

Indoor Localization with Aircraft Signals*

Manuel Eichelberger
ETH Zürich
manuelei@ethz.ch

Simon Tanner
ETH Zürich
simtanner@ethz.ch

Kevin Luchsinger
ETH Zürich
kevinlu@ethz.ch

Roger Wattenhofer
ETH Zürich
wattenhofer@ethz.ch

ABSTRACT

The standard method for outdoor localization is GPS, because it is globally available, relatively accurate and receivers are inexpensive. However, GPS does not work well indoors due to low signal strength.

We explore a new localization approach, which uses the same principle as GPS localization, but employs signals transmitted by aircraft. Compared to GPS, aircraft signals are strong and can be received indoors. Our prototype implementation achieves a user localization accuracy of approximately 25 meters.

CCS CONCEPTS

• **Information systems** → **Spatial-temporal systems; Global positioning systems;**

KEYWORDS

ADS-B, aircraft signal, air traffic control, indoor localization, multi-lateration, secondary surveillance

ACM Reference format:

Manuel Eichelberger, Kevin Luchsinger, Simon Tanner, and Roger Wattenhofer. 2017. Indoor Localization with Aircraft Signals. In *Proceedings of SenSys '17, Delft, The Netherlands, November 6–8, 2017*, 14 pages. DOI: <https://doi.org/10.1145/3131672.3131698>

1 INTRODUCTION

The British Science Association has recently called GPS technology the #1 invention “that changed the world”.¹ Originally developed by the US military, GPS is now used in a majority of mobile devices. GPS has enabled a multitude of applications, which 50 years ago must have sounded like magic. However, GPS also has a major drawback — its satellite signals cannot be received indoors.²

*The authors of this paper are alphabetically ordered.

¹Richard Gray. Top 10 ‘inventions’ that changed the world. The Daily Telegraph. March 13th 2009. <http://www.telegraph.co.uk/news/science/4981964/Top-10-inventions-that-changed-the-world.html>

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SenSys '17, Delft, The Netherlands

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. 978-1-4503-5459-2/17/11...\$15.00

DOI: <https://doi.org/10.1145/3131672.3131698>

The problem is intrinsic. A GPS satellite gets its energy from a dual solar array, which (depending on the generation of the satellite) generates about 400-2900 W of power,³ about the consumption of a GSM base station. With an altitude of about 20,000 km, this relatively weak signal barely makes it to earth, already the free space path loss is in the order of 180 dB.⁴ If only GPS satellites flew lower and had more power!

In contrast, airplanes and other aircraft are flying at an altitude below 13.7 km. They also have ample power, e.g. a Boeing 747 has an average power consumption of 140 MW, leaving enough power for communication.

For safety reasons, airplanes and helicopters repeatedly transmit their location, pretty much like GPS satellites. These so-called ADS-B signals are strong enough to be received indoors, even with cheap hardware. But are these air traffic control signals *precise* enough to not only locate the aircraft but *any* mobile device?

As air traffic control signals have not been designed for indoor localization, we have to deal with three challenges:

- (1) Aircraft do not fly on an orbit; aircraft do not have accurate predetermined flight paths and unexpected changes to their route are always possible, for instance due to a holding pattern when approaching a crowded airport.
- (2) Aircraft are not uniformly distributed in the sky. This is in stark contrast to GPS satellites, which cover the sky in a regular pattern in order to maximize user localization performance.
- (3) Aircraft position signals are not precise: an aircraft has an unpredictable delay between learning its position from the GPS satellites and retransmitting this position;⁵ unlike GPS satellites with their atomic clocks, aircraft transmissions do not include time information; some aircraft in fact do not even include position information.

But there are good news as well. Although aircraft do not fly on an orbit (1), passengers and crew certainly do not appreciate abrupt flight path changes. Also, even though the aircraft positions are not optimized for the localization of users on ground (2), but rather for air traffic safety, at least in urban areas there are more aircraft

²Also, multipath effects are detrimental for GPS signal reception and positioning accuracy. The official GPS website by the US government has a list of causes for inaccurate position estimation, which includes the “urban canyon” effect in inner cities and also less frequent causes: <http://www.gps.gov/systems/gps/performance/accuracy/#problems>

³Data for newer satellite models: http://www.au.af.mil/au/awc/awcgate/smc-fs/gps_fs.htm

⁴See, e.g., <http://gpsinformation.net/main/gpspower.htm>.

⁵Uncompensated latency of up to 0.6 s: <https://www.law.cornell.edu/cfr/text/14/91.227#e>

available than satellites.⁶ A large number of received signals hopefully reduces the statistical uncertainty of the position estimation from noisy measurements (3).

So, do these good news compensate for the bad news above? It turns out, we cannot simply use the aircraft signals. Instead, we propose to install a few receivers in known positions. Not many receivers are needed, in principle already a single receiver is enough to serve a large metropolitan area or even a small country. These receivers will calibrate the received aircraft messages, so that any mobile receiver can deduce its position without any additional infrastructure. More formally, our localization system consists of

- a network of receivers with known positions, which we call *ground stations*
- a receiver whose position should be determined – we refer to this receiver as the *handset*
- a *server*, which connects the ground stations and the handset

The ground stations receive aircraft signals, decode messages, precisely time their arrival times and send the messages together with the time stamps to the server. Given enough records for the same message, the server can then determine the message's transmit position and time. In practice, most aircraft in Europe and a growing number of aircraft in other countries already send their position [6, 11] so the only unknown is the transmit time, in which case one ground station is sufficient to determine this transmit time.

To localize a handset, its received messages are also sent to the server like the messages received at the ground stations. The server then matches the handset messages with the corresponding messages from the ground stations for which the transmit position and time are known. A multilateration approach will reveal the localization solution, which consists of the handset position and time.

We use receivers which consist of a 7 cm long antenna, a USB software-defined radio receiver dongle and a Raspberry Pi 3 board. The total cost of such a receiver is less than \$100. In the future, such a receiver design could easily be integrated into a smartphone. A 7 cm antenna is small enough to fit inside the casing, and current smartphones have more processing power than the Raspberry Pi 3.⁷ Our signal reception and decoding software is a modified version of the open source project *dump1090*,⁸ which we enhanced with more accurate time resolution. With these cheap receivers, messages sent from aircraft as far as 250 km away can be received.

The GPS Performance Standard [13] by the US government currently lists a worst-case horizontal GPS accuracy better than 17

meters in 95% of all cases. Depending on the quality of the receiver and whether different advanced correction methods are available, the horizontal accuracy can be substantially better, in the order of 3-7 meters. Usually, indoor localization methods attempt to be even *more* accurate, as for instance people want to precisely know in which aisle of the shopping mall they are.

In terms of accuracy, our method cannot compete with these indoor localization techniques, and does not even achieve the typical GPS outdoor accuracy. Our prototype implementation has a median error of about 25 meters. On the plus side, it works both indoors *and* outdoors!

Even though 25 meter accuracy is not exciting, 25 meters is still better than *nothing at all*, and various applications do not need a very precise position. For instance, our method may tell you in which building you took some photograph. Also, our method can help you catch the next bus from your current location, since no precise position is necessary to determine the closest bus stop and look up the corresponding timetable. Moreover, one can automatically log the working time of employees by just knowing an approximate indoor position. Due to the higher signal strength of the aircraft signals, our system also works in urban canyons, where GPS receivers may not be able to detect the signals from the GPS satellites, because urban canyons behave similar to indoor environments. Generally, our system works well when the application needs both indoor and outdoor locations, or when dedicated indoor localization infrastructure is too costly to deploy. Finally, and probably most importantly, our method may serve as a more accurate initial guess for Assisted GPS (A-GPS) than currently provided by cellular networks. Having a better estimate of the receiver position and time can speed up the initial GPS fix, especially when a maximum likelihood method like Collective Detection [2, 3] is applied.

Another upside of our method is that it is pretty much independent from additional infrastructure, as for real-time tracking only an Internet connection is necessary to synchronize information with the ground stations. It is also possible to post-process data, in which case no infrastructure is needed at all. The latter approach has for instance been successfully applied to GPS localization [25].

Since our method and implementation presented in this paper are a proof of concept, future improvements might yield more accurate positioning using the same aircraft signals and could thus allow such a localization system to be used for even more indoor applications.

Our method is one of the first basically infrastructure-free indoor localization methods. As discussed in Chapter 6, the accuracy of our method is comparable to LTE positioning, which is a competitor to infrastructure-free indoor positioning due to the wide availability of cellular networks. These networks are designed for communication purposes and therefore cover an area with at least one antenna, but often not more, in order to save costs. Moreover, cellular networks use transmit antennas located close to the ground. For positioning, it is therefore often not possible to receive signals from a sufficient number of antennas. In contrast, aircraft are high up in the sky and can thus provide a better area coverage for instance in cities, forests and mountains.

⁶Note that at night, the frequency of received aircraft signals is substantially lower than during daytime. But since sleeping people do not need indoor localization, we do not consider the resulting degradation of positioning accuracy to be a problem in practice.

