Single GNSS Antenna Heading Estimation

Fabian Rothmaier, Stanford University
Yu-Hsuan Chen, Stanford University
Sherman Lo, Stanford University
J. David Powell, Stanford University

BIOGRAPHIES

Fabian Rothmaier is a PhD candidate at the GPS Laboratory at Stanford University. He received his B. Engr. degree from the University of Applied Sciences Bremen, Germany in 2015 and his M. Sc. degree from Stanford University in 2017.

Yu-Hsuan Chen is a research engineer at the Stanford GPS Laboratory. He received his Ph.D. in Electrical Engineering from National Cheng Kung University, Taiwan in 2011.

Sherman Lo is a senior research engineer at the GPS Laboratory at Stanford University. He received his Ph.D. in Aeronautics and Astronautics from Stanford University in 2002. He has and continues to work on navigation robustness and safety, often supporting the FAA. He has conducted research on Loran, alternative navigation, SBAS, ARAIM, GNSS for railways and automobiles. He also works on spoof and interference mitigation for navigation. He has published over 100 research papers and articles. He was awarded the ION Early Achievement Award.

J. David Powell is a Professor Emeritus of Aeronautics and Astronautics and of Mechanical Engineering at Stanford University.

ABSTRACT

Heading is an important state in modern vehicle control, especially in applications like ships in currents, aircraft in crosswinds, helicopters or Vertical Take Off and Landing (VTOL) aircraft. And yet there is only limited means to measure a vehicle's heading on earth, and each approach has its drawbacks and limitations. Magnetometers for example are susceptible to electromagnetic interference and are of limited use in polar regions. Alternatives such as gyrocompass systems or GNSS antenna arrays come with a significant price tag and / or complexity.
This paper presents a heading estimation procedure using inexpensive hardware and without the limitations of a system based on measuring magnetic field strength: the combination of a single element Dual Polarization GNSS Antenna (DPA) that delivers a low frequency, low accuracy absolute heading measurement and a MEMS rate gyro.
Both sensors are characterized, and their measurement models justified. Two algorithms are then described to fuse the measurements from both sensors in a Kalman Filter. Simulation results outline the technique's performance, depending on the quality of the rate gyro used and the DPA's measurement characteristics. Test results from a dynamic driving test compare the heading estimate using the current implementation of the DPA to the estimate based of a magnetometer combined with the same rate gyro.
The DPA based result is comparable to the magnetometer-based result in heading rate, but shows lower accuracy in absolute heading. Pareto curves based on simulation results outline how this limitation could be overcome in the future.

INTRODUCTION

A vehicle's heading with respect to some reference direction has been an essential part of navigation since the first explorers set sail beyond the visible horizon. In the age of GNSS navigation, where the determination of one's velocity vector azimuth is easily determined, it may seem that the need for one's heading has less relevance than it did through history.
While that is partially true, there is a definite value in knowing heading for vehicles that are stationary, for ships in currents, for aircraft in cross winds, for helicopters and Vertical Take Off and Landing (VTOL) aircraft, and for high performance race cars - where the heading vector is not necessarily the same as the course over ground or no velocity exists. Hence, in today's

modern aircraft and vessel control, heading is a heavily used state. And yet, there is only limited practical alternatives to measure heading on earth. This project introduces a new technique, based on a single GNSS antenna and a low-cost gyro. An overview of the current alternatives is given in [1].
Providing this capability in a small form factor with very limited extra hardware comes with a few drawbacks compared to GNSS antenna arrays that can measure their attitude with high precision. Due to the antennas implementation using a single integrated circuit, it can currently only measure heading with an update rate below 0.5 Hz and with significant noise. Every measurement further comes with a 180 degree ambiguity, leading to a 180 degree ambiguity in the calculated heading.
The following paragraph gives an overview of alternatives for heading estimation. The next section describes the sensors used in the approach presented by this paper, followed by a description of the dynamic models and Kalman Filter integration modes in the next section. It is followed by a section describing simulation results, which are completed by hardware implementation test results. The paper concludes with a summary of the results and some closing thoughts.
Unless stated otherwise, heading and all other angles about the vertical axis are defined to be positive in clockwise direction starting from true north. Throughout all equations in the next sections, unless stated otherwise, a variable with a hat "^" shall denote a (noise affected) measurement or estimate. Vector quantities are expressed by bold characters.

Heading Estimation Solutions

A magnetic compass, magnetometer or flux valve, possibly augmented with gyroscopes, is used in cost sensitive applications or where magnetic heading is desired. Combined with means of positioning and a database of magnetic field variation values, heading with respect to true north can be derived as is done in [2] and [3]. If combined with accelerometers to measure gravity and vehicle accelerations, solving Wahba's problem based on the magnetic field and gravity vectors results in an attitude solution [4].
An alternative is a high-performance gyrocompass that measures attitude with respect to the earth's rotational axis directly. In combination with a position source and a variation database, the magnetic heading can be derived. This solution requires navigation-grade gyrocompass systems for a reliable heading information [5].
Another option is measuring the vehicles orientation through multiple GNSS antennas. [6], [7], [8] have shown promising results using this approach using as few as two GPS antennas and augmenting the attitude solution with rate gyros. Similar to gyro compassing, this results in a heading measurement with respect to true north.

All those options have considerable drawbacks. Magnetic devices are based on measuring the magnetic field around the sensor. The trend towards electric means of propulsion, as well as the increased use of electronics in vehicles for navigation, automation and operator comfort distorts the electromagnetic field. Therefore, extra care has to be taken to shield the magnetometer or place it appropriately on the vehicle. In addition, the high inclination of magnetic field lines and the wandering magnetic north pole further limits their use especially in polar regions. Melting icecaps enabling increased traffic in extreme latitudes as well as a more rapidly changing earth magnetic field incentivize the search for alternative solutions [9]. High performance gyrocompass systems come with a significant price tag, weight and size making their use infeasible outside commercial aircraft, military equipment or single applications like research vehicles.
Using multiple GNSS antennas for the vehicle's orientation requires a carefully calibrated setup and sufficient space on the vehicle. This method has significantly higher complexity and cost than the approach presented in this paper.
[10] presents an interesting approach of attitude determination using a single GPS antenna, but limits the results to pitch and roll and requires an accurate dynamics model of the vehicle.

TWO COMPLEMENTORY ANGLE SENSORS

The heading estimation approach presented is based on a single GNSS antenna and a rate gyro. The next section introduces the Stanford Dual Polarization Antenna (DPA), the GNSS antenna used. Originally developed for interference mitigation, it generates bearing measurements to the GNSS satellites in view at low accuracy and low frequency. These can be used for an absolute heading measurement and are then augmented with the angular rate measurement of a rate gyro introduced in the section after.

