Low-cost angle of arrival-based auxiliary navigation system for UAV using signals of opportunity

Adrian Winter
Institute for Engineering Cybernetics
NTNU
Trondheim, Norway
adrian.winter@ntnu.no

Nadezda Sokolova
Sustainable Communication Technologies
Sintef Digital
Trondheim, Norway
nadia.sokolova@sintef.no

Aiden Morrison
Sustainable Communication Technologies
Sintef Digital
Trondheim, Norway
aiden.morrison@sintef.no

Oliver Hasler
Institute for Engineering Cybernetics
NTNU
Trondheim, Norway
oliver.hasler@ntnu.no

Tor Arne Johansen
Institute for Engineering Cybernetics
NTNU
Trondheim, Norway
tor.arne.johansen@ntnu.no

***Abstract*—Many modern devices rely on satellite navigation for reliable, precise and ubiquitous localization. While this enables technological advances that have been unthinkable a few decades ago, it also has clear implications on safety and resilience. There have been many real-world incidents where through deliberate or accidental interference, as well as simply by design, satellite navigation has been shown to be unreliable. Consequences of encountering interference can range from anywhere between minor inconveniences to catastrophic failure with loss of life. For unmanned aerial vehicles (UAVs), the consequences will often be significant economical damages or injury, but the risk is not high enough to justify expensive backup navigation systems that are rarely used.**

This paper investigates a sensor build on low-cost commercial off-the-shelf (COTS) hardware that can be used in costal areas if localized global navigation satellite systems (GNSS) interference occurs. This sensor works by measuring the incoming angles of radio signals sent by maritime vessels as part of the Automatic Identification System (AIS) and triangulating the receiver's position.

The measurement precision of the hardware combination is evaluated from experiments with real-world signals. It is then shown in simulations that this achievable performance can lead to a significant improvement of the vehicle's state estimation compared to pure dead reckoning in GNSS-outage scenarios.

***Index Terms*—Navigation, Signals of Opportunity, UAV**

I. INTRODUCTION

It is well known that many of today's technological achievements are only possible due to the ubiquitous availability of low-cost, high-accuracy positioning, navigation and timing (PNT) services. The applications go far beyond simple navigation aids for vehicles and include safety critical use cases such as autonomous flight.

The default solution to these challenges is to use Global Navigation Satellite Systems (GNSS), namely the United States' GPS, the European Galileo, Chinese Beidou and Russian Glonass constellations. These navigation systems are extremely useful and offer multiple advantages such as low-cost receivers, high precision, world-wide absolute positioning and timekeeping using purely passive reception without the need for external information or two-way communication. The latter implies that low power receivers are possible and that no subscription fees must be paid for the base service.

However, the drawbacks of GNSS are also well known. They include degraded performance to the point of inoperability indoors and in certain environments such as deep urban canyons, as well as susceptibility to accidental or deliberate attacks, such as jamming or spoofing. Many different techniques to counter these challenges have been developed and deployed. These countermeasures are not perfect, however, and offer improvements only in certain circumstances. A severe limitation of alternatives to GNSS is their often prohibitively high cost, as many approaches require specialized and expensive hardware.

One such approach is to leverage dissimilar alternative localization techniques, such as optical navigation, odometry, or radio navigation. These techniques, too, have their respective drawbacks. For example, visual navigation or LIDARs do not perform well in self-similar environments such as in tunnels or over the sea. Local radio navigation has multiple advantageous properties compared to satellite navigation. Most notably, the power level at the receiver antenna is typically many orders of magnitudes higher, making jamming and spoofing attacks infeasible or at least easy to detect. However, local radio navigation usually requires dedicated infrastructure such as VORs, UWB base stations, or others, and is not commonly available for most users [1], [2]. Many previously used radio navigation systems such as LORAN-C have also been phased out or are being scaled down considerably.

One approach to leverage the advantages of terrestrial radio signals without relying on a dedicated network is to use so-called Signals of Opportunity (SoOp), also called opportunistic navigation. The idea for SoOp-based navigation is to use radio

The research has been funded by the Research Council of Norway through IKTPLUSS grant 288634 (Sintef Digital) and SFF AMOS grant 223254.

signals that are available without the user's intervention, but are *not* optimized or intended for navigation by the operators. Examples include TV or radio networks, cellular networks or decentralized sources such as WiFi.

Within this paper, a SoOp-based system is investigated. The exploited signals are the periodic position reports sent by many small boats and almost all large ships[1] as part of AIS [3]. The sensor measures the incoming bearing of these signals and triangulates the sensor's position. A more detailed discussion about underlying assumptions will follow.

Every section of this paper is divided into two subsections: first, the physical sensor system is described and evaluated, and afterwards the integration of the sensor into an actual airborne platform is discussed. The first part of the paper builds on measurements taken from real transmissions in the field, whereas the second part employs software simulations.