⁷OpenBenchmarking.org, which is one of the standard benchmark sites for Linux systems, shows benchmark results comparing the Raspberry Pi 3 to the ODROID-C2 development board, which features an ARM Cortex-A53 SoC, also found on current low-cost phones: <https://openbenchmarking.org/result/1603051-GA-ODROIDPI362>. Even this low-cost SoC, which is the core of the Qualcomm Snapdragon 410, 610 and 615 [http://www.anandtech.com/show/7573/qualcomm-announces-snapdragon-410-based-on-64bit-arm-cortex-a53-and-adreno-306-gpu, http://www.anandtech.com/show/7784/snapdragon-610-615-qualcomm-continues-down-its-64bit-warpath-with-48core-cortex-a53-designs] beats the Raspberry Pi's performance in all tested workloads.

⁸<https://github.com/mutability/dump1090>

Our paper is organized as follows. First, in Section 2 we give a short introduction to air traffic surveillance systems and the multilateration technique. Then, we dive into the details of our method in Section 3. We present our implementation in Section 4, followed by our results in Section 5. To the best of our knowledge, there is no closely related prior work to our approach. We discuss various weakly related existing indoor localization methods and air traffic surveillance systems in Section 6. The paper is concluded in Section 7.

2 BACKGROUND

2.1 Air Traffic Surveillance

Primary Radar. Air traffic surveillance has emerged as a fundamental requirement for air traffic control. The classic method for localizing aircraft is radar technology. Radar systems send out powerful pulses of electromagnetic signals. When these pulses are reflected at an object, the back-scattered energy can be detected. Given the angle of arrival and the round trip time, the position of that object can be determined. Because the energy of the reflected signal is much weaker than the pulse transmitted from the radar installation, the energy of the radar pulse needs to be high in order to be able to detect the reflected signal. Due to the resulting high energy consumption, radar antennas are tied to the ground as fixed facilities. Today, this classic radar technique is termed the primary surveillance system.

Secondary Surveillance. The messages we employ for our method are part of the secondary surveillance system for air traffic control. Unlike the primary radar, aircraft actively send messages in the secondary surveillance system. Two types of secondary surveillance techniques exist: Aircraft can be “interrogated” by ground stations and respond appropriately, or they simply transmit messages periodically. The former is denominated *Secondary Surveillance Radar* and the latter is called *ADS-B*. Messages may or may not include various types of information, for instance the current position or velocity of the aircraft.

In contrast to the primary radar, the received power of the secondary surveillance signals is higher, because the signals are actively sent by the aircraft. The secondary surveillance system can therefore achieve wider ranges at the same transmission power. Another advantage of the secondary surveillance system is that not only ground control stations are aware of aircraft and their position in the sky, but all aircraft in the sky can receive the messages sent by other aircraft. This latter functionality is supported by having two separate transmitters at the bottom and the top of the aircraft hull [16, Paragraph 3.1.2.10.4].

Secondary Surveillance Radar. The *secondary surveillance radar (SSR)* is an established secondary surveillance system, which detects and identifies aircraft and receives barometric pressure measurement data, from which the flight level of aircraft can be determined. The radio signal uplink from the ground to the aircraft uses a carrier frequency of 1030 MHz and the downlink is at 1090 MHz. The SSR has multiple *interrogation modes* with different message formats. The most important modes are A, C and S. To Mode A interrogations, aircraft reply with their identity. Mode C is used to gather aircraft altitude calculated from the barometric pressure

and Mode S supports different message formats, including ADS-B messages [37].

ADS-B. A modern secondary surveillance technology is the *Automatic Dependent Surveillance – Broadcast (ADS-B)*. In contrast to the secondary surveillance radar, aircraft with ADS-B transponders send out messages periodically, that is without interrogation. ADS-B is a *dependent* surveillance system because ground stations depend on the aircraft sending out data. ADS-B messages can contain a variety of information. For our system, mainly the position and the velocity of the aircraft are of interest. The position data is always derived from a GPS receiver in the aircraft. Also information used for collision avoidance is transmitted over ADS-B.

ADS-B signals can be sent over different physical links, but the most commonly used is the 1090 MHz channel, because many aircraft already have a transponder for Mode C installed, which operates at this frequency. Mode S comprises a so-called *Extended Squitter* message format, which allows including ADS-B messages in Mode S packets. According to [11] and [6], at the moment most airliners in Europe and a growing number in North America are equipped with an ADS-B transponder.

The data is transmitted at 1 MBit/s with pulse position modulation. Each packet starts with a preamble of 8 bits followed by 56 or 112 bits of data. The data contains the message type, aircraft address, the actual information (depending on the type) and parity bits for error correction [36]. ADS-B messages are sent in response to an *interrogation* by a ground surveillance station and are additionally also transmitted approximately every second, by choosing a random time interval between 0.8 and 1.2 seconds between two transmissions [16, Paragraph 3.1.2.8.5.2].

2.2 Multilateration

Multilateration is an established technique and used in many radio navigation systems such as GPS and LORAN-C. It is also called *time difference of arrival (TDOA)* which stems from its usage of measurements of signal arrival time differences from multiple transmitters. Broadcast signals from different transmitters at known positions are received and timestamped. This way the position difference of the senders to the receiver can be calculated. For each pair of transmitters, a hyperboloid describes the possible locations of the receiver. With multiple pairs of transmitters, the intersection of those hyperboloids can be calculated where the receiver should be located. The system of equations is normally solved using a least squares approach. But other methods can be used. For instance Collective Detection (CD) can tolerate more noise than the least squares method and has been successfully applied to the GPS localization problem [3].

Compared to GPS, our method employs signals transmitted from aircraft instead of satellites. Since aircraft travel at an altitude of around 10 km instead of over 20,000 km of the GPS satellites, the transmitters of aircraft are much closer to the user on the ground and give stronger received signals. Therefore, our method is more suitable for indoor localization than GPS.

Timing. For the TDOA method, the transmission time of the messages has to be known. Normally, this is achieved with time synchronized transmitters that include their current time in the

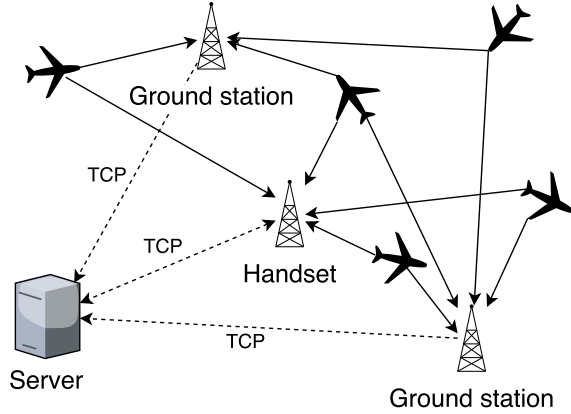


Figure 1: System consisting of a server, ground stations at known locations, a handset at unknown location and aircraft. The ground stations and handsets determine the receive time of messages from aircraft to then calculate the send time of the messages and the location of the handset.

messages. None of the SSR or ADS-B messages contain a time stamp. Since the messages travel at the speed of light, even if an aircraft is 300 km away, the propagation time will only introduce a delay of one millisecond. Assuming the delay in the aircraft from the generation of the position to the transmission event is small or compensated, the position of the aircraft will have changed by only a few meters by the time the signal arrives at the ground. For Mode C messages, the height will probably still be the same to within the measurement error. Therefore, for air traffic control, send time stamps are not necessary. But this timing is not good enough for positioning a user, since every millisecond of time error also alters the measured distance to the aircraft by 300 kilometers. Therefore, we use a ground station with known position to determine the transmit time of messages. Multiple ground stations can be used to increase the range of the system. The details of our method are explained in the next section.

3 METHOD

In this section, we show how the challenges explained in Section 1 can be addressed. The main idea of our method is to replace GPS satellites with aircraft in order to receive stronger signals, which are more suited for indoor reception than GPS signals. As aircraft messages do not include a time stamp, we also solve a time synchronization problem using a small number of ground stations with known positions. A significant fraction of aircraft transmit their position – which might or might not be accurate. For aircraft with legacy systems (SSR modes A and C), which do not provide their position, our infrastructure also determines the transmit positions of the aircraft messages. In the end, the system we present provides users with a system similar to GPS, but with aircraft instead of satellites. The user localization is done using multilateration like in GPS.

Figure 1 shows the concept of our system. The ground stations and the handset send the recorded ADS-B messages with the corresponding time stamps to a server. The server collects the messages




Figure 2: Handset multilateration: The handset receives messages from different aircraft that are then used for multilateration.

and matches the messages from the handset to those from the ground stations and computes send times of the messages and the position of the handset by solving least squares problems using the relative timing of these messages and their transmit position. The localization method can be partitioned into two steps:

- Calculation of message send times and aircraft positions (if not given in messages): Matching messages received at multiple receivers are used to calculate the clock offset and drift of the receivers, the message send times and optionally the aircraft positions.
- User localization: Multilateration using messages with now known send times.

These individual steps are explained in detail in the next sections.

As explained in the introduction, our localization system consists of three hardware components: unsynchronized ground stations that receive ADS-B or SSR messages from aircraft, a user handset and a server that collects all the received messages.

