

Crowdsourced Multilateration: An Empirical Evaluation of `mlat-client/server` in Switzerland

Martin Strohmeier¹, Michael Greub³, and Matthias Schäfer^{4*}

¹ armasuisse Science + Technology, ² OpenSky Network, ³ ETH Zurich, ⁴ Hochschule RheinMain

*Corresponding author: strohmeier@opensky-network.org

Abstract

Crowdsourced multilateration (MLAT) is widely deployed and used operationally, yet little is documented about the end-to-end behaviour of the canonical open-source `mlat-client/server` stack. We evaluate OpenSky's test MLAT deployment in Switzerland: 27.8 million position estimates accepted by the server and 27.7 million MLAT positions stored in OpenSky's Trino database, collected from 20 May to 20 July 2026. A configuration snapshot immediately after the study contains 26 connected receivers, including 19 in Switzerland. In `mlat-client/server`, Automatic Dependent Surveillance-Broadcast (ADS-B) position reports synchronise the receiver clocks. They also serve us as an operational accuracy reference.

For the 19,862,378 position estimates with a qualifying ADS-B reference, the median horizontal disagreement is 106.0 m and the 95th percentile 654.5 m. Geometry and solution type also matter: three-receiver altitude-constrained position estimates account for 15.0% of position estimates but 46.8% of disagreements above 5 km. In our deployment, MLAT is available within two seconds for 36.4% of ADS-B position reports over Switzerland, falling to 18.1% below 3 km.

Keywords: multilateration; crowdsourced sensing; ADS-B; synchronisation; OpenSky

Abbreviations: ADS-B: Automatic Dependent Surveillance-Broadcast, MLAT: multilateration, GDOP: geometric dilution of precision, HDOP: horizontal dilution of precision, NIC: navigation integrity category, NAC: navigation accuracy category

1. Introduction

Automatic Dependent Surveillance-Broadcast (ADS-B) made large-scale open aircraft surveillance practical and widely used for open research [1]. Many crowdsourced networks also use multilateration (MLAT): reception-time differences at separated receivers constrain a transmitter's position. MLAT can locate Mode S aircraft without a usable ADS-B position. Wide-area multilateration (WAM) is established surveillance technology [2, 3, 4]. Here we study the open-source `mlat-client/server` implementation used by volunteer and commercial networks alike.

Today, crowdsourced MLAT primarily fills gaps in live aircraft positions. This need decreases as ADS-B equipage grows. However, archived MLAT data also offer valuable research opportunities. As we will show, around 10% of aircraft in our study were not providing ADS-B positions and are thus excluded from measurement studies such as measuring climate impact [5]. This share can vary considerably with the weather: general aviation aircraft often carry no ADS-B, and better weather brings more general aviation traffic and therefore a higher share of aircraft without ADS-B. Use of MLAT data to supplement such studies requires that researchers understand their accuracy, timestamps, gross errors, and availability. These properties depend on selective forwarding, clock alignment, grouping, acceptance, and downstream transport, as well as on the solver itself.

We audit OpenSky's first test deployment of a standard `mlat-client/server` setup and evaluate its output. We release the full data of this study on Zenodo at <https://zenodo.org/records/23675290>.

Contributions. This paper makes three main contributions:

- We explain how the default `mlat-client/server` stack selects observations, synchronises clocks and accepts positions, tracing the changes that permit routine MLAT of ADS-B-equipped aircraft.
- We deploy a multilateration server with 26 connected receivers, including 19 in Switzerland, for real-world evaluation. The prototype has since been opened and accepts connections from global ADS-B feeders at <https://opensky-network.org/feed/mlat>.
- We evaluate accuracy and availability in a large operational dataset in Switzerland, examining receiver geometry, solution type, altitude and reported uncertainty.

2. Related work

Research on MLAT addresses both the positioning algorithms and the conditions under which their measurements are obtained. Yaro compares reference-station selection techniques in simulation, showing that the choice of stations used to form reception-time differences affects position accuracy [6]. Figuet et al. combine machine learning with triplet-wise and all-in-view multilateration, using OpenSky data to study how receiver count affects accuracy and coverage [7]. Sonnleitner and Hobiger investigate wide-area multilateration with unsynchronised low-cost ADS-B receivers [8]. Their setting makes clock alignment a central concern, as it is in the volunteer deployment studied here. Osypiuk and Surma analyse localisation uncertainty when altitude is supplied as an additional constraint, including the influence of station geometry [9]. We examine how constrained solutions perform in operational output, where altitude, geometry and observation selection vary together.

Timing measurements also support verification of aircraft reports. Baeye et al. implement statistical location verification in OpenSky and evaluate it on eight months of Radarcape data, detecting anomalous trajectories and sensor or transponder problems [10]. Olive et al. examine noise, uncertainty and quantisation in ADS-B and Mode S data and describe trajectory filters implemented in the traffic library [11].

Schäfer and Jonáš use selected (professional, non-crowdsourced) high-performance MLAT as a reference for ADS-B [12]. Their comparison associates positions with the same airborne downlink format 17 (DF17) transmission and restricts altitude and geometry to obtain a precise reference. In our work, the roles are reversed. Our evaluation complements these studies by following a deployed open-source system from observation selection through position estimation to downstream storage.

