WMPS: A Positioning System for Localizing Legacy 802.11 Devices

Pierluigi Gallo[1], Domenico Garlisi[1], Fabrizio Giuliano[1], Francesco Gringoli[2], and Ilenia Tinnirello[1]

[1] Universit`a degli Studi di Palermo / Italy pierluigi.gallo@dieet.unipa.it
[2] Universit`a degli Studi di Brescia / Italy

* Corresponding Author: Pierluigi Gallo

Received July 30, 2012; Revised September 20, 2012; Accepted August 19, 2012; Published October 31, 2012

***Abstract*:** The huge success of location-aware applications has called for the rapid development of an alternative positioning system to the global positioning system (GPS) for indoor localization based on existing technologies, such as 802.11 wireless networks.
This paper proposes the Wireless MAC Processor Positioning System (WMPS), which is a localization system running on off-the-shelf 802.11 Access Points and based on the time-of-flight ranging of users' standard terminals. This paper proves through extensive experiments that the propagation delays can be measured with the accuracy required by indoor applications despite the different noise components that can affect the result: latencies of the hardware transreceivers, multipath, ACK jitters and timer quantization. Key to this solution is the choice of the Wireless MAC Processor architecture, which enables a straightforward implementation of the ranging subsystem directly inside the commercial cards without affecting the basic DCF channel access algorithm.
In addition to the proposed measurement framework, this study developed a simple and effective localization algorithm that can work without requiring any preliminary calibration or device characterization. Finally, the architecture allows the measurement methodology to be adjusted as a function of the network load or propagation environments at the run time, without requiring any firmware update.

***Keywords*:** Indoor positioning system, WiFi-based localization, Time of arrival

1. Introduction

The proliferation of location-aware applications is prompting hand-set manufacturers and mobile operators to develop new mechanisms to extend or even outperform the Global Position System (GPS). Assisted-GPS, which is a technique for speeding up the detection of the satellite constellation by retrieving fresh information from 802.11 or 3G infrastructures [1, 2], has paved the way for new systems where the terminal receives so much information from the Network that it can determine its position without the need for GPS. In this context, large players, such as Google and Apple, have recently started mapping Base Station IDs and WiFi network names (beacons) to their corresponding geographical position, allowing localization in urban canyons and in GPS hostile places, such as airports, museums and most indoor environments. Unlike GPS, however, the accuracy of such systems is low [3], and the possibility of refining the positioning using inexpensive and widespread technology, such as 802.11, is very interesting.

Although several commercial systems based on WiFi have been proposed [4, 5], one of the most critical aspects affecting their accuracy depends on the limitations of current implementations in collecting physical parameters, such as the received signal strength or the propagation delay, capabilities that were not included in the original 802.11 Standard. For this reason, recent studies [6, 7] reported ways of improving the localization accuracy by relying on customized measurement instruments (e.g., using FPGA and SDR boards for acquiring the 802.11 signals), and/or complex statistical analysis of long measurement runs [8, 9]. Unfortunately, such customizations offer solutions that are not flexible enough for supporting alternative measurement methodologies that might be required by new applications. This situation can be changed without the need for hardware refinement

simply by reconsidering the software that drives these systems. Modern cards are equipped with very powerful logics and sensing mechanisms, and current issues in collecting accurate localization information are due mostly to the interface available at the driver level rather than to the hardware itself. For example, on one side, some hardware signals are not exposed directly to the driver, whereas the access rate to the exposed signals depends on the host operating system. Starting from these considerations, we define a Positioning System that is *flexible* in terms of the measurement methodology and data processing, and can be *deployed* in off-the-shelf network adapters due to the Wireless MAC Processor [10] architecture. This approach to wireless card programmability has recently illustrated how to use programmable state machines profitably to implement complex MAC protocols on top of a MAC Engine, where the underlying machine code implements only basic operations, such as frame transmission and reception. By adding some measurement primitives triggered by specific events to the legacy DCF state machine, this paper proves the effectiveness of a localization solution based on time-of-flight (i.e. propagation delay) ranging.

The contribution of this paper is two-fold. First, the different noise components of time-of-flight measurements (latencies of the hardware transceivers, multipath, ACK jitters, timer quantization, etc.) and their effects on the ranging accuracy are analyzed in detail. Time-of-flight ranging was also used in reference [11], but the present study is able to perform the measurements by exploiting different triggering events (frame start, frame end, preamble start, etc.) internal to the card and by avoiding preliminary characterization of the devices. Second, an overall localization system was designed, in which an high-level service deployed on the 802.11 infrastructure (namely, the APs) can run-time exploit the low-level measurement functionalities. This paper also proposes a simple and effective adaptation of the Bancroft's localization scheme [12], which is a well known solution for the localization problem devised to avoid any preliminary calibration of both the anchor nodes and target devices. These results show absolute positioning errors lower than 2.5 m in outdoors (and lower than 5 m in 80% of the tested indoor positions), as well as the capability to track the targets moving at pedestrian speeds.

2. Related Work

A number of technologies and approaches have been proposed to solve the critical issues of ranging in indoor environments. Among these, the solutions based on 802.11 appear quite promising because they can rely on pervasive Access Point deployments and user interfaces. Although the initial proposals mainly considered the power received as the main ranging parameter [13], the use of propagation delays was recently proposed [8, 9, 11]. Propagation delays are linearly dependent on the distance, whereas the relationship between the received power and the distance is often quite specific to the propagation environment, thus requiring long calibration phases. On the other hand, propagation delays track the distance of the radio signal, which can be different from the actual distance because of reflections. Moreover, they are normally difficult to measure because off-the-shelf

802.11 cards can only measure the time of arrivals at the driver level with a resolution of 1 μs, which corresponds to a distance quantization error of 300 *m* (unsuitable for indoor localization). The clocks used in commercial cards are also of limited quality and their clock drifts cause significant errors. Two-way measurements are often used to solve this issue because they do not need to synchronize the clocks (used for timestamps) of independent nodes [16].

