DOI: 10.13009/EUCASS2022-7461
9TH EUROPEAN CONFERENCE FOR AERONAUTICS AND SPACE SCIENCES (EUCASS)

GNSS-Denied Navigation using Direction of Arrival from Low-Cost Software Defined Radios and Signals of Opportunity

Adrian Winter[⋆]*, Nadezda Sokolova*[⋆†]*, Aiden Morrison*[†]*, and Tor Arne Johansen*[⋆]

[⋆] *Center for Autonomous Marine Operations and Systems, Department of Engineering Cybernetics*
Norwegian University of Science and Technology (NTNU), 7034 Trondheim, NORWAY
[†] *Sintef Digital, 7034 Trondheim, NORWAY*

adrian.winter@ntnu.no, nadia.sokolova@sintef.no,
aiden.morrison@sintef.no, tor.arne.johansen@ntnu.no

Abstract

This paper describes a novel navigation system for outdoor navigation in conditions where reliable satellite navigation cannot be assumed. It is build around inexpensive of-the-shelf hardware and could be used with several different signal types, allowing flexibility in usage. The system is currently in a proof-of-concept stage, and this paper shows that there are promising preliminary results.

1. Introduction

It is well known that Unmanned Aerial Vehicles (UAVs) today strongly depend on Global Navigation Satellite Systems (GNSS) such as GPS, Galileo, Glonass or Beidu for position and sometimes attitude estimation. Unless specifically designed otherwise, e.g. to enable indoor or confined space operations, GNSS is used by almost any unmanned vehicle of any size, civilian or military to aid the Inertial Measurement Unit (IMU), such that accurate position and attitude estimation becomes possible.

1.1 Background

Due to the low transmission power and long distance between satellite and receiver ($\approx 20.000\,\text{km}$), the power levels at the receiver antenna are extremely low, which makes GNSS particularly susceptible to accidental or deliberate radio frequency interference (RFI). Such RFI could be caused by malfunctioning equipment such as TV antenna amplifiers [1] or incorrect use of testing equipment [2]. More often, however, RFI is associated with deliberate, malicious attacks such as jamming and spoofing. Jamming is a denial of service attack, where GNSS reception is severely degraded or impossible in a certain area. Jamming can be done with simple, cheap ($\lessapprox \$10$) devices that can be plugged into a car's cigarette lighter socket, or can be done using professional or military equipment. Jamming is simpler to perform and arguably not as dangerous as spoofing, because it is relatively simple to detect by a user – the GNSS reception under a jamming attack is poor or unavailable, and this will likely cause a warning of some sort. Spoofing, on the other hand, intends to supply a receiver with a valid, but incorrect, GNSS message and is thus much more difficult to detect. From a user perspective and if done carefully, there may be no indication at all that the system is providing incorrect data. This could for example be used to "hijack" vehicles [3].

As of today, there are three basic approaches to mitigate the risk of GNSS outages for low-cost UAVs:

1. Simply accept short-term outages, and ensure that the system is performing satisfactorily for a certain time, after which the GNSS outage should have ended. If it has not ended by that time, then there is a significant risk of loss of the vehicle. This is often still acceptable, as GNSS outage events are relatively rare and often very localized [4], and the low probability of loss and relatively low cost when a loss of a vehicle does occur mean that this risk can be accepted.

2. Use sophisticated receivers with features such as multi-frequency, multi-constellation reception or beamforming and nulling to reduce the risk of being affected by an attack. The idea behind beamforming and nulling is that

DOI: 10.13009/EUCASS2022-7461

the satellites always are over the horizon and geometrically spread, while a typical spoofer would be on or under the horizon, and all faked signals would come from the same position. Antennas that direct gain to the known azimuths and elevations of visible satellites (beamforming) and potentially reduce gain in the direction of a suspected RFI source (nulling) can allow operations even with a malicious actor present. However, hardware capable of these techniques is neither cheap nor common.

3. Use additional, GNSS-independent sensors. Navigation without GNSS has been the norm until recently, and especially in manned aviation only within the last few years has GNSS become anything more than an informative but not flight-critical system. Any vehicle capable of indoor or underwater use must also be capable of navigating without GNSS. However, all of these methods have some sort of limitation, included but not limited to requirements regarding ground visibility, trackable terrain features, specialized, pre-deployed equipment, knowledge about the local terrain, maximum distance of features (e.g. caused by limited range of LIDAR sensors). High-quality inertial reference systems that allow GNSS-free flight for much longer than consumer hardware are prohibitively expensive and often subject to export restrictions. Often these sensors are precise enough for long-distance navigation along specified routes, but not precise enough for navigation close to the ground.

