

Distributed Ground Sensor Fusion Based Object Tracking for Autonomous Advanced Air Mobility Operations

Thomas Lombaerts* Keerthana Kannan† Evan Kawamura‡
NASA Ames Research Center, Moffett Field, CA 94035

Chester Dolph§
NASA Langley Research Center, Langley, VA 94035

Vahram Stepanyan¶ George Gorospe|| Corey Ippolito**
NASA Ames Research Center, Moffett Field, CA 94035

Autonomous operations are a crucial aspect in the context of Advanced Air Mobility and other emerging aviation markets. In order to enable this autonomy, an accurate and detailed understanding of the positions of the various vehicles in the air is necessary. Full localization independent of on-board sensors makes the system suitable for noncooperative vehicles. This paper focuses on the object tracking part that relies on distributed ground-based sensor fusion, considering specific properties and limitations of different sensor types. Results show satisfactory performance in nominal scenarios with full coverage. Dropouts of individual sensors affect the accuracy of the tracking results, which agrees with expectations for partial coverage, when full localization is not achievable anymore. Finally, a study is performed to identify which parameters have the largest impact on the fit error.

I. Introduction

RECENTLY, there have been numerous developments in the area of advanced air taxi operations. Combined with other novel transportation concepts such as autonomous aerial cargo delivery etc., they form the Advanced Air Mobility (AAM) ecosystem that will help to meet the challenging mobility needs of an increasing population density in metropolitan areas, taking into account the limited capacity of the currently available transportation infrastructure. To realize the full vision of AAM, technology will be needed to allow a few operators to operate many vehicles (m:N). As most m:N implementations will require highly autonomous vehicles, the current state-of-the-art will have to advance for on-board autonomous functionality. One of the key areas requiring further technology development is sensor-based perception techniques. For this perception, techniques are being developed to address problems such as GPS-free Alternative Positioning, Navigation and Timing (APNT) [1]. Advanced sensing and tracking algorithms will be necessary from the ground in an urban environment, where traditional traffic tracking infrastructure such as air traffic radars will be useless. This infrastructure will be needed for traffic monitoring, traffic guidance, as well as spacing and separation in order to increase overall capacity and efficiency in the airspace. In this context, ground-based perception systems, capable of APNT, are critical to safe autonomous operations. This APNT relates to all moving objects in the sky, which can be cooperative (aircraft, AAM vehicles, etc.) as well as noncooperative (smaller drones, birds, ...). The distinction between both lies in the fact that there is some form of data exchange (ADS-B or similar) with the former category in order to coordinate path planning and separation. This exchange of information is nonexistent for the latter category, which means that other sensors will be required to detect and track these objects independently from the ground. Ground-based object detection and tracking are two essential prerequisites for the purpose of traffic monitoring, and also to provide some off-board redundancy for detect and avoid (DAA). For multi-target tracking, there is an additional challenge of observation-to-track association.

* Aerospace Research Engineer, KBR, Intelligent Systems Division, Mail Stop 269-1, AIAA Associate Fellow, email: thomas.lombaerts@nasa.gov

† Software Engineer, KBR, Intelligent Systems Division, Mail Stop 269-1, AIAA member, email: keerthana.kannan@nasa.gov

‡ Computer/GNC Engineer, Intelligent Systems Division, Mail Stop 269-1, AIAA member, email: evan.kawamura@nasa.gov

§ Aerospace Engineer, Aeronautics Systems Engineering Branch

¶ Principal Scientist, KBR, Intelligent Systems Division, Mail Stop 269-1, AIAA senior member, email: vahram.stepanyan@nasa.gov

|| Senior Research Scientist, Intelligent Systems Division, Mail Stop 269-1, email: george.e.gorospe@nasa.gov

** Aerospace Scientist, Intelligent Systems Division, Mail Stop 269-1, email: corey.a.ippolito@nasa.gov

A. Focus of this paper

This paper focuses on the single object tracking part that relies on adaptive multi-sensor fusion in a distributed ground-based sensor setup, taking into account specific properties and limitations of different sensor types. The sensor fusion is done centrally by means of a Kalman Filter, relying on the raw measurements from the network of non-collocated ground-based sensors. Some other approaches include local processing at the individual sensor or node of sensors level, after which the different local estimates are integrated by some form of consensus filter, for example using weighted averages. The sensor fusion strategy used here is adaptive, so that lower sampling rates of specific sensors are taken into account, as well as loss of signal (for example because the object is outside the field of view or the useful range for a particular sensor location and viewing angle, or because of a sensor malfunction resulting in a sensor signal dropout), and invalid data. This capability is achieved by incorporating valid flags in the weighting matrix of the Kalman filter. This paper also explores which sensor properties and/or network characteristics have a dominant impact on the accuracy of the estimate, such as the latency and the number of sensors.