3. Background: The crowdsourced MLAT implementation

Oliver Jowett, known as mutability, developed `mlat-client` and `mlat-server` in 2015 [13]. He joined FlightAware in April of that year, and PiAware included an MLAT client by May 2015 [14, 15]. The software was subsequently developed publicly in multiple forks. ADS-B Exchange and the community aggregator adsb.fi distribute forks of `mlat-client` through their feeder installation scripts, while adsb.lol publishes deployment configurations for `mlat-server` [16, 17, 18]. These public sources document the wide adoption of the general software family. Flightradar24 also operates MLAT. Its public description does not identify this codebase as the implementation [19].

We audited the Wiedehopf-derived OpenSky server fork at revision 7e85f5bc (19 December 2025) [20]. For comparison, we also examined the differences to the original mutability/`mlat-client` revision fe70767b (version 0.2.13, 23 October 2023) and mutability/`mlat-server` revision bee59806 (17 February 2016) [21, 22]. Fig. 1 shows the information and control flow of the `mlat-client/server` stack. The following subsections follow its components.




Figure 1. Information and control flow of mlat-client and mlat-server in this study (schematic). D1 and D2 are the two evaluated MLAT datasets.

3.1 Receiver and mlat-client

Each receiver runs a Mode S decoder, which timestamps every message with the local receiver clock, and the mlat-client, which selects the messages sent to the server. The mlat-client reports the aircraft it receives every 30 s and removes an aircraft after 60 s without observation. It forwards messages of an aircraft only after ten observations and only if the server has requested that aircraft. It sends two kinds of messages. For every DF17 airborne position message with altitude and a navigation uncertainty category (NUCp) of at least 6, it sends a sync pair: the latest even and odd position messages of that aircraft, if they are at most 5 s apart. Ordinary Mode S messages are forwarded as MLAT messages, but this is suppressed for 30 s after a locally observed ADS-B position. The software therefore treats MLAT as a gap filler for aircraft without ADS-B.

3.2 Receiver interest and clock synchronisation

The mlat-server uses the reports of the mlat-clients to request traffic from relevant individual receivers. Two reporting receivers establish interest in an aircraft, but do not establish whether a position can be solved.

ADS-B-equipped aircraft also act as moving clock references for relative time synchronisation, a concept described for example in the original OpenSky paper or the OpenSky data competition on this topic [1, 23] but also various other applications [24]. Given the position reported in a sync pair and the known receiver coordinates, the server calculates the expected propagation-time difference between two receivers for the same transmission. Subtracting this from the observed timestamp difference gives an estimate of the relative clock offset. Repeated observations fit a relationship between two receivers: an offset and a relative clock-rate drift. Each receiver therefore need not have absolute GPS synchronisation.

3.3 Grouping, normalisation and solving

The grouping step collects identical copies of a message from all receivers for 0.9 s. MLAT messages enter it directly. Sync pairs provide a second route: in this fork, the even DF17 message of every sync pair enters the grouping step when at least three receivers track the aircraft. Absent from the original upstream server, this route was introduced in the ADS-B Exchange/Wiedehopf lineage in 2022. For our comparison in this paper, we make one single change to the stock image above and remove the restriction that prevents aircraft with a recent ADS-B position to be multilaterated. For aircraft with ADS-B, the even DF17 messages therefore make up almost all position estimates, because the mlat-client holds back their other Mode S messages.

The usable pairwise clock models form a graph. For each grouped message, the deployed normaliser composes clock transformations along a minimum spanning forest and retains the largest connected component with enough receivers. A receiver that receives a message does not necessarily contribute to the position estimate. Nominal timing models also differ: Radarcape declarations use a 1 GHz timestamp model and 15 ns jitter, while the unknown/dump1090 class, typically Raspberry Pis with RTL-SDRs, uses 12 MHz and 500 ns. Normalisation also checks physical feasibility. Timestamps are split at 2 ms gaps, each receiver appears at most once per cluster, and pairwise delays must satisfy a bound based on receiver separation. Poor synchronisation or stale clock pairing prevents admission.

The solver estimates a weighted nonlinear least-squares problem. The state contains three Earth-centred Earth-fixed (ECEF) position coordinates and a common range offset. Four independent receiver observations can locally determine these four unknowns. With only three receivers, the barometric altitude reported in Mode S supplies a fourth, weighted observation. Barometric altitude differs from geometric height by a weather-dependent offset, which enters the position as a systematic error. This permits a constrained solution with non-degenerate geometry, without guaranteeing uniqueness or accuracy. Receivers less than 1 km apart supply separate timestamps but count as one distinct site for this rule. The solver compares the pressure-derived altitude with WGS84 ellipsoidal height.

3.4 Aircraft tracker and outputs

The server keeps a state for each aircraft (aircraft tracker in Fig. 1). It caches the last plausible altitude report for at most 45 s. The standard error of a cached altitude starts at 250 ft and grows by 70 ft/s. A Kalman filter combines successive position estimates of the aircraft and estimates its speed, heading and position uncertainty. These Kalman estimates are not part of the outputs used in this study.

The server writes each accepted position estimate with its covariance to a pseudorange JSON file (D1). It also re-encodes each position estimate as DF18 extended squitter frames in Beast format, because Beast carries only Mode S frames. The OpenSky readsb/Kafka/Trino pipeline reads these frames and produces the analytical position table *readsb_mlat_sv* in OpenSky's Trino interface (D2).¹ Neither output records the full number of attempted or rejected solves.