With this proposed system, the messages from the aircraft can be used for localization without large infrastructure costs. The ground stations and handsets do not need synchronized clocks and work with inexpensive software-defined radios (SDRs). Each receiver only needs little processing power to decode the messages and forward them to the server.

3.1 User Localization

Let us first discuss the handset localization assuming that the transmit position and transmit time of the received messages is known.

The calculation of the position of the handset is similar to the multilateration in global navigation satellite systems, like GPS. Instead of satellites with known position and time of signal transmission, we use aircraft and ADS-B messages transmitted by them. Figure 2 shows the concept. For each received message, we can create a pseudorange equation

$$\|P_H - P_j\|_2 + c\Delta t_H \stackrel{!}{=} c(t_{j,H}^r - w_{j,H} - t_j^t)$$

where P_j is the position from where the message was sent, t_j^t the send time, P_H the handset position, Δt_H the clock offset of the

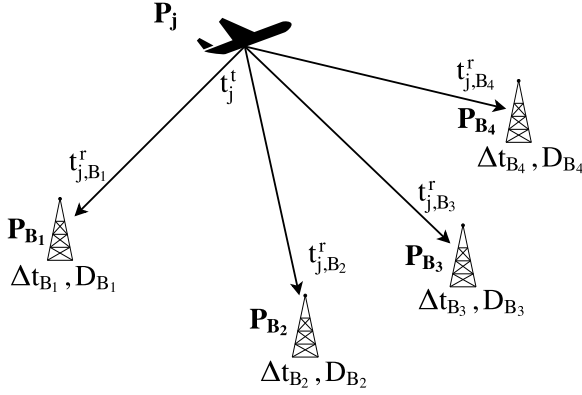


Figure 3: Calculation of message send time and aircraft position: Multiple ground stations receive the messages from the aircraft and can calculate the send time of the messages and optionally the position of the aircraft.

handset to the ground stations and $t_{j,H}^r$ the receive time of the message at the handset with the noise $w_{j,H}$. The handset position and clock offset can be found by solving a least squares problem.

$$e_j(\mathbf{P}_H, \Delta t_H) = \frac{1}{c} \left\| \mathbf{P}_H - \mathbf{P}_j \right\|_2 + \Delta t_H + t_j^t - t_{j,H}^r + w_{j,H}$$

$$(\widehat{\mathbf{P}}_H, \widehat{\Delta t}_H) = \arg \min_{(\mathbf{P}_H, \Delta t_H)} \sum_j e_j(\mathbf{P}_H, \Delta t_H)^2$$

This cost function is minimized using a non-linear least squares solver. At least four points \mathbf{P}_j are needed to estimate the handset position and time offset. We used the *Levenberg–Marquardt algorithm* which combines the *Gauss-Newton method* with *gradient descent* [14].

Dilution of Precision (DOP) is a concept originating from GPS, but it can also be applied to our problem. When doing multilateration, the geometric distribution of the used points (for instance GPS satellites or aircraft and the handset) influences the quality of the position estimation. The position estimation error is the product of the signal arrival time measurement error and the DOP value. A rule of thumb is that the larger the volume spanned by the aircraft, the better the localization precision [22]. And the receiver should be close to the aircraft. We can leverage other advantages of aircraft signals, like higher signal strength or a higher number of available aircraft to reduce the measurement errors and get good localization accuracy.

Due to the aircraft geometry, the vertical positioning uncertainty is often larger than the horizontal uncertainty, because most aircraft are at a low elevation angle, and therefore a vertical position change in the user position results in a smaller change in the measured distance to the aircraft than a horizontal movement. Therefore, it is useful to determine the handset height using a different method, for instance with a barometric pressure sensor. This can increase the positioning accuracy significantly, as shown in Section 5.6.

3.2 Transmission Time and Location

To calculate the position of the user, the transmission time t_j^t of the message at the aircraft has to be known. By using a ground station

with a fixed location, the time-of-flight from the aircraft to the ground station can be calculated to then compute the transmission time.

To reduce the error of the time stamps and increase the covered area, multiple ground stations can be used. But the message time stamps at the different ground stations are not synchronized. Moreover, due to small deviations of the receiver oscillators, the sampling rate is varying and not equal at different ground stations. In order to use the time-stamped messages from the different ground stations for the user localization, we have to compensate the clock offsets and drifts.

The receive time stamp t_{j,B_i}^r consists of the message send time at the aircraft (in global time) t_j^t plus the clock offset Δt_{B_i} and clock drift D_{B_i} of the ground station relative to station 1 plus the time-of-flight of the signal from location \mathbf{P}_j to \mathbf{P}_{B_i} .

$$t_{j,B_i}^r - w_{j,B_i} - \frac{1}{c} \left\| \mathbf{P}_{B_i} - \mathbf{P}_j \right\|_2 = t_j^t + \Delta t_{B_i} + D_{B_i} (t_j^t - t_1^t) \quad (1)$$

This system of equations can be solved using a linear least squares solver to compute the transmission times t_j^t of the messages in the synchronized time.

As the clock drift rate and offset may change over time, the synchronization has to be repeated regularly. When a few messages have been received at multiple ground stations, Equation (1) is solved. Additionally new clock offsets and drifts for the ground stations are calculated. Now, only the handset has a clock offset which is the same to all ground stations.

Not all messages sent by the aircraft contain the position. To be able to also use those messages in the localization of the handset, we have to determine the aircraft position \mathbf{P}_j at the time when the message was sent. In this case, Equation (1) is not linear anymore due to the distance term. To calculate \mathbf{P}_j , we therefore use a non-linear least squares solver.

Ground Station Requirement. Before going on, let us discuss the ground station requirement. By the end of the year 2019, all aircraft will be required to support ADS-B [6] and therefore send their position in their messages, and many aircraft already do it today, as mentioned in Section 2.1. Since signals can be received several hundred kilometers from an aircraft, in principle one ground station is enough to cover a large urban area or a small country. Using multiple ground stations has several benefits:

- Measurement errors can be reduced.
- In mountainous regions, a ground station should either be deployed on the highest mountain in order to receive aircraft signals from all directions or several receivers can be used, which all see a part of the sky. Note that the different parts need to have some overlap in order to be able to synchronize the ground stations' times.
- Less messages are lost and therefore a larger number of the messages of a handset can be matched by the server with a message from a ground station. This increases the number of measurements per time and thus allows tolerating faster moving receivers because a sufficient number of messages for positioning can be collected at the handset during a shorter time.



Figure 4: Hardware of ground station: Raspberry Pi 3 with FlightAware Pro Stick Plus and antenna. The Raspberry Pi decodes the messages and sends them to the server via LAN or WLAN.

The number of ground stations for a full deployment of our system is orders of magnitude less than for other localization techniques, like for instance cellular network or WiFi based localization. By leveraging already existing ADS-B receiver network communities (cf. Section 6.2), deploying a receiver network for our proposed method might be possible by simply providing updated software for these receivers.

Lastly, note that handsets do not need to be in line-of-sight of any ground station since the time synchronization of aircraft messages and the position computation of a handset are decoupled.

4 IMPLEMENTATION

4.1 Receiver

To receive the messages from the aircraft, we use software-defined radios (SDRs). For this project USB dongles based on the RTL2832U chipset are used, which were originally designed to be DVB-T TV tuners. Therefore, the SDRs are inexpensive. In a second iteration we used FlightAware Pro Stick Plus which is a USB dongle based on the same chipset and has been optimized for the reception of ADS-B. It contains a band-pass filter and an amplifier. These SDRs provide a stream of complex samples (I/Q samples) at a rate of 2.4 MHz. The power consumption of the SDR is at most 1.5 W.⁹

For the processing of the data, each SDR is connected via USB to a Raspberry Pi 3 model B, a single-board computer. The operating system is Raspbian, which is based on Debian. The Raspberry Pi consumes a maximum of 6.7 W under stress.¹⁰

The detection of the messages is done with dump1090-mutability. It is able to run on the Raspberry Pi 3 without dropping samples. This software generates time stamps with a resolution of 83.3 ns. In order to improve the time stamp quality further, we added local

upsampling of the messages and GPU_FFT is used to compute the FFT for the convolution on the GPU.

After detecting the messages, the program sends the detected messages with the corresponding time stamps to the server where they get collected and the localization is performed.

Currently the ground stations and the handset are based on identical hardware configurations.

Barometric Pressure Sensor. Since the ground stations and the handsets have a similar altitude and the aircraft are close to the horizon, the vertical dilution of precision is much higher than the horizontal dilution of precision. A solution is to equip the handset with a barometric pressure sensor. The altitude value provided by the sensor can be used to perform two-dimensional multilateration which reduces the number of needed aircraft by one. The sensor has to be calibrated with the barometric pressure on sea level (QFF) which is weather-dependent. This value can be obtained from a weather API or can be calculated at a point with known altitude.

4.2 Server

Figure 5 shows the overview of the server application. The server consists of four different threads that are connected via message queues.

The *collector* thread accepts TCP-connections from the receivers and parses the received messages.