In references [8, 9], it was shown that the limited resolution of the delay measurements can be overcome by statistical means, such as stochastic resonance. Unfortunately, statistical techniques require the observation of a large number of samples, which prevent the tracking of moving objects. A different approach was proposed in [6, 7], where dedicated mechanisms for measuring the propagation delay were designed using either external hardware [6] or software-radio [7]. The solutions are quite expensive because these measurements need to be taken at each anchor in an indoor environment. To reduce the number of independent measurements required for localization, [14] proposed to perform indirect measurements of four-way propagation delays. An alternative solution based on cheap Atheros cards was described in [11]. Owing to the availability of a well documented open-source driver, the measurement resolution was improved using the host CPU for opportunistically polling some card registers. In particular, these registers account for the cumulative time intervals in which the medium has been sensed idle and busy, and are updated at the card internal clock, i.e., at 44 MHz. Reference [15] reported that for typical indoor environments, even in presence of a strong multipath (without line of sight propagation), the average error on the propagation delay is 1 clock cycle at most (at 44 MHz). Such an error can be compensated for if multiple reference points with independent propagation conditions are available for trilateration.

Apart from the measurement resolution, another critical aspect of delay-based ranging solutions is the need to estimate and compensate for hardware-dependent measurement delays, which are added to the propagation delays due to clock drifts [17], transceiver latencies and jitters in the ACK frame scheduling, i.e. in the SIFS interval. References [18, 19] and [20] proposed a characterization of the hardware additional delay to identify the fingerprinting of different cards. In reference [11], such a hardware additional delay was assumed to be bi-modal and related to the quality of the received signal. By jointly measuring received power levels and propagation delays, the ranging was then improved by compensating for the estimated hardware delay.

3. Ranging

This paper proposes a solution for performing propagation delay measurements with a resolution of 1/88

MHz using ultra-cheap 802.11 cards. Rather than modifying the card driver and polling the card registers as proposed in reference [11], the solution is completely *internal* to the card, thereby avoiding the uncertain delays introduced by the host operating systems. The approach was enabled by the Wireless MAC Processor architecture recently proposed in [10], whose implementation has also been released in [22] for the Broadcom AirForce54G cards. As described in the next subsection, this architecture can integrate a (programmable) measurement framework to the MAC operations without any firmware update.

3.1 Platform

The Wireless MAC Processor (WMP) architecture enables the execution of different MAC schemes, which are specified in terms of programmable state machines, on the same hardware. Indeed, MAC protocols can be described effectively in terms of the state machines provided that a complete list of actions (such as transmit a frame, set a timer, build an header field, switch to a different frequency channel, etc.), events (channel up/down signals, indication of reception of specific frame types, expiration of timers, enqueueing of a new packet, etc.) and verifiable conditions (frame address, medium status, RTS thresholds) are available.

The proof-of-concept of such a vision was obtained by replacing the firmware of a commercial card (for which an open firmware was available) with a new firmware implementing an executor of generic state machines and a core list of events, conditions and actions [10]. Because the events exposed as a programming interface also include the hardware signals involved in a propagation delay measurement, this paper proposes adding a few more actions (devised to perform ranging estimates) to the MAC state machine. Apart from the solution described in 3.2, the approach is intrinsically programmable and several methodologies, such as multiple handshake delays, can be supported by specifying different events for activating and stopping the time measurements.

3.2 Methodology

In this solution, the propagation delay measurements are carried out during the medium access operations and can work on normal DATA/ACK frame handshakes between a target station, i.e. a station to be localized, and an anchor station, i.e. a reference station whose location is known.

Measurement Operations are quite simple: a timer can be activated immediately after a DATA frame transmission and stopped at the start of the ACK reception. Both the MED START and MED END events (signaling the start and end of channel activity) are available on the WMP programming interface. Because the MED END events are captured more precisely than the MED START events, which might be delayed in the case of low power quality [11], it was decided to modify the previous procedure slightly by stopping the timer at the end of the ACK reception (as shown in Fig. 1). This results in $t_{MEAS}(d) = 2 \cdot t_P(d) + T_{SIFS} + T_{ACK} = q + m \cdot d$, i.e. it is linearly dependent on the radio distance, where $t_P(d)$ is the propagation delay, $m = 2/c$, and c is the speed of the light. Timing errors due to the latency required for activating (t_{SET}) and stopping the timer (t_{STOP}) can be considered as additive noise components.

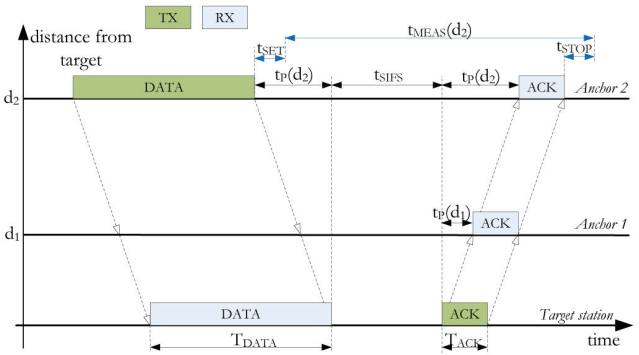

Fig. 1. Two-way delay measurements: triggering events and timing errors.

In principle, the station performing the measurements, i.e. the ranging station, can be either the target or the anchor station, provided that it supports the WMP architecture. This suggests that the ranging station can work on frame handshakes that begin with its own DATA frame transmissions, whereas the other station (sending the ACK replies) can be a legacy station. This type of measurements is called active measurements. As shown in the figure, in this paper, the delay measurements were performed by the anchor stations to have a system that can work with every (legacy) target station. The only requirement for the target stations is to be associated with the anchors of the ranging system thereby enabling the ACK replies to the DATA frames. The figure also shows that the anchor stations can work in passive mode, as in the case of anchor 1, by activating and stopping the timer during a frame handshake beginning from an another station. This solution allows the measurement overheads to be limited by exploiting traffic frames for ranging or by multiplying the number of independent measurements carried out on a single frame handshake.

Measurement Resolution obviously depends on the clock of the ranging stations. The internal clock of the card used for the WMP implementation was confirmed to works at 88 MHz. The resulting quantization error on $t_{MEAS}(d)$ was 1000/88 ns, which corresponds to a distance of 3.4/2 = 1.7 m at the speed of light. Such an error was considered acceptable for common localization applications. The measurement errors (due to quantization and latencies in the timer management) and physical errors could also be distinguished.