Each dataset records a different time. A D1 record carries the server clock at the moment the first copy of the message in the solved cluster arrived from an mlat-client. This is not the transmission time. It is later by the time the mlat-client waits before sending the message and by the network delay to the server. The Beast frames carry a fixed MLAT marker instead of a receiver timestamp, and D2 records the time of OpenSky's database. A D3 record carries the earliest reception time of the ADS-B position report at OpenSky.

For even DF17 messages, each mlat-client sends the message twice, with the preceding and with the

¹Discussing the exact behaviour of readsb is out of scope of this paper.

following odd message. If every receiver of the cluster sent both copies, the server keeps the later one. The D1 time can then lag the transmission by the interval to the next odd message, typically about 0.5 s. This is consistent with the shift of 0.180 s that best aligns D1 with ADS-B (Appendix Appendix B). We did not measure the share of affected position estimates.

4. Evaluating the quality of mlat-server/client

4.1 Study datasets

Table 1. Datasets covering 20 May to 20 July 2026 inclusive (UTC).

	D1: mlat-server raw	D2: OpenSky Trino MLAT	D3: OpenSky Trino ADS-B
Rows	27,842,111	27,666,844	1,509,890,842
Aircraft addresses	13,821	13,809	15,740

We build three datasets covering 20 May to 20 July 2026 (UTC, 62 days) (Table 1). D1 contains the accepted position estimates from the mlat-server logs, including the original server timestamps, receiver coordinates, solution covariance and available altitude information. D2 contains the MLAT positions stored in OpenSky’s readsb_mlat_sv table after going through a Kafka/readsb pipeline, with their database timestamps and coordinates.

D3 contains decoded DF17 ADS-B positions from position_data4 in OpenSky’s Trino interface. We include the Swiss viewport (45.7 to 47.95° N, 5.7 to 10.7° E), together with reports from the same aircraft within 60 s of any D1 or D2 position, including reports outside the viewport. We retain reception timestamp, aircraft address, coordinates, altitude, quality indicators and raw message.

4.2 Matching MLAT results to ADS-B tracks

Our deployment multilaterates ADS-B-equipped aircraft, so that their position estimates can be matched with their ADS-B position reports. We use DF17 position reports from D3 with valid coordinates and navigation integrity category (NIC) of at least seven as references.

The mlat-server does not multilaterate every ADS-B position message. We searched for the position message behind each D1 position estimate. Across all 27,842,111 D1 position estimates, only 73,187 have a position message from the same aircraft within 0.5 ms. Most D1 position estimates of ADS-B aircraft are themselves even DF17 position messages. The D1 timestamp is the arrival time at the server, not the transmission time (Section 3.4), so only these few match within 0.5 ms. We therefore interpolate between ADS-B position reports to obtain a reference at the timestamp of each D1 record.

For each D1 position estimate, we select the nearest qualifying ADS-B reference from the same aircraft at or before its timestamp and at or after it. Both references must be within two seconds, have distinct timestamps, and imply a speed no greater than 500 m/s. We linearly interpolate latitude and longitude at the MLAT timestamp, taking the shorter longitude difference, and calculate horizontal disagreement using the WGS84 inverse geodesic [25]. Figure 2 illustrates this comparison.

Under these rules, not every position estimate has an ADS-B reference. We found references for 69.41% of D2 in the Trino database and 71.34% of D1.

The ADS-B reference is not fully independent of the MLAT calculation. The server can use an ADS-B message to adjust receiver clocks and then estimate that same message’s position through MLAT. In a snapshot, 20 of the 21 aircraft with an accepted result had an ADS-B message as their latest MLAT target. The pseudorange records do not identify the multilaterated message, so we cannot determine how often this happens across the full dataset. The study fork also logs every sync pair and MLAT




Figure 2. Matching an MLAT position to an ADS-B reference track.

message it receives from the mlat-clients, which can identify the multilaterated message in future deployments. 174 175

Our results in the next section describe how closely MLAT and ADS-B agree in this deployment. 176 They may overstate the accuracy attainable for aircraft without ADS-B, whose messages do not 177 provide the same synchronisation information. 178

4.3 Measuring receiver geometry and dilution of precision 179




Figure 3. Network-potential GDOP from all 19 Swiss snapshot receivers at 11 km ellipsoidal height. Triangles mark receivers. LV95 axes preserve regional geometry. Colours span the 1st to 99th percentiles on a logarithmic scale. This is not the geometry of an individual position estimate. 180

Geometric dilution of precision (GDOP) measures how receiver geometry amplifies measurement 180

errors when estimating position and the common range offset. We calculate GDOP from the receiver and aircraft coordinates in the Earth-centred, Earth-fixed (ECEF) frame. For receiver i at $\mathbf{r}_i$ and aircraft at $\mathbf{p}$, each row of the geometry matrix H is

$$H_i = \begin{bmatrix} \frac{(\mathbf{p} - \mathbf{r}_i)^T}{\|\mathbf{p} - \mathbf{r}_i\|}, & 1 \end{bmatrix}, \quad \text{GDOP} = \sqrt{\text{tr}[(H^T H)^{-1}]}.$$

Figure 3 shows the GDOP for a message received by all 19 Swiss receivers from an aircraft at 11 km. Median GDOP across this map is 7.12. This is a lower bound. A message reaches only a subset of the receivers, usually three to five, and each subset has a higher GDOP. Each D1 record is a single position estimate, so its error follows the GDOP of its own subset. A tracker that combines position estimates changes this relation.