A. Signals of Opportunity-based Navigation

Signals of Opportunity-based navigation is an old technique whose roots can be traced back to the beginning of radio communications: if one considers the use of direction finding of non-directional beacons (NDBs) as signals of opportunity, then this has been used for aircraft navigation since the 1930's. While NDBs are dedicated infrastructure and thus do not fulfill the definition of SoOps, direction finding can also use other transmitters such as medium wave (AM) radio stations.

Due to this legacy, there has also been significant research in this field. To the author's knowledge, there are four commonly used basic techniques for stand-alone (i.e. without using base stations or other reference receivers) navigation. A short overview will be given here:

1) Angle of Arrival (AoA): The idea of Angle of Arrival-based localization is to measure the angle of incoming signals using either a physically rotating directional antenna, or an antenna array. This technique has few requirements regarding the source signal but does require special hardware. Modern high-throughput communication systems such as 5G often employ antenna arrays for their multiple input multiple output (MIMO) transfers, making this technique interesting for those devices. Disadvantages of this technique are its susceptibility to multi-path reflections [4], as well as the increasing inaccuracies over longer distances. This approach is also commonly called "Direction of Arrival" (DoA).

To use this technique, it is necessary to know or estimate the position of the transmitter and to be able to identify the transmitter. That is, if multiple transmitters are active, these must be distinguishable.

2) Received Signal Strength Indication (RSSI): For this technique the reception strength of incoming signals is measured and compared to either a model of the physical area or to a database of pre-determined measurements that were taken from a known location. Since most receivers are able to provide this measurement, there are very low hardware requirements. There are also very few requirements regarding the signal itself, making this technique adequate for a large group of signals. In particular, a synchronization between base stations is not necessary, enabling the ubiquitous WiFi routers in residential areas to be used.

For this method to provide accurate results, either a high-quality database of high-resolution measurements must be accessible, or a very good model of radio frequency (RF) propagation in that area must be built. The location of the transmitters must be known or estimated from measurements. Any deviations from the model or changes to the environment will likely result in degraded measurements [5].

3) Time Difference of Arrival (TDoA): This technique requires that a signal is sent simultaneously (or with a known and constant delay) from two different locations. The time offset between reception of the signals can be measured, and after multiplication with the speed of light yields a distance differential. This differential is a measurement how much nearer/farther the first transmitter is compared to the second one. On a two-dimensional surface, this is represented by a hyperbola.

Since the transmitters must be synchronized, it is only possible with some source signals. Many large-scale networks, including digital radio [6], [7] and TV as well as telecommunication networks such as the LTE mobile phone networks, are synchronized and can be used for TDoA navigation. The receiver must be able to measure the time differential, but the transmitter and receiver clocks need not be synchronized. This implies relatively low hardware requirements for the receiver, but high requirements for the transmitter network.

4) Time of Arrival (ToA): Conceptually, this technique is very simple: the time of reception of a signal with known time of transmission is measured, and the difference between transmission and reception, again multiplied by the speed of light, yields the distance to the transmitter. By repeating this with multiple transmitters, the receiver position can be uniquely calculated.

Unfortunately, this has very high requirements regarding the signal itself, as the signal must be sent at a precise time and with a structure that makes precise measurement of the transmission time possible. The receiver must either have a very high quality local clock (in practise this will likely be an atomic clock) or must co-estimate the time with the position. The latter approach is the one used by GNSS.

Due to these requirements, only very few source signals are possible for ToA localization. Nonetheless, significant scientific work has been done using the technique, for example using LTE cellular networks [8].

5) Other Methods: Since there is no universally agreed on definition of Signals of Opportunity, there are other methods that can be seen as SoOp-based navigation. For example, if base stations can be used, then analogue signals such as FM radio stations with poor frequency and phase stability become usable, and if user-side transmission is allowed then probing WiFi routers using round-trip time measurements [9] allows significantly better results in dense urban environments.

---

[1] Within this paper, the terms "boat", "ship", "vessel" and "beacon" are used interchangeably and all refer to a transmitter that sends valid AIS position reports. This includes ships, but also buoys or costal base stations

B. Related work

The most closely related work has been published by Baine and Gross [10]. Their work has a very similar premise to this work, however it is mainly focused on the algorithm as opposed to implementation. Additionally, instead of AIS they used the analogous Automatic Dependent Surveillance-Broadcast (ADS-B) transponder system for aircraft as data source. ADS-B is conceptually similar to AIS, but has somewhat different challenges to solve, such as the imprecision caused by latency: A latency of 1 second between GPS fix and position report broadcast will cause errors of less than approximately $20\,\mathrm{m}$ in boats, but up to approximately $250\,\mathrm{m}$ for subsonic airliners at cruise speed. They later expanded this work to three dimensions [11], still verifying the algorithm in simulations only.