Physical errors are caused by differences in the actual relationship between delay and distance and the expected one. A number of factors can cause such a misalignment. First, because of the multipath, the measured delay can refer to the strongest radio path rather than to the direct one, resulting in a higher estimated distance. Second, because the measurement refers to two-way propagation delays, it is affected by errors on the ACK scheduling time. In fact, the standard defines a tolerance on the SIFS

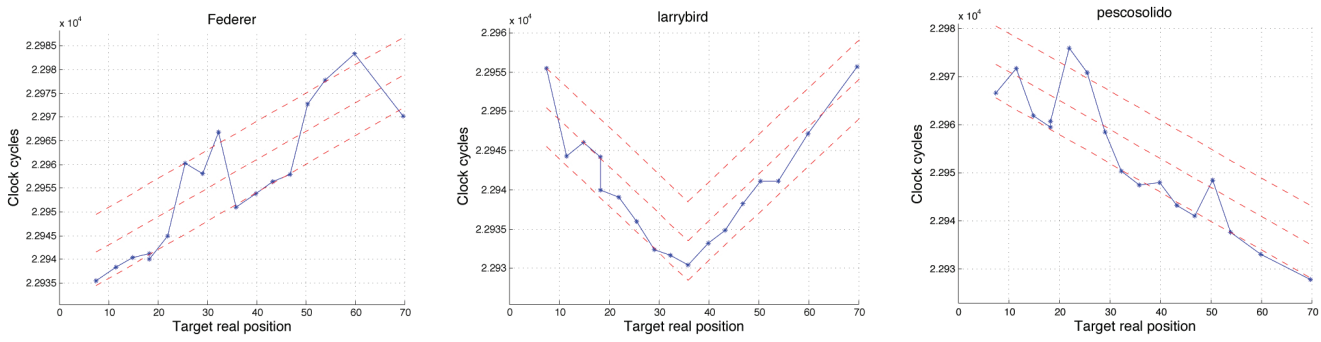

(a) (b) (c)

Fig. 2. Delay measurements carried out by three independent anchors while a target station moves along a 70 m.

scheduling (up to $1\mu s$), which is reflected into a measurement error that can be systematic or not for the target stations produced by different vendors. Finally, also for anchor nodes, the physical implementation of the transreceiver introduces some latencies in revealing the MED_START and MED_END events, which might have a random component. These latencies can be higher in passive measurements because in this case, the receiver latency affects both the stopping and starting of the timer.

3.3 Experimental results

As a first ranging experiment, three different anchor points (called federer, larrybird and pescosolido) were set up in a long corridor of approximately 70 m at the University of Brescia. Two anchor nodes were placed at the corridor edges, whereas a third is placed precisely in the middle. All the anchors were equipped with an Airforce54G wireless interface running the WMP and loaded with the modified DCF state machine. The target station used the same wireless adaptor but with the original DCF firmware.

Fig. 2 shows the measurements carried out independently by each anchor node. Each measurement was averaged by considering a large number of measured samples (e.g. 1000 different samples). The delay measurements are expressed in terms of clock cycles, which should be linearly dependent of the radio distance with a slope equal to $2/c \cdot 88MHz = 0.5871$ clock cycle/m. Indeed, such a linear relationship works well when the target is close to the anchor node (within a distance of approximately 20 m), whereas it can deviate from the linear relationship for larger distances. For example, in Fig. 2(a), it is evident that the delay versus distance relationship switches between two different lines (the bottom one, from 5 m to 20 m and from 35 m to 50 m, and the upper one from 20 m to 35 m and from 50 m to 60 m), which have identical slopes and different constant terms. Such a phenomenon may due to the multipath, because the dominant path of the received signal can change as the target node moves along the corridor. On the other hand, the phenomenon is evident for large distances, where the difference between the direct and reflected path should be, in percentage, less evident. For example, when the target node moves from 32.5 m to 35 m, the federer anchor node measures a delay difference of approximately 18 clock cycles which corresponds to a path difference of approximately 60 m. At a corridor width and height of approximately 5 m, such a difference cannot be justified by a single reflection.

For a better understanding of the multipath effect, some experiments were also carried out outdoors, where it is reasonable to assume that the dominant path is the one in direct line of sight. Fig. 3 shows the results of different measurement campaigns performed in the outdoor court of the University of Brescia with a single anchor point. The timing measurements were mapped to distance estimates by considering a reference distance of 5 m (for evaluating the q coefficient).The distance range was limited to 25 m, because the phenomena of interest were observed, even in this reduced range. Despite the experiment being performed outdoors, the expected linear relationship between the delays and distance were affected by fluctuations of approximately 6 clock cycles (i.e. 20 m) even at short distances. The same experiment was repeated using a directional Yagi antenna with horizontal polarization for the anchor node and a dipole antenna with the same polarization for the target one. The use of a directive antenna on the anchor node has an impressive effect in improving the ranging performance, as shown in Fig. 3(b). Such a result might be due to the improved quality of the ACK signal (at the anchor node) that reduces the latency on the MED_END events. On the other hand, in Fig. 3(c), where the Yagi antenna is polarized vertically while maintaining the horizontal orientation of the dipole, the ranging was still better than in the omni-directional case, with the exception of a single point where there was again an error of approximately 60 m.

Based on these results, it was suspected that the main cause of the so called physical errors is the variability of the received signal levels in the subsequent measurement samples rather than their absolute value because such a variability has some effects on adjusting the anchor transceiver.

Note that systematic errors on the SIFS timings are not a problem because they are eliminated by the localization scheme as a constant additional bias (similar to the clock bias in GPS), whereas SIFS random jitters might originate positioning errors because they vary at each frame handshake started by different anchors.