5. Results

Figure 4 compares MLAT and ADS-B tracks in three selected examples. A straight cruise and a turn show close agreement. A take-off shows a gap in the MLAT output at low altitude.




Figure 4. Three selected five-minute comparisons of D1 position estimates (orange circles) with ADS-B position reports (D3, blue). The maps show every D1 position estimate in each interval. The error plots include only position estimates with an eligible ADS-B match, with both counts given in each panel. Map style adapted from traffic [26].

5.1 Overall agreement of MLAT with ADS-B

We first examine D1 with the timestamps recorded by the mlat-server. The 19,862,378 D1 position estimates with an ADS-B reference have a median horizontal disagreement of 106.0 m, a 95th percentile of 654.5 m and a 99th percentile of 1.577 km.

The median alone does not describe how often a position estimate differs substantially from its reference. Of the matched D1 position estimates, 51.83% exceed 100 m, 8.18% exceed 500 m, 2.201% exceed 1 km and 0.306% exceed 5 km. These correspond to 10,294,411, 1,624,967, 437,150 and 60,805 position estimates, respectively. All of them remain in the analysis.

For the positions in the downstream Trino database (D2), the disagreements exceed 100 m and 500 m more often, whereas the proportions above 1 km and 5 km are much closer. Interpreting this difference requires accounting for the timestamps and for the different positions included in the two datasets. Table 2 lists all results.

Table 2. Horizontal disagreement with ADS-B for all D1 and D2 positions with an eligible reference, over 20 May to 20 July 2026. P95 is the 95th percentile. The last four columns give percentages above each threshold.

Dataset	Matched positions	Median (m)	P95 (m)	> 100 m (%)	> 500 m (%)	> 1 km (%)	> 5 km (%)
D1: <i>mlat-server</i>	19,862,378	106.0	654.5	51.83	8.18	2.201	0.306
D2: Trino	19,202,951	239.7	765.4	93.52	13.66	2.567	0.218

5.2 Influence of the OpenSky pipeline time offset




Figure 5. Horizontal disagreement with ADS-B before and after subtracting 0.915 s from D2 timestamps. Each curve includes all positions with an eligible reference at its comparison time. The arrow connects the two D2 medians.

D2 timestamps are later than D1 by a median of 0.915 s, added by OpenSky’s downstream pipeline. This is the median over 27,656,495 pairs of D1 and D2 records matched by aircraft address within 10 m and 2 s. The 5th and 95th percentiles are 0.456 and 0.936 s. After subtracting the median from every D2 timestamp, D2 agrees with ADS-B as closely as D1: the 19,798,558 corrected positions with an eligible reference have a median of 100.6 m and a 95th percentile of 641.2 m, compared with 239.7 and 765.4 m before (Fig. 5). The rest of this paper uses D1. Shifting the D1 timestamps 0.180 s earlier improves the agreement further (Appendix Appendix B). We keep the original D1 timestamps.

5.3 Which solutions account for large disagreements?

Large disagreements are not evenly distributed across solution types. Three-receiver position estimates with an altitude constraint make up 15.0% of the D1 position estimates with an ADS-B reference, but contribute 46.8% of the 60,805 disagreements above 5 km. Their share increases with the

distance threshold: they account for 17.2% of disagreements above 100 m, 20.6% above 500 m and 27.8% above 1 km. These position estimates therefore contribute disproportionately to the largest differences from ADS-B.

The frequency within each group. The share of large disagreements contributed by a group differs from the frequency of a large disagreement within that group. Table 3 and Figure 6 distinguish the two.

Table 3. Solution types among the 19,862,378 D1 position estimates with an ADS-B reference. The last two columns distinguish the percentage above 5 km within each group from its share of all 60,805 disagreements above 5 km.

Solution type	Positions	P95 (m)	> 5 km within group	Share of all > 5 km
3 receivers + altitude	2,984,192	865.2	0.953%	46.8%
≥ 4 receivers + altitude	5,662,169	475.3	0.438%	40.7%
No altitude constraint	11,216,017	683.3	0.068%	12.5%




Figure 6. Solution types among the D1 position estimates with ADS-B references. Left: each group's share of all matched position estimates and of disagreements above each threshold. Right: the percentage exceeding each threshold within each group. Each position estimate belongs to exactly one group.

Position estimates without an altitude constraint make up 56.5% of the matched D1 position estimates. They contribute 58.2% of disagreements above 100 m and 63.6% above 500 m, but only 12.5% above 5 km. Within this group, 0.068% exceed 5 km, compared with 0.438% for constrained position estimates using four or more receivers. The latter group has a lower 95th percentile but a higher proportion above 5 km.

These differences do not isolate the effect of altitude assistance. The groups also differ in receiver geometry, aircraft altitude, synchronisation and reference availability. Groups with very few position estimates, particularly at high receiver counts, cannot support firm conclusions about rare large disagreements.

Receiver count and degrees of freedom. Among D1 position estimates with an ADS-B reference, the median and 95th percentile are 131 and 865 m for three receivers, 87 and 517 m for six, and

56 and 403 m for ten. For zero degrees of freedom, they are 146 and 873 m, compared with 63 and 399 m for four degrees of freedom. We next examine whether the arrangement of the receivers also matters when their number is held fixed.

