Synchronicity: Pushing the Envelope of Fine-Grained Localization with Distributed MIMO

Jie Xiong
University College London
j.xiong@cs.ucl.ac.uk

Kyle Jamieson
University College London
k.jamieson@cs.ucl.ac.uk

Karthikeyan Sundaresan
NEC Laboratories America
karthiks@nec-labs.com

Abstract

Indoor localization of mobile devices and tags has received much attention recently, with encouraging fine-grained localization results available with enough line-of-sight coverage and enough hardware infrastructure. Synchronicity is a location system that aims to push the envelope of highly-accurate localization systems further in both dimensions, requiring less line-of-sight and less infrastructure. With Distributed MIMO network of wireless LAN access points (APs) as a starting point, we leverage the time synchronization that such a network affords to localize with time-difference-of-arrival information at the APs. We contribute novel super-resolution signal processing algorithms and reflection path elimination schemes, yielding superior results even in non-line-of-sight scenarios (with one to two walls separating client and APs). We implement and briefly evaluate Synchronicity on the WARP hardware radio platform using standard 20 MHz wireless LAN channels.

Categories and Subject Descriptors

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

Keywords

TDoA; Synchronization; Distributed MIMO; Location tracking

1. INTRODUCTION

Recently, indoor wireless LAN-based localization systems have broken the meter accuracy barrier both for WiFi devices [1, 8] and RFID tags [7], but to achieve these results, they need many APs and antennas, an RF environment without too many obstacles blocking the line-of-sight from client to AP, or a combination of both.

The most promising approaches to indoor localization involve analyzing the signals clients send jointly: both in the spatial domain, analyzing the angle-of-arrival (AoA) of the signal to the access point, and in the time domain, analyzing the time-of-arrival of different parts of the signal. But time-of-arrival analysis has a particular challenge: for a typical 802.11g WiFi channel with only 20 MHz bandwidth, the signal is sampled once every 50 nanoseconds, during

Physical layer	Bandwidth	Resolution
802.11g WiFi	20 MHz	15 m
802.11n WiFi	40	7.6
802.11ac WiFi	< 160	> 1.9
Ultra-wideband	> 500	< 60 cm

Table 1: Popular physical layers used in localization, their frequency bandwidth, and naïve spatial resolution—the resulting distance light travels between sampling instants at that bandwidth.

which light travels a full 15 meters, as shown in the first row of Table 1. As the next rows of the table show, later 802.11n and 802.11ac standards enhance this resolution, but still achieve just 1.9 meters, limiting the utility of time-of-arrival analysis on its own for the purpose of accurate localization. Even ultrawideband systems, whose expensive high-speed A/D converters sample at a rate of at least 500 MHz, achieve just 60 cm spatial resolution.

However, two recent advances in knowledge are changing the above landscape of time-of-arrival analysis:

1. First, a new class of array signal processing algorithms, informally known as “super-resolution” algorithms [2, 5], can look deeper into the received signal to overcome the naïve resolution limit of Table 1.
2. Second, recent work making distributed MIMO wireless LANs practical [4] raises the possibility of leveraging the fine-grained time synchronization such systems achieve to measure the time-of-flight difference between APs, thus computing *time-difference of arrival* (TDoA) in the distributed MIMO LAN.

Synchronicity is a system that leverages both above observations to achieve high localization accuracy in a typical, relatively narrowband (20 – 40 MHz) distributed MIMO wireless network (our techniques are also applicable to Distributed Antenna Systems). Synchronicity breaks through the 15 meter resolution barrier of the 20 MHz wireless channel illustrated in Table 1, to achieve localization accuracy on the order of one meter with just *four single-antenna APs*, as shown experimentally in §3. A rough system sketch (see Figure 1) highlighting our novel idea contributions, is as follows:

Synchronicity first measures time-difference of arrival of a client’s transmission at pairs of APs in the network. In order to do this, Synchronicity combines elements of two state-of-the-art super-resolution algorithms, looking at the correlation between incoming signals on different subcarriers as improvements of MUSIC [6] do [2], and performing generalized eigenanalysis (as “matrix pencil” schemes do [5]), crucially, just on data identified as coming from the signal. This allows Synchronicity to achieve a higher accuracy in the impulse response profile at one AP than either prior super-resolution algorithm. To synchronize time across APs, Synchronicity uses

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).
HotWireless '14, September 11, 2014, Maui, Hawaii, USA.
ACM 978-1-4503-3076-3/14/09.
<http://dx.doi.org/10.1145/2643614.2643619>.