1.2 Problem Description and Scope

The task and scope of this paper is to develop a navigation system that is capable of bridging medium-term (i.e. a few minutes) GNSS outages in conjunction with other on-board navigation sensors. Mitigation of sophisticated attacks such as directed spoofing or wide-area GNSS jamming are *not* within scope. The assumed attack scenario is a low-power jammer with a range of a few kilometers.

1.3 Signals of Opportunity-based Navigation

This paper proposes a novel navigation system, based on Direction of Arrival (DoA) measurement of signals of opportunity (SoOp). Using signals of opportunity, that is signals that are available independent of the mission and outside of the operator's control, provides the benefit that no mission preparation is necessary. Regulatory approval for usage of radio transmitters can be difficult or expensive, if not outright impossible to get, whereas legal operation of pure receivers is unproblematic in most jurisdictions. Providing ground based equipment may be expensive, possibly require personnel to protect it, takes time to deploy and collect, and limits the operational range to where the equipment has been pre-deployed – just to name a few potential challenges when not using SoOps.

However, using SoOps is not without challenges in itself. A typical issue is that the signal structure is usually not *designed* to be particularly suitable for navigation, and therefore may perform poorly simply because of the signal itself. For example, long-term frequency and phase stability of the carrier wave is often not necessary for radio communication, but may be paramount to certain navigation techniques. Similarly, an unpredictable synchronization mismatch of 1 ms may not be noticeable to a typical user of a communications signal, but would render any time-of-arrival based navigation system completely useless, as it would introduce errors of 300 km. Additionally, the physical distribution of transmitters may prove to be poor. Many navigation techniques ideally require a wide distribution of many transmitters, so a low area of uncertainty can be achieved in two or three directions. The navigation system proposed in this paper would, for example, only be able to deduce that one is "somewhere along a line", if the transmitters and receiver are all on a line. Alternatively, the achievable measurement accuracy may require relatively close-by signals, especially if angles are measured. This may render a system based on a sparse long-range network, such as (relatively) low-frequency broadcast towers, infeasible.

This can be summarized in a set of requirements to identify potential signals:

- Suitable signal structure, in particular with respect to long-term phase stability or timing characteristics
- Favorable geometry
- Known transmitter frequency and signal structure
- Known transmitter location
- Regular or continuous transmission
- Each measurement must be attributable to transmitter, e.g. by distinct transmitter frequencies
- Favourable frequency (range, antenna size, line-of-sight transmission)

DOI: 10.13009/EUCASS2022-7461

Note that not all of these characteristics are required for all SoOp-based navigation techniques. Many techniques may only require a subset of this list.

Discounting navigation techniques with purely passive receivers but dedicated infrastructure (e.g. GNSS, VOR, LORAN) are from a user perspective are similar to SoOp-navigation. Other works related to SoOp-based navigation includes exploiting the Single-Frequency Network (SFN) for TV broadcast (DVB-T) [5] or Radio broadcast (DAB) [6], cellular networks (LTE, 5G) [7], [8] or received signal strength (RSSI). These are all promising ideas with their own strengths and weaknesses and may even be complimentary to the idea proposed here.

2. Methodology: The Proposed System

The system works by measuring incoming signals' angle of arrival relative to the vehicle (bearing). Assuming a favorable geometry, three signals are enough to calculate the vehicle's 2D-position and heading.

While it's theoretically also possible to estimate the incoming elevation of the signals, and therefore the altitude of the vehicle, it seems unlikely that the achievable measurement accuracy will allow calculating any usable altitude information – especially considering that the transmitters are likely all close to the horizon, giving a very poor geometry. In addition, measuring altitude is rather trivial using a barometric sensor if the GNSS outage time is short enough that time-invariant air pressure can be assumed. Considering the limited range and endurance of many UAVs as well as the assumed use cases, this limitation does not appear to be of relevance in practice and only becomes relevant after roughly an hour or flight distances of more than 100 km. The rest of this paper will therefore purely focus on 2D position estimation and assume that the altitude is known.