The Stanford Dual Polarization Antenna

The Stanford Dual Polarization Antenna (DPA) depicted in Figure 1 is a single GNSS patch antenna that is compliant with the ARINC 743 form factor specification for GNSS antennas on aircraft [11]. It is capable of measuring the approximate azimuthal Direction of Arrival of satellite signals, denoted bearings $\boldsymbol{\theta}$ [12]. With two feeds and a hybrid coupler, the DPA detects the phase difference between Right Hand Circular Polarized (RHCP) and Left Hand Circular Polarized (LHCP) components of each GNSS signal. An approximate bearing can be calculated from this phase difference. The interested reader is referred to the explanations in [11]–[13] for more detail on the antenna. Within the scope of this work, it is sufficient to know that the DPA delivers a noisy bearing to each tracked GNSS satellite every 2.56 seconds. Figure 2 shows a histogram of the θ measurement errors made during a stationary test on the rooftop of the Durand building on Stanford Campus. The histogram shows over 8000 measurements from all elevations for GPS and GLONASS satellites. Given that Figure 2 depicts noisy, real world data, the measurement error is further modelled and assumed to be identically independently distributed (IID) according to a Normal distribution with a standard deviation of 20 deg as given by Equation (1).

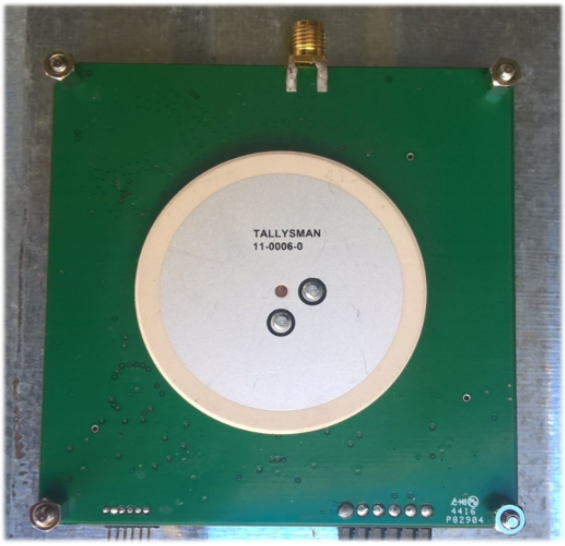

Figure 1: The Stanford DPA

$$\varepsilon_\theta \sim \mathcal{N}(0, \sigma_\theta^2) \qquad \text{with} \qquad \sigma_\theta = 20\ deg \tag{1}$$

There are two implementation strategies that could be followed when using bearing measurements for heading estimation.

i. The bearing θ_i to the ith satellite is related to the satellite's azimuth ϕ_i and the antenna or vehicle heading ψ as described in Eq. (2). The azimuths $\boldsymbol{\phi}(t)$ are known from the antenna's position and the ephemeris message. One possible approach to get from N bearing measurements, one $\hat{\theta}_i$ to each satellite, to a heading estimate $\hat{\psi}$ is to find the maximum likelihood heading estimate (MLHE) that minimizes the measurement error of all bearing measurements. Since the bearing measurements are IID, the MLHE $\hat{\psi}$ at each epoch t is the solution to the (weighted) least squares problem given by Equation (3).

$$\theta_i = \phi_i - \psi \qquad , \qquad i = 1, \dots, N \tag{2}$$

$$\hat{\psi}(t) = \underset{\psi(t)}{\operatorname{argmin}} \left\| diag(\boldsymbol{w}) \cdot \left(\hat{\boldsymbol{\theta}}(\boldsymbol{t}) - \left(\boldsymbol{\phi}(\boldsymbol{t}) - \mathbf{1}_{\boldsymbol{N}} \psi(t) \right) \right) \right\|^2 \tag{3}$$

With the optional weighting vector $\boldsymbol{w}$ representing individual $\sigma_{\theta,i}$ (see the notion on $\sigma_{\theta,i}$ in the next paragraph), the vector of bearing measurements at time t $\hat{\boldsymbol{\theta}}(\boldsymbol{t})$, the vector of satellite azimuths $\boldsymbol{\phi}(\boldsymbol{t})$ and a N x 1 vector of ones for N tracked satellites at time t, $\mathbf{1}_{\boldsymbol{N}}$. The diag() command here creates a diagonal matrix from a vector. Other cost functions such as the L1 norm and Huber penalty function have been experimented with by the author, but the above least squares formulation has led to equal or better results and is therefore used in the loosely coupled implementation of the Kalman Filter described in the next section. It is worth noting that due to the wrapping of angles, Equation (3) is not necessarily convex and can have local minima. But since the domain of the problem is limited to [0 2π], a solution can nevertheless be easily found using a combination of the optimization techniques described in [14] or Matlab's fminbnd function. Figure 3 shows the heading values computed according to (3) to be approximately zero-mean Gaussian with standard deviation σ_{DPA} = 3.4 deg, based on the same dataset used for Figure 2. Taking the least squares average at each epoch has reduced the standard deviation by more than a factor of 5 compared to the individual bearing measurements to each satellite.

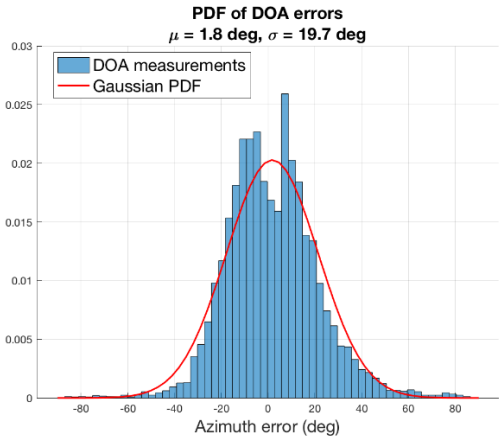

Figure 2: Bearing error histogram during a stationary rooftop test (deg)