Figure 1: Synchronicity high-level view. Each pair of AP antennas (A1, A2, and A3) measures time difference of arrival (TDoA) to generate a possible location whose locus is a hyperbola. A third AP generates another hyperbola, and hence the intersection is identified.

existing methods [3] as a starting point, and leave improvements to time synchronization functionality as near-term future work.

Next, Synchronicity analyzes the output of the above step (the MUSIC *pseudospectrum*) to identify useful spectra, retrieving accurate information even when the overall spectrum is not accurate due to the resolution limit of MUSIC. Synchronicity also potentially eliminates multipath propagation which is the major challenge for indoor localization. Here novel peak classification algorithms identify the accurate direct-path peak on the spectrum for further localization processing.

Finally, Synchronicity compares TDoA readings across pairs of APs in the network in order to estimate and refine the mobile client's location. Here a novel algorithm based on the classical triangle inequality discards measurements that must have arisen from a multipath reflection. Then clustering, outlier rejection, and averaging complete the processing chain, yielding the location estimate from a mobile client's transmission. Synchronicity does not need any offline training, calibration or decoding of the packet: the preamble of a *single* packet is enough for it to work well.

In the remainder of this short paper, we sketch the design of Synchronicity, supplementing our description with microbenchmark experiments to justify key design choices. A brief WARP-based experimental evaluation comparing Synchronicity with super-resolution MUSIC in a cluttered working office environment follows, and finally the paper concludes.

2. DESIGN

As mentioned above, Synchronicity operates in three stages. In this section, we expand on the design of each stage in turn, motivating design choices with microbenchmarks conducted over cabled RF links. In §3, we evaluate the design over the indoor wireless channel. We assume AP locations are known and pre-configured.

2.1 Super-resolution spectral processing

In the first stage, Synchronicity measures the time of arrival (ToA) of a client's transmission at one AP. This is done with the following super-resolution processing on the received signal. We begin with the MUSIC algorithm [6]. Previous work on angle-of-arrival has applied MUSIC in the spatial (bearing to AP) domain [8], but we note that MUSIC can also be applied in the time domain [2].

The multipath indoor radio propagation channel is usually modeled as complex impulse responses as:

$$h(t) = \sum_{k=1}^D \alpha_k \delta(t - \tau_k) \quad (1)$$

where D is the number of paths. α_k and τ_k are the complex attenuation and propagation delay of the k th path. Processing starts with

sub-carrier channel response of Equation 1 in the frequency domain:

$$x[n] = H[f_n] + w[n] = \sum_{k=1}^D \alpha_k e^{-j2\pi(f_0 + n\Delta f)\tau_k} + w[n] \quad (2)$$

where $w[n]$ denotes additive white noise with mean zero and variance σ_w^2 . f_n and Δf are the carrier frequency and subcarrier bandwidth. Taking the DFT of the received 64-sample long training symbols in the preamble we obtain the subcarrier frequency response $x[n]$. In 802.11a/g, 52 out of 64 subcarriers are used and we employ all of them for Synchronicity. The correlation matrix is defined as:

$$\mathbf{R}_{xx} = E\{x[n]x[n]^*\} \quad (3)$$

The *time steering vector* $\mathbf{a}(\tau)$ is defined as:

$$\mathbf{a}(\tau) = \begin{bmatrix} 1 \\ \exp(-j2\pi\tau\Delta f) \\ \vdots \\ \exp(-j2(M-1)\pi\tau\Delta f) \end{bmatrix} \quad (4)$$