2.1 Use-case and requirements

Previous work by Morrison et al. [9] has shown that most GNSS RFI incidents are short-term and local, and likely caused by car-based equipment – although the latter statement is speculation, as it is very difficult to find and interrogate users of such illegal equipment. This is also the scope for the research presented here. While large-scale, long-term GNSS outages are not impossible and have been reported [10], [11], this is out of scope at the moment. It is also assumed that GNSS is available during initialization of the system.

This scope allows limited flight under RFI influence. Such flights could occur during flights where one or multiple UAVs are used for localization of RFI sources [12], a task that the authorities are very interested in to enforce legal use of licensed frequency bands. This also gives additional requirements to a navigation system:

- small and light-weight, so can be carried by a small UAV
- low power consumption, to not decrease flight endurance significantly
- relatively cheap as to not increase the hardware cost by an unreasonable sum.

All of these requirements are fulfilled by the $\approx$ \$200 multi-channel coherent Software Defined Radio (SDR) receiver *KerberosSDR*, which essentially consists of four RTL-SDRs tied to a common clock. Refer to section 2.5 for a more detailed description of the SDR. In addition, antennas and cabling contribute to the weight, size and cost budget.

2.2 Principle

Assuming known transmitter location(s), ideal measurements, and non-singular geometry, the following information can be derived:

1. Using one beacon and no heading: No information can be derived
2. Using one beacon and heading: the UAV's position can be localized anywhere on a semi-infinite line
3. Using two beacons and no heading: the UAV's position can be localized anywhere on an ellipse
4. Using two beacons and heading: the UAV's position can be uniquely identified
5. Using three beacons: The UAV's position and heading can be calculated

It is reasonable to assume that any modern UAV has a magnetic sensor available on-board, however magnetic sensors are not perfect. They can be influenced by on-board equipment (especially motor torque dependent magnetic fields), and magnetic navigation becomes unreliable at high latitudes [13]. Whether or not this is a restriction in practice depends much on the hardware configuration, the effort spent on compass calibration, as well as the earth's magnetic field (and therefore, implicitly, the given mission).

DOI: 10.13009/EUCASS2022-7461

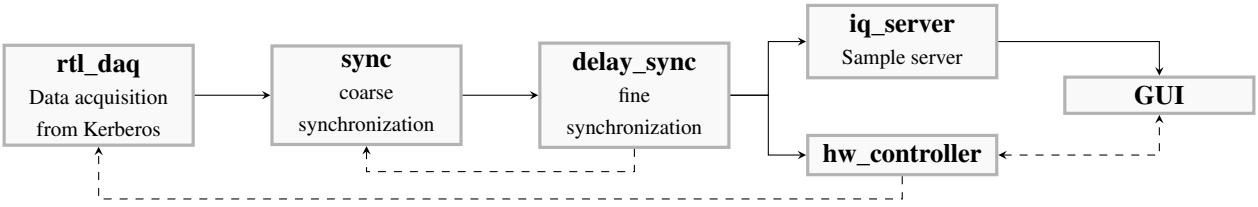

Figure 1: Simplified flowchart of pre-processing

2.3 Suitable Signals of Opportunity

As mentioned in section 1.3, the signals should have several of characteristics to be suitable for Direction of Arrival-based navigation. In particular, there must be multiple uniquely identifiable, transmitters with a suitable geometry and a known location. In addition, from the requirements due to the platform (section 2.1), the measurement must be taken with a system that is not too large or heavy. KerberosSDR uses phase differences for DoA-estimation, and therefore requires an antenna array with an antenna spacing d that satisfies $0.3\lambda \lessapprox d < 0.5\lambda$, where λ is the wavelength of the incoming signal. This implies a carrier frequency of $\gtrapprox 100\,\text{MHz}$ – else, the antenna array would be either physically too big or the antennas too close to each other, resulting in poor measurement accuracy.

While there are certainly several potential signals available, all of the requirements are fulfilled by the Automatic Identification System (AIS) used mainly for collision avoidance in ships and boats. AIS is a system, where the ship's current position (amongst other information) is broadcast without encryption on 161.975 MHz and 162.025 MHz. Thanks to Time Division Multiple Access (TDMA) schemes, it's simple to associate the transmitter to the DoA measurement, since there's (almost) always only one transmitter simultaneously sending. The transmission powers are relatively high and far above the noise floor, which makes filtering simple - a lowpass filter is sufficient. In coastal areas, there are typically many boats simultaneously available, although this obviously very much depends on the area where the mission is supposed to be flown. The main disadvantage of using AIS signals is that it itself depends on GNSS - the broadcast information is always taken from the ship's GNSS system. However, as described in section 2.1, this is likely not a critical problem in our scope, as almost all GNSS interference events are localized and therefore will only affect a few transmitters.