The *planettracker* thread performs the multilateration for messages that do not contain the position of the aircraft as described in Section 3.2. Furthermore, all messages are Kalman filtered using a constant velocity model. The model is updated with the received velocity and position messages and the positions calculated with the multilateration. The hash of the messages and the message time stamps are then used to match the messages received at different ground stations and at the handset. The messages then get forwarded to the *synchronizer* thread.

The *synchronizer* thread performs the time synchronization of the ground stations. It looks for corresponding messages that have been received at multiple ground stations and uses them to calculate the time offset and clock drift of the ground stations and the transmission time stamps of the messages. When the transmission time of a message, that has also been received at the handset, has been calculated, it gets forwarded to the *localizer* thread.

The *localizer* thread calculates the handset position using the messages with the calculated transmission time stamp. For each handset, the messages are accumulated in a queue. As soon as it contains enough messages (at least four different aircraft), the multilateration of the handset is performed. Handsets with known altitude can be located using two-dimensional multilateration and therefore only need three received messages.

5 RESULTS

To evaluate our method, we deployed six ground stations in a region approximately 110 km in diameter. The positions of the ground stations and the handset can be seen in Figure 6. Placing ground stations outdoors, for instance on the roofs of buildings or on hills, would be beneficial in order to maximize received signal strength, number of received messages and observed unique aircraft, as well as to reduce multipath errors. However, building weatherproof

⁹<http://de.flightaware.com/adsb/prostick/>

¹⁰<https://www.raspberrypi.org/help/faqs/#powerReqs>

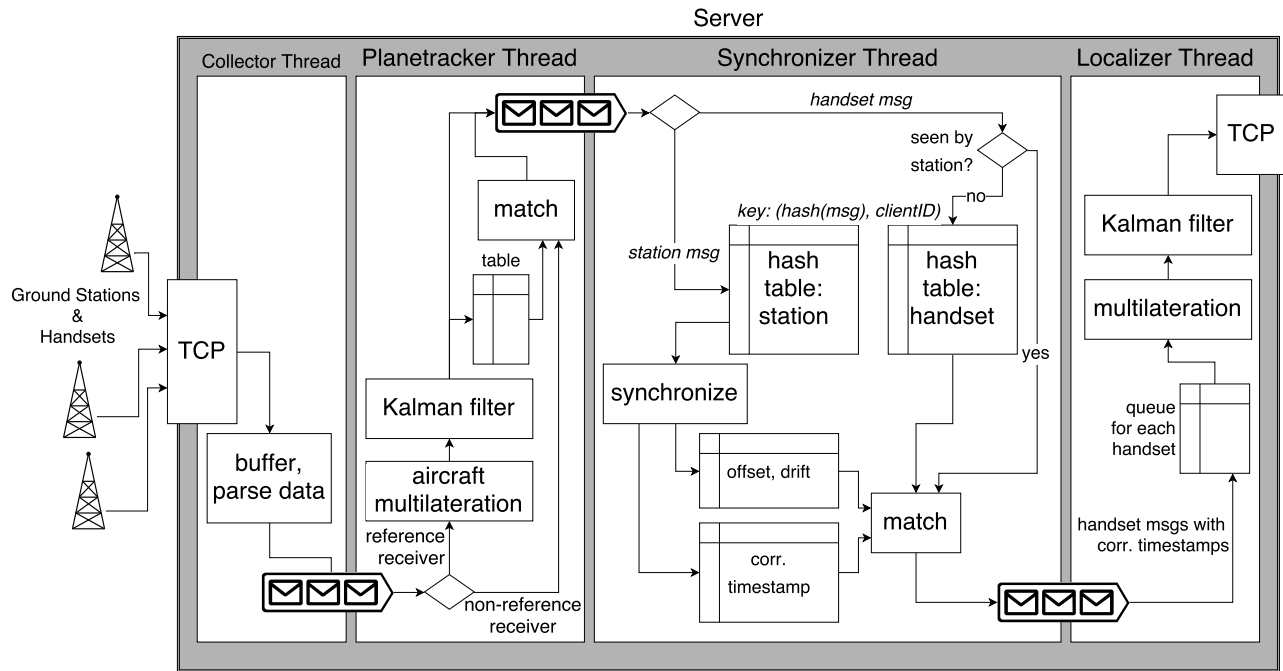


Figure 5: Server application: The server receives the messages from the ground stations and handsets, matches the messages from the different receivers. Then the clock offset and drift between the ground stations is computed with messages from different receivers. In the last stage, the multilateration of the handsets is performed.

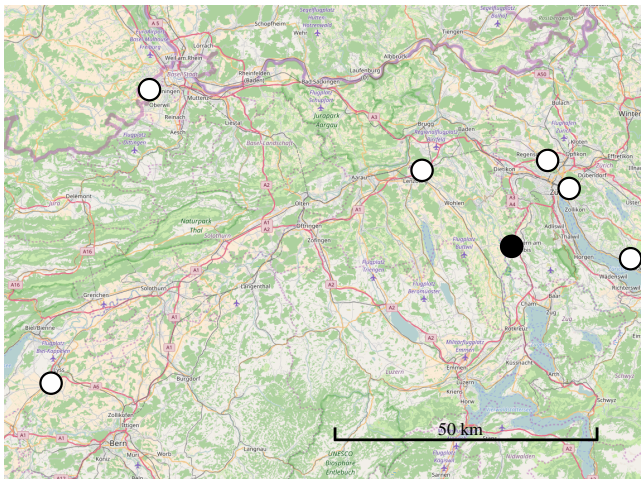


Figure 6: Location of the ground stations (white circles) and a handset (black circle) for the evaluation. The ground stations span over a region approximately 110 km in diameter.

cases was outside of the scope of this work. Therefore, the ground stations for our evaluation are placed inside buildings. Note that our preliminary test setup uses ground stations with known, but inaccurate positions, which have been estimated manually by locating the receivers on a map, therefore introducing several meters of error.

In our setup, the received signal is sampled at 2.4 MHz and up-sampled by a factor of 25 at the ground stations and the handset. The handset's height is determined using a barometric pressure sensor, so only the latitude and longitude of the handset are computed based on the data decoded from the aircraft signals.

Our evaluation should be considered as a mostly qualitative analysis, since we could not yet test our system performance extensively. Doing so would require measurements spanning a long time, maybe even a year, due to the variability of air traffic. Not only is there a daily cycle of flight patterns and air traffic density, but for instance different wind directions influence the routing of aircraft on different days and there are even seasonal differences, as during holidays, the number of passenger flights increases significantly. Also, aircraft geometry depends on the flight routes assigned in different regions of the world and is not mostly uniform, like for instance the GPS satellites with their nicely distributed orbits. Furthermore, the landscape around the receiver also influences the reception of aircraft signals, for instance by blocking the line of sight between a handset and some aircraft. Further variable parameters, like the geographic distribution of the ground stations, also influence the performance of our system. Finally, note that all these parameters are not independent, which makes it even more challenging to evaluate our system thoroughly.

5.1 Reception Quality

The maximum range within which aircraft signals can be received heavily depends on the antenna characteristics and placement. With our ground stations, which were all located indoors, messages from

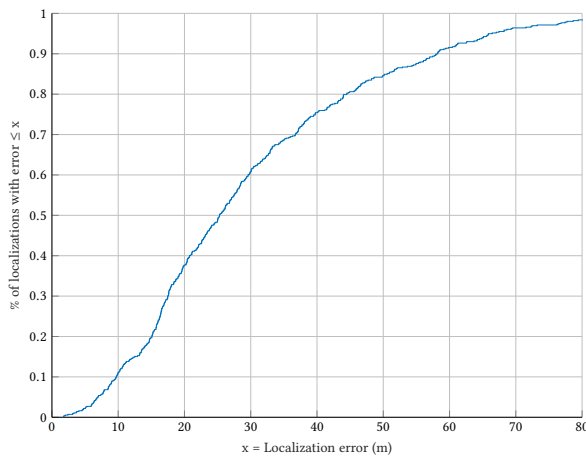


Figure 7: Cumulative distribution function of the positioning errors of an outdoor handset using data from six ground stations. The handset altitude is known and the horizontal coordinates are determined using the aircraft signal measurements from the ground stations.

approximately 190 km away can be received. When placing the antenna on a roof, the range increases to about 250 km. Comparing the indoor and outdoor cases, we observe that indoors the number of unique aircraft from which signals are received decreases by a quarter. Note that our ground stations have cheap passive antennas. With more expensive active antennas, ideally mounted on a roof, the received signal strength should be higher, resulting in increased signal reception range and more received messages.

5.2 Localization Accuracy

Figure 7 shows the localization accuracy of our method using our six deployed ground stations and a handset outdoors. The median error between the computed positions and the ground truth is 25.3 meters and the maximum error is 118.6 meters. The ground truth was estimated using Google Maps,¹¹ the error of the ground truth should be less than 3 meters.

The results of our measurements are approximately normally distributed. An example distribution of the computed positions around the ground truth can be seen in Figure 8.