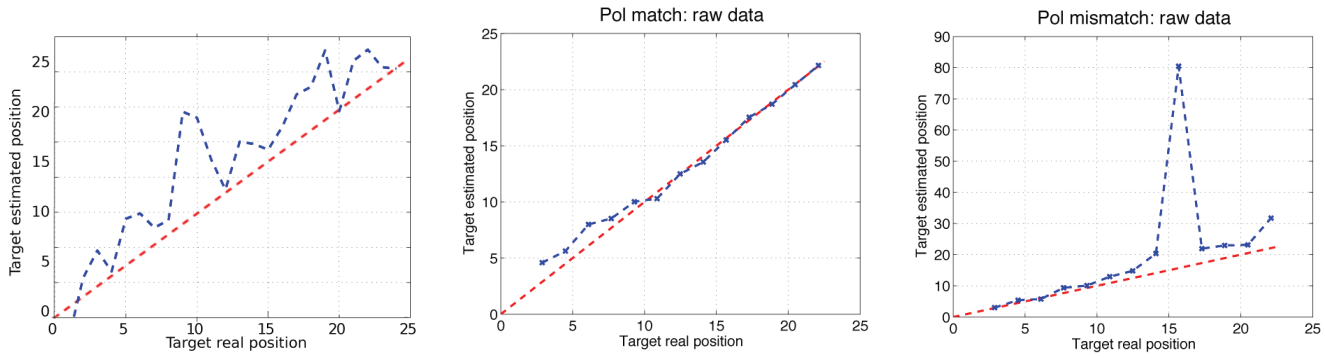

Fig. 3. Distance estimate in an outdoor scenario under different propagation conditions (omindirectional antenna (a), directional antenna pointed (b) and not pointed (c) to the target.

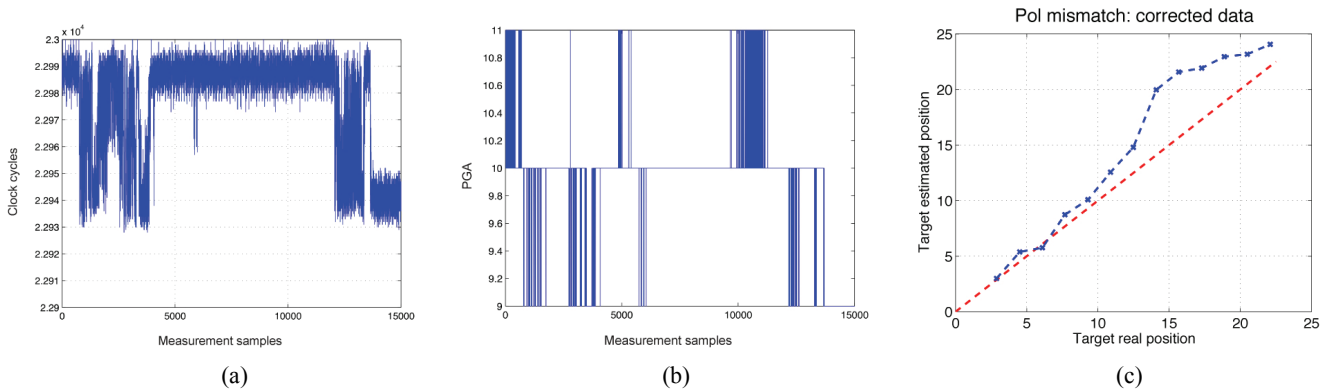

Fig. 4. Effects of AGC on the two-way delay measurement: temporal trace (a), quantized programmable gain (b), and distance estimate with corrections (c).

4. Dissecting Receiver Noise Components

The ranging results presented in the previous section were obtained by considering a long acquisition interval of the measurement samples at each ranging position. The goal to identify the stationary errors by attempting to filter as much as possible the random measurement noise due to quantization and timing inaccuracies. Despite the long measurement runs, systematic errors where found, whose origin did not appear to be related to multi-path propagation. Fig. 4(a) shows a temporal trace of the delay measurements (in clock cycles) for a run of 15000 consecutive handshakes captured at the same distance of 15 m. The figure suggests that apart from the measurement noise, there is a non-stationary noise component which makes the average values oscillate between the two values. A similar phenomenon was also reported in reference [11], and justified in terms of the different latencies of detection mechanisms implemented in an Atheros card for revealing strong and weak packets.

Starting from these considerations, this study conducted further tests to analyze the impact of the latencies introduced by the anchor transceiver, as well as by the jitter of the ACK scheduling times. The characterization of these noise components is important for refining the measurement methodology, providing distance estimates with the desired confidence and tuning opportunistically the observation interval.

4.1 AGC

The first element that this study attempted to quantify was the impact of the Automatic Gain Control (AGC) on the detection of measurement events. The AGC adjustments due to sudden variations in the received signal quality can introduce measurement delays. The WMP architecture implemented on the Broadcom open firmware allows the time-varying programmable gain introduced by the receiver amplifier to be read in a 4 bit quantization scale.

Fig. 4(b) shows the quantized value of the programmable gain amplifier (PGA) for the same experiment plotted in Fig. 4(a). The shortest delay measurements corresponded to a value of 9, whereas the highest delays were obtained when the PGA value was 10 or 11. In several other experiments, it was found consistently that a high value of the PGA always corresponds to additional measurement latency. This latency could be mapped to a significant distance error (up to 60 m in Fig. 3(c)) when the expected delay versus distance relationship was tuned for a short reference distance for which the latency was absent. By correlating the delay measurements with the PGA values, it was possible to correct such an additional latency and produce

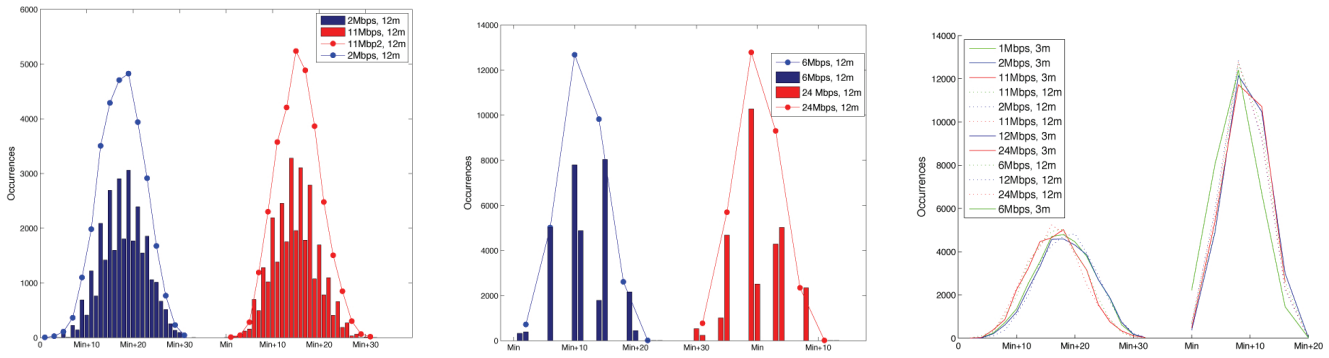

(a) (b) (c)

Fig. 5. Distributions of the measurement samples collected under 802.11b (a) and OFDM (b) modulation schemes and at different distances (c).

a more reliable distance estimate, as shown in Fig. 4(c).