Chen et al. investigated the inverse problem: They localized airplanes using multiple synchronized ground based antenna arrays [12]. This is useful for spoofing protection, but of limited use to the aircraft itself. They also required GPS-disciplined rubidium oscillators, and while they describe it as "relatively low cost", this approach will likely be prohibitively expensive for a backup system for low-cost UAVs.

II. METHODOLOGY

The navigation sensor employs Angle of Arrival (AoA) measurements from multiple, geometrically spread sources as measurements to a Kalman filter based navigation state estimator. Triangulation with AoA measurements only yields 2D information, unless elevation measurements are also available. The four-element linear antenna array used for this project does not allow elevation measurements, it is instead assumed that a barometric altimeter is available. Due to the low price of barometric altimeters and their common availability, this condition should be fulfilled by most UAVs. Due to the short and local flights of UAVs, the potential change in ambient air pressure is small, making the measurements of a barometric altimeter precise and exhibiting limited and slowly varying drift after initialization.

In addition it is assumed that a magnetic compass is available, which is also a reasonable assumption for most UAVs, although accuracy may be limited by calibration and at high latitudes. As described more thoroughly in the previous paper on this sensor [13], the system also yields heading information if enough beacons are in range. However, since magnetic compasses are so wide-spread and any dissimilar measurement increases robustness, it is assumed that one is available.

Further it is assumed that the estimators have converged to the correct state before start of flight. This is a reasonable assumption to make because any flight will have a known start position and an initialization phase. The navigation system is also not expected to be used as a *replacement* for GNSS but rather as a backup if an outage occurs *during* a flight. The estimators will have converged during initialization and under GNSS-aided initial flight.

A. KerberosSDR measurements

KerberosSDR is a semi-commercial development platform and essentially consists of four RTL-SDR low-cost receivers tied to a common oscillator. It has since been replaced by the five-channel *KrakenSDR*, which can be bought for about 400 €.

AIS signals were chosen because they have multiple desirable characteristics:

- Frequency of $162\,\mathrm{MHz}$ is low enough for RTL-SDR and high enough to have manageable antenna array sizes
- High transmitter power ($5\,\mathrm{W}$ to $25\,\mathrm{W}$)
- Source position known, because it is part of the unencrypted transmission payload
- High geometric diversity and multiple signal sources in the relevant environment
- Only one transmitter active on each channel due to Time Division Multiple Access (TDMA) scheme[2]
- Low signal bandwidth

RTL-SDRs are low-cost software defined radio (SDR) receivers with a relatively low sample rate of about 2 MS/s, a maximum carrier frequency of about 2 GHz and a sample resolution of 8 bit, built around the TV receiver chips RTL2832U (from which they derive their name) and R820T2. Sample rate and frequency bands have clear implications on which signals can be received, and exclude the license-free ISM band on $2.4\,\mathrm{GHz}$ as well as higher-bandwidth transmission such as mobile phone networks. The receivers are perfectly adequate to receive both AIS channels simultaneously and with high reception quality.

AIS is not the only potential source signal fulfilling these criteria, the system could be adapted to other sources such as ADS-B as well.

From the University's premises, there is a clear view of Trondheimsfjorden, a bay of the Norwegian Sea. The fjord is enclosed on all sides by mountains of $300\,\mathrm{m}$ to $600\,\mathrm{m}$ altitude that obstruct view of the open ocean. The area visible from the University is approximately 15 by 15 km in size[3] and includes a signal-rich harbor.

In summer 2022 a dataset was recorded from the roof of the University for publication at a conference. Unfortunately at this point in time it was not possible to investigate the fidelity of the measurements in detail due to unresolved software issues. For that reason, this previous paper only included very heavily selected data that did show that the repetition precision to the same source was high, but no statement about the system performance could be made [13].

These software issues have since been resolved and the results yielded from the same data set are described in the next section.

---

[2] The Self Organizing Time Division Multiple Access (SOTDMA) and Carrier Sense Time Division Multiple Access (CSTDMA) scheme should ensure that collisions do not occur or if they do that a transmitter backs up. However in certain circumstances, such as caused by hidden nodes, multiple transmitters could be active simultaneously.

[3] The fjord is much larger, but the part that is both visible and has significant maritime traffic is relatively small

B. Performance Verification Simulations

To investigate the feasibility of using such a sensor system, the approximate minimal required measurement accuracy and precision need to be investigated. Obviously, this heavily depends on the requirements of the mission, and therefore no pass/fail threshold can be defined without also defining a platform and a mission. It is instead investigated how the results depend on the measurements, allowing an informed decision on whether the sensor may be appropriate for each individual application.