Suppose D signals $s_1, \dots, s_D$ arrive at different time $t_1, \dots, t_D$ at M ($M > D$) frequency domain subcarriers. The array correlation matrix $\mathbf{R}_{xx}$ at AP will then have M eigenvalues associated respectively with M eigenvectors $\mathbf{E} = [\mathbf{e}_1 \ \mathbf{e}_2 \ \dots \ \mathbf{e}_M]$. The eigenvalues are sorted in non-decreasing order, the smallest $M - D$ eigenvalues correspond to the noise while the next D eigenvalues correspond to the D incoming signals. Based on this process, the corresponding eigenvectors in $\mathbf{E}$ can be classified as noise or signal:

$$\mathbf{E} = \begin{bmatrix} \underbrace{\mathbf{E}_N}_{\mathbf{e}_1 \ \dots \ \mathbf{e}_{M-D}} & \underbrace{\mathbf{E}_S}_{\mathbf{e}_{M-D+1} \ \dots \ \mathbf{e}_M} \end{bmatrix} \quad (5)$$

We refer to $\mathbf{E}_N$ as the *noise subspace* and $\mathbf{E}_S$ as the *signal subspace*. The MUSIC time of arrival (ToA) spectrum is then computed as:

$$P(\tau) = \frac{1}{\mathbf{a}(\tau)^H \mathbf{E}_N \mathbf{E}_N^H \mathbf{a}(\tau)} \quad (6)$$

2.1.1 Algorithm

Synchronicity includes a new super-resolution scheme which employs the signal subspace $\mathbf{E}_S$ instead of the noise subspace $\mathbf{E}_N$, then applies the matrix pencil [5] method to obtain the ToA information. The traditional Matrix Pencil scheme is applied to the raw channel response data. However, we notice that with Matrix Pencil applied to the noiseless signal-space eigenvectors, we are able to achieve even better performance than either MUSIC or traditional Matrix Pencil although the computational load is a little bit higher. We define two matrices $\mathbf{E}_{S1}$ and $\mathbf{E}_{S2}$ using the first $M - 1$ rows and the 2nd to M th rows of $\mathbf{E}_S$ respectively:

$$\mathbf{E}_{S1} = [\mathbf{I}_{M-1}, \mathbf{0}_{M-1,1}] \mathbf{E}_S \quad (7)$$

$$\mathbf{E}_{S2} = [\mathbf{0}_{M-1,1}, \mathbf{I}_{M-1}] \mathbf{E}_S \quad (8)$$

where $\mathbf{I}$ and $\mathbf{0}$ are the identity matrix and the zero matrix respectively. We then obtain ToA information by finding the generalized singular values for the matrix pencil $\mathbf{E}_{S1}$ and $\mathbf{E}_{S2}$.

2.2 Spectrum accuracy determination

We now describe the processing that happens on the ToA profile computed in §2.1. When the lengths of a line-of-sight path and a reflected path are too close to each other, MUSIC is unable to resolve the two signals correctly in the time domain on the pseudospectrum. This leads to either inaccurate pseudospectrum peak positions or multiple peaks merged into one. However, Synchronicity leverages the insight that we can still retrieve useful and accurate information from these *inaccurate* pseudospectra.

Figure 2: Testing MUSIC’s ability to resolve two paths with decreasing path length difference. The ground truth path length delays are denoted with dotted lines.

Figure 3: Peak error (translated and shown in meters) on the pseudospectrum when two signals are too close for MUSIC to resolve.

2.2.1 Microbenchmark characterization

MUSIC super-resolution capability. To determine when line-of-sight and reflected paths are too close to each other for MUSIC to resolve at 20 MHz, we use RF splitter-combiners to split a signal into two, delay one branch, and then combine the other branch and the delayed copy, thus simulating multipath propagation with one reflection path in a controlled manner. We use different cable lengths¹ to control the relative path lengths, and attenuators to control the respective path signal strengths to the same level.