B. Literature survey

As already introduced before, most literature considers local processing at the individual sensor or node of sensors level, after which the different local estimates are integrated by some form of consensus filter, for example using weighted averages. [2] considers a distributed sensor network, consisting of a set of processing nodes collecting data from sensors. The nodes communicate with each other, so that each node updates its local estimate with information received from the other nodes. [3] reviews developments in distributed sensor networks from four aspects: network structure, data processing paradigm, sensor fusion algorithm with emphasis on fault-tolerant algorithm design, and optimal sensor deployment strategy. [4] describes the solution to a real-time distributed resource allocation application involving distributed situation assessment. The solution decomposes the environment into a number of sectors, where individual sensor nodes in a sector are specialized dynamically to address different parts of the objective to track objects moving in the environment in real-time, given uncertainty and constraints on sensor loads, communication, power consumption, action characteristics, and clock synchronization. [5] investigates a binary decentralized detection problem in which a network of wireless sensors provides relevant information about the state of nature to a fusion center. The fusion center attempts to accurately reconstruct the state of nature. Optimal sensor configuration and capacity constraints are considered. The gain offered by having more sensors exceeds the benefits of getting detailed information from each sensor. The method of [6] consists of two parts. First a distributed sensor fusion algorithm integrates two different distributed consensus filters to achieve cooperative sensing among sensor nodes. This fusion algorithm has two phases. In the first phase, the weighted average consensus filter is developed, which allows each sensor node to find an estimate at each time step. In the second phase, the average consensus filter is used to allow each sensor node to find a confidence of the estimate at each time step. The final estimate is iteratively updated during the movement of the sensors via weighted average. [7] studies the problem of location detection for commercial aircraft and methods to improve location detection over any terrain the flight path traverses. The proposed techniques are based on the internet-of-things model for aircraft, where the aircraft can communicate with each other within a certain range. Distributed algorithms detect location using such methods that work effectively when the aircraft is outside the range of radar and on an oceanic route. Using the proposed methods, the precise location of all aircraft can be tracked to a higher degree.

The approach followed in this paper is centralized and does not rely on middle layer sensor nodes. Although resulting in an overall longer communication path, it removes adverse effects caused by the consensus filter.

C. Paper Structure

The structure of this paper is as follows. Sections II and III describe respectively the scenario and the relevant sensor characteristics which have been used to evaluate the performance of the developed object tracker. Section IV elaborates how these distributed sensor data are combined in a Kalman Filter based object tracker setup. Section V shows some results for four different scenarios. The first one is a flight across the bay, the second one is a similar set up but where the cameras have a reduced visual range, the third one is a landing scenario and the final one is a brief study that focuses specifically on the impact of sensor latency on the accuracy of the estimate. Finally, Section VI gives some overall conclusions and recommendations based on this work.

II. Scenarios

There are two main scenarios in this publication, both are in the San Francisco Bay Area. A longer flight from the East Bay to downtown San Francisco is the first one, where a few distributed sensors are strategically located along the flight path in order to achieve full coverage. The second scenario focuses specifically on the landing portion of the aforementioned scenario, which is near the Fifth and Mission parking structure downtown San Francisco. For this second scenario, a few distributed sensors are strategically placed at shorter range around the vertiport and along the approach path.

A. First scenario: flight across the bay

The first scenario considered in this publication to evaluate the centralized object tracking capabilities of the network of distributed ground sensors is as follows. A 6 person quadrotor eVTOL vehicle, illustrated in Fig. 1 and discussed in detail in previous work [8–10], flies from middle harbor shoreline park near Oakland (labeled MHP) over the San Francisco Bay to the parking structure near Mission and Fifth Avenue in downtown San Francisco (FMG), as illustrated in Fig. 2. The data for this scenario is obtained from a high fidelity simulation in a custom NASA flight software called ‘Reflexion’ and visualized by a commercial software called X-plane [11]. In addition to the takeoff and landing locations, several additional sensors are located along the trajectory of the vehicle, namely on the East tower of the Bay Bridge (Oakland span) (labeled BBE), Alameda Naval Station (ANS), the antenna tower on the top of Yerba Buena Island (YBA), the West tower of the Bay Bridge (BBW), and the Salesforce Tower (SFT). An overview of the general outline of the scenario and the sensor locations is given in Fig. 3. Table 1 gives detailed coordinates and viewing directions per camera location. The yellow marker is the common reference point (origin) for the local East-North-Up reference frame.