Since NTNU is located in coastal areas and the affiliated research group AMOS (Centre for Autonomous Marine Operations and Systems) mainly researches maritime applications, AIS was selected as an initial research focus for this work. However, other promising signals, in particular the very similar ADS-B system used by aircraft, could be investigated in a later step.

2.4 Processing Chain

The processing toolchain as it exists now is a non-real time capable software collection built around the KerberosSDR system and its driver and analysis software `heimdall_daq_fw`. For a simplified flowchart, where all deactivated blocks are omitted, see fig. 1. The other software used for KerberosSDR, `krakensdr_doa`, is in this context purely used for command and control of the underlaying toolchain. In particular, its DoA functionalities are currently not used. Since this part of the software is not directly relevant to this work, it will not be discussed further.

The last step of the multi-process toolchain `heimdall_daq_fw` is modified so that all received and partially processed packets are saved to a SQLite database for post-processing. For simplifying the development process, real-time integration was not yet done, but no performance issues or other potential showstoppers are identified or expected.

This part of the toolchain is responsible for sample synchronization, that is ensuring both coarse (sample-level) and fine (subsample-level) alignment between the channels. After all those steps are performed, the samples are saved.

The post-processing toolchain again consists of several independent steps/modules, whose task is to identify and decode the payload message, and then to calculate the direction of arrival of each message. For a flowchart of these steps, see fig. 2. The modularity decreases performance and user-friendliness, mainly because some steps are performed several times, but it allows for detailed debugging and is therefore currently preferred.

The first step of the post-processing toolchain is to identify the burst boundaries. For this, the first step is to pre-filter the relatively large databases for promising sections. This is done by concatenating two data packets (who each have a length longer than the AIS signal burst) and trying to decode the AIS payload message. If this decoding is unsuccessful, the packet can be rejected, as being able to decode the payload message is a requirement for all following steps.

DOI: 10.13009/EUCASS2022-7461

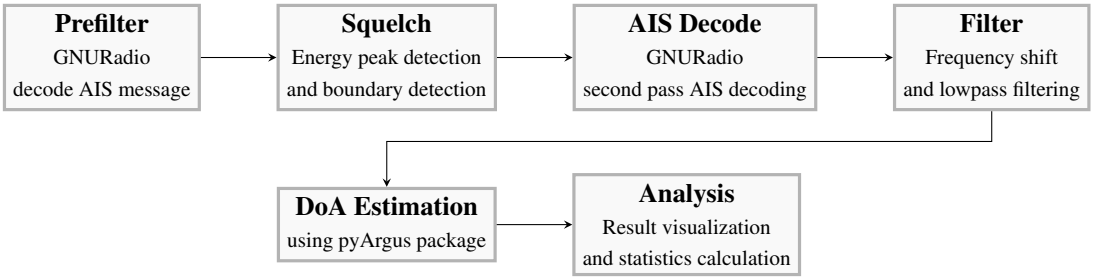

Figure 2: Flowchart of post-processing

The second step consists of localizing and "cutting" the burst, which is essentially a squelch algorithm. Using scipy's `Hilbert()` function, the envelope of the signal is calculated. The convolution of the envelope and a rectangular test window with the burst's known length gives a triangular scalar function as a result, whose maximum value is the centre of the burst. This relatively simple method works only if the signal length is known with sufficient accuracy (as otherwise the peak of the triangular function is no longer sharp), but the assumption of a known pulse length holds true for the overwhelming majority of the bursts. It also only works if the signal is clearly above the noise level, but again this assumption should not be a limitation in practice, as AIS uses high signal powers and far-away signals provide little information for navigation due to high errors.

After the burst is cut, it's sent though the AIS decoder once more. This step is done for two reasons: first, to ensure that the resulting, cut burst can still be decoded. Second, if the original, uncut concatenated packet contained more than one burst, there might be cases in which the cut packet and the decoded payload do not match. This is not unlikely to happen, as each concatenated packet has a length of 64 ms, whereas each AIS burst is 26 ms in length. Considering that the original burst also contains channels A and B, which can be used simultaneously, up to four AIS bursts could be within one original packet.

The decoded burst contains the information about which channel the message was sent on. Using this information, the burst is then shifted by 25 kHz in the appropriate direction and lowpass-filtered. This step is necessary because the RTL-SDR's oscillators are set to 162.0 MHz, and the message is sent on either 161.975 MHz or 162.025 MHz. In other words, the message at this point is still on some intermediate frequency.