Decreasing the path length difference from 13.5 m gradually to 2.7 m results in the MUSIC pseudospectra shown in Figure 2. We see from the figure that MUSIC is able to resolve both paths quite accurately when their lengths are sufficiently different, but once the distance between the two signals is decreased to around six meters, MUSIC is not able to generate accurate pseudospectra anymore. This fundamental limit to all the super-resolution scheme is determined by the sampling rate (bandwidth).

Effect of path attenuation. Most of the time, the direct path and reflection path signals don’t have the same amplitude. We observe that even when the two peaks are too close for MUSIC to resolve, the larger peak on the pseudospectrum corresponding to the larger signal is still quite accurate compared to the smaller peak.

We demonstrate this observation by the following experiments: we fix the distance between the direct path signal and reflected signal

¹Because of lower transmission speed in cable, we translate cable length to equivalent air propagation distance. (The delay of a 1.8 m RG-58 cable is equivalent to 2.7 m propagation delay in the air).

Figure 4: Peak deviations when two peaks merged into one with different relative strengths

as 5.4 meters with the help of different length cables. Then we tune the relative signal strength ratio of the direct path and reflection path from 3 dB to 9 dB and show the peak position error in Figure 3. We see clearly that when the direct path signal is stronger, although MUSIC is not able to resolve both of them correctly, the error of the direct-path peak is quite small (less than 0.5 meter). On the other hand, the smaller reflection path peak has a much larger error, so we can still extract relatively accurate information from the MUSIC spectrum in these scenarios as we only care the direct-path peak.

Next, we show that when the two signals are even closer (separated by 2.7 meters) and the two peaks actually merge into one as shown in Figure 2 bottom subplot, as long as the direct-path signal is stronger, the error is still small. We vary the relative strength ratio from 22 dB to −7 dB and show the error results in Figure 4. We can see that the error is well under one meter as long as the direct-path signal is stronger. The error increases significantly when the reflection path is stronger.

2.2.2 Algorithm (Spectrum identification)

As shown in the preceding microbenchmarks, useful and accurate information can still be retrieved even when MUSIC fails to resolve all the signals correctly as we only care about the direct-path peak. A critical step for Synchronicity is to identify useful pseudospectra; the algorithm is as follows:

Step 1. We take the first two peaks on the pseudospectrum as input to the algorithm. If the two peak positions are separated by less than half the sample period, they fall into the zone that MUSIC is not able to resolve accurately and we proceed to Step 2.

Step 2. Next, we compare the relative amplitudes of the two peaks. From the above microbenchmarks, we know that as long as the direct path signal is stronger than the reflection path signal, the direct-path peak position will be accurate. So if the amplitude of the first peak is greater than the second peak, we mark the spectrum as useful. Otherwise it’s discarded. Referring to Figure 5, we identify the spectrum in the upper figure as useless and the lower as useful.

Step 3. Then we check whether the first peak is a single-signal peak or a merged peak. A single-signal peak is much thinner compared with a merged peak. The difference is apparent when we compare the higher peak in Figure 5 (lower) with the peaks in Figure 6. We measure the peak lobe width at 90% of the peak amplitude and compare it with a predefined value W_l (This value is stable at a particular bandwidth) to make the decision.² If it’s a single-signal peak, we identify this spectrum as useful and keep it for localization. If it’s a merged peak, we proceed to the next step.

Step 4. Now we check the direction of the peak’s skew. We measure the width of the lobe with respect to the peak position as shown in Figure 6. Here the width of the green part is much larger than

²We will investigate using the second moment (variance).

Figure 5: Single-signal peaks. Identifying useful ToA spectra by comparing the amplitudes of the first two peaks.

Figure 6: Merged-signal peaks. Classifying useful spectra by the direction (earlier or later) that a merged peak is skewed towards.

the purple part so it's identified as useless. The lower plot shows a spectrum skewing earlier in time (merged towards the direct-path peak). In this case the peak has a reasonably small error, and can thus still be employed for localization.³

Step 5. After the above steps, only the useful peak remains. Then we calculate the relative ToA corresponding to the useful peak from the generalized singular values with outputs derived from Equations 7 and 8 and thus the TDoA between a pair of APs.