Fig. 1 6 passenger quadrotor eVTOL vehicle used for the scenario



Fig. 2 3D view of the flight trajectory overlaid over the bay area of San Francisco




Fig. 3 Map view of the flight track from Middle Harbor Shoreline Park to Fifth and Mission Garage

In this scenario, there can be either a camera or a combination of camera and radar located in each sensor location. Both have an identical field of view, which is 60 deg horizontally and 33.75 deg vertically (which makes the aspect ratio 9/16). The cameras have typically a range of 4,000 m and the radars have only half that range, namely 2,000 m. Heavier ground based radars may have a better range, but these conservative values are taken to reflect low cost sensors with more limited performance in an urban environment. For this scenario, the viewing directions of the individual cameras and radars are set up in such a way as to optimize global coverage over the trajectory of the vehicle. The coverage of the individual sensors is illustrated by means of sensor cones in Fig. 4. Cameras provide 2D info only by means of azimuth and elevation. The radars provide 3D info by adding range information to the two aforementioned sensed variables. Sec. III discusses the sensor characteristics together with the synthetic measurements of these sensors, perturbed by these disturbances.

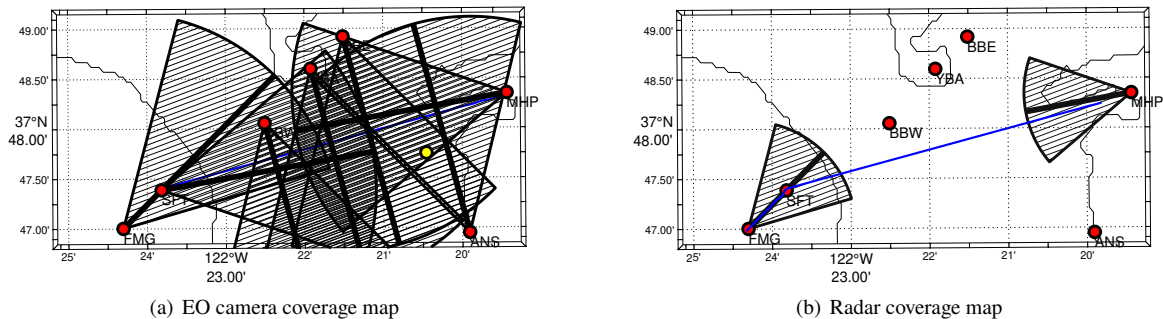



Fig. 4 Sensor coverage maps for all the distributed sensors in the scenario

Each sensor location as shown in Fig. 4 has a camera but only the departure and arrival points have cameras as well as radars.

Fig. 5 shows the number of cameras that are simultaneously observing the target at any given time. For successful triangulation, in absence of a radar measurement, at least two different camera observations are necessary. The significant overlap of coverage provides redundancy in case of sensor failures or degraded measurements, etc.

Table 1 Distributed camera locations and angles across the bay

Camera Name	Latitude (°)	Longitude (°)	Abs. Alt. (m)	Heading (°)	Elev. (°)
Mission Harbor Shoreline Park (MHP)	37.8058404	-122.3238733	4.53	260	0
Tower of Bay Bridge East Span (BBE)	37.815309	-122.358504	160	165	0
Yerba Buena Island Antenna Tower (YBA)	37.809956	-122.365332	104	165	0
Alameda Naval Station (ANS)	37.782494	-122.331785	0	345	0
Tower of Bay Bridge West Span (BBW)	37.800911	-122.375117	160	165	-9
Salesforce Tower (SFT)	37.789782	-122.396968	326	80	0
Fifth & Mission Garage (FMG)	37.783379	-122.405077	26	45	0

**Fig. 5 Number of cameras simultaneously observing the target****B. Second scenario: Landing in downtown San Francisco**

The second scenario focuses on the landing portion of the aforementioned flight over downtown San Francisco down towards the Fifth and Mission parking structure, as illustrated in Fig. 6. A few distributed sensors are strategically placed at shorter range around the vertiport and along the approach path. The rooftop of the parking structure is the only location with a collocated camera and radar (FMG), looking towards the bay along the approach path as also shown in the previous scenario in Fig. 4. In addition, other cameras are placed on the tower of the San Francisco Chronicle building (SFC) to the West, the rooftop of the San Francisco International hotel (SFI) to the South, the high rise building at Market and Fourth to the North, and the SF Marriot Marquis to the Northeast, and finally the Salesforce Tower farther to the East. Due to the orientation in the urban environment, some cameras have a much shorter range than others, as illustrated in Fig. 7. Specific location and orientation details per camera are given in Table 2.