5.3 Indoor vs. Outdoor Accuracy

We conducted experiments to evaluate the accuracy of our method indoors. Surprisingly, the accuracy indoors is extremely close to the accuracy outdoors. The median error is only 5.6 % larger indoors, and the standard deviation increases by 14.7 % compared to the outdoor case. The cumulative distribution function looks almost identical to the outdoor case.

5.4 Number of Ground Stations

To test the influence of different parameters, we conducted additional experiments. First, we tested the accuracy using different numbers of ground stations. The results are shown in Figure 9.

¹¹<https://maps.google.com>




Figure 8: Distribution of positioning errors in latitude and longitude direction. The handset altitude was known. These results were obtained using six ground stations and a handset outdoors. The plot shows that the positioning solutions are approximately normally distributed.




Figure 9: Absolute positioning error outdoors for different numbers of ground stations. As expected we do not see large differences of the localization error since one ground station is sufficient to calculate the send time of an aircraft message.

We do not see a large increase in the positioning accuracy when more ground stations are used. This is what we expected, since one ground station is sufficient to calculate the send time of a message from an aircraft. One reason for the slightly better results with 5 and 6 ground stations might be that errors in the calculated clock offsets and drifts of the ground stations are averaged out.

5.5 Number of Observed Aircraft

Further, we tested the influence of the number of unique aircraft used for a positioning solution. Figure 10 shows the results from this experiment. When using different numbers of aircraft, the localization error does not change. As in GPS with more satellites,

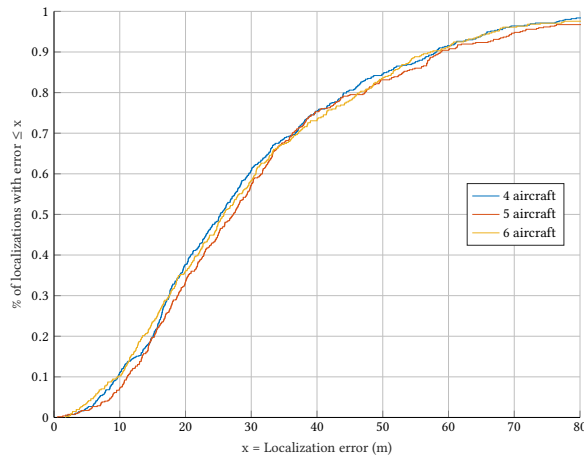


Figure 10: Absolute positioning error for different numbers of aircraft used for multilateration. The localization error does not decrease with more aircraft used in the multilateration. We assume that the synchronization error increases because messages have to be combined over longer periods of time.

we expected the localization error to decrease with more aircraft. However, when not many aircraft are within range of the handset and the receivers, messages from a longer time period might have to be combined. This can lead to a larger synchronization error between the ground stations.

5.6 Known vs. Unknown Altitude

Another interesting experiment evaluates the benefit of using a barometer for the receiver’s height estimation. Figure 11 shows that the spatial positioning error is much higher than the horizontal one with known receiver height. However, an increased error is expected due to an added degree of freedom of the solution. Also note that most aircraft are at a low elevation angle from the receiver position. This results in a badly conditioned problem for the height estimation. A simple way to see this is that a height change of the handset does not influence the arrival time measurements as much as a change in the horizontal plane. Since we conducted this test early in the development of our method, some improvements present in the final system described in this paper were still missing. Therefore, the horizontal positioning accuracy is also worse than the results shown before. The conclusion of this experiment is that a barometer is a crucial feature of a handset multilateration method based on aircraft signals. Adding the barometer reduces the median error by 52%.

5.7 Upsampling

As explained in Section 4.1, we upsample the received signal by a factor of 25. Figure 12 shows the achieved localization accuracy compared to that which results when using the time stamps of the standard dump1090 software, which determines the phase of the messages by correlating with five fixed patterns. As we understood from analyzing the source code, these patterns used in dump1090

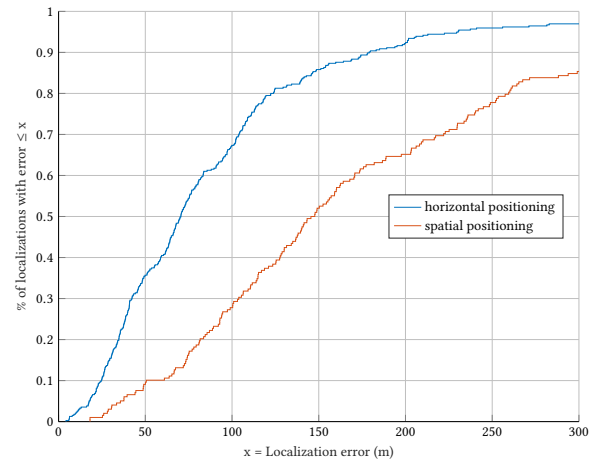


Figure 11: Positioning error with known handset altitude (horizontal positioning) compared to unknown handset altitude (spatial positioning). These results are from a preliminary experiment with a version of our implementation, which did not include all the improvements present in the final version.

correspond to the expected sample values when shifting the signal by multiples of a fifth of a sample duration. Therefore, the accuracy should be the same as for a five times upsampled signal. However, we observe that the five possible phases do not appear equally likely, which we interpret as an implementation error. Therefore, the results using the dump1090 time stamps might be somewhat worse than theoretically achievable using an optimal such phase estimation technique. Note that like in the previous experiment, also the localization results of this test are less accurate than those of the current system, because the results were derived using a preliminary implementation. Figure 12 demonstrates that the positioning error decreases when upsampling of the received signal is performed. However, the necessary computation power increases substantially. Due to performance limitations of our prototype ground stations featuring a Raspberry Pi 3, we were not able to continuously use 50-fold upsampling. In fact, even when upsampling by a factor of 25, it happens every few hours in expectation that the processor is overloaded. This manifests itself in dropped samples, because buffers are emptied slower than samples are recorded. The reason why this happens infrequently and unpredictably is that the number of received messages depends on the number of aircraft in the range of the receiver and therefore varies considerably. When the processor has a lot of load, it overheats and at 80°C, the clock rate of the CPU is automatically reduced. The result is that the processor is even less able to cope with all the incoming messages. This problem could probably be resolved by installing a heat sink on the processor. However, if regions are covered by multiple ground stations, short outages of one ground station for a few seconds or minutes can easily be tolerated. If this is not the case and only one ground station is available, its upsampling factor can also be reduced a bit, without resulting in a large increase in positioning error, as Figure 12 shows.




Figure 12: Positioning error based on 25-fold and 50-fold upsampled time stamps, as well as the default time stamps output by *dump1090*. These results were obtained using two ground stations and also an old version of our implementation, like the results from the barometer experiment.

5.8 Error Sources

Multiple possible error sources of the localization exist. The multi-path effect caused by signal reflections from buildings and ground has a big impact, since the ground stations are placed in offices and residential buildings and do not have direct line of sight in all directions. Accurate time stamps are essential for the multilateration. As shown in Section 5.7, the time stamps can be improved using upsampling. Since the ground stations do not have synchronized clocks, the estimation of the clock offset and drift has to be repeated regularly. If only few aircraft messages are received, they have to be collected over a longer period of time during which the synchronization error accumulates and the time stamps become less accurate. Note that the opposite case, when many aircraft send messages, which then collide and therefore cannot be decoded, is not a problem. In this case, due to the large number of aircraft, also many messages can successfully be decoded. The extreme case in which basically all messages are lost should not occur in practice, since this would also be a threat to air traffic surveillance and guidance. In practice, up to 50 % message loss is possible [35, Figure 3b]. The time stamps are more precise if the signal-to-noise ratio (SNR) is high and therefore the correlation gives a easily distinguishable peak for the message arrival time. But this correlation can also deteriorate because of the frequency shift caused by the Doppler effect. Small errors are also introduced by the ground station positions. Depending on the location of the aircraft and the handset, a bad dilution of precision can occur which increases the localization error. Additional errors are introduced by the uncompensated latency in the aircraft between the reception of the GPS position and the transmission of the ADS-B message. Also the GPS receiver and the transponder have a position offset depending on the aircraft model.




Figure 13: Distribution of aircraft multilateration errors using six ground stations. Errors are computed as the difference between the determined positions and the positions reported by the aircraft themselves in the transmitted messages.

5.9 Using Messages from Multilaterated Aircraft

As described in Section 3.2, the system can also compute the position of the aircraft that do not transmit their position. Especially smaller aircraft are not equipped with ADS-B capable transponders. According to the OpenSky Report 2016 [31], about 70% of transponders support ADS-B messages. ADS-B will be mandatory by 2020. In the report it is also mentioned that 26% of received Mode S messages are ADS-B messages. The most frequent types are altitude replies, which account for 35% of all messages. As a result, only 2D multilateration has to be performed, since the altitude is already known.

Figure 13 shows the result of the aircraft multilateration performed with six ground stations. For the evaluation of the position accuracy, the multilateration has been performed on messages from aircraft that also transmit their position. Therefore, part of the error is also due to unknown errors in the transmitted aircraft positions. The calculated positions have a median error of approximately 300 m. With a median error of 25 m of the handset multilateration when using the aircraft with known position, the calculated aircraft positions can not be used to get more accurate handset positions. But in cases with too few ADS-B equipped aircraft in range, the aircraft multilateration could still be used to get a rough position estimate of the handset.