To this end, an in-house navigation toolbox, written in MATLAB, is used with simulated data. This allows comparison between estimation and truth as well as artificially changing measurement fidelity to investigate the influence of measurement errors on estimation errors.

This in-house toolbox is a modular estimation tool, which implements an Error State Kalman Filter (ESKF), also called Multiplicative Extended Kalman Filter (MEKF), and several sensor models. The sensor data for this project is generated from flight simulations. ArduPilot's integrated flight simulator was chosen, because it is easy to use and offers sufficient fidelity for the intended use case. To simplify all further steps, a flat earth Cartesian North-East-Down (NED) coordinate system is used. Given the range of the sensor and the UAV of a few kilometers, the errors introduced by the linearisation are negligible.

The flight simulator output is a noise-free ground-truth data file which is subsequently loaded into MATLAB. From the noise-free state a sensor model is derived, where sensor imperfections can be modelled. Currently, additive Gaussian noise with a configurable standard deviation is used.

Simulated bearing measurements are generated from a configurable number of beacons, which can be placed freely around the initial starting position of the UAV. Each beacon sends a signal after a certain configurable time, similar to transmission frequencies used in AIS. The beacons are currently simulated to be stationary and their position reports are noise-free, but this can easily be changed in the code. Noise in the position reports has a linear influence on the solution, and since typical GNSS errors of a few meters is smaller than the other errors, it is not expected that this will cause significant deterioration of the solution. The physical geometry of the ships could influence the solution significantly – refer to Section V-B for further discussion.

Using simple planar geometry, the measurement is calculated and corrupted by noise:

$$\beta_t^i = \arctan_2(\Delta E^i, \Delta N^i) + \Psi_t + \mathcal{N}(\mu, \sigma^2) \tag{1}$$

where β_t^i is the body-fixed bearing to the i^{th} boat at time step t. ΔE and ΔN are the distances between boat and sensor in north and east direction, respectively. Ψ_t is the current heading of the vehicle, $\mathcal{N}(\mu, \sigma^2)$ represents the added noise in bearing measurements, modelled as Gaussian noise with the mean μ and standard deviation σ.

This also yields the measurement function $\hat{y}_t^i = \beta_t^i = h(\vec{x})$, which is the base for calculating the measurement Jacobian $H(\vec{x})$ that is required for the ESKF. Since the Jacobian is dependent on the state parametrisation $\vec{x}$ and thus implementation-dependent, it will not be repeated here.

C. Conditioning of angular measurements

It is important to ensure that the non-linear and periodic angular measurements are introduced into the ESKF correctly. In general, it is not valid to treat angular measurements in the same way as linear cartesian coordinates. For example, it is not valid to calculate the arithmetic mean of angles – however if it can be ensured that the angular differences are small and in the same or adjacent quadrants, then the measurements can be treated as any other linear measurement.

This significantly simplifies the entire filtering algorithm, because no special handling of any filtering steps is necessary if angular measurements are available. The assumption of small errors should always be valid for a correctly working Kalman filter, barring some special cases such as a close pass of the UAV to the beacon. Since such edge cases are solvable and have no impact on the fundamental findings, they are currently ignored and the linearized approach is chosen.

To ensure that two angular measurements are in adjacent quadrants, a conditioning algorithm [14] is executed before any angular differences, in particular the innovation step of the Kalman filter, are calculated. The difference δ between two angles α and β is calculated using the formula

$$\delta = ((\alpha - \beta + \pi) \mod 2\pi) - \pi \tag{2}$$

and the angle β is modified to β', which is then used instead of β in all subsequent steps:

$$\beta' = \alpha - \delta. \tag{3}$$

III. RESULTS

A. KerberosSDR

The dataset was recorded from the Institute's roof in Trondheim, Norway, on the 16th of June 2022. The recording consists of twelve minutes of data, and includes ca. 2300 AIS transmissions. After rejecting the transmissions that were physically obstructed, far away or with otherwise poor numerical confidence, 655 acceptable transmissions were left. This results in approximately one usable transmission per second.

The true azimuth angle (i.e. bearing with respect to north) between antenna and transmitter is known, because both the location of the receiver and transmitter are known during the experiments. The receiver position can be measured using a GNSS receiver or a map, and the transmitter position is broadcast by the vessel. This assumes that the transmitter noise, caused by GNSS measurement errors or other effects, is not dominating. This is a reasonable assumption, because the *minimal* distance between transmitter and receiver is $1.8\,\mathrm{km}$, and for the measurement error to reach 1° – which is the

TABLE I
RESULTS OF AOA MEASUREMENTS

Description	σ [°]	Max error [°]
Unfiltered Set	20	107
Filtered Set	9.5	70
Mild Pruning (4% removed)	4.9	17
Heavy Pruning (33% removed)	2.1	4.5

measurement's resolution – the GNSS measurement error would need to be at least $\arctan(1°) \cdot 1800\,\mathrm{m} = 80\,\mathrm{m}$. This is an order of magnitude higher than what can reasonably be expected from a GNSS measurement.