The second to last step consists of sending the bursts to the DoA estimation algorithm. Currently, a library developed and maintained by the lead developers of the KerberosSDR, `pyArgus` is used, but since the data simply consists of four channels of complex samples, most DoA estimation implementations should be usable.

As the last step, an evaluation of all previous steps is done. Since it is known where the signals were recorded from during development, the azimuth between antenna and received signal can be easily calculated and, barring small measurement errors and latencies from the boats, seen as ground truth. This ground truth measurement can then be compared to the DoA measurements.

At many points in the toolchain, there are tap-offs or plotting functionalities to help with debugging.

2.5 Used hardware

As already mentioned before, the test system is built around the KerberosSDR four-channel coherent SDR system. KerberosSDR is a development platform essentially consisting of four RTL-SDR modules connected to a common oscillator, which ensures phase-coherent sampling. There is also a configurable noise source and an integrated USB hub, the former being necessary for the initialization process. RTL-SDR refers to the very common software defined radios built around the Realtek RTL2832U chip.

KerberosSDR is connected using USB 2.0 to a Raspberry Pi 4 running Ubuntu Linux. Connected to the Raspberry Pi is a USB 3.0 SSD, which is fast enough to save the recorded data.

For signal reception, four Glomex RA-111 AIS *rubber ducky* antennas are connected. These are monopole antennas, i.e. they require a groundplane. At the moment, this groundplane is realized by connecting four copper wires of appropriate length to the base of the antenna.

2.6 Forward-backward-ambiguity

Currently, a linear antenna array consisting of four antennas is used. Linear antenna arrays cannot distinguish whether signals come from in front or behind the array. However, it should be easy to reject the unlikely signal based on the

DOI: 10.13009/EUCASS2022-7461

current estimate of position and attitude of the UAV with respect to the beacon, so that the issue of aliasing is currently assumed to be a solvable problem.

An alternative to accepting aliasing and rejecting the unlikely signal is to use a circular array, which is identical to a square array for four antennas. Both hardware and software support circular arrays and the authors will investigate the implications of changing to a circular array in the future. It seems likely that the added information comes at a cost of reduced precision, but this can currently neither be confirmed nor disproved.

2.7 Measurement and Integration

The measurement chain provides two simultaneous "measurements": the estimated bearing to the signal, and the AIS payload. As mentioned before, the AIS payload contains several fields, of which mainly the position is relevant for this work. The associated measurements must then be fused into the existing state filter. To do this, there are two fundamental approaches:

The first one can calculate a 2D position and potentially heading based on the measurements alone, and fuse this into the existing Kalman filter. This would be analogous to a loosely-coupled GNSS measurement integration. The advantage of this approach, as with any loosely-coupled system, is that it is significantly simpler to implement. Loosely coupled sensor integration can be seen as independent modules, where it is irrelevant for the platform *how* the measurements were calculated. Another added benefit of this approach is that it's trivial to compare the GNSS and SoOp-measurements, which can be a simple cross-check to verify the integrity of GNSS.

The second approach is similar to the tightly-coupled navigation approach. In this case, the bearing measurements are directly integrated into the state filter. The advantages and disadvantages are therefore also similar to tightly coupled GNSS navigation: the integration is more difficult, however the achievable performance will most likely be higher, as more information can be processed.

In either case, the main mathematical challenge stems from the non-linearity of the measurements. Some additional challenge will be caused by the fact that each individual measurement only gives a line, not a point, as a potential position candidate. Each individual measurement yields an underdetermined system of equations. A naive solution to this is to simply collect at least two measurements from different sources, whose intersections then yield a unique solution.

However, the proposed navigation system is *not* an instance of the "typical" bearings-only navigation problem [14], because the entire position can be observed without additional assumptions or maneuvering by the platform – instead, the observability is ensured by tracking bearings to several targets.

2.7.1 Measurement function

The measurement function used in a typical Kalman filter approach is dependent on the state vector used for state estimation. Assuming a simplified state vector only consisting of position and heading

$$\vec{x} = (x^U, y^U, \psi)^T \tag{1}$$

and a position of the i-th boat

$$\vec{b_i} = (x_i^B, y_i^B)^T \tag{2}$$

where the first of the entries (x^U, x^B respectively) is the position northwards, the second entry (y^U, y^B) is the position east and the heading ψ is the compass direction relative to true north. Additionally, a Cartesian flat earth frame of reference is assumed. This assumption is not restrictive in this application due to the local nature of both the UAV and signals combined with the relatively low expected accuracy. By converting from a geodetic frame to a local frame the notation can be *significantly* simplified.