We did not reach a conclusion yet, why these errors are much larger than the handset multilateration errors. Different additional error sources are possible. Among them are:

- Large dilution of precision due to bad geographical distribution of the ground stations.
- Inaccurate ground truth of the position received from the aircraft: unpredictable delay from receiving position to transmitting it and unknown position offset between the GPS antenna and the aircraft transponder.

- Accumulation of synchronization errors between the ground stations.

6 RELATED WORK

Our method employs aircraft messages for the purpose of indoor localization. To the best of our knowledge, nobody has done this before. In the following, we discuss existing work in the two independent fields of indoor localization (Section 6.1) and air traffic surveillance and control (Section 6.2).

6.1 Indoor Localization

Much of the research on indoor localization focuses on providing accurate localization, for instance to room level or even sub-meter accuracy. The cost factors to get so accurate are

- (I) the installation of dedicated Infrastructure, like for instance one beacon in each building up to several in each room;
- (T) a Training or initialization phase to gather data which is necessary for the subsequent localization;
- (E) the usage of Expensive user equipment.

Most methods do not have all three of these drawbacks, but at least one. In contrast, our method is basically infrastructure free, does not need training and receivers are cheap. Simply adding a small antenna can turn a cheap smartphone into a user handset. Our method requires only one or a few ground stations for a region, which can be hundreds of kilometers in diameter. Therefore, our method is suitable to be used on a global scale. In contrast to GPS, which is also a global positioning system, our method also works indoors. To get rid of all the cost factors listed above, we give up some accuracy. As outlined in the introduction, various applications exist for which localization precision is not essential. In other words, our method fills the void between cheap, global and easy to use outdoor localization such as GPS and precise, but local or expensive, indoor localization.

A plethora of indoor localization methods exist and different classifications are possible. For instance, Seco et al. [32] classify the methods into four categories: geometry-based, minimization of a cost function, fingerprinting and Bayesian. Another classification can be made based on the type of employed signals. For indoor localization, signals in basically the whole range of the electromagnetic spectrum up to light have been used. A third possibility is to distinguish the methods by their kind of fundamental measurements, which include received signal strength (RSS), time of arrival (TOA), time difference of arrival (TDOA) or angle of arrival (AOA). An overview of these localization techniques is given by Liu et al. [24]. Our method uses the TDOA technique, which is also called *multilateration*.

Here, we will discuss different indoor localization systems based on the employed signals' type. For each category, we indicate in parentheses which drawbacks the method has, using the letters from the list above.

WiFi. (T) WiFi signals are popular for indoor localization, because WiFi hotspots are widely in use. Therefore, no dedicated infrastructure, like beacons, is necessary for WiFi localization methods. In a survey by Liu et al. [24], many types of wireless indoor

localization methods are compared and WiFi based approaches generally have an accuracy of a few meters. This finding is confirmed by more recent results from the annual *Microsoft Indoor Localization Competition*.¹² WiFi localization methods require a training phase in which either the positions of WiFi access points are determined or fingerprints at different locations are gathered. Furthermore, infrastructure changes have to be detected and the database needs to be updated accordingly.

Ultrasound. (I) In contrast to WiFi based localization, which is infrastructure free, ultrasound based methods require dedicated hardware. However, ultrasound systems are relatively inexpensive and have proven to be very accurate compared to many other indoor localization methods. For instance the SmartLOCUS [5] system and the SpiderBat [27] system achieve centimeter-level accuracy. Still, ultrasound techniques are not in much favor any more, because the necessary signal strength for distances more than a few meters is high and these systems are prone to ambient noise, like for instance jingling keys. Also, ultrasound systems raise concerns about animal health compatibility – mostly pets like dogs and cats.

Light. (T,E) The most accurate results in the Microsoft Indoor Localization Competition are achieved by laser- and camera-based methods.¹³ The best system achieves an accuracy of 5 cm using two lasers and multiple high-end cameras. However, this system costs a quarter million dollars. Still, even cheap cameras today have a high number of pixels, resulting in a fine resolution when used for positioning. For instance, smartphones featuring Google's *Project Tango* hardware¹⁴ have an accuracy of a few centimeters [15]. However, the relative error of approximately 4% compared to the operating range is rather large. Also, because the maximum range of these phones' localization systems is limited to 4 m,¹⁵ and because the compute requirements are high, mapping of large rooms or even a whole building takes a long time. After the training phase, in which a building is mapped, users can localize themselves by matching their locally detected features with those in the model of the building. Thereby, they establish the link between their local coordinate frame and a global coordinate system used for mapping the building. Like this, users learn in which room of a building they are.

The widespread use of light emitting diodes (LED) and the miniaturization of processors has opened a new field of visible light communication and localization techniques. Pathak et al. [28] give an extensive overview of current methods. Among the advantages of visible light indoor localization they list the large installation base of LEDs in buildings, which helps canceling out noise in measurements. For localization, approaches based on RSS [23] and AOA [21] have been used. RSS based approaches with light can achieve sub-meter accuracy [23] and therefore are one order of magnitude more accurate compared to WiFi RSS techniques [9]. Analyzing a camera image to perform AOA localization is even more accurate and can yield a positioning error of 10 cm [21]. Sub-centimeter visible

¹²<https://www.microsoft.com/en-us/research/event/microsoft-indoor-localization-competition-ipsn-2014/#official-results>

¹³<https://www.microsoft.com/en-us/research/event/microsoft-indoor-localization-competition-ipsn-2016/#official-results>

¹⁴Currently, the *Lenovo Phab 2 Pro* and the *Asus ZenFone AR* are available.

¹⁵Google's *Tango* developer guide lists a range of 0.5 to 4 m: https://developers.google.com/tango/overview/depth-perception#usability_tips

light localization has also been demonstrated, by using multiple receivers [38]. Similar to methods leveraging WiFi base stations, also LED transmitters' positions need to be learned in a training phase. In contrast to WiFi signals, light does not penetrate walls, which can be a benefit or a downside. On the one hand, interfering multipath signals from neighbouring rooms, which introduce errors, are eliminated, but on the other hand, multiple LEDs have to be installed in every room.

Bluetooth. (T,I) Another type of signal used for indoor localization is Bluetooth. Bluetooth is similar to WiFi in that both systems share the 2.4 GHz frequency band. Compared to WiFi, which can take tens of seconds for identifying base stations, faster response times can be achieved with Bluetooth [7, 26]. This is important, because at walking speed, the set of visible beacons can change quickly. Not being able to use signals from intermittently visible transmitters is detrimental for the localization accuracy.

Due to the protocol specifications, Bluetooth devices have to be paired before user data can be exchanged. However, it has been shown that only using publicly announced device names and received signal strengths of Bluetooth beacons is enough to achieve a localization error of less than 3 meters [18].

While most Bluetooth RSS localization approaches achieve an accuracy of multiple meters, more elaborate approaches, for instance ones using Neural Networks can achieve sub-meter accuracy [1]. Accuracy using trilateration can be even better with an error less than 0.5 m [29].

RFID. (I) There exists also work on localization with RFID tags. RFID tags come in two flavors: Active tags have an internal battery and passive tags do not. The latter therefore have limited capabilities. Since even active tags only have extremely limited energy, RFID tags can only communicate over short distances and are mostly useful to identify the proximity of objects. Still, various more elaborate RFID localization schemes exist, applying RSS, TOA or TDOA. Bouet and Dos Santos [4] provide an overview of the work on RFID localization. RFID tags are cheap, but due to the short range, many tags have to be deployed for a localization system serving a whole building.

Sensor Fusion. (not stand-alone) Sensor assisted localization methods are particularly favored in smartphone applications, because basically all of these devices feature an inertial measurement unit (IMU) comprising an accelerometer, a gyroscope and a compass. For instance, the accelerometer can be used to estimate the movement speed and the compass can provide the orientation of the device. Based on this principle and by counting steps of a person, pedestrian localization systems have been developed [17]. Although using the sensor data alone over long time periods is not accurate due to accumulating errors, the sensors can bridge gaps in the operation of another localization system. This technique is called *dead reckoning*. For example, cars driving into a tunnel will lose signals from GPS satellites, but based on measurements of the current driving speed, their position can be estimated until the tunnel ends [34]. Although IMUs are cheap and require no infrastructure, they are not suitable as stand-alone localization and navigation systems due to the drift of the estimated position. Therefore, continuous recalibration using a second localization system is necessary. Nevertheless, by relaxing

the requirements from absolute three-dimensional positioning to only detecting floor level changes, IMUs alone *can* provide the basis for sufficiently accurate results over multiple hours [39].