Since the absolute orientation with respect to north of the antenna array is not known, as there is no magnetometer used in the setup, the measurements are adjusted such that the errors become zero-mean. Essentially, this is estimating a constant bias that is then removed. The difference between these corrected measurements and the true azimuth is the measurement error.

Table I lists the results of these measurements. The standard deviation is significantly influenced by a few bad measurements, as can be seen when the worst measurements are removed as outliers. Two different thresholds for rejection are shown, a mild pruning at the 2^{nd} and 98^{th} percentile (i.e. removing 4% of all measurements), as well as more aggressive pruning of the top and bottom sixth (i.e. the 16^{th} and 83^{rd} percentile), resulting in the removal of a third of all measurements. Purely for comparison, the results for the entire dataset without any rejected measurements are also included.

Figure 1 shows the errors as histogram, together with the borders of the pruning percentiles.

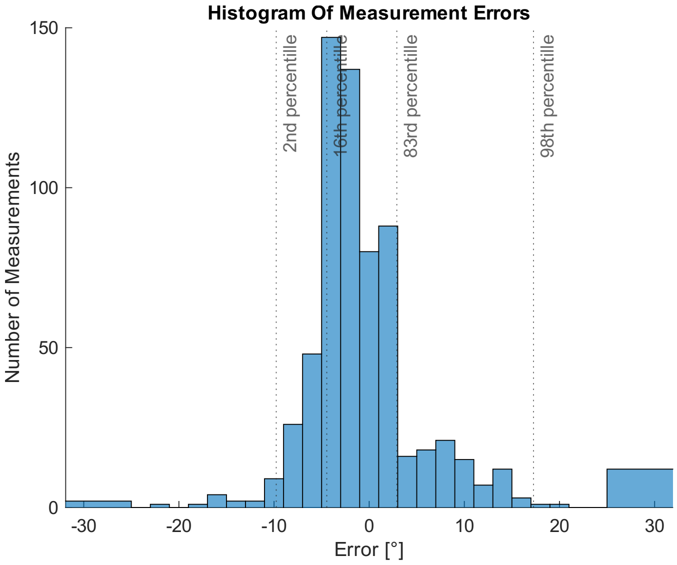

Fig. 1. Histogram of Measurement Errors

B. Performance Verification Simulations

As described in the previous section, simulations were performed to evaluate whether or not the expected accuracy and precision can realistically improve the state estimate compared to dead reckoning. Multiple simulations were performed with differing measurement noise in the AoA-measurements but otherwise identical conditions. For comparison, simulations without position aiding, that is dead reckoning with IMU, altimeter and compass measurements, are included. The dead reckoning simulations too were run with the same tuning parameters. The dead reckoning delibertely was not tuned perfectly, as the intention of this comparison is to compare *imperfect* dead reckoning to the aided solution.

Figure 2 shows an overview of the scenario including the location of the beacons. Figure 3 shows a short section of the flight path, where the impact of degraded measurement performance can be seen clearly.