2.3 Multi-AP data fusion

In the last stage of TDoA localization, the measured TDoAs are converted into corresponding distance differences and used to estimate the mobile's location. Each pair of APs yields one TDoA estimate in the shape of half a hyperbola. Thus three APs are able to localize the client at the intersection of two hyperbolas⁴.

Occasionally, the direct-path signal may be totally blocked, with only reflection signals existing.⁵ We propose the following two algorithms to handle this very challenging scenario.

³We will investigate using the third moment (skewness).

⁴If there is no intersection, we discard data from that triplet of APs.

⁵Please note "non-line-of-sight" does not necessarily mean the direct path is blocked. At 2.4 GHz, the WiFi signal is able to penetrate several walls before becoming undetectable.

Figure 7: Using triangle inequality to identify direct path blockage.

Figure 8: Employ clustering and outlier rejection to remove non-accurate estimates.

2.3.1 Triangle inequality

As shown in Figure 7, when both AP antennas are able to receive direct-path signals from the client, we can deduce that the path difference is always smaller than the distance between the two APs.

$$d1 + dc > d2 \Rightarrow d2 - d1 < dc \quad (9)$$

However, when $d2$ is totally blocked and only reflection paths (dr) exist, the path length difference may exceed the distance between the APs. Whenever we detect $dr - d1 > dc$, dr must be a reflection path. We then exclude this AP from localization. Note that it's possible that with direct path blocked, the path difference is still smaller than dc and not detected. However, by checking many AP pairs, we reduce the possibility of missing a reflection.

2.3.2 Clustering and outlier rejection

Clustering and outlier rejection schemes further reduce the error caused by 100% blockage of the direct path signal and errors from other sources. This is based on the fact that the direct path signals of multiple APs will localize the clients close to the true source location while reflection path signal will localize the client at random locations. As shown in Figure 8, APs A1, A2, A3 and A4 all have direct path signals while A5 has direct path signal blocked. The estimates with any three APs from A1, A2, A3 and A4 will be around the true source S . For AP A5, because of the blockage, the signal appears to emanate from virtual source VS . So a location estimate from A1, A2 and A5 'E125' will be far away from the true location, and can be easily detected and removed. We leave implementation and evaluation of this post-processing step for future work.

3. EVALUATION

We implement Synchronicity on the WARP platform. In our current evaluation, we use four antennas attached to the WARP to serve as four synchronized APs and distribute them with equal-length cables in a large, cluttered 20×30 m office (Figure 9). We place the clients at 16 randomly-chosen locations denoting their positions as red dots on the floor plan. Half of the chosen clients locations are not in the same room as the APs, with at least one to two walls in between.

Figure 9: Floor plan of our evaluation environment with APs marked as black numbers and client locations as red dots.

Figure 10: Comparison of MUSIC and Synchronicity with/without spectrum identification (SI) and triangle inequality (TI).

We compare the following schemes:

1. MUSIC, picking the first peak in each spectrum as the direct-path peak, deriving TDoA information from the MUSIC spectrum.
2. MUSIC with spectrum identification (SI, §2.2.2) and the triangle inequality (TI, §2.3.1), deriving TDoA information from the MUSIC spectrum.
3. **Synchronicity**, including SI and TI, deriving TDoA information from generalized singular values (§2.1.1).
4. **Synchronicity with neither SI nor TI**, but deriving TDoA information from generalized singular values.