Using simple trigonometry, the azimuth between UAV and beacon can be calculated:

$$\chi_i = \text{arctan2}(\Delta E, \Delta N) \tag{3}$$

where $\text{arctan2}(x, y)$ follows the north-clockwise convention and ΔE, ΔN is the distance of the beacon from the UAV in east and north direction, respectively.

Subtracting the current heading angle ψ from the azimuth angle will give the expected measurement, i.e. the bearing:

$$\theta = \chi_i - \psi \tag{4}$$

$$= \text{arctan2}(\Delta E, \Delta N) - \psi \tag{5}$$

$$= \text{arctan2}(y_i^B - y^U, x_i^B - x^U) - \psi \tag{6}$$

DOI: 10.13009/EUCASS2022-7461

This measurement function can then be used for the filtering or estimation process, potentially using linearization or other additional steps if necessary for the filter implementation.

3. Results

To verify that the system is working, it has been placed on the institute's roof and two datasets of approximately 15 minutes each have been recorded. For the first recording, the antenna array was facing towards a base station on a nearby mountain, whereas for the second recording the antenna was turned by about 60°.

This base station, MMSI 2573325, was chosen mainly because of its location. There is no unobstructed line-of-sight to other permanent transmitters on the nearby fjord, while the base station is located on a hill. Together with the relatively high receiver antenna location on top of a building, which in turn is on top of another hill, it is reasonable to assume that factors such as multipath due to an obstructed direct signal are not dominating.

The simple fact that the AIS payload can be decoded does not guarantee a high-quality signal. AIS is designed to work well with degraded communication channels, in particular by using robust Gaussian Minimum Shift Keying (GMSK) and very low data rates of 9600 baud. This allows decoding of signals coming from transmitters approximately 60 km away that are obstructed by mountains, however it is reasonable to assume that these signals have very poor signal quality and can only be used with caution within this project.

Similar problems exist with using signals from the nearby harbor area. The coast is only approximately 2–3 km from the institute's building, however the entire area between the University and the shore is urban. Together with potential reflections from the surrounding mountains, it can be assumed that the received signals from boats are also degraded.

These suboptimal conditions make it difficult to identify issues within the system, as it is difficult to say exactly where a measurement error was introduced. This is the main reason why heavy data pruning is used for this proof of concept.

Table 1 lists the results if only one target, the base station on the nearby hill, is used. *Mean* is the mean value of the errors (i.e. measurement minus ground truth). Ideally, this should be equal to the yaw angle of the antenna array, since the antenna array is not oriented towards north. The table lists four levels of data pruning:

Orientation 1		Orientation 2	
Mean [°]	Std. dev. [°]	Mean [°]	Std. dev. [°]
+3	0.5	-27	0.9

Table 1: Comparison of mean and standard deviation for different pruning levels

A discussion why only one target is included as of now, will follow in the next section.

4. Discussion

The obvious point to address in the discussion is that the results are preliminary in the sense that only parts of the system have been verified to be working. There are no known fundamental issues, but whether or not the system will be able to work in practice is to be seen. Many issues have been identified and most of them are fixed, but some are still open to an extend where a sound assessment of the performance is currently not possible. Currently, analyzing the measurements from more than one boat results in unacceptable errors, which is why these results are not included here.

However, as can be seen in table 1, the measurement precision (given as standard deviation) is good if only one target is used. The mean is a proxy measurement for the orientation and difficult to assess for its validity, as the true orientation of the antenna array is unkown.

Obviously the approach of simply discarding any measurements that do not fit is not a valid approach, but it *does* yield an upper bound of the expected accuracy as well as an indication that the system *should* be able to perform well enough, if all the persisting issues are addressed.

4.1 Error Sources

Error sources can broadly be sorted into three categories:

1. Propagation and reception errors
2. Pre-processing and recording errors

DOI: 10.13009/EUCASS2022-7461

3. Post-processing errors

A description of these three categories follows. The description of potential error sources is not complete, especially since there might be un-identified issues still.

4.1.1 Propagation and reception errors

These errors include all errors that are introduced between the transmitter and the receiver firmware. The authors do not have equipment to inject synthetic messages into the RF frontend, so real transmissions from AIS transmitters have to be used.