These considerations also justify our previous findings regarding the impact of directional antennas on the ranging accuracy. Indeed, a directional antenna allows a lower PGA value when correctly pointed, or a consistently high PGA value when non-pointed. Ranging errors emerge only when the distance vs. delay relationship is derived under a given PGA condition and maintained for all values. Although these results were obtained for this reference acquiring card, i.e. the Airforce54G one by Broadcom, it is reasonable to assume that the conclusions can be generalized, i.e. the amplifier latencies depend on the amplification entity and need to be compensated for by producing a reliable ranging mechanism, e.g. by updating periodically the bias of the distance function, as described in 5.2.

4.2 Transreceiver latencies

Fig. 4(a) shows that for a given latency of the AGC block (e.g. from sample 5000 to sample 10000), there are some other measurement fluctuations. The measurement distributions obtained for a fixed PGA value on a long run of 30000 measurement samples were analyzed to determine which of the random latencies discussed in 3.2 is mainly responsible for these fluctuations and understand how to limit or filter such a noise component.

Fig. 5 summarizes these results for various 802.11b and OFDM modulation schemes while the target was 12 m and 3 m distant from the reference anchor. The ACK frames were sent at the data rate. Rather than plotting the actual clock values, which obviously depend on the employed data rate, to simplify the comparison of the curves, the delay distribution were expressed as a function of the minimum clock value of 99% of the samples. The number of occurrences for each integer value in the range [*Min, Min* +Δ] is indicated by the bars in the figure, which refer to 2 Mbps and 11 Mbps in (a) and 6 Mbps and 24 Mbps in (b).

The first interesting observation that emerges from the figure involves the quantization and timing errors. Although the timer available on the anchor has a quantization error of 1 clock cycle (1/88MHz), the bars plotted in Fig. 5(a) appear to suggest a further uncertainty of an additional cycle, which is likely to be due to the WMP firmware execution (a clock cycle is required to respond to the events and starting or stopping the timer). In other words, a given delay resulting from the propagation delay and transreceiver/ACK latency is mapped into two consecutive bars because of the timer management. A similar phenomenon is presented in Fig. 5(b), where there is another quantization effect of 4 clock cycles (from the start of a group of two bars to the next one) which is justified as a quantization of the transceiver capability to detect the MED_START and MED_END events in the case of OFDM modulations.

The figure also plots the discrete curve obtained by summing the values of the consecutive bars, to allow an easier a comparison between different distributions. Fig. 5(c) shows many overlapped distributions obtained for different distances and rates. For OFDM modulations, the delay spread of 99% of the samples was lower (approximately 20 clock cycles) than that for the 802.11b (approximately 30 clock cycles). Moreover, the shape of the curve was quite insensitive to the distance and to the specific modulation and coding scheme. These observations suggest that the delay measurements performed on OFDM frames have greater *precision* than the ones performed on 802.11b frames (i.e. they have a lower variability). On the other hand, because of the receiver implementation, the measurements performed on the OFDM packets exhibit a quantization of 4 clock cycles, which generally reduce the measurement *accuracy* (i.e. how the estimated range is close to the actual one).

4.3 ACK jitters

The spread of the previous distributions is given by the sum of two independent latencies: the latency due to the anchor transreceiver in revealing the measurement events, and the errors due to the target node in scheduling the ACK frames. Indeed, the standard allows a jitter of 1 μs on SIFS scheduling (i.e. 88 clock cycles), which might completely impair any ranging mechanism based on the propagation delay measurements. To identify the impact of the ACK jitters, an experiment was carried out with two

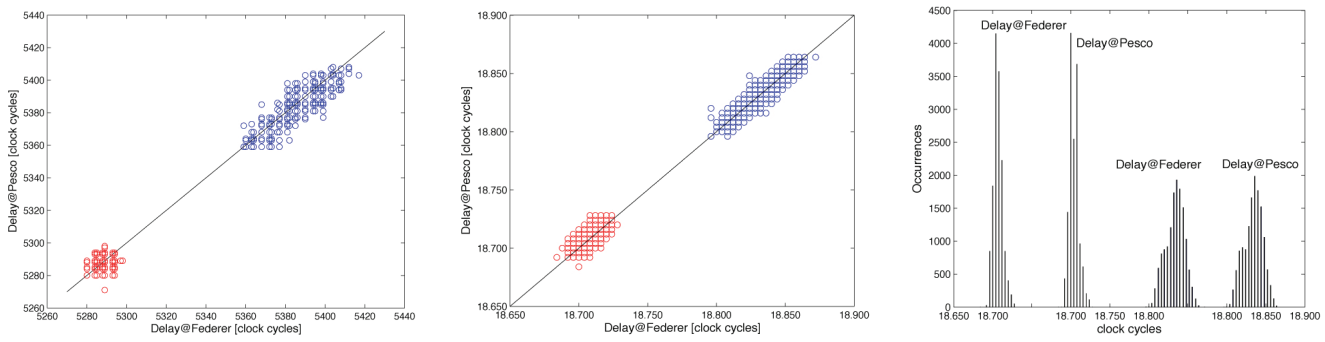

(a) (b) (c)

Fig. 6. Effects of ACK jitters introduced by two different targets (blue points for an iphone, red points for a macbook): correlated delays samples of two independent anchors at 6 Mbps (a) and 11 Mbps (b) and the overall delay distributions (c).

independent anchors performing delay measurements in *passive* mode on the same frame handshakes. The two anchors were placed at a similar distance from the target (actually, there was a 2 m of difference, which is within the clock cycle precision). They registered the sequence number of the data frames for each corresponding measurement sample. The measurements carried out by the two anchors were then correlated based on these sequence numbers. The experiment was then repeated for two different targets (namely, an iphone and a mackbook) and at two different rates (6 Mbps and 11 Mbps).