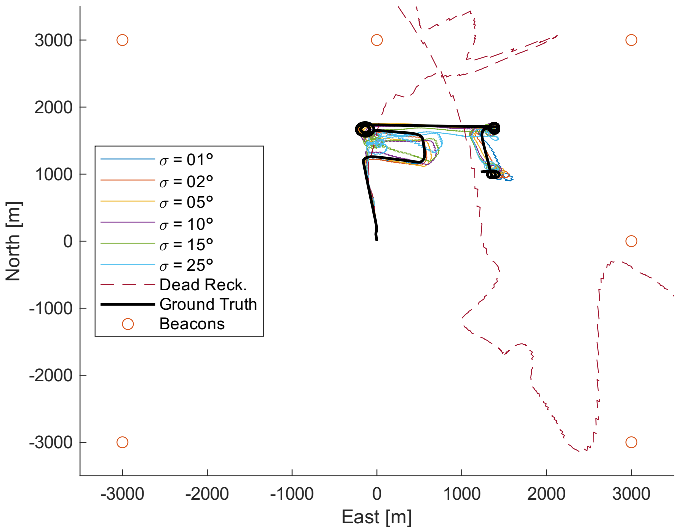

Fig. 2. Overview of the flight scenario

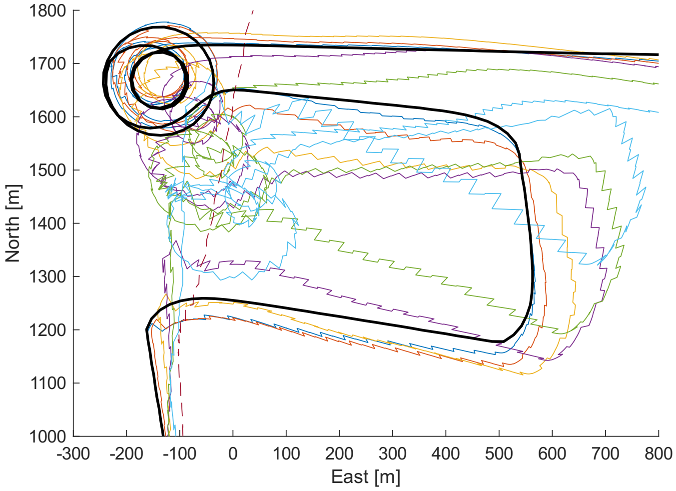

Fig. 3. Detail of the flight scenario

Additionally, different metrics were calculated (see Table II). These metrics are the Root Mean Square Error (RMSE) for the entire flight path, the maximum error during the flight, as well as the time until the error exceeds the threshold of 50 m and 200 m.

TABLE II
STATISTICS OF FLIGHT SIMULATIONS

σ	**RMSE**	**Max. Err.**	**Err $>$ 50 m**	**Err >200 m**
1 °	48 m	175 m	65 s	N/A
2 °	47 m	144 m	61 s	N/A
5 °	70 m	143 m	18 s	N/A
10 °	107 m	219 m	16 s	133 s
15 °	139 m	272 m	16 s	26 s
25 °	201 m	476 m	16 s	25 s
Dead R.	3318 m	5358 m	16 s	25 s

IV. DISCUSSION

A. KerberosSDR

The measurements described in Section III-A were taken in a complex environment: the University is about 2 km from the coastline within a city. Therefore, a significant part of the signals were obstructed by buildings, likely creating a complex multi-path environment. There is also surrounding mountains and hills, which too will cause some reflections of any signal. Therefore it can be assumed that the measurement fidelity is certainly not a best case scenario.

For practical reasons, the measurements are taken from a known stationary position, allowing simple comparison with the ground truth. There are multiple potential error sources that could be introduced by dynamic antenna movements. Doppler shifts can be ruled out to cause noticeable problems, as the lateral and rotational velocities achievable for UAVs is slow compared to the speed of light. Yawing while a AIS message is transmitted will result in a non-constant AoA. However, since the AIS burst time is very short (26 ms), and each AIS message is handled individually from all others, the yaw angle can be seen as constant for each transmission. Noise, latency between position fix and AIS transmission and other errors stemming from the transmitter side will influence the UAV the same way as a stationary test setup. RF propagation effects such as multipath and fading will be different and time-varying in flight, but due to the fact that the current test setup with the receiver close to the ground and in proximity to buildings and other structures, it seems unlikely that the reception quality will be worse for a vehicle in free flight. However, it cannot be ruled out entirely.

The measurements were also filtered based on geometry and a performance metric. Filtering by geometry is necessary because the high signal powers and robust signal structure of the AIS standard allow reception of boats several dozens of km away, despite being obstructed by multiple mountain ranges. These measurements were rejected, as there is no reasonable expectation to receive realistic measurements in this case. This is also a simple form of spoofing rejection, if a transmitter sends a unreasonably far away position report, such as (0°N, 0°E) during a hardware test.

Additionally, the employed AoA algorithm, MUSIC [15], provides a confidence value. All measurements with a low confidence were rejected. This could happen for example due to low signal quality despite geometric closeness, or due to incorrect filtering and data association in the processing toolchain. The toolchain assumes that only one transmitter is active at any time on either AIS channel, but in certain circumstances this assumption is violated and can lead to unexpected behaviour. The current implementation of the tool chain is also not very sophisticated with respect to burst begin and end detection, so it can happen that two consecutive bursts are evaluated together, again leading to unexpected behaviour. Since such cases are relatively rare *in the Trondheimsfjord area* where there is only moderate maritime traffic, these disadvantages are accepted – but they can and should be addressed if the system is supposed to be used in busier areas.

After filtering the measurements in the described way, the system shows very promising performance. The majority of measurements are within a few degrees of the true value and the relatively poor standard deviation of 10° is dominated by a few outliers. Removing a small amount of outliers – something that can be done in practice as well, if updates to the Kalman filter that seem too unlikely are rejected – significantly improves the standard deviation, as the majority of errors is well below 10°.

B. Performance Verification Simulations

The goal of the simulations was to show that even with moderately precise measurements, a significant improvement of the state estimate can be achieved compared to a baseline without using these measurements (dead reckoning). The results presented in Section III-B show that this goal was indeed reached, and even a relatively poor AoA measurement provides significant improvement of the navigation solution.