Due to the location of the receiver antenna array in a wide trough-like valley, i.e. surrounded by mid-size mountains of $\lessapprox 600\,\mathrm{m}$ in elevation and "behind" a city, it is very likely that the experiments have been performed in a rich multipath environment. The effects of multipath are very difficult to predict and to what extend it can and does degrade the reception is currently not known.

The other suspected error source is the antenna array itself. Currently, four independent short helical *"rubber ducky"* antennas are used with four angled dipoles each as ground plane. Using three or more dipoles as a groundplane for a whip antenna is commonly done [15], but whether or not these four antennas and sixteen groundplane dipoles influence eachother in a detrimental way is not known to the authors. There are, however, indications that this might indeed happen, as there are strange variations in relative received signal strengths between the antennas as well as a general underestimation of the received DoA, which could be caused by antenna behaviour that significantly deviates from an ideal dipole.

4.1.2 Pre-processing and recording errors

All software errors that happen during the actual data recording fall in this category. Broadly speaking these are bugs in `heimdall_daq_fw`. Misconfiguration of the toolchain during recording could also be sorted in this category, for example by setting the receiver gains to an inappropriately low or high value.

This part of the toolchain is difficult to debug. While the toolchain *is* modular in the sense that each task has its dedicated process that broadly speaking only does one single thing, the toolchain is heavily inter-dependent and it's difficult to pause the execution of any one process or run reduced subset of processes for debugging purposes. The toolchain is under active development by the project maintainers, and bugs have been identified and fixed by both the project maintainers and the authors of this paper, but due to the fact that the software only very recently left beta-testing it is not inconceivable that there are some bugs still left.

4.1.3 Post-processing errors

All bugs that are within the custom toolchain developed for this project operating on the recorded data from the previous steps fall in this category. This includes incorrect decoding of the payload message, matching DoA measurements to the payload message as well as the actual DoA estimation. Due to the fact that this part of the toolchain is developed by a single person and not cross-examined/tested by other users, it is likely that there are many bugs left in this section in particular.

This part of the toolchain is relatively easy to debug, as each subsequent step is done independently of the previous one, so operating on a subset of data or pausing execution is easily possible. The fact that this part need not run in real time also significantly simplifies the debugging process, since time-consuming but useful steps such as plotting each burst are feasible.

4.2 Known Issues

The standard deviation of the errors in the previous section are clearly unacceptably high. However, the errors are not random. This gives reason to believe that there is a relatively simple fundamental problem with the processing steps, which for example introduces a scaling error, where all calculated DoA-values are multiplied with a constant factor. This could, for example, be caused by incorrect assumptions about the (electrical) distance between the antennas or a simple bug. Adding to this theory is the very *low* standard deviation if only one target is considered.

While the system in its current state is clearly not usable with errors that large and obvious, the authors believe that with some further tuning it should be possible to reduce these issues to an acceptable level.

While the phase-amplitude plots of some bursts look exactly as expected, that is with an roughly equal amplitude for all four channels and a phase difference between adjacent channels that is equal for all four antennas, other bursts look considerably worse. Sometimes, the reception power levels are vastly different between the four channels, and

DOI: 10.13009/EUCASS2022-7461

sometimes the phase differences are unequal. Sometimes the phase differences between adjacent channels is higher than 120°, which is the highest value physically possible.

All of this are likely symptoms of a deeper, as-of-yet unidentified issue.

4.3 Further Research

Section 4.2 mentioned several problems *within the existing software*, but there are also several open problems to investigate that go beyond the tasks that have already been started. These tasks need to be done before operational capabilities are established and before flight tests can be done:

The current setup is not fully real-time capable, as the post-processing steps require around 1 second per each identified 25 ms burst. This does not mean that a speedup of a factor of 40 is required, because data is not transmitted continuously. However, for actual use the system still requires a significant speedup. Much of this speedup can likely be achieved by optimizing the existing toolchain, for example by avoiding unnecessary initialization and destruction of processes and by avoiding multi-pass decoding of messages. Additionally, it might be possible to decrease the sample rate or to decimate the data stream, however this will likely decrease the overall system performance.

The integration of the measurements into an actual flight platform (see section 2.7) is still open. Only when this works, a substantiated statement about the required measurement accuracy will be possible. This refers to both the hardware integration, e.g. mounting the antennas to the platform, and to the software integration.