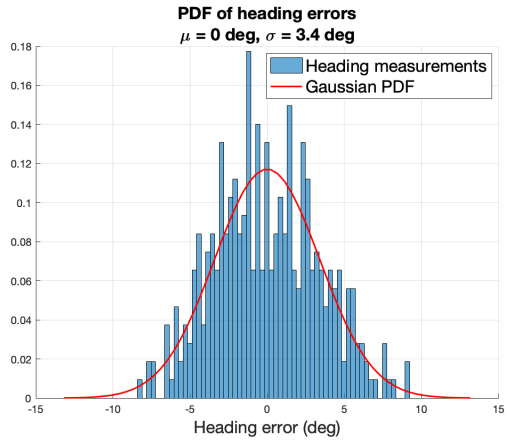

Figure 3: DPA heading error histogram during a stationary rooftop test (deg)

ii. The bearing measurements $\hat{\theta}_i$ can also be used in a Kalman Filter directly, as the bearing θ_i to the ith satellite is related to the satellite's known azimuth ϕ_i and the antenna or vehicle heading ψ as described in Equation (2). Eq. (2) is rewritten as Eq. (4) to phrase N conditionally independent heading measurements $\hat{\psi}_i$, one for each of the N tracked satellites. The measurement noise v_θ is zero-mean Gaussian with covariance matrix R_θ (Equation (*5*)).

$$\hat{\psi}_i = \phi_i - \hat{\theta}_i + v_{\theta_i} \quad , \quad i = 1, \dots, N \tag{4}$$

$$\boldsymbol{v_\theta} \sim \mathcal{N}(0, R_\theta) \quad \text{with} \quad R_{\theta_{[ij]}} = \begin{cases} \sigma_{\theta,i}^2 & i = j \\ 0 & otherwise \end{cases} \tag{5}$$

where $\sigma_{\theta,i} = \sigma_\theta$, the standard deviation defined in Equation (1). A more precise definition of individual $\sigma_{\theta,i}$ is possible based on information from the phase difference detection procedure of the DPA but beyond the scope of this work. This approach is implemented in the tightly coupled version of the KF described later.

The Rate Gyro

A lot of prior work exists on modelling rate gyros, [5] or [15] for example. This work will closely follow the nomenclature and setup of [15] which is further used in [3].
After removing the constant gyro bias or drift by calibration with a static test, following the cited work the gyro's measurements can be described as Gaussian Normally distributed about the heading rate $\dot{\psi}$ plus a time varying bias $b(t)$:

$$\hat{\dot{\psi}}_{gyro}(t) \sim \mathcal{N}(\dot{\psi}(t) + b(t), \sigma_{gyro}^2) \tag{6}$$

The variance σ_{gyro}^2 is the power spectral density of the white noise that is the driver behind the rate gyro phenomenon called angle random walk (ARW). σ_{gyro} is in $\mathrm{deg/s}/\sqrt{Hz}$ and read off of the Allan Deviation plot or given on sensor specification sheets for integration times of 1 sec [16]. The bias b can be modelled as an exponentially correlated Gaussian random process [15]:

$$\dot{b}(t) = -\frac{1}{\tau} b(t) + w_b \tag{7}$$

with the driving process noise w_b:

$$w_b \sim \mathcal{N}(0, Q_{w_b}) \quad \text{with} \quad Q_{w_b} = 2\,\frac{\sigma_{w_b}^2}{\tau} \tag{8}$$

for the bias instability standard deviation σ_{w_b} from specification sheets and correlation time $\tau = 1000s$.

KALMAN FILTER INTEGRATION

With the sensor measurements described in the previous sections, we want to track the heading ψ, the heading rate $\dot{\psi}$ and the gyro bias b in the following state vector:

$$\boldsymbol{x} = \begin{bmatrix} \psi \\ \dot{\psi} \\ b \end{bmatrix} \tag{9}$$

As outlined in the previous sections, multiple different modes of integration are possible to fuse heading and heading rate measurements. This work follows the integration logic described in [17] of loosely and tightly coupled integration.
Both implementations follow the same principle of integrating the rate gyro measurement for a heading solution when no absolute measurement from the DPA is available and updating the integrated estimate when an absolute measurement becomes available. The KF runs at the same update rate as the rate gyro. The vehicle is assumed to undergo no pitch or roll changes. Expanding this work to three-dimensional attitude solutions is reserved for future work as outlined in the last section of this paper.

Loosely coupled integration

Equations (10) through (12) describe the linear measurement equations without and with available DPA measurements in the loosely coupled integration. The gyro measurement is as described by Equation (6), the sum of the heading rate $\dot{\psi}$ and bias b with Normally distributed measurement noise with covariance σ^2_{gyro}. The DPA measurement is the solution to Eq. (3) with Normally distributed measurement noise with covariance σ^2_{DPA}. Both measurements are considered conditionally independent given the state $\boldsymbol{x}$.

$$\boldsymbol{y_t} = C\boldsymbol{x_t} + \boldsymbol{v} \quad \text{with} \quad v \sim \mathcal{N}(0, R) \tag{10}$$

$$C_{noDPA} = \begin{bmatrix} 0 & 1 & 1 \end{bmatrix} \quad \text{and} \quad R_{noDPA} = \sigma^2_{gyro} \tag{11}$$

$$C_{DPA} = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \end{bmatrix} \quad \text{and} \quad R_{DPA} = \begin{bmatrix} \sigma^2_{gyro} & 0 \\ 0 & \sigma^2_{DPA} \end{bmatrix} \tag{12}$$

The gyro bias term b is therefore estimated in the KF and is taken into account in each rate measurement, as illustrated schematically in Figure 4.