Fig. 6 summarizes the results of these experiments. The first two subfigures were obtained by plotting a point cloud, where each point (x_i, y_i), refers to the measurement x_i and y_i performed respectively by the anchor node federer and pesco, respectively, on the same *i*-th frame handshake. Obviously, in the case of perfect ACK scheduling, a square shaped cloud would expected because for each noise value introduced by the transceiver of a given anchor, the second anchor noise would have been completely independent (between the maximum and minimum most likely values). On the other hand, the figure shows that the two measurements are strongly correlated, i.e. when an anchor produces a given measurement, the one produced by the other anchor spreads on a smaller set of values than the overall possible values. Although the vertical (horizontal) points collected for a given $x(y)$ values can be related to the transceiver noise, the overall spread of the measurement values is related to the ACK jitters. The phenomenon is evident for both the rates (Fig. 6(a) and 6(b)) and targets, with an evident lower accuracy of the iphone (blu points) in ACK scheduling.

Although the point spread appears to be much higher than in Fig. 5, a closer look at the number of occurrences shows that most of the measurement samples (90%) fall in a range of 30 clock cycles at 11 Mbps and 20 clock cycles at 6 Mbps. The range is slightly wider for the 99% of the samples, because the passive measurements refer to different events than the active ones. Fig. 6(c) shows the overall distributions of the samples.The two independent anchors (by the same vendor) have almost identical measurement distributions.

5. Localization System

The proposed localization system was designed by exploiting the previous findings on the time-of-flight measure-ment accuracy and precision. The system includes the following components: i) the WLAN infrastructure, which is made up of Access Points supporting the WMP paradigm and the (programmable) measurement functionalities, and equipped with directional antennas; ii) the target node, which is equipped with a legacy 802.11 interface; and iii) a localization service running both on the target node and Access Points, for activating, collecting and processing the ranging measurements. Although the measurements are entirely demanded by the WLAN infrastructure, which is also responsible for the simple pre-filtering operations, the localization algorithm runs at the target node. The same data frames sent by the Access Points for measuring the two-way propagation delays can be used to send the previous filtered measurements to the target node.

5.1 Measurements

Once the target node is associated to the WLAN infrastructure and begins the localization service, each Access Point in visibility (receiving a signal from the target node at a value higher than a minimum threshold) continuously performs the delay measurements. While the serving Access Point can exploit downlink traffic frames for starting the two-way measurements, the other Access Points necessarily add a traffic overhead for sending their ranging data frames. Although for space reasons this paper does not describe the state machine responsible for managing the ranging frames, is is believed that such a mechanism can be modified slightly (due to the WMP flexibility) according to the network load conditions and the number of targets to be localized simultaneously. For example, to prioritize the ranging frames, it is possible to schedule these transmissions with shorter backoff values or to send back-to-back a multiple number of ranging handshakes spaced of a SIFS time. Because of the receiver

For the iphone case, the measurement spread is approximately 70 clock cycles, which is still compatible with the 1μs SIFS accuracy.

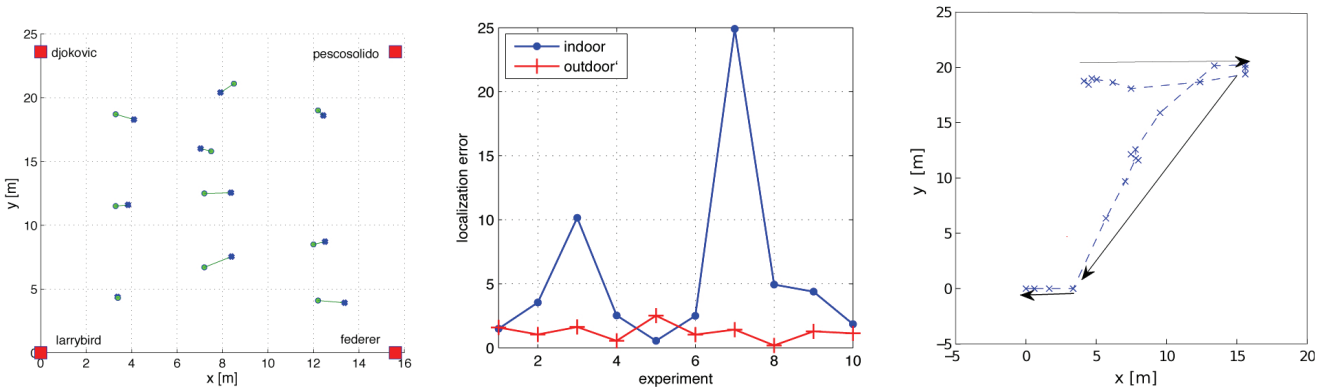

(a) (b) (c)

Fig. 7. Localization experiments: actual position (green points) and estimated ones (blue points) in an open court, errors in different positions for indoor and outdoor experiments (b), dynamic tracking (c).

random latencies, it is required to use a number of measurements between 20 and 100 for averaging before applying the localization algorithm. Because accuracy is more important for this application than precision (being more anchors available for independent ranging), we prefer to use 802.11b modulation schemes for the ranging frames, even though OFDM high rates can reduce the channel resource consumption.

For a rough evaluation of the channel resource consumption due to the measurements, the minimum channel time for an handshake performed at 11Mbps with the shortest possible data frame can be considered to be equal to 425 μs, thus resulting in a 0.425 20 4=34 ms minimum channel time (4 being the minimum number of anchor points required by the localization algorithm).

5.2 Localization Algorithm

The positioning problem can be faced by the solutions already established for GPS, by considering the peculiarity of this system. For each anchor point a_i, the target node periodically receives the information t^i_{MEAS} on the estimated two-way delay, whereas the information on the anchor position (x_i, y_i) can be considered to be known to all the nodes because the anchors are static and this information can be broadcast periodically in the beacon frames. The time measurements can be converted to pseudo-distances by considering the expected m slope (i.e. 0.5871 clock cycle/m) of the time vs. distance relationship. Anchor clocks are obviously non synchronized and are generally of a poor quality (typically, with an error of 25 ppm). On the other hand, the lack of synchronization is not a problem for two-way measurements that do not require timestamp comparisons, whereas the different speeds of two anchor clocks correspond to a much lower error than the quantization error for the typical indoor distances.