5.4 Receiver geometry and stricter selection

Geometry at a fixed receiver count. We use all eligible D1 position estimates for the geometry and uncertainty calculations. To examine receiver arrangement while keeping receiver count and altitude assistance fixed, we compare 2,348,697 D1 position estimates with an ADS-B reference that use exactly four distinct receivers, four measurement rows and no altitude constraint. Median GDOP in this group is 16.41. From the lowest to the highest GDOP fifth, median disagreement rises from 112.8 to 290.1 m, while the proportion above 1 km rises from 0.65% to 11.59% (Figure 7). Receiver geometry therefore remains associated with disagreement within this more comparable set of position estimates.

The comparisons cover 10,495 aircraft and 99,555 aircraft-days. The Spearman correlation between GDOP and disagreement is 0.267. A fit on logarithmic axes gives a slope of 0.418, with an aircraft-day bootstrap 95% interval from 0.416 to 0.420.




Figure 7. Disagreement with ADS-B and GDOP for 2,348,697 D1 position estimates using exactly four receivers and no altitude constraint. Left: number of position estimates per hexagon. Right: median and 95th-percentile disagreement in five equally sized GDOP groups. The comparison describes accepted position estimates, not a causal effect of geometry.

5.5 Position quality and availability at low altitude

For each eligible ADS-B position report over Switzerland (D3), we look for a D1 position estimate of the same aircraft within one, two or five seconds of its original timestamp. This measures availability whenever ADS-B shows that an aircraft is present. Of all reports, 36.38% have a D1 position estimate within two seconds. Below 3 km, this falls to 18.07%. At 3 to 6 km, it reaches 54.24%. Fig. 8 shows availability per 0.1-degree grid cell, with higher values over the northern, less mountainous parts of Switzerland.

For D1 position estimates with an ADS-B reference, median disagreement varies only from 93.9 to 123.0 m across altitude bands (Figure 9). Below 3 km, the 2,530,262 comparisons have a median of 93.9 m, slightly lower than the 113.6 m median among 9,479,145 comparisons at 9 to 12 km. Yet 1.205% of the position estimates below 3 km disagree by more than 5 km, compared with 0.057% at 9 to 12 km.




Figure 8. OpenSky ADS-B position reports over Switzerland (left) and the percentage with a D1 position estimate for the same aircraft within ± 2 s of its original timestamp (right).

Radio visibility, receiver geometry and the positions accepted by the server also vary with altitude. All reference altitudes used in these comparisons are barometric. Geometric-altitude messages provide no qualifying reference pairs for D1, so we cannot compare the two altitude types. We assess horizontal disagreement within each altitude band and do not evaluate vertical accuracy.




Figure 9. Horizontal disagreement with ADS-B by aircraft altitude for D1 position estimates, using their original timestamps. Lines show the median, 95th and 99th percentiles on a logarithmic scale. Position estimates without an eligible ADS-B reference or decoded reference altitude are omitted.

Update rate. Across all 704.1 million OpenSky ADS-B position reports over Switzerland, availability increases from 23.06% within one second to 51.75% within five seconds. With the two-second window, daily availability ranges from 25.28% to 41.23%, with a median of 37.44%. The overall value of 36.38% weights days by their number of reports, whereas the median gives each day equal weight.

5.6 mlat-server uncertainty analysis

mLat-server itself offers uncertainty information, which estimates how sensitive a calculated position is to errors in the measurements. The *mLat-server* combines receiver geometry, estimated

reception-time errors and, when used, the assumed altitude error. Its solver returns a covariance matrix: numbers describing the estimated uncertainty in three spatial directions and how errors in those directions relate to each other. It estimates this covariance from how the weighted least-squares fit changes when the position is moved slightly. Weakly constrained directions receive larger uncertainty estimates.

Larger reported uncertainty generally corresponds to larger disagreement among D1 position estimates with an ADS-B reference (Figure 10, left). Here, uncertainty is the horizontal standard deviation in the direction with the largest estimated spread. Assuming Gaussian errors, we can also calculate a region around each position that should contain the aircraft with a chosen probability (right). The observed proportions of ADS-B references inside these regions match the expected percentages only in part. The nominal 95% and 99% regions contain 89.7% and 92.2% of the ADS-B references, respectively, but the nominal 50% region contains 76.2%.




Figure 10. Estimated uncertainty and disagreement with ADS-B for D1 position estimates. Left: median and 68th-percentile disagreement within groups of similar estimated horizontal uncertainty. Both axes are logarithmic. Right: expected percentages of position estimates and observed percentages of ADS-B references inside the uncertainty regions, in steps of one percentage point. The diagonal indicates agreement between the two.

5.7 Additional coverage of non-ADS-B aircraft provided by MLAT

We count how many aircraft MLAT adds to those with DF17 ADS-B positions inside Switzerland, as distinct aircraft addresses over the full study period from 20 May to 20 July 2026.

ADS-B provides positions for 13,211 addresses and MLAT for 13,518, with 11,924 present in both. Combining the two therefore adds 1,594 addresses to those observed through ADS-B alone, an increase of 12.1% over the study period.

6. Discussion

6.1 Position quality is sufficient for research applications

Our evaluation shows that crowdsourced MLAT can provide positions close to the ADS-B reference across much of the observed traffic. For *mlat-server* positions with an ADS-B reference, the median horizontal disagreement is 106.0 m. This quality is sufficient for many research applications. From our results in the previous section, we argue that making MLAT even more useful for research also requires sharing how a position was obtained, including receiver count and altitude assistance.

However, this incurs serious storage trade-offs when scaling, for potentially only very few relevant applications.