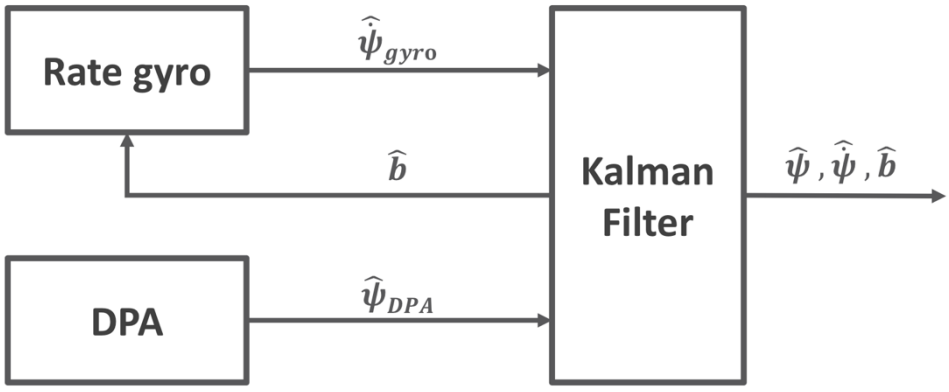

Figure 4: Loosely coupled integration with gyro bias estimation

Tightly coupled integration

Two aspects differentiate the tightly coupled from the loosely coupled integration. First, instead of a single heading measurement from the DPA, N heading measurements $\hat{\psi}_i$ are now fed directly into the KF, each a function of the measured bearing and azimuth corresponding to each individual satellite as we've shown in Eq. (4).
The second aspect is to mitigate a problem inherent to the DPA measurements. The N individual heading measurements $\widehat{\boldsymbol{\psi}}(\boldsymbol{t})$ are taken throughout an approximately 3.5 second long measurement period before becoming available every 2.56 seconds. Each measurement is a snapshot of the heading at the exact epoch it was taken. Any change in the antenna's heading between the measurement's epoch and the time it becomes available biases the measurement, similar to the blur occurring on a long exposure photo of a moving object. The corresponding exact measurement epoch is related to the measured phase difference and therefore known to the receiver. We integrate the rate gyro measurements backwards to each measurement epoch to correct the measurements for the relative heading change between the time the measurement was taken and the time it becomes available. During this short time frame the gyro measurements are an order of magnitude more accurate than the noise in $\widehat{\boldsymbol{\psi}}(\boldsymbol{t})$ and can be used for this correction without problem. The corrected heading measurements correspond to what the measurement would have been, had it been taken at the time when it becomes available to the filter. Figure 5 shows an example of the effect. In blue we show the measured gyro rate integrated backwards from the epoch the DPA measurements become available at t = 292.2 sec, indicating a heading change by almost 70 deg during the measurement cycle. The original measurements are shown in black and the corrected measurements in red, each at the time they correspond to. We notice a slope in the original measurements fairly parallel with the blue line, their bias. The corrected measurements are much closer to representing the same value, the bias has been largely removed.

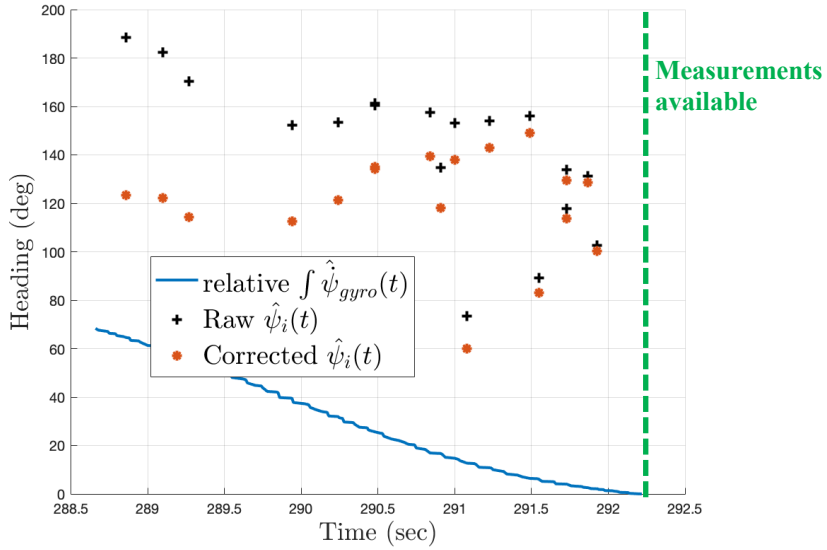

Figure 5: Heading measurement bias effect due to heading changes. The blue line indicates the heading during the measurement period relative to the measurement epoch. Black crosses indicate the original measurements, red circles the corrected values.

Analyzing test data, it becomes clear that the measurements further do not only become available with a delay but were already taken with a delay. The heading measurements obtained for the DPA correspond to a drop in the Carrier to Noise (CN0) ratio reported by the GNSS receiver [12]. The authors suspect that this CN0 value is likely passed through a low pass or moving average filter within the used ublox M8 receiver module, leading to a delay in the detected CN0 drop. We integrate the measured gyro rotation backwards by another 0.8 seconds to compensate each measurement for this delay. In Figure 6 we show the error of raw and corrected DPA measurements as well as the angular rate as measured by the gyro during phases of constant heading and rapid heading changes. We notice an increase in error whenever a heading change is present. The correction of the measurements reduces the effect but does not mitigate it completely. Heading changes lead to an increase in measurement covariance despite the correction. The faster changes starting around 280 seconds lead to a larger increase.

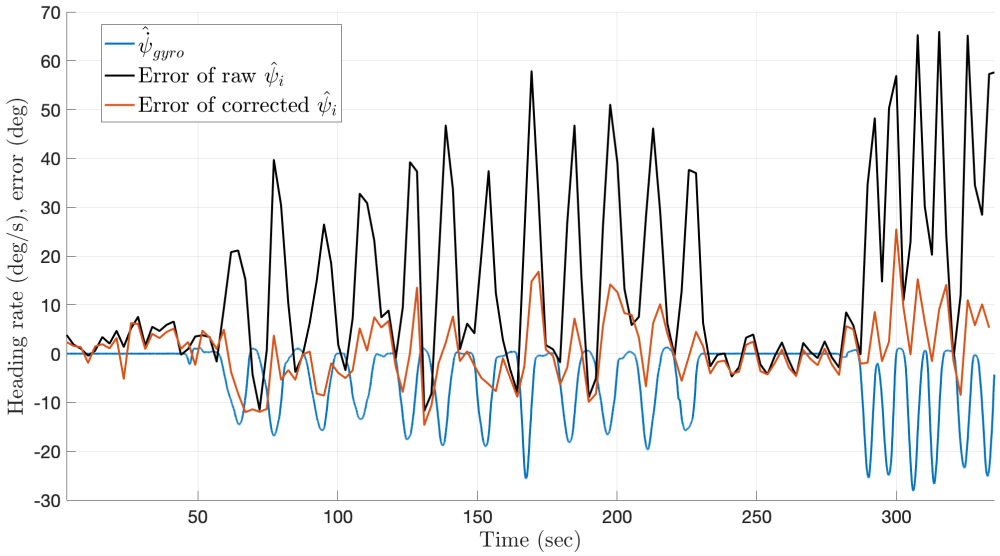

Figure 6: Heading rate and error of DPA heading measurements

The measurement equation without available DPA measurements is the same as for the loosely coupled integration and as described in (10) and (11). In case of an available DPA measurement, $\boldsymbol{y(t)} \in R^{N+1}$ with output and covariance matrices as follows:

$$C_{DPA} = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \\ & \dots & \\ 1 & 0 & 0 \end{bmatrix} \quad \text{and} \quad R_{DPA} = \begin{bmatrix} \sigma^2_{gyro} & 0 & \dots & 0 \\ 0 & \sigma^2_{\theta,1} & \dots & 0 \\ & & \dots & \\ 0 & \dots & 0 & \sigma^2_{\theta,N} \end{bmatrix} \tag{13}$$

Figure 7 illustrates the setup schematically.