Basic Scheme. The information relative to each anchor node can be organized in a vector $\mathbf{a_i} = [x_i, y_i, d_i]$, where d_i is the pseudo-distance to the target (i.e. t^i_{MEAS} / m), and the unknown information relative to the target node node in a vector $\mathbf{u} = [x, y, b]$, where b is the fictitious distance bias due to the measurement methodology (i.e. because even with a 0 propagation delay, t_{MEAS} includes the SIFS time, the ACK duration and the latencies introduced by the receiver). With such a formalization, the positioning problem corresponds precisely to the GPS case, provided that the bias b depends on the target only (and not on the anchors). By imposing this on each anchor node, the sum of b and the distance between (x, y) and (x_i, y_i) is equal to d_i, it is possible to find u, without the need for a preliminary evaluation of b. This suggests that there is no need to measure the specific SIFS value of the devices produced by different vendors as considered in reference [11]. An efficient implementation for solving the previous system of equations is to apply the Bancroft's algorithm that transforms the nonlinear problem to a linear algebra problem.

Possible Generalizations. The previous scheme is based on the simplifying assumption that the bias bdepends on the target node only. As discussed in 3.2, this is true for most delay components included in t_{MEAS}, which contain the ACK duration and the actual SIFS time. On the other hand, the latency introduced by the anchor receiver is only partially corrected by the AGC delay compensation; this latency can vary according to the propagation conditions, resulting anchor-dependent and potentially, for mobile targets, even time-varying. When calibration is not available or not desired, it is possible to apply iteratively some adjustments. For example, the target node can cooperate with the network infrastructure by sending back the results of the previous positioning estimates. Based on these estimates and its own measurements, each anchor can opportunistically translate the delay vs. distance line to compensate for its specific bias.

5.3 Experiment Results

Fig. 7 shows some of results of the localization experiments. The performance of the localization system was analyzed in different (fixed) target locations (a-b) and under target mobility (c). Fig. 7(a) shows the actual

The delay samples were filtered using a simple moving average filter. Before being passed to the filter, each measurement sample was correlated with the corresponding PGA value, in order to remove the additional AGC delay when necessary.

coordinates of the target (green points) as well as the estimated ones, with four anchor points located at the vertices of a square in an open court at the University of Brescia. The propagation environment has a few obstacles, whereas the anchor antennas are all oriented towards the internal space. Under these (almost ideal) conditions, 100 different measurement samples were used to produce a single pseudo-distance and the basic localization schemes were run. The figure clearly shows that the positioning works very well in all the different locations. This is quite impressive, considering that no preliminary calibration was performed for the anchors and the target node. Although the figure refers to a target adapter of a given brand, i.e. a Broadcom card, a similar performance was obtained with the adaptors of different brands.

Fig. 7(b) plots the errors (in terms of the distance from the actual position) for each of the positions indicated in the map, as well as for other indoor experiments. For the outdoor experiments, an error $\leq 1.5m$ can be quantified in 9 positions out of 10; it becomes $\leq$ 3m by using a measuring window of 20 samples rather than 100. Indoor experiments were carried out by placing two anchor nodes in a corridor and two other anchor nodes in a perpendicular one, in order to create conditions for path reflections. The basic localization scheme still provided good results (with an error $\leq 5m$) in 8 positions out of 10. On the other hand, there were two cases in which the positioning was poor. This was attributed to the different bias experienced by each anchor that is not addressed by the Bancroft's scheme.

Finally, Fig. 7(c) shows a run-time experiment, during which the target position was tracked according to the path shown in the figure. The tracking was based on a measuring window of 100 samples and on the basic localization scheme because the experiment was performed outdoors. The ressults were considered satisfactory.

6. Conclusions

Current solutions for indoor localization based on 802.11 do not meet the set of conflicting requirements, such as high precision, fast convergence, and minimal environmental calibration. This limits the utilization of WLAN infrastructures to assisting existing navigation systems or providing localization services for indoor areas. Indeed, off-the-shelf devices offer limited access to the low-level hardware signals that can be exploited for ranging, resulting in the need to adopt simple ranging schemes based mainly on the received signal strength indicator, or deploying customized anchor nodes based on dedicated hardware or SDR.

This study examined the potentialities arisen by the availability of advanced API (accessing the hardware signals) for developing localization functionalities on future wireless cards on top of which a localization positioning system can be built. In particular, this paper discussed how this API can be used to implement a ranging solution based on time-of-flight measurements. After dissecting the benefits and limits of the approach, some interesting conclusions were drawn. The multipath has a negligible impact (in terms of the final localization error) on this type of measurements. The use of directional antennas on the anchors can improve the performance of the ranging system. Higher amplifier gains correspond to higher receiver latencies that can be compensated for by providing reliable delay measurements. Localization experiments exploiting the proposed measurement framework showed absolute errors $\leq 2.5m$ in an outdoor environment and $\leq 5m$ in 80% of the indoor positions tested, as well as the ability to track targets moving at pedestrian speeds.

Acknowledgement

This work has been carried out in the frame of the EU FP7-FLAVIA project, contract number 257263.

References

[1] A. Varshavsky, E. de Lara, J. Hightower, A. LaMarca, and V. Otsason, "GSM indoor localization", Pervasive and Mobile Computing 3(6), pp. 698-720, Dec. 2007. Article (CrossRef Link)

[2] H. Lim, L.C. Kung, J.C. Hou, and H. Luo, "Zero-configuration indoor localization over IEEE 802.11 wireless infrastructure", Wireless Networks, 16(2), pp. 405-420, Feb. 2010. Article (CrossRef Link)

[3] P802.11v -Amendment to Standard for Information Technology Telecommunications and information exchange between systems Local and Metropolitan networks specific requirements Part II: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: IEEE 802.11 Wireless Network Management. Article (CrossRef Link)

[4] Awiloc positioning system by Fraunhofer Institute, http://www.iis.fraunhofer.de/en/bf/ln/technologie/rssi.

[5] Cisco Wireless Control System, http://www.cisco.com/en/US/products/ps6305/index.html.

[6] M. C. Adell, "Contributions to TOA-based location with WLAN, Ph.D. dissertation, Universitat Politecnica de Catalunya, Departament dEnginyeria Telematica, Barcelona, Spain, Jan. 2010.