Fig. 8 shows the number of cameras that are simultaneously observing the target at any given time. The significant overlap of coverage in short range provides more redundancy compared to the previous scenario.

Table 2 Distributed camera locations and angles around landing area, [12]

Camera Name	Lat (°)	Long (°)	Abs. Alt. (m)	Heading (°)	Elev. (°)	Range (m)
Fifth & Mission Garage (FMG)	37.783379	-122.405077	26	45	16.875	1,000
San Francisco Chronicle (SFC)	37.782546	-122.406520	34	45	16.875	1,000
San Francisco Intercontinental (SFI)	37.781944	-122.404722	110	345	-15	300
Marriot Marquis Hotel (MMH)	37.785278	-122.404167	135	165	-45	300
Market & Fourth Street (MFS)	37.784684	-122.405732	76	165	-10	300
Salesforce Tower (SFT)	37.789782	-122.396968	326	225	-10	1,200




Fig. 6 3D view of the landing trajectory overlaid over the financial district area of San Francisco, [12]




Fig. 7 Camera coverage maps for the landing scenario




Fig. 8 Number of cameras simultaneously observing the target for the landing scenario

III. Sensor Characteristics

There are two types of sensors that are considered in the aforementioned scenarios, namely electro-optical camera's and radar. Each sensor type has different sensed variables, together with their respective sensor characteristics, such as limitations on field of view and useful range, noise, latency, update rate and resolution. Some characteristic values from the literature of state-of-the-art sensor hardware [13] are given in Table 3.

Table 3 Sensor characteristics

sensor	EO camera	radar
sensed variables	azimuth [deg] elevation [deg]	azimuth [deg] elevation [deg] range [m]
limits on field of view (16/9)	az: 60deg el: 33.75deg	az: 60deg el: 33.75deg
max useful range	4,000m	2,000m
min useful range	0m	20m
noise sigma	$9.1 \cdot 10^{-3}$ deg	az: 1deg el: 3deg range: 3.25m
latency	0.1s	0.1s
update rate	30fps	10Hz
resolution	$1.3 \cdot 10^{-3}$ deg	az/el: 1deg range: 3.25m

A. Sensor readings for flight across the bay

Fig. 9 shows the synthetic measurements for the flight across the bay, provided by the different distributed cameras and radars with the disturbances included, such as noise and resolution. Noise clearly has a larger impact on the radar measurements.

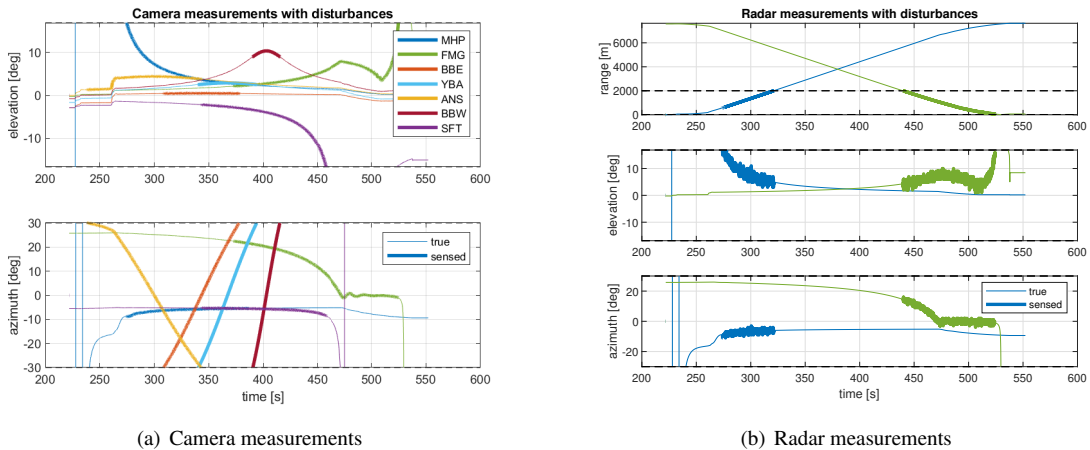


Fig. 9 Simulated camera and radar measurements with disturbances for across the bay scenario

B. Sensor readings for landing in downtown San Francisco

Fig. 10 shows the synthetic measurements for different distributed cameras only in the landing scenario. The long range cameras in FMG and SFC observe the vehicle first between $t=490$ s and $t=500$ s. The shorter range cameras perpendicular on the flight path observe the vehicle starting around $t=510$ s. Touchdown happens near $t=550$ s, after which all observations remain constant over time.