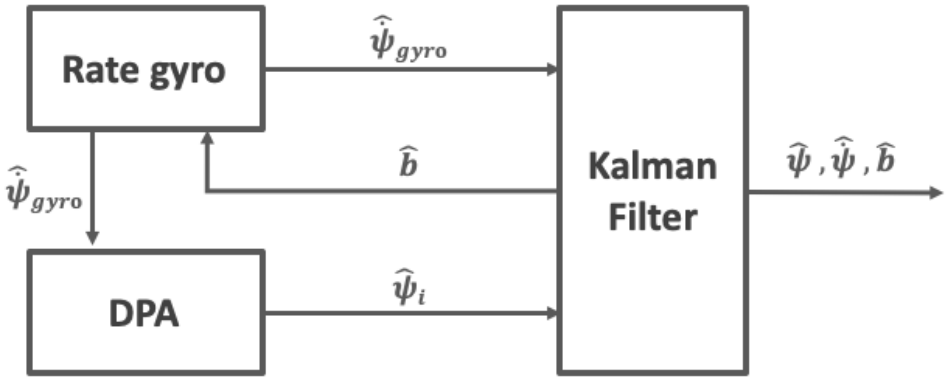

Figure 7: Tightly coupled integration with gyro bias estimation and bearing correction

TEST RESULTS

The sensor setup described in the previous sections was tested both in simulation and in driving tests mounted on a vehicle. Note that in simulations, the control input u is known to the filter, where as in driving tests it is not. Throughout the remainder of the paper, the vehicle is assumed to have negligible pitch and bank angles. Expanding this analysis to three-dimensional attitude is reserved for future work.

Dynamics

The same straight forward dynamics are assumed for the driving test and used in simulation, linearized according to the methods described in [18] for a time step δt and given by Equations (14) - (16). The simulation assumes a rigid body with inertia I about the vertical axis and applies an input torque $u_t = I \,/\, 100 \sin(t)$. The process noise w is zero-mean Gaussian with covariance Q given by (16).

$$\boldsymbol{x_{t+1}} = A\boldsymbol{x_t} + Bu_t + \boldsymbol{w} \quad \text{with} \quad \boldsymbol{w} \sim \mathcal{N}(0, Q) \tag{14}$$

$$A = \begin{bmatrix} 1 & \delta t & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 - \dfrac{\delta t}{\tau} \end{bmatrix} \quad \text{and} \quad B = \begin{bmatrix} 0.5\,\delta t^2 \,/\, I \\ \delta t \,/\, I \\ 0 \end{bmatrix} \tag{15}$$

$$Q = diag(\delta t^2 [0.0001 \quad 0.001 \quad Q_{w_b}]) \tag{16}$$

The same dynamics equation is used in the prediction steps of the Kalman Filter for both the simulation and the driving test. The only difference is that since the input is unknown in the second case, u is set to 0 in the predict steps of the KF and the process noise Q inflated.

Simulation Results

We conduct simulations to evaluate the performance differences from using gyros of varying quality. We expose the system to a small sinusoidal input $t/100$ during all simulations. The measurement delay effect described in the previous section is not modelled and the loosely coupled implementation is used in all simulations. Figure 8 shows the resulting error for each state variable and their 2σ confidence values for a low-quality industrial grade rate gyro with angle random walk of $10\ \mathrm{deg/h/\sqrt{Hz}}$ and bias instability of $100\ \mathrm{deg/h}$ over a simulation period of 1000 s. Since the system's C and R matrices are time variant, we can not solve Ricatti's equation for a steady state covariance. We can however clearly see the convergence from Figure 8, reaching 2σ bounds of around $[1.1\ \mathrm{deg} \quad 0.05\ \mathrm{deg/s} \quad 0.004\ \mathrm{deg/s}]$ for the three state variables after 1000 sec. It should not be surprising that the error bound and error values for heading ψ and heading rate $\dot{\psi}$ differ similarly as the respective measurements from DPA and rate gyro, with a significant reduction each. The heading error from this KF implementation is an order of magnitude better than the standalone DPA measurements.

Of course, this performance is strongly dependent on the grade of rate gyro chosen. But what kind of gyro leads to which performance? The Pareto curve shown in Figure 9 attempts an answer, by showing resulting 95% or 2σ values for the heading state ψ after 1000 sec for a range of gyro bias instabilities and angle random walk values. Values on the order of 1-2 degrees seem very well reachable with very little influence of the angular random walk.

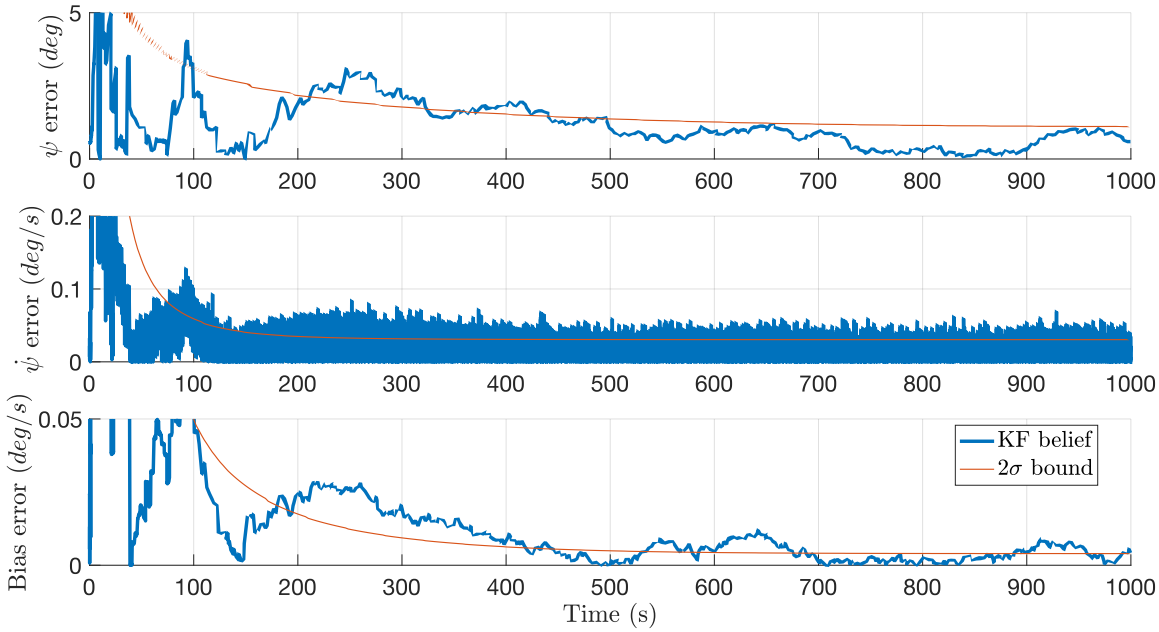

Figure 8: Errors and 2σ confidence of the loosely coupled implementation for a low-grade gyro with ARW 10 deg / h / $\sqrt{Hz}$ and bias instability 100 deg / h in simulation