6.2 Comparison with Schäfer and Jonáš (2025)

Schäfer and Jonáš evaluate ADS-B using a dense German network of more than 170 homogeneous GPS-synchronised receivers, with most evaluated transmissions received by 50 or more receivers [12]. Their evaluation matches the exact DF17 transmission and requires flight level 200 or above, HDOP below 0.5, and systematic filtering of cell-hours. These rules retain 7,506,442 of 28,981,127 candidate reports. They report a cross-track standard deviation of 3.24 m and an absolute cross-track 95th percentile of 6.17 m.

Effect of ADS-B latency on synchronisation. Schäfer and Jonáš find a median ADS-B latency of 73.7 ms, which displaces a reported position mainly along track. The server takes the position in a sync pair as the true transmitter position. For all 14,891,360 matched D1 position estimates with reference speeds of 50 to 350 m/s and decoded altitude at both ADS-B endpoints, evaluated against all 87 Swiss synchronisation edges, a reference age of 74 ms produces a pairwise range error with a median of 4.1 m and a 95th percentile of 26.3 m. This error enters the clock models. We do not measure its effect on the final position error.

6.3 Validating aircraft positions in security applications

A major application of MLAT data is the validation of aircraft positions during suspected or known GNSS interference or ADS-B anomalies such as spoofing [10]. Since MLAT estimates a position from signal arrival times, it offers a way to compare what an aircraft reports with where its transmissions appear to originate. In a crowdsourced system, this check depends on ADS-B itself, because ADS-B positions synchronise the receiver clocks.

6.4 Inclusion of non-ADS-B aircraft into research data

Beyond checking existing positions, MLAT could make position-based research more complete by adding aircraft for which receivers hear Mode S transmissions but obtain no usable ADS-B position. MLAT increased the number of aircraft addresses seen in Switzerland by 12.1% over the study period (Section 5.7). In a separate global count for 20 July 2026, 40,547 Mode S addresses had at least ten conventional Mode S records, a DF11 reply and more than 30 minutes between their first and last observations. Of these, 3,877 (9.6%) had no decoded DF17 ADS-B position. Positions from DF18 messages (non-transponder ADS-B, TIS-B and ADS-R) do not count as ADS-B positions here. This gives an upper limit on the observed address population that could gain positions through MLAT. It does not establish continuous tracking or sufficient receiver coverage to locate every address.

6.5 Computational cost

The computational requirements are modest at the observed load. We operated the MLAT server on a Hetzner CX33 virtual machine at a hosting cost below €5 per month. In the 21 July snapshot, the server container used 23.45% of one CPU core and 150.2 MiB of memory while reporting 4,147 messages per second from 26 receivers. This suggests that central MLAT processing is affordable for similar regional research deployments, although peak demand and the resources needed for larger receiver networks remain to be measured.

6.6 Future research directions

The source-code analysis points to four directions for further work: 1) The *mlat-server* currently favours larger, more recent groups of receivers. We would investigate whether including geometry

in this choice improves positions without losing too much coverage. 2) Re-solving the same observations with geometric or weather-corrected altitude would help establish how much pressure-altitude assistance contributes to horizontal disagreement. 3) On the computational side, `leastsq` is called without explicit derivatives, suggesting an analytical derivative implementation as a candidate optimisation. 4) Schäfer and Jonáš find that ADS-B position errors lie mostly along track and recommend using only reports from aircraft with a low range rate relative to the receiver: slow aircraft, or aircraft moving nearly perpendicular to the line of sight at sufficient distance [12]. The server could apply the same rule when it selects the sync pairs for clock synchronisation.

7. Conclusion

In this paper, we evaluated the open-source `mlat-client/server` stack in Switzerland, combining basic source-code analysis with 27.8 million accepted position estimates collected over 62 days. The D1 position estimates with an ADS-B reference have a median horizontal disagreement of 106.0 m and a 95th percentile of 654.5 m. At low altitude, large disagreements are more frequent, while a smaller proportion of ADS-B position reports have a corresponding MLAT position estimate.

The modest computational load in the recorded server snapshot encourages further use and evaluation of this approach. OpenSky has since moved beyond the research setup examined in this paper. It made MLAT available to all interested feeders globally², and the output of its `mlat-server` is available in its Trino database interface.

Author statements and reproducibility

We thank the OpenSky receiver community and the open-source MLAT contributors. No funding or conflicts of interest are reported.

The reproduction package available at <https://github.com/openskynetwork/mlat-paper> contains the D1 to D3 analysis code, figures, plotted data, numerical claim inventory and hash manifests. All D1 to D3 analyses use the full 62-day data release, subject to their stated matching and selection conditions. The data is available at Zenodo at <https://zenodo.org/records/23675290>.

The authors are responsible for study design, validation and conclusions. Specifically, M.St. was responsible for draft, analysis, revision and supervision. M.G. was responsible for setting up the operational `mlat-client/server` test environment and early data analysis. M.Sc. was responsible for supervision, revision and the analysis of the system architecture. M.St. and M.G. were responsible for deployment of the `mlat-clients`. OpenAI Codex assisted with source inspection, analysis code, and figures. Scientific Agent Skills guided visualisation [27].