[7] D. Humprey, M. Hedley, "Super-Resolution Time of Arrival for Indoor Localization", in Proc. IEEE ICC '06, 2008. Article (CrossRef Link)

[8] A. Gunther and C. Hoene, "Measuring round trip times to determine the distance between WLAN nodes, in Proc. Networking 2005, Waterloo, Canada, 2005. Article (CrossRef Link)

[9] K. I. Ahmed and G. Heidari-Bateni, "Improving two-way ranging precision with phase-offset measurements, in Proc. IEEE GLOBECOM 06, 2006. Article (CrossRef Link)

[10] I. Tinnirello, G. Bianchi, P. Gallo, D. Garlisi, F. Giuliano, F. Gringoli, "Wireless MAC Processors: Programming MAC Protocols on Commodity Hardware," in Proc. IEEE INFOCOM '12, March 2012. Article (CrossRef Link)

[11] D. Giustiniano, S. Mangold "CAESAR: Carrier Sense-Based Ranging in Off-The-Shelf 802.11 Wireless LAN," in Proc. ACM CoNEXT '11, New York, NY, USA. Article (CrossRef Link)
[12] S. Bancroft, "An Algebraic Solution of the GPS Equations," IEEE Transactions on Aerospace and Electronic Systems, vol. AES-21, no.1, pp.56-59, Jan. 1985. Article (CrossRef Link)
[13] P. Bahl, V.N. Padmanabhan, "RADAR: an in-building RF-based user location and tracking system," in Proc. IEEE INFOCOM '00, vol. 2, pp. 775-784. Article (CrossRef Link)
[14] C. Hoene and J. Willmann, "Four-way TOA and software-based trilateration of IEEE 802.11 devices, in Proc. IEEE PIMRC, Cannes, Sep. 2008. Article (CrossRef Link)
[15] M. Ciurana, F. Barcel´o, S Cugno, "Multipath profile discrimination in TOA-based WLAN ranging with link layer frames," in ACM WiNTECH '06, Los Angeles, CA, USA, pp. 73-79. Article (CrossRef Link)
[16] S. A. Golden and S. S. Bateman, "Sensor measurements for Wi-Fi location with emphasis on time-of-arrival ranging, IEEE Transactions on Mobile Computing, vol. 6, no. 10, 2007. Article (CrossRef Link)
[17] R. Krishna and C. Hoene, "Calculating relative clock drifts using IEEE 802.11 beacons, in Proc. IEEE INDICON, Ahmedabad, India, Dec. 2009. Article (CrossRef Link)
[18] Bartlomiej Sieka, "Active Fingerprinting of 802.11 Devices by Timing Analysis," in Proc. IEEE Consumer Communications and Networking Conference, 2006. Article (CrossRef Link)
[19] Günther Lackner, Udo Payer and Peter Teufl, "Combating Wireless LAN MAC-layer Address Spoofing with Fingerprinting Methods," International Journal of Network Security, Vol.9, No.2, PP.164-172, Sept. 2009. Article (CrossRef Link)
[20] Guenther Lackner, Peter Teufl, "IEEE 802.11 Chipset Fingerprinting by the Measurement of Timing Characteristics," Australasian Information Security Conference, Perth, 2011. Article (CrossRef Link)
[21] Open firmware for WiFi networks, http://www.ing. unibs.it/openfwwf/.
[22] WMP Wireless MAC Processor

Pierluigi Gallo has been an Assistant Professor at the University of Palermo since November 2010. He graduated with distinction in Electronic Engineering in July 2002 and worked at CRES (Electronic Research Center in Sicily) untill 2009. His work there was dedicated to QoS in IP core routers, IPv6 network mobility andWireless Networks. His research activity has focused on wireless networks at the MAC layer and 802.11 extensions, localization based on the time of arrival and cross layer solutions. P. Gallo has contributed to several national and European research projects: ITEA-POLLENS (2001-2003) on a middleware platform for a programmable router; IST ANEMONE (2006-2008) about IPv6 mobility; IST PANLAB II on the infrastructure implementation for federating testbeds; ICT FLAVIA (2010-2013) on Flexible Architecture for Virtualizable future wireless Internet Access.

Domenico Garlisi was born in Italy on December 8, 1984. He received his Masters Degree in Telecommunication Engineer-ing from the University of Palermo, Palermo, Italy in 2010. He is currently a Ph.D student in the Department of Telecommu-nication, University of Palermo, 2012. He is also working as collaborator researcher for CNIT (Consorzio Interuniversitario per le Telecomunicazioni), in the section of European Project FLAVIA (http://www.ict-flavia.eu).

Fabrizio Giuliano was born in Italy on December 22, 1982. He received his Engineering Masters Degree in Telecommuni-cations from the University of Palermo, Palermo, Italy in 2009. He is currently a Ph.D student in Department of Telecommu-nication, University of Palermo, 2012. He is also working as collaborator researcher for CNIT (Consorzio Interuniversitario per le Telecomunicazioni) involved in European Project ICT-FLAVIA.

Francesco Gringoli has been an Assistant Professor of Telecommunications at the Dept. of Information Engineering of the University of Brescia, Italy, since 2005. He received his Laurea degree in Telecommunications in 1998 and Ph.D. in Information Engineering in 2002. His current research interests include performance evaluation and medium access control in Wireless LANs, statistical classification of network traffic, and high-speed packet inspection. He has participated is several European and National Projects and he is currently involved in the European Project FLAVIA and National PRIN IMPRESA. His research projects have been financed by major telecommunication companies such as Alcatel-Lucent, Cisco Systems and Juniper Networks.

Ilenia Tinnirello has been an Assistant Professor at the University of Palermo since January 2005. She received her Laurea degree in Electronic Engineering and the Ph.D. in Communications in April 2000 and February 2004 respectively. She was also a Visiting Researcher at Seoul National University, Korea, in 2004, and Nanyang Technological University of Singapore in 2006. Her research has focused mainly on wireless networks particularly on multiple access algorithms with quality of service provisioning, cross-layer interactions between access solutions and physical layer; and mobility management and load balancing in wireless packet networks. I. Tinnirello has been involved in several international and national research projects, which include the bilateral Italian-Korean research project INFINITY and the national research project PRIN MIMOSA, as a scientific coordinator for the research unit of Palermo, and the European project FP7 FLAVIA, as technical coordinator.