It is noteworthy that improving the rate gyro beyond a bias instability of 20 deg/h does not seem to notably improve the result, a performance level that is delivered by industrial grade MEMS gyros for around $20 [19]. At this point the low-quality absolute measurement updates are the limiting factor. This again prompts a question: how much does a varying measurement quality from the DPA influence the resulting heading covariance? Figure 10 shows a second set of Pareto Curves, this time for varying accuracies and frequencies of the DPA-based measurement combined with the rate gyro used in the simulation depicted in Figure 8. The red circle indicates the current specifications of the DPA. The displayed simulation results indicate that for increasing accuracy of the measurements, update frequency becomes less critical. If the measurement covariance R_{DPA} were in the future limited to $(0.5\ \text{deg})^2$ at constant measurement update rate, the 2σ confidence bounds would be reduced down to around 0.3 deg.

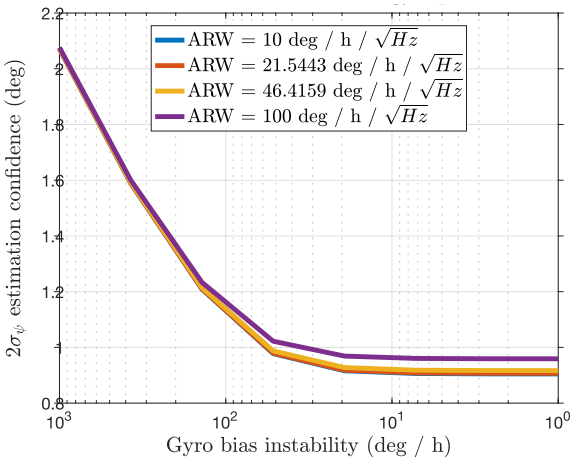

Figure 9: 2σ confidence depending on the gyro quality

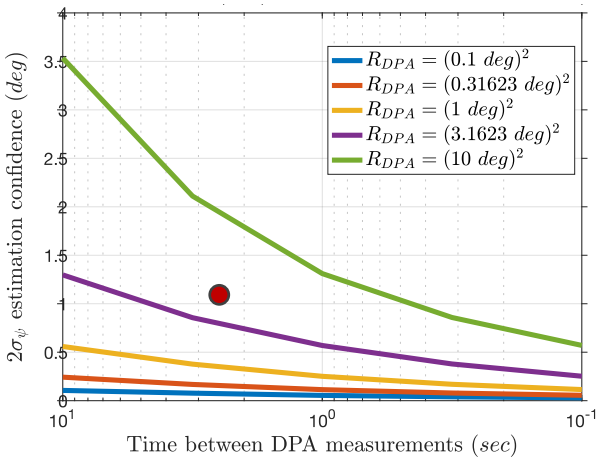

Figure 10: 2σ confidence depending on the absolute measurement quality (gyro ARW 10 deg/h/$\sqrt{Hz}$, bias instability 100 deg/h)

Driving Test

To confirm the performance of the DPA-gyro-combination determined in simulation and test the setup's performance in the case of unknown control inputs, the antenna and a STMicroelectronics iNEMO 6 Degree of Freedom (DoF) inertial measurement unit (IMU) including magnetometer has been mounted to the roof of a car for a dynamic test. Figure 11 shows the test setup. The DPA, marked by the red arrow, is mounted in the back of the car with the iNEMO IMU contained in the box underneath. From static tests the IMU rate gyros are characterized with an angle random walk of $10\ \mathrm{deg/h/\sqrt{Hz}}$ and bias instability of $5\ \mathrm{deg/h}$ according to the procedures described in [16]. The magnetometer is characterized with a zero-mean Gaussian Normally distributed error with variance of $(0.5\ \mathrm{deg})^2$. No magnetic disturbance was expected or noticed during the test, making the magnetometer a fairly reliable heading source. Three Kalman Filters are run in parallel: one based on rate gyro and magnetometer measurements from the iNEMO IMU, as well as the loosely and tightly coupled KF implementations based on rate gyro and DPA described previously. We use the KF estimate based on gyro and magnetometer measurements from the iNEMO IMU as a truth reference.

Figure 11: Driving test setup. The DPA is the aft antenna (red arrow), placed on an aluminum ground plane for better reception. The IMU is in the box underneath.

The test drive was a repeated rectangle ground track, on top of a parking structure on Stanford campus. Heading changes were applied between seconds 60 and 230 and from 280 seconds onwards as shown in Figure 6. Figure 12 shows an excerpt of the reference KF using the magnetometer as well as two DPA-based heading measurements during the drive. Both measurements are the solutions to eq. (3). The loosely coupled (l. c.) heading ones are calculated directly from the bearing measurements. The tightly coupled (t. c.) values were then corrected for the rotation induced bias as described in the previous section. Figure 13 shows an excerpt of the heading solutions from the three run Kalman Filters: magnetometer and rate gyro, loosely and tightly coupled DPA and rate gyro implementation. We clearly see a delay in the red loosely coupled measurements in Figure 12. It corresponds to the delayed availability of the DPA measurements discussed in the section on the tightly coupled implementation. We drove repeated left turns leading to a heading value reducing over time. The delay causes the raw measurements used in the loosely coupled implementation to be too large. The error leads to a bias of the estimates from the loosely coupled filter. We notice an offset of the red curve in Figure 13 compared to the blue reference showing this effect. The correction we computed in the tightly coupled filter seems to mitigate this issue to a large extend, both the measurements in Figure 12 and the filter estimate in Figure 13 (yellow curves) show much less of an error.