It is very important to note that the numerical values must not be interpreted as absolute result, but should only interpreted relative to each other. This is because the assumptions made in the simulations are not necessarily realistic, for example noise values and distribution might be different.

Based on the simulations, it becomes immediately clear that dead reckoning using only a magnetic compass, an altimeter and the IMU is not sufficient. This is despite the fact that the simulations have no artificially added noise apart from unavoidable quantification noise in the IMU measurements. This does not come as a surprise, as it is well known that inertial navigation alone can only be used for certain, very specific applications and usually only when relying on expensive, high quality IMUs.

Clearly there is little potential to use this approach as a high-precision localization system. However, for many applications a low-precision backup to GNSS is still useful. Such applications include

- Integrity Monitoring
- Fault bridging, assuming short-term outages
- Flight out of GNSS-denied area
- Return to remote pilot or save flight termination zone
- Dissimilar redundancy for low-accuracy applications, e.g. in conjunction with visual navigation

The system also cannot be a backup system in case of large-scale GNSS outages or attacks, since it indirectly relies on satellite navigation systems: Although the AIS standard has some provisions for non-GNSS position measurements, all vehicles usually rely on GNSS as their data source. In case of a large-scale outage or a spoofing attack, the primary data source thus would no longer be correct. In practice, this is not necessarily a severe restriction, because most RFI incidents are short-term and localized [16].

Further investigation of influences such as the number of beacons, the distance between beacons and UAV, the AIS transmission frequency, the geometric distribution or other interference sources could be made here. This will yield a better understanding of dominating error sources and potential challenges for real flights, however both due to time constraints and due to limited applicability to the results to real world flights this has not been done yet.

V. OPEN PROBLEMS AND FURTHER OUTLOOK

To use this navigation sensor in-flight, some open problems still need to be addressed.

A. Bank angle and non-planar measurements

Thus far, it was assumed that all signals are received from within the antenna array plane. For low flight altitudes and a horizontal antenna array (i.e. no bank or pitch angle), this assumption holds true. However if the UAV passes close-by over a beacon or if significant bank or pitch angles are achieved, then the measurement must be corrected for this effect.

Tilting the antenna array with respect to the incoming wave front effectively shortens the array in the direction of the wave. In case of a linear array an underestimation of the angle of incidence will result, whereas a square array array will become rhombic. There are two potential approaches to this problem: either the current state of the UAV, which includes the pitch and bank angles, could be fed back to the DOA estimation algorithm. Using the (estimated) orientation with respect to the incoming wave front, the dimensions of the antenna array could be corrected. This would potentially require an iterative correction algorithm, as the measurement result feeds back into the correction.

Alternatively, at least for the linear antenna array, it would be possible to not compensate for the array's tilt angle during the DOA estimation steps, but rather consider all possible solutions for DOA and tilt angle combinations that result in the given measurement. However, this would introduce another degree of freedom into the estimation algorithm and thus increase uncertainty.

B. Beacon position estimation

During the simulations, it was assumed that the position of the transmitter is exactly known. In reality, there will be at least three error sources that make this assumption invalid:

1) Measurement noise of the vessel's GNSS receiver
2) Physical distance between measured position and AIS transmitter antenna
3) Latency between position fix and AIS broadcast

The first error source is straightforward. The vessel's measurement error will typically be less than approx. 10 m and therefore not dominant compared to other error sources. However, during a spoofing or jamming scenario, this measurement error could become large.

The second error is a systematic problem for larger vessels. The AIS standard provisions for the physical dimensions of boats by broadcasting data about the relative dimensions of the vessel measured from the GNSS antenna. This information is contained in AIS message type 5 and broadcast every 6 minutes [3]. However, the distance between GNSS receiver antenna and UHF transmitter antenna is not broadcast, as in normal use cases this information is irrelevant. For navigation purposes, however, this offset could become noticeable: Given that ships can be about 400 m in length, the measurement error could also become that large assuming a worst case scenario, that is GNSS antenna mounted at the stern and UHF antenna at the bow or vice-versa.

The third error can be relevant for fast ships that simultaneously have a long latency between position fix and transmission. The AIS data packet contains a timestamp with second-precision of the time where the positioning fix was created. It should be investigated if this timestamp can be used to monitor whether or not this latency is likely to be a problem in practice, and if it can be used to mitigate the influence of the latency.

By co-estimating the vessel's state, that is position, velocity, heading, as well as potentially the offset between the antennas, it might become possible to mitigate these three error sources. However, tracking the vessel will also introduce additional uncertainty and latency, so a thorough investigation of advantages and drawbacks needs to be done before this can be considered.