²<https://opensky-network.org/feed/mlat>

References

- [1] Matthias Schäfer, Martin Strohmeier, Vincent Lenders, Ivan Martinovic, and Matthias Wilhelm. “Bringing Up OpenSky: A Large-Scale ADS-B Sensor Network for Research”. In: *Proceedings of the 13th International Symposium on Information Processing in Sensor Networks*. IEEE, 2014, pp. 83–94. DOI: 10.1109/IPSIN.2014.6846743.
- [2] International Civil Aviation Organization. *Guidance Material on Surveillance Technology Comparison*. Tech. rep. ICAO Asia and Pacific Office, 2007. URL: https://www.icao.int/sites/default/files/APAC/Documents/edocs/CNS/gmst_technology.pdf (visited on 07/22/2026).
- [3] Federal Aviation Administration. *ADS-B Wide Area Multilateration (WAM)*. 2022. URL: https://www.faa.gov/air_traffic/technology/adsb/atc/wam (visited on 07/22/2026).
- [4] Regina Kaune, Christian Steffes, Sven Rau, Wolfgang Konle, and Jürgen Pagel. “Wide Area Multilateration Using ADS-B Transponder Signals”. In: *15th International Conference on Information Fusion*. 2012, pp. 727–734.
- [5] Junzi Sun, Luis Basora, Xavier Olive, Martin Strohmeier, Matthias Schäfer, Ivan Martinovic, and Vincent Lenders. “OpenSky Report 2022: Evaluating aviation emissions using crowd-sourced open flight data”. In: *2022 IEEE/ALAA 41st Digital Avionics Systems Conference (DASC)*. IEEE, 2022, pp. 1–8.
- [6] Abdulmalik Shehu Yaro. “Aircraft Position Estimation Comparison of Multilateration System Lateralization Algorithms with Different Reference Selection Techniques”. In: *ELEKTRIKA – Journal of Electrical Engineering* 18.1 (2019), pp. 16–21. DOI: 10.11113/elektrika.v18n1.133.
- [7] Benoit Figuet, Raphael Monstein, and Michael Felux. “Combined Multilateration with Machine Learning for Enhanced Aircraft Localization”. In: *Proceedings* 59.1 (2020), p. 2. DOI: 10.3390/proceedings2020059002.
- [8] Clemens Sonnleitner and Thomas Hobiger. “Wide-Area Multilateration Airspace Surveillance with Unsynchronized Low-Cost ADS-B Receivers Using Time Difference of Arrival Observations”. In: *NAVIGATION: Journal of the Institute of Navigation* 72.3 (2025). DOI: 10.33012/navigation.704.
- [9] Rafał Osypiuk and Filip Surma. “Uncertainty of Aircraft Localization with Multilateration and Known Altitude”. In: *Electronics* 14.12 (2025), p. 2420. DOI: 10.3390/electronics14122420.
- [10] Lukas Baege, Patrick Schaller, Vincent Lenders, and Martin Strohmeier. “Scaling the Timing-Based Detection of Anomalies in Real-World Aircraft Trajectories”. In: *Journal of Open Aviation Science* 1.2 (2023). DOI: 10.59490/joas.2023.7205.
- [11] Xavier Olive, Jan Krummen, Benoit Figuet, and Richard Alligier. “Filtering Techniques for ADS-B Trajectory Preprocessing”. In: *Journal of Open Aviation Science* 2 (2024). DOI: 10.59490/joas.2024.7882.
- [12] Matthias Schäfer and Petr Jonáš. “ADS-B Positional Accuracy and Anomalies: A Comprehensive Analysis Using High-Resolution MLAT Data”. In: *2025 Integrated Communications, Navigation and Surveillance Conference (ICNS)*. IEEE, 2025. DOI: 10.1109/ICNS65417.2025.10976935.
- [13] Oliver Jowett. *Comment on MLAT software authorship and deployment*. Post 357, 16 September 2015, under the username obj. 2015. URL: <https://forum.flightradar24.com/forum/radar-forums/flightradar24-feeding-data-to-flightradar24/9031-archived-beta-test-mlat-software-for-raspberry-pi/page24> (visited on 10/19/2026).
- [14] FlightAware. *ADS-B Network News: September 2025*. Published 30 September 2025. 2025. URL: <https://blog.flightaware.com/ads-b-network-news-september-2025> (visited on 10/19/2026).
- [15] FlightAware Developers. *PiAware changelog*. Historical entries for May 2015. URL: https://github.com/flightaware/piaware_builder/blob/master/changelog (visited on 10/19/2026).
- [16] ADS-B Exchange. *Feed client installation and update script*. Selects the ADS-B Exchange mlat-client repository. URL: <https://github.com/ADSbexchange/feedclient/blob/master/update.sh> (visited on 10/19/2026).