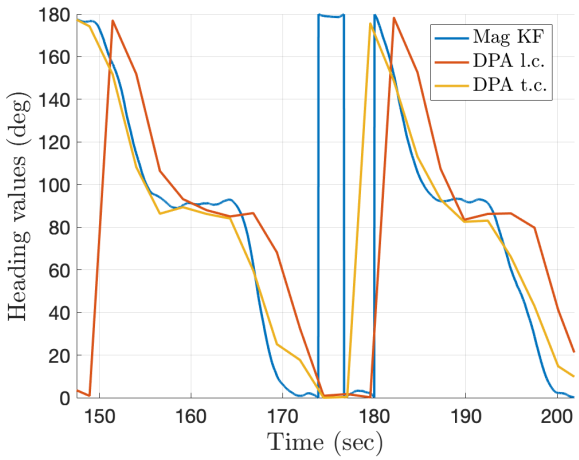

Figure 12: DPA heading resulting from bearings used in loosely and tightly coupled implementation

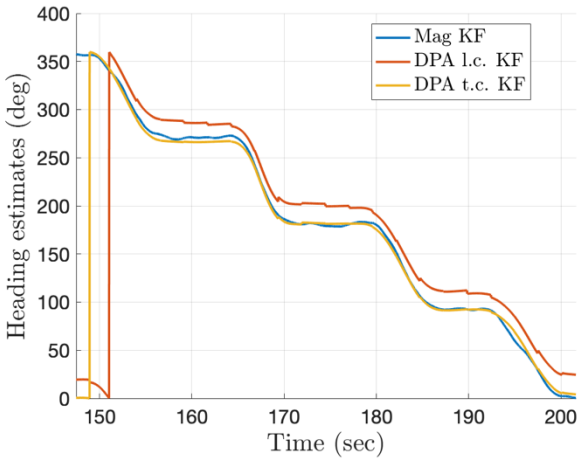

Figure 13: Excerpt of resulting KF headings: Magnetometer-based, loosely coupled and tightly coupled DPA-based implementation

We summarize the results of the entire driving test in Figure 14. We depict the 2σ error bounds of all three states, together with the same bounds obtained from the reference KF based off the magnetometer and rate gyro measurement. We consider the difference between the belief values of either DPA based KF and the reference as the estimation error ϵ and depict it as well in Figure 14. Since the input u is unknown to the filter, we inflate the process noise value for $\dot{\psi}$ to respect expected angular acceleration with 1σ of 20 deg / s^2. Based on the results shown in Figure 6 we increase the measurement covariance σ_{DPA}^2 in the loosely coupled (l. c.) implementation to $(20\text{deg})^2$ and triple the measurement standard deviation $\sigma_{\theta,i}$ used in the tightly coupled (t. c.) implementation.

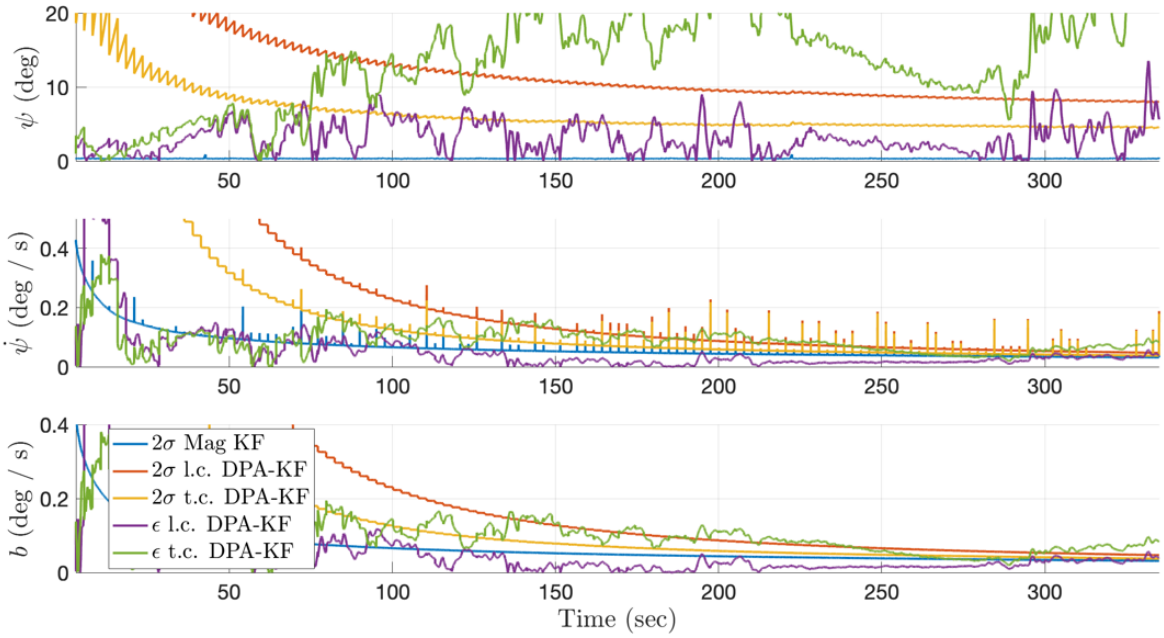

Figure 14: 2σ confidence of the state vector elements heading ψ, heading rate $\dot{\psi}$ and bias b for the magnetometer-based, loosely and tightly coupled KF implementations during the dynamic driving test

The l. c. error exceeds its 2σ bound not surprisingly. The measurement error observed in Figure 6 and Figure 12 is an offset rather than a Gaussian error and therefore leads to an erroneous belief of the Kalman Filter. The tightly coupled implementation shows significantly smaller errors thanks to the removed measurement bias. Only during the fast heading changes towards the end of the test drive does the error seem to systematically exceed the bound.

The results demonstrate that the biggest degradation of using the DPA for absolute heading reference in the Kalman Filter instead of a magnetometer is in the absolute heading value. The main reasons for this are the low frequency, low accuracy and delay-affected measurements of the DPA. The accuracy could be improved through a higher update rate or measurement quality of the antenna.
Aside from inflated errors and error bounds in ψ, spikes in the bound of $\dot{\psi}$ are noteworthy. They appear during phases of coasting on a low-quality DPA measurement while waiting for an update from the gyro. The large process noise covariance used in the filter then immediately inflates the covariance of the belief.
After $300\ \mathrm{sec}$, 2σ bounds of around $4.5\ \mathrm{deg}$ are reached using the tightly coupled implementation despite inflated process noise and measurement covariance. Judging from Figure 6 and Figure 14 these bounds are satisfied for heading changes of up to $15\ \mathrm{deg/sec}$. As mentioned before, these values are strongly dependent on the update quality and frequency of the DPA. In magnetometer-denied environments or where highly precise heading measurements are not required, the DPA-gyro combination builds a very cost- and space efficient alternative.