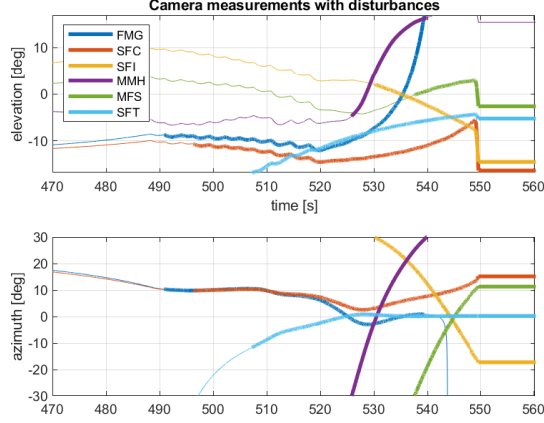


Fig. 10 Simulated camera measurements with disturbances for landing scenario

IV. Object Tracker Setup

The actual object tracker is based on a conventional Kalman Filter as elaborated in detail in Ref. [8, 14]. The filter gain calculation is modified in order to incorporate valid flags for the different sensor measurements, which makes this method adaptive and robust against invalid data, loss of signal between the different distributed sensors, outliers and lower sampling rates. This new implementation integrates the measurements from EO cameras and radars in the different locations, whenever each of them is available. This modular architecture is fairly straightforward to extend so that other sensors and sensor locations can be incorporated as well. Fig. 11 illustrates the general overview how the Kalman Filter is set up. The prediction step assumes constant velocity for calculating the predicted state $\hat{\mathbf{x}}_k$ based on the previous state \mathbf{x}_{k-1} . The state consists of six elements, namely the three position coordinates and the three velocity coordinates, all in ENU coordinates (East–North–Up) with respect to a predefined fixed referenced point. Predicted measurements $\hat{\mathbf{z}}_{k,i}$ are calculated for each sensor location i , which are a function of the predicted state $\mathbf{h}_i(\hat{\mathbf{x}}_k)$ and the sensor location coordinates. These predicted measurements consist of the measured variables that are provided by the relevant sensors for each location, namely range, azimuth and elevation from the radar (when available) as well as azimuth and elevation from the EO camera. By comparing the actual measurement $\mathbf{z}_{k,i}$ with the predicted measurement $\hat{\mathbf{z}}_{k,i}$ for each sensor location i , one can calculate the innovation $\varepsilon_i = \mathbf{z}_{k,i} - \hat{\mathbf{z}}_{k,i} = \mathbf{z}_{k,i} - \mathbf{h}_i(\hat{\mathbf{x}}_k)$. Main drivers for the innovation are errors in the predicted state (due to the constant speed assumption) and sensor disturbances. Next, the information from all individual sensor locations and sensors is stacked together in one large innovation array ε . As last but not one, the Kalman gain matrix \mathbf{W} is calculated by means of the predicted covariance $\hat{\mathbf{P}}_k$, measurement matrix $\mathbf{H} = \frac{d\mathbf{h}(\mathbf{x})}{d\mathbf{x}}$, and the innovation covariance matrix \mathbf{S} . This Kalman gain matrix \mathbf{W} is weighted with valid flags to deal appropriately with the sensor measurements as explained previously. The corrected state \mathbf{x}_k is calculated by means of predicted state $\hat{\mathbf{x}}_k$, the innovation ε and the Kalman gain matrix \mathbf{W} via the following equation: $\mathbf{x}_k = \hat{\mathbf{x}}_k + \mathbf{W} \cdot \varepsilon$. Finally, the corrected covariance \mathbf{P}_k is calculated. The complete procedure is described in more detail below.