GSM. (T) All the mentioned localization techniques above only provide a local coordinate system by themselves. When setting up the infrastructure or in the training phase, the local coordinates of the system in the building have to be related to a global coordinate system, like for instance GPS coordinates. For this, some reference points need to be matched. For example, positions outside the building that have already been mapped to the established local system coordinates, can be linked to the reported location of a GPS receiver. Our method only needs the coordinates of a ground station for each region, which due to the relaxed accuracy requirements, can even be estimated by picking the position manually on a map. Thus, our system does not require relating a coordinate frame for each building to global coordinates. In this respect, the localization techniques most similar to our method are those based on GSM signals. These techniques leverage signals sent by cellular network antennas, which can be received over distances up to 35 kilometers [8]. Therefore, in theory, the positions of only relatively few antennas need to be known. Because the available frequency spectrum is limited, the number of simultaneous users per cell is bounded. Thus, in order to serve more users, the signal strength of practically all GSM antennas is intentionally set much lower than the maximum. With this measure, cells at distances closer than 35 km can reuse the same frequency spectrum. Normally, cells have a diameter of only a few kilometers [30]. Aircraft signals can be received over distances of up to 400 km, limited by the curvature of the Earth, which is an increase of two orders of magnitude compared to these practical GSM cell sizes. And note that the covered area increases quadratically with the diameter. This means that our method needs far less ground/base stations than GSM-based localization schemes. A problem of the GSM cells being designed for small overlap between neighboring cells is that at a given position, signals from less than the required four antennas might be available. In this case, no position can be determined. The accuracy of GSM localization is 50 to 150 m outdoors [20], which is a factor of 2 to 6 less accurate than our method. Methods leveraging the wider band LTE channels can achieve better accuracy than GSM-based methods. The Cramér-Rao bound on the accuracy is 20 to 40 meters [12], but practical systems usually do not get close to this theoretical limit.

GSM localization methods do not need additional infrastructure, as cellular antennas are already widely available all around the globe. However, to learn the positions of the cellular antennas, a training phase is required, because this information is not publicly available with good enough accuracy.

6.2 Air Traffic Surveillance and Guidance

Various projects exist which localize and track aircraft using multilateration based on received ADS-B messages. Using a network of ground stations, which are time-synchronized using GPS receivers, a median aircraft localization error of 128 meters has been achieved [33]. Another test series has shown a horizontal positioning error of 127 meters 95% of the time [10]. In simulation, it has been shown that the horizontal aircraft localization error could be

as low as 11 meters if the geographic distribution of the ground stations is good [19].

Our system also employs a network of ground stations which determine the time and optionally the position of message transmissions by aircraft, but our ground stations do not need to be time-synchronized. However, we do not stop here, but use the determined message transmission positions and times to perform multilateration of a *handset*. To the best of our knowledge, nobody has done this before.

Aircraft Signal Receiver Networks. Even though aircraft tracking systems only need a relatively sparse network of ground stations, it is nevertheless a practical challenge setting up a global ground station infrastructure. At least two companies have taken part in this endeavor by leveraging the participation of hobbyists. Note however that these companies do not localize users.

FlightAware¹⁶ is a large system of aircraft signal receivers organized by a company with the same name. The network mostly consists of ground stations installed in private homes, sending their data to a FlightAware server. The company provides online instructions, software and hardware to set up a ground station.¹⁷ The ground stations decode aircraft messages and send them together with a time stamp to the company's servers. The aircraft positions are shown on an online map¹⁸ or can be accessed through an API by registered users who contribute data. For tracking aircraft which do not send their position, the servers apparently multilaterate the aircraft based on the time stamps associated with the received messages. However, the exact method is not disclosed to the best of our knowledge and therefore the accuracy is unknown. Based on the fact that the time stamps sent by the ground stations are less accurate than those in our method, we are relatively certain that the accuracy is not good. As we show in Section 5.9, even with our accurate time stamps from upsampled signals, determining aircraft positions accurately is difficult. For the purpose of displaying aircraft positions on a map, the positions do not need to be determined accurately. Therefore, for FlightAware's business case, this data is good enough. However, the more stringent accuracy requirements of user localization ask for more accurate aircraft positions. The advantage of FlightAware is their large user base, which gives their network good coverage of many regions. But compared to our method, FlightAware does not localize users.

Flightradar24¹⁹ is another provider of a website displaying current aircraft positions on a map. This website integrates a bunch of additional information about aircraft, such as the aircraft model and technical data, and for commercial aircraft additionally the flight number, origin, destination and the current delay. Flightradar24 also collects data from ground stations installed by hobbyists, but the company also has an active program, in the course of which receivers are sent out to those interested persons, which increase their regional coverage the most.²⁰ The receivers, which are distributed by the company are expensive and feature a GPS receiver

for accurate timing of received messages. Therefore, these receivers have more accurate time stamps than our proposed method, at the trade-off of much higher costs. Although we could not find data on the accuracy of multilateration of aircraft using these receivers, due to the synchronized clocks, the results should be good. However, not all of the ground stations in the Flightradar24 network are such expensive receivers. Also Flightradar24 does not localize users.

Although we are not able to use ground station networks as dense and large as those of FlightAware and Flightradar24, we can still test our method relatively well with a low number of receivers, because aircraft signals can be received over several hundred kilometers. Also, as discussed in the next section, many aircraft send their position, so one ground station receiving a message is enough, because the transmit time is the only variable.

While FlightAware and Flightradar24 do not provide historical data, OpenSky [31] is an effort to do so, mainly for research purposes. This project also relies on volunteers deploying ground stations and sharing the gathered data. One problem of OpenSky is that the necessary hardware costs about € 700, which is a significant entry barrier. In contrast, our ground stations cost less than \$ 100. The project website claims that the network comprises "hundreds of receivers".²¹

7 CONCLUSION

We have shown that using aircraft signals to localize users is a viable approach, even when the receiver is indoors. Our method fills a gap between globally available outdoor localization and accurate but expensive or cumbersome indoor localization. A few ground stations are enough to serve a region several hundred kilometers in diameter, which makes our method basically infrastructure free.

To better understand the possibilities and limitations of localization using aircraft signals, a thorough evaluation of the influence of various parameters on the performance, as outlined in Section 5, is necessary. For instance, the time of day influences the density of air traffic, and it would be interesting to determine the effects on the area coverage, since even a low number of aircraft can be sufficient for localization.

Plenty of accuracy improvements to our prototype system are possible:

- Using more advanced signal processing to more precisely detect signal arrival times and detect more signals per time.
- Improvements to the RF chain, such as employing antennas designed for ADS-B.
- Applying enhancements to the position estimation, such as selecting only "good" measurements for the least squares computation, computing a weighted least-squares solution, applying multipath mitigation techniques or using a different localization algorithm such as a maximum likelihood approach.
- Precisely localizing ground station positions.
- Choosing an optimal placement of the ground stations, to reduce error sources such as multipath and to maximize received signal strength, number of received messages and observed unique aircraft.

¹⁶<https://flightaware.com/>

¹⁷USB SDR ADS-B receiver dongles: <http://flightaware.com/adsb/prostick/>. PiAware ADS-B ground station software setup guide: <http://flightaware.com/adsb/piaware/build>

¹⁸Live map of aircraft currently in the sky: <https://flightaware.com/live/>

¹⁹<https://www.flightradar24.com/>

²⁰Flightradar24.com – Apply for receiver: <https://www.flightradar24.com/apply-for-receiver>

²¹The OpenSky Network – Services: <https://opensky-network.org/services>

In the future, a larger number of aircraft will be equipped with ADS-B transponders due to regulatory requirements and growing air traffic, which will increase the availability of our proposed localization method and also improve its accuracy due to more possibilities for error correction.

The presented handset design shows that our method could be integrated in a smartphone. The only additional hardware required in a smartphone is a small antenna, which easily fits into such a form factor, and a few components for the RF front end. Given the usual level of system-on-chip integration, this should be an inexpensive addition.