One issue not addressed so far is the fact that the DPA's measurements come with a 180 degree ambiguity [12]. A computed heading will therefore have the same limitation, just like the heading estimates depicted in Figure 12. If the antenna is mounted on a vehicle, this could be mitigated by for example simply using the heading closer to the GPS course over ground direction, the solution that is used in the result depicted in Figure 13. A Multi-Hypothesis Kalman Filter or Particle Filter could be employed for more traceable solutions detecting the correct of two heading values based on the vehicle's GPS velocity vector.

SUMMARY AND CONCLUSIONS

The Stanford DPA is a dual feed patch GNSS antenna that generates approximate measurements of the antenna's bearing to GNSS satellites. This paper introduces two Kalman Filter algorithms to fuse the DPA's bearing measurements with heading rate measurements from a rate gyro for heading estimation. We present simulation and dynamic driving test results that are an order of magnitude more accurate than standalone heading estimates from the DPA. We present Pareto curves that outline the system's performance for different gyro grades. We show that an industrial grade rate gyro with a bias instability of 20 deg / h is sufficient for best system performance.
This paper presents that the DPA aided by an industrial grade MEMS rate gyro through a tightly coupled Kalman Filter delivers heading with an accuracy of around 2 deg ($1\ \sigma$) for heading rates of up to $15\ \mathrm{deg/sec}$ based on experimental evaluations. We examine the DPA's measurement quality depending on heading rate and show their close correlation.
This approach already outperforms antenna arrays and gyrocompass systems in price, complexity and form factor without the limitations of devices based on the earth's magnetic field. Given their low turning rate, the heading solution presented in this paper could be an interesting alternative for vessels and aircraft, especially if operation in high latitudes or large electromagnetic interference due to e.g. electric propulsion prohibits the use of a magnetometer.

The presented results show a close correlation between the measurement accuracy and heading rate. Introducing a state-based covariance matrix that models this dependency is the first recommendation for future work. Improving the update rate and measurement quality of the DPA is then the main future task to improve the accuracy of this approach. The use of a Software Defined Radio (SDR) could alleviate or remove the DPA measurement's time delay issue completely.
A more precise (nonlinear) dynamics model of the moving platform could further be employed to aid the (Extended) Kalman Filter and lead to more precise results.

The IMU used for the driving test was a 6 Degree of Freedom (DoF) unit. A logical next step would be to expand the filter architecture to encompass all six axes to examine the DPA's performance in this setup. It could potentially create a 6DoF attitude and position measurement unit at the size of and price tag close to a GPS antenna and independent of magnetic field measurements.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the Federal Aviation Administration and the Stanford Center for Position Navigation and Time for supporting this research. We would also like to thank especially Dr. Yu-Hsuan Chen for his support during the test campaign.

DISCLAIMER

The views expressed herein are those of the authors and are not to be construed as official or any other person or organization.

REFERENCES

[1] K. Gade, "The Seven Ways to Find Heading," *J. Navig.*, vol. 69, no. 05, pp. 955–970, 2016.
[2] D. Gebre-Egziabher, R. C. Hayward, and J. D. Powell, "A low-cost GPS/inertial attitude heading reference system (AHRS) for general aviation applications," in *IEEE Position, Location and Navigation Symposium*, 1998, pp. 518–525.
[3] D. Gebre-Egziabher, R. C. Hayward, and J. D. Powell, "Design of multi-sensor attitude determination systems," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 40, no. 2, pp. 627–649, 2004.
[4] D. Gebre-Egziabher, G. H. Elkaim, J. D. Powell, and B. W. Parkinson, "A gyro-free quaternion-based attitude determination system suitable for implementation using low cost sensors," in *IEEE PLANS*, 2000, pp. 185–192.
[5] V. M. N. Passaro, A. Cuccovillo, L. Vaiani, M. De Carlo, and C. E. Campanella, "Gyroscope technology and applications: A review in the industrial perspective," *Sensors (Switzerland)*, vol. 17, no. 10, 2017.
[6] C. E. Cohen, "Attitude determination using GPS: Development of an all solid-state guidance, navigation, and control sensor for air and space vehicles based on the global positioning system," Stanford University, 1992.
[7] I. Bar-Itzhack, P. Montgomery, J. Garrick, I. Bar-Itzhack, P. Montgomery, and J. Garrick, "Algorithms for attitude determination using GPS," in *Guidance, Navigation, and Control Conference*, American Institute of Aeronautics and Astronautics, 1997.
[8] R. C. Hayward, A. Marchick, and J. D. Powell, "Single Baseline GPS Based Attitude Heading Reference System (AHRS) for Aircraft Applications," in *Proceedings of the American Control Conference*, 1999, pp. 3655–3659.
[9] B. Carey, "Moving Target," *Aviation Week & Space Technology*, pp. 26–28, Mar-2019.
[10] R. P. Kornfeld, R. J. Hansman, and J. J. Deyst, "Single-antenna GPS-based aircraft attitude determination," *Navig. J. Inst. Navig.*, vol. 45, no. 1, pp. 51–60, 1998.
[11] E. McMilin, "Single Antenna Null-Steering for Gps & Gnss Aerial Applications," Stanford University, 2016.
[12] Y.-H. Chen, S. Lo, A. Perkins, F. Rothmaier, D. M. Akos, and P. Enge, "Demonstrating Single Element Null Steering Antenna Direction Finding for Interference Detection," *Proc. 2018 Int. Tech. Meet. Inst. Navig.*, pp. 240–259, 2018.
[13] Y. Chen, F. Rothmaier, D. M. Akos, S. Lo, and P. Enge, "PCB Implementation of a Single Null-steering Antenna and its Anti-spoofing / jamming Testing," *Proc. 2017 Int. Tech. Meet. Inst. Navig.*, 2017.
[14] M. J. Kochenderfer and T. A. Wheeler, *Algorithms for Optimization*, 1st ed. Cambridge, Massachusetts: The MIT Press, 2018.
[15] D. Gebre-Egziabher, "Design and performance analysis of a low-cost aided dead reckoning navigator," Stanford University, 2004.
[16] Institute of Electrical and Electronics Engineers, *IEEE Standard Specification Format Guide and Test Procedure for Single- Axis Interferometric Fiber Optic Gyros*. 2008.
[17] C. Jekeli, *Inertial Navigation Systems with Geodetic Applications*. 2012.
[18] G. Franklin, J. D. Powell, and M. Workman, *Digital Control of Dynamic Systems*, 3rd ed. Half Moon Bay, CA: Ellis-Kagle Press, 1998.
[19] A. M. Fitzgerald, "Recent Innovations in MEMS Sensors for PNT Applications," 2017.