3.1 End-to-end location accuracy

We show location accuracy results in Figure 10. The CDF is composed of 160 measurements, 10 at each location measured at different times. With only a 20 MHz sampling rate, we are able to achieve around 1.6 meters median accuracy, significantly better than both the naïve spatial resolution of 15 m and MUSIC’s spatial resolution of approximately 6.5 m. With Synchronicity’s SI scheme, the modified version of MUSIC is also able to achieve around 1.75-meter location accuracy. Note that we include all the schemes described in §2 except clustering and outlier rejection, simply averaging all the location estimates across each triplet of APs. With more APs, we expect cluttering and outlier rejection to further refine the accuracy of Synchronicity’s location estimate. We plan to run more through evaluations at higher bandwidth and with outlier rejection to improve Synchronicity’s accuracy in near-term future work.

If we exclude our spectrum identification and triangle inequality schemes, how will MUSIC and Synchronicity perform? From Figure 10, we can see that the median error is around 2.8 m and 3.6 m for Synchronicity and MUSIC respectively. We notice that MUSIC without TI & SI has a particularly long tail mainly caused by not detecting a blockage of the direct client-AP path, and employing inaccurate peak without SI when the sampling rate is not high enough to accurately resolve all signals in the spectrum.

3.2 Impact of time-synchronization error

In our current setup, the four antennas serving as APs are fully synchronized and connected to a WARP. In a distributed MIMO system, there will be time synchronization errors between APs, leading to a

performance degradation for Synchronicity. We evaluate the effect of time synchronization error on Synchronicity in Figure 11. We borrow the time synchronization error data from SourceSync [3] and incorporate it into our TDoA estimation for Synchronicity. We can see that, with 5 ns and 10 ns 95th percentile time synchronization error, Synchronicity is still able to perform reasonably well, achieving a median localization error below two meters. With 20 ns time synchronization error, a larger location error is found. However, as long as we can keep this error below 10 ns, location accuracy remains mostly unaffected. We expect to minimize this synchronization error in our future work when we push Synchronicity to work on larger bandwidth for greater accuracy.

Figure 11: Synchronicity performance in the presence of inter-AP time-synchronization error.

4. CONCLUSION

Synchronicity is an indoor localization system that pushes the accuracy envelope with very few antennas and APs, and relatively narrowband wireless channels. Synchronicity is able to achieve 1.6 m median accuracy at 20 MHz bandwidth. We expect to realize centimeter accuracy localization at the 160 MHz bandwidth of 802.11ac. We also leave the use of multiple antennas at each AP in Synchronicity to further improve accuracy as future work. Multiple antennas at each AP can bring angle-of-arrival information as well as enhancing ToA information at the AP, thus potentially improving the location accuracy performance.

5. ACKNOWLEDGEMENTS

The research leading to these results has received funding from the European Research Council under the EU’s 7th Framework Programme (FP/2007-2013), ERC grant agreement n° 279976. Jie Xiong is supported by a Google European Doctoral Fellowship in Wireless Networking.

6. REFERENCES

- [1] K. Joshi, S. Hong, and S. Katti. PinPoint: Localizing Interfering Radios. In *Proc. of NSDI*, 2013.
- [2] X. Li and K. Pahlavan. Super-resolution TOA estimation with diversity for indoor geolocation. *IEEE Trans. on Wireless Comms.*, 3(1), 2004.
- [3] H. Rahul, H. Hassanieh, and D. Katabi. Sourcesync: A distributed wireless architecture for exploiting sender diversity. In *SIGCOMM*, 2010.
- [4] H. Rahul, H. Hassanieh, and D. Katabi. JMB: Scaling wireless capacity with user demands. In *SIGCOMM*, 2012.
- [5] T. Sarkar and O. Pereira. Using the matrix pencil method to estimate the parameters of a sum of complex exponentials. *IEEE Antenna and Propagation Magazine*, 37(1), 1995.
- [6] R. Schmidt. Multiple emitter location and signal parameter estimation. *IEEE Trans. on Antennas and Propagation*, AP-34(3):276–80, 1986.
- [7] J. Wang and D. Katabi. Dude, where’s my card? RFID positioning that works with multipath and non-line of sight. In *SIGCOMM*, 2013.
- [8] J. Xiong and K. Jamieson. ArrayTrack: A fine-grained indoor location system. In *NSDI*, 2013.