid indoor positioning system to practice," in *Proceedings of the 24th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 2015, pp. 938–948.

[15] "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," pp. 1–2793, 2012.

[16] openFWWF, "Open firmware for wi-fi networks," http://www.ing.unibs.it/openfwwf/.