VI. CONCLUSION

Within this paper, it was shown that the proposed backup navigation sensor shows potential for the intended use case, such as short term interference bridging or flight to safe areas in case of GNSS interference events.

It was shown using simulations that bearing measurements to known beacons can significantly improve the vehicles state estimation compared to dead reckoning, even if the precision of the measurements is poor. It was also shown using practical experiments in a complex environment that a measurement precision that is much better than assumed for the simulations can be achieved.

VII. ACKNOWLEDGEMENTS

The author would like to thank Torleiv H. Bryne for his significant work with the initial development and as project lead of the navigation toolbox, as well as all other colleagues that contributed to the current state of this software.

The author would also like to thank his colleagues Praveen Jain, Joseph Garrett, Dennis Langer and Mika Okuhara for their expertise whenever a question during development arose.

REFERENCES

[1] M. L. Sollie, K. Gryte, T. H. Bryne, and T. A. Johansen, “Outdoor Navigation Using Bluetooth Angle-of-Arrival Measurements,” *IEEE Access*, vol. 10, pp. 88 012–88 033, 2022.

[2] K. Gryte, T. H. Bryne, S. M. Albrektsen, and T. A. Johansen, “Field Test Results of GNSS-denied Inertial Navigation aided by Phased-array Radio Systems for UAVs,” in *2019 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, 6 2019, pp. 1398–1406. [Online]. Available: https://ieeexplore.ieee.org/document/8798057/

[3] “Technical characteristics for an automatic identification system using time division multiple access in the VHF maritime mobile band (Recommendation ITU-R M.1371-4),” International Telecommunication Union, Tech. Rep., 2005. [Online]. Available: https://www.itu.int/rec/R-REC-M.1371/en

[4] O. Hasler, T. H. Bryne, and T. A. Johansen, “Robust Phased Array Navigation in Reflective Prone Areas,” in *submitted to: European Control Conference*, Bucharest, 2023.

[5] “Moved a Wi-Fi router? That could mess with an iPhone’s location — Macworld.” [Online]. Available: https://www.macworld.com/article/559641/wifi-iphone-location.html

[6] O. Lepistö, S. Tanner, and R. Wattenhofer, “DAB+ Positioning,” Ph.D. dissertation, ETH Zürich, Zürich, 4 2021.

[7] D. Palmer, “Position estimation using the Digital Audio Broadcast (DAB) signal,” Ph.D. dissertation, University of Nottingham, Nottingham, 8 2011. [Online]. Available: http://eprints.nottingham.ac.uk/12456/ http://eprints.nottingham.ac.uk/12456/1/539180.pdf

[8] K. Shamaei, J. Khalife, and Z. M. Kassas, “Performance Characterization of Positioning in LTE Systems,” in *Proceedings of the 29th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2016)*, 11 2016, pp. 2262–2270. [Online]. Available: https://www.ion.org/publications/abstract.cfm?articleID=14732

[9] “Wi-Fi location: ranging with RTT — Android Developers.” [Online]. Available: https://developer.android.com/guide/topics/connectivity/wifi-rtt

[10] N. A. Baine and R. A. Gross, “Algorithm for geodetic positioning based on angle-of-arrival of automatic dependent surveillance-broadcasts,” in *2018 IEEE/ION Position, Location and Navigation Symposium, PLANS 2018 - Proceedings*. Institute of Electrical and Electronics Engineers Inc., 6 2018, pp. 990–1001.

[11] ——, “Algorithm for Three-Dimensional Geodetic Positioning Based on Measurements of Automatic Dependent Surveillance-Broadcast Signals,” in *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. IEEE, 4 2020, pp. 1381–1392. [Online]. Available: https://ieeexplore.ieee.org/document/9109834/

[12] Y. H. Chen, S. Lo, D. M. Akos, G. Wong, and P. Enge, “A testbed for studying automatic dependent surveillance broadcast (ADS-B) based range and positioning performance to support alternative position navigation and timing (APNT),” *26th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS 2013*, vol. 1, pp. 263–273, 2013.

[13] A. Winter, N. Sokolova, A. Morrison, and T. A. Johansen, “GNSS-Denied Navigation using Direction of Arrival from Low-Cost Software Defined Radios and Signals of Opportunity,” in *EUCASS*, 2022.

[14] Tom J Nowell, “How can I find the smallest difference between two angles around a point?” [Online]. Available: https://stackoverflow.com/a/7869457/6561822

[15] R. Schmidt, “Multiple emitter location and signal parameter estimation,” *IEEE Transactions on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, 3 1986.

[16] A. J. Morrison and N. Sokolova, “Multi-band Multi-site GNSS RFI Monitoring Results after a Year of Operation,” *European Journal of Navigation*, vol. 21, no. 3, pp. 4–15, 2021.