Adapting the system to three dimensions might pose challenges too. In section 2 it was mentioned that the geometry of low-altitude UAV and boats close to the horizon imply that the problem can be seen as two-dimensional and the third dimension can be ignored. This simplification becomes invalid in two cases: if roll- or pitch angle of the platform are not zero (i.e. if the platform is tilted with respect to earth), and if the UAV is close to and above the signal source. Out-of-plane signals effectively shorten the distance between the antennas, which in return has an influence on the calculated DoA. It is likely that this influence is different if a linear or circular antenna array is used, and it is likely that this can be mathematically compensated for using the known attitude and relative geometry. Nonetheless, the effects of this need investigating.

5. Acknowledgments

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

References

[1] J. R. Clynch, A. A. Parker, R. W. Adler, W. R. Vincent, P. McGill, and G. Badger, *The Hunt for RFI: Unjamming a Coast Harbor*, Jan. 2003. [Online]. Available: https://www.gpsworld.com/the-hunt-rfi/.

[2] L. Scott, *Spoofing Incident Report: An Illustration of Cascading Security Failure*, 2017. [Online]. Available: https://insidegnss.com/spoofing-incident-report-an-illustration-of-cascading-security-failure/.

[3] T. Humphreys, *Humphreys Research Group Successfully Spoofs an $ 80 million Yacht at Sea*, Jul. 2013. [Online]. Available: https://www.ae.utexas.edu/news/humphreys-research-group.

[4] N. Gerrard, A. Rødningsby, A. Morrison, N. Sokolova, and C. Rost, "GNSS RFI monitoring and classification on Norwegian highways - An authority perspective," *Proceedings of the 34th International Technical Meeting of the Satellite Division of the Institute of Navigation, ION GNSS+ 2021*, pp. 864–878, 2021. DOI: 10.33012/2021.17952.

[5] D. Serant, O. Julien, C. Macabiau, *et al.*, "Development and validation of an OFDM/DVB-T sensor for positioning," *Record - IEEE PLANS, Position Location and Navigation Symposium*, pp. 988–1001, 2010. DOI: 10.1109/PLANS.2010.5507273.

[6] D. Palmer, "Position estimation using the Digital Audio Broadcast (DAB) signal," 2011. [Online]. Available: http://eprints.nottingham.ac.uk/12456/.

[7] J. A. del Peral-Rosado, J. A. López-Salcedo, G. Seco-Granados, *et al.*, "Software-defined radio LTE positioning receiver towards future hybrid localization systems," *31st AIAA International Communications Satellite Systems Conference, ICSSC 2013*, 2013. DOI: 10.2514/6.2013-5610. [Online]. Available: https://arc.aiaa.org/doi/10.2514/6.2013-5610.

DOI: 10.13009/EUCASS2022-7461

[8] A. Dammann, R. Raulefs, and S. Zhang, *On prospects of positioning in 5G; On prospects of positioning in 5G*. 2015, ISBN: 9781467363051. DOI: 10.1109/ICCW.2015.7247342.

[9] 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.

[10] Inside GNSS, *Lessons to be Learned from Galileo Signal Outage*, Oct. 2019. [Online]. Available: https://insidegnss.com/lessons-to-be-learned-from-galileo-signal-outage/.

[11] C. Kråkenes, *Norwegian Armed Forces creating GPS jamming alert system*, Jun. 2020. [Online]. Available: https://www.nrk.no/tromsogfinnmark/norwegian-armed-forces-creating-gps-jamming-alert-system-1.15073878.

[12] N. Ahmed, A. Winter, and N. Sokolova, "Low Cost Collaborative Jammer Localization Using a Network of UAVs," *IEEE Aerospace Conference Proceedings*, vol. 2021-March, Mar. 2021, ISSN: 1095323X. DOI: 10.1109/AERO50100.2021.9438441.

[13] *Magnetic field in the Arctic regions*. [Online]. Available: https://www.geomag.nrcan.gc.ca/mag_fld/arctics-en.php.

[14] N. Gordon and M. Pitt, "A COMPARISON OF SAMPLE BASED FILTERS AND THE EXTENDED KALMAN FILTER FOR THE BEARINGS-ONLY TRACKING PROBLEM," 1998.

[15] Electronics Notes, *Antenna Ground Plane: theory & design*. [Online]. Available: https://www.electronics-notes.com/articles/antennas-propagation/grounding-earthing/antenna-ground-plane-theory-design.php.