- [17] adsb.fi. *adsb.fi feeder installation and update script*. Selects the adsbfi/mlat-client-adsbfi repository. URL: <https://github.com/adsbfi/adsb-fi-scripts/blob/master/update.sh> (visited on 10/19/2026). 418 419 420
- [18] adsb.lol. *Public MLAT server deployment configuration*. Container and command configuration. URL: <https://github.com/adsblol/infra/blob/main/bases/mlat-server/base/mlat-server.yaml> (visited on 10/19/2026). 421 422 423
- [19] Flightradar24. *Multilateration (MLAT)*. URL: <https://www.flightradar24.com/how-it-works/mlat> (visited on 10/19/2026). 424 425
- [20] Fabian Stettler and contributors. *mlat-server-wiedehopf-wipro*. Deployed base revision 7e85f5bcab9a844d0d08f3a02a5adb58226ca626; accessed 21 July 2026. Dec. 19, 2025. URL: <https://github.com/fabian-stettler/mlat-server-wiedehopf-wipro>. 426 427 428
- [21] Oliver Jowett and contributors. *mlat-client*. Version 0.2.13. Revision fe70767be859100176983b948140046b6ecdd34a; accessed 21 July 2026. Oct. 23, 2023. URL: <https://github.com/mutability/mlat-client>. 429 430 431
- [22] Oliver Jowett and contributors. *mlat-server*. Revision bee59806af1fda9728489279a96f0f5ac56abfe7; accessed 21 July 2026. Feb. 17, 2016. URL: <https://github.com/mutability/mlat-server>. 432 433
- [23] Martin Strohmeier, Mauro Leonardi, Sergei Markochev, Fabio Ricciato, Matthias Schäfer, and Vincent Lenders. “In Pursuit of Aviation Cybersecurity: Experiences and Lessons From a Competitive Approach”. In: *IEEE Security & Privacy* 21.4 (2023), pp. 61–73. 434 435 436
- [24] Manuel Eichelberger, Kevin Luchsinger, Simon Tanner, and Roger Wattenhofer. “Indoor localization with aircraft signals”. In: *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems*. 2017, pp. 1–14. 437 438 439
- [25] Charles F. F. Karney. “Algorithms for Geodesics”. In: *Journal of Geodesy* 87.1 (2013), pp. 43–55. DOI: 10.1007/s00190-012-0578-z. 440 441
- [26] Xavier Olive. “traffic, a Toolbox for Processing and Analysing Air Traffic Data”. In: *Journal of Open Source Software* 4.39 (2019), p. 1518. DOI: 10.21105/joss.01518. 442 443
- [27] Timothy Kassis, Vinayak Agarwal, Yuhuan He, Darshil Patel, and Aubrey M. Brueckner. *Scientific Agent Skills: A Library of Procedural Knowledge for Research Agents*. 2026. DOI: 10.48550/arXiv.2609.00065. arXiv: 2609.00065. 444 445 446

Appendix A. Common-period research datasets

All results have been computed from D1 to D3. The separate reproduction package records the analysis code, derived tables, figures and the source of each updated numerical claim.

Record structure and extraction. We retain the original D1 JSON records alongside a Parquet representation for analysis. D2 also retains the source partition hour. For D3, the timestamp is the earliest reported reception time. Altitude-type and surface-status flags and the source partition hour are also retained. The compact D3 export omits per-receiver reception lists and other fields unused by our analyses, while preserving every selected row and its retained values.

We merge overlapping reference intervals for each aircraft so that the D3 selection does not duplicate records. A position-table row may combine receptions from several sensors and is not necessarily a unique transmission. Any deduplication remains an explicit analysis step.

Table A.1 lists whole hours for which completed source queries return no D3 records. These are gaps in the selected ADS-B references, rather than unfinished downloads. For example, on 25 May from 21:00 to 22:00 UTC, D1 and D2 contain MLAT observations but D3 contains no selected reference records.

Table A.1. Whole hours with no D3 records found in completed source-count queries. D1 and D2 counts show the MLAT observations available during the same intervals. All 1,488 hours have been counted.

Interval (UTC, end exclusive)	D1 rows	D2 rows	D3 rows
25 May 21:00–25 May 22:00	9 250	9 224	0
06 Jun 19:00–06 Jun 20:00	20 100	20 065	0

Reproducibility and sharing. The artefact contains the data, dataset definitions, extraction code, SQL queries, source schemas, query identifiers, daily and hourly counts, and file checksums.

Appendix B. Sensitivity to a constant timestamp correction

We tested whether a further change to the recorded time improves agreement with ADS-B. We chose one constant correction per dataset to minimise median disagreement over 20 May to 19 June 2026, then applied it without further adjustment to 20 June to 20 July. Coordinates and ADS-B matching rules remained unchanged. During fitting, we used every position with an eligible reference at every tested correction, retaining 40.8% of the originally matched D1 positions and 41.5% of D2. This comparison therefore favours positions with consistently available references.

Table B.1. Timestamp corrections evaluated on 20 June to 20 July 2026. Each dataset uses the same positions across its rows. D1 and D2 have separate comparison populations. Negative adjustments move the comparison time earlier. The 0.915 s correction is the previously measured median timestamp difference between D1 and D2.

Dataset	Adjustment (s)	Positions	Median (m)	P95 (m)
D1: mlat-server	0	10,067,612	108.4	655.3
D1: mlat-server	−0.180	10,067,612	104.7	640.7
D2: Trino	0	9,330,947	240.1	764.5
D2: Trino	−0.915	9,330,947	103.5	640.8
D2: Trino	−1.050	9,330,947	100.8	630.5

The fitted corrections give a modest further improvement (Table B.1). Both the median and 95th percentile improve on all 31 validation days, including the comparison with the previous D2 correction. The minimum is fairly flat: tested corrections from 0.11 to 0.25 s for D1 and from 0.99 to 1.10 s for D2 are within 1 m of their respective fitting minima. These ranges are not confidence intervals.

There is also a cost in reference availability. Across all validation positions, the fitted corrections yield 4.8% fewer eligible matches for D1 than its original timestamps, and 4.3% fewer for D2 than the 0.915 s correction. The fixed-population comparison in the table accounts for this change in selection. These results measure relative alignment with ADS-B rather than an absolute clock delay. The main analyses retain the original D1 timestamps.