REFERENCES

- [1] Marco Altini, Davide Brunelli, Elisabetta Farella, and Luca Benini. 2010. Bluetooth Indoor Localization with Multiple Neural Networks. In *5th IEEE International Symposium on Wireless Pervasive Computing (ISWPC)*. IEEE, 295–300.
- [2] Penina Axelrad, Ben K. Bradley, James Donna, Megan Mitchell, and Shan Mohiuddin. 2011. Collective Detection and Direct Positioning Using Multiple GNSS Satellites. *Navigation: Journal of the Institute of Navigation* 58, 4 (2011), 305–321.
- [3] Pascal Bissig, Manuel Eichelberger, and Roger Wattenhofer. 2017. Fast and Robust GPS Fix Using One Millisecond of Data. In *Proceedings of The 16th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN 2017)*. 11. DOI: <https://doi.org/10.1145/3055031.3055083>
- [4] Mathieu Bouet and Aldri L. Dos Santos. 2008. RFID Tags: Positioning Principles and Localization Techniques. In *1st IFIP Wireless Days (WD 2008)*. IEEE, 1–5.
- [5] Cyril Brignone, Tim Connors, Geoff Lyon, and Salil Pradhan. 2005. *SmartLOCUS: An autonomous, self-assembling sensor network for indoor asset and systems management*. Technical Report. Mobile Media Syst. Lab., HP Laboratories, Palo Alto, CA.
- [6] Pete Bunce. 2016. The Clock is Ticking on ADS-B Equipage. (April 2016). Retrieved April 11, 2017 from <http://www.aviationpros.com/article/12180962-the-clock-is-ticking-on-ads-b-equipage>
- [7] Sudarshan S. Chawathe. 2009. Low-Latency Indoor Localization Using Bluetooth Beacons. In *12th International IEEE Conference on Intelligent Transportation Systems (ITSC 2009)*. IEEE, 1–7.
- [8] Mike Y. Chen, Timothy Sohn, Dmitri Chmlev, Dirk Haehnel, Jeffrey Hightower, Jeff Hughes, Anthony LaMarca, Fred Potter, Ian Smith, and Alex Varshavsky. 2006. Practical Metropolitan-Scale Positioning for GSM Phones. In *International Conference on Ubiquitous Computing (UbiComp 2006)*. Springer, 225–242.
- [9] Krishna Chintalapudi, Anand Padmanabha Iyer, and Venkata N. Padmanabhan. 2010. Indoor Localization Without the Pain. In *Proceedings of the sixteenth annual international conference on mobile computing and networking (MobiCom 2010)*. ACM, 173–184.
- [10] Anastasios Daskalakis and Patrick Martone. 2003. A Technical Assessment of ADS-B and Multilateration Technology in the Gulf of Mexico. In *Proceedings of the 2003 IEEE Radar Conference*. IEEE, 370–378.
- [11] Jason Davidson. 2013. ADS-B Requirements Coming Into Effect. (September 2013). Retrieved April 11, 2017 from <http://www.universalweather.com/blog/2013/09/ads-b-requirements-coming-into-effect/>
- [12] José A. del Peral-Rosado, José A. López-Salcedo, Gonzalo Seco-Granados, Francesca Zanier, and Massimo Crisci. 2012. Achievable Localization Accuracy of the Positioning Reference Signal of 3GPP LTE. In *International Conference on Localization and GNSS (ICL-GNSS)*. IEEE, 1–6.
- [13] United States of America Department of Defense. 2008. *Global Positioning System Standard Positioning Service Performance Standard (GPS SPS PS), 4th Edition*. Technical Report.
- [14] Burton S. Garbow, Kenneth E. Hillstrom, and Jorge J. More. 1980. Documentation for MINPACK Subroutine LMDIF – Double precision version. (March 1980). Retrieved April 10, 2017 from <https://www.math.utah.edu/software/minpack/minpack/lmdif.html>
- [15] Eberhard Gülich. 2016. Investigations on Google Tango Development Kit for Personal Indoor Mapping. In *The 19th AGILE International Conference on Geographic Information Science (AGILE 2016)*. 3. https://agile-online.org/Conference_Paper/cds/agile_2016/posters/102_Paper_in_PDF.pdf Poster Abstract.
- [16] International Civil Aviation Organisation (ICAO). 2014. *Aeronautical Telecommunications – Surveillance and Collision Avoidance Systems. Annex 10 to the Convention on International Civil Aviation. Volume IV*. (July 2014).
- [17] Antonio R. Jimenez, Fernando Seco, Carlos Prieto, and Jorge Guevara. 2009. A Comparison of Pedestrian Dead-Reckoning Algorithms using a Low-Cost MEMS IMU. In *IEEE International Symposium on Intelligent Signal Processing (WISP 2009)*. IEEE, 37–42.
- [18] Troy A. Johnson and Patrick Seeling. 2012. Localization Using Bluetooth Device Names. In *Proceedings of the thirteenth ACM international symposium on Mobile Ad Hoc Networking and Computing (MobiHoc 2012)*. ACM, 247–248.
- [19] Regina Kaune, Christian Steffes, Sven Rau, Wolfgang Konle, and Jürgen Pagel. 2012. Wide Area Multilateration using ADS-B Transponder Signals. In *15th International Conference on Information Fusion (FUSION)*. IEEE, 727–734.
- [20] Tomislav Kos, Mislav Grgic, and Gordan Sisul. 2006. Mobile User Positioning in GSM/UMTS Cellular Networks. In *48th International Symposium ELMAR-2006 focused on Multimedia Signal Processing and Communications*. IEEE, 185–188.
- [21] Ye-Sheng Kuo, Pat Pannuto, Ko-Jen Hsiao, and Prabal Dutta. 2014. Luxapose: Indoor Positioning with Mobile Phones and Visible Light. In *Proceedings of the 20th annual international conference on Mobile computing and networking (MobiCom 2014)*. ACM, 447–458. DOI: <https://doi.org/10.1145/2639108.2639109>
- [22] Richard B. Langley. 1999. Dilution of Precision. *GPS World* 10, 5 (May 1999), 52–59.
- [23] Liqun Li, Pan Hu, Chunyi Peng, Guobin Shen, and Feng Zhao. 2014. Epsilon: A Visible Light Based Positioning System. In *Proceedings of the 11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2014)*. 331–343.
- [24] Hui Liu, Houshang Darabi, Pat Banerjee, and Jing Liu. 2007. Survey of Wireless Indoor Positioning Techniques and Systems. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 37, 6 (2007), 1067–1080.
- [25] Jie Liu, Bodhi Priyantha, Ted Hart, Heitor S. Ramos, Antonio A.F. Loureiro, and Qiang Wang. 2012. Energy Efficient GPS Sensing with Cloud Offloading. In *Proceedings of the 10th ACM Conference on Embedded Networked Sensor Systems (SenSys 2012)*. ACM, 85–98.
- [26] Nicholas Mair and Qusay H. Mahmoud. 2012. A collaborative Bluetooth-based approach to localization of mobile devices. In *8th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom)*. IEEE, 363–371.
- [27] Georg Oberholzer, Philipp Sommer, and Roger Wattenhofer. 2011. SpiderBat: Augmenting Wireless Sensor Networks with Distance and Angle Information. In *Proceedings of the 10th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN 2011)*. IEEE, 211–222.
- [28] Parth H. Pathak, Xiaotao Feng, Pengfei Hu, and Prasant Mohapatra. 2015. Visible Light Communication, Networking, and Sensing: A Survey, Potential and Challenges. *IEEE Communications Surveys & Tutorials* 17, 4 (2015), 2047–2077.
- [29] Aswin N. Raghavan, Harini Ananthapadmanaban, Manimaran S. Sivamurugan, and Balaraman Ravindran. 2010. Accurate Mobile Robot Localization in indoor environments using Bluetooth. In *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 4391–4396.
- [30] Moe Rahnama. 1993. Overview Of The GSM System and Protocol Architecture. *IEEE Communications Magazine* 31, 4 (April 1993), 92–100.
- [31] Matthias Schäfer, Martin Strohmeier, Matthew Smith, Markus Fuchs, Rui Pinheiro, Vincent Lenders, and Ivan Martinovic. 2016. OpenSky Report 2016: Facts and Figures on SSR Mode S and ADS-B Usage. In *35th Digital Avionics Systems Conference (DASC)*. IEEE/AIAA, 1–9. DOI: <https://doi.org/10.1109/DASC.2016.7778030>
- [32] Fernando Seco, Antonio R. Jiménez, Carlos Prieto, Javier Roa, and Katerina Koutsou. 2009. A survey of Mathematical Methods for Indoor Localization. In *IEEE International Symposium on Intelligent Signal Processing (WISP 2009)*. IEEE, 9–14.
- [33] Rudolf Seller and Ádám Szüllő. 2013. Wide Area Multilateration Demonstration System. In *21st International Conference on Applied Electromagnetics and Communications (ICECom)*. IEEE, 1–5.
- [34] Salman Ali Shaikat, Khalid Munawar, Muhammad Arif, Aamer Iqbal Bhatti, Umar Iqbal Bhatti, and Ubaid M. Al-Saggaf. 2016. Robust vehicle localization with GPS dropouts. In *6th International Conference on Intelligent and Advanced Systems (ICIAS)*. IEEE, 1–6.
- [35] Martin Strohmeier, Matthias Schäfer, Vincent Lenders, and Ivan Martinovic. 2014. Realities and Challenges of NextGen Air Traffic Management: The Case of ADS-B. *IEEE Communications Magazine* 52, 5 (May 2014), 111–118.
- [36] Junzi Sun, Jacco Hoekstra, and Joost EllerBroek. 2017. ADS-B Decoding Guide. (2017). Retrieved April 10, 2017 from <https://adsb-decode-guide.readthedocs.io/en/latest/content/introduction.html>
- [37] Phil Vabre. 2010. Air Traffic Services Surveillance Systems, Including An Explanation of Primary and Secondary Radar. (2010). Retrieved April 11, 2017 from <http://www.airwaysmuseum.com/Surveillance.htm>
- [38] Se-Hoon Yang, Hyun-Seung Kim, Yong-Hwan Son, and Sang-Kook Han. 2014. Three-Dimensional Visible Light Indoor Localization Using AOA and RSS With Multiple Optical Receivers. *Journal of Lightwave Technology* 32, 14 (2014), 2480–2485.
- [39] Haibo Ye, Tao Gu, Xiaorui Zhu, Jinwei Xu, Xianping Tao, Jian Lu, and Ning Jin. 2012. FTrack: Infrastructure-free Floor Localization via Mobile Phone Sensing. In *IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 2–10.