A. State Prediction $\hat{\mathbf{x}}_k$

The one step ahead prediction of the state vector is split up as follows:

$$\hat{\mathbf{x}}_k = \begin{bmatrix} \hat{\mathbf{x}}_{k_1} \\ \hat{\mathbf{x}}_{k_2} \end{bmatrix} \quad (1)$$

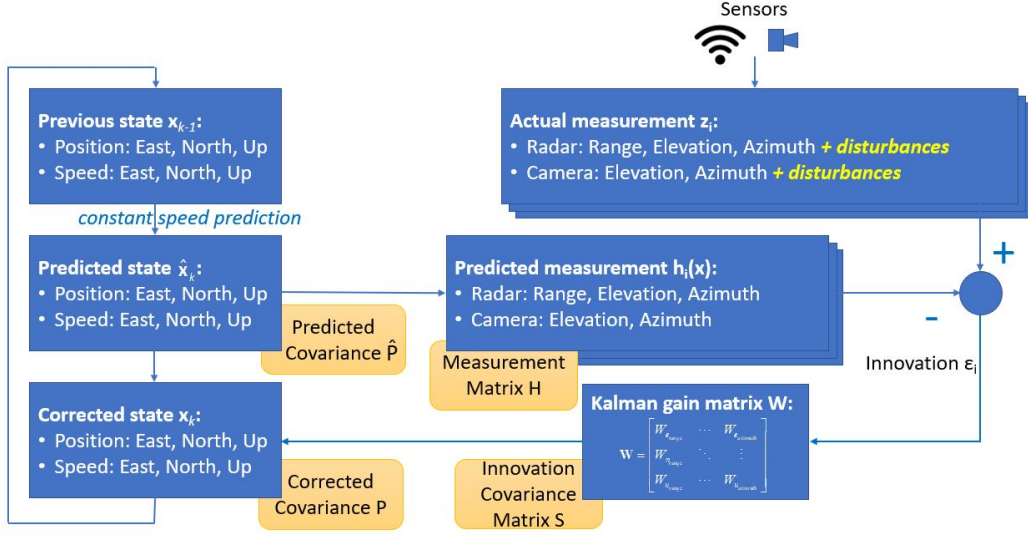


Fig. 11 General overview of the Distributed Sensor Kalman Filter setup

With the position in the upper half of the vector:

$$\hat{\mathbf{x}}_{k_1} = \begin{bmatrix} \hat{E} \\ N \\ U \end{bmatrix}_k = \begin{bmatrix} E \\ N \\ U \end{bmatrix}_{k-1} + \Delta t \cdot \begin{bmatrix} \dot{E} \\ \dot{N} \\ \dot{U} \end{bmatrix}_{0_{k-1}} \quad (2)$$

And the velocity in the lower half of the vector:

$$\hat{\mathbf{x}}_{k_2} = \begin{bmatrix} \hat{E} \\ N \\ U \end{bmatrix}_k = \begin{bmatrix} \dot{E} \\ \dot{N} \\ \dot{U} \end{bmatrix}_{0_{k-1}} \quad (3)$$

Note that the coordinates are East–North–Up with respect to a predefined fixed reference point and that the velocity is assumed to be close to constant. Moreover:

$$\begin{bmatrix} \dot{E} \\ \dot{N} \\ \dot{U} \end{bmatrix}_{0_{k-1}} = \mathbf{L}^\top \cdot \begin{bmatrix} \dot{E} \\ \dot{N} \\ \dot{U} \end{bmatrix}_{T_{k-1}} \quad (4)$$

where \mathbf{L}^\top is the rotation matrix from the target ENU-frame to the reference point ENU-frame, as elaborated in the Appendix. Next step is to rotate from ENU to ECEF reference frame:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix}_{0_{\text{ECEF}}} = \mathbf{T} \begin{bmatrix} E \\ N \\ U \end{bmatrix}_{0_{k-1}} \quad (5)$$

where \mathbf{T} is the rotation matrix from the ENU reference frame to the ECEF reference frame, as defined in the Appendix. The predicted target position is converted straight from ECEF to WGS-84 (i.e. latitude, longitude and altitude) via the transformation elaborated in the appendix, resulting in the predicted target position coordinates $\left[\phi_t \quad \lambda_t \quad h_t \right]_{\text{pred}}^\top$.

B. State Error Covariance Prediction $\hat{\mathbf{P}}_k$

$$\hat{\mathbf{P}}_k = \mathbf{F}\mathbf{P}_{k-1}\mathbf{F}^\top + \mathbf{Q} \quad (6)$$

with:

$$\mathbf{F} = \left[\begin{array}{ccc|ccc} 1 & 0 & 0 & \Delta t & 0 & 0 \\ 0 & 1 & 0 & 0 & \Delta t & 0 \\ 0 & 0 & 1 & 0 & 0 & \Delta t \\ \hline & & \mathbf{0}_{3 \times 3} & 1 & 0 & 0 \\ & & & 0 & 1 & 0 \\ & & & 0 & 0 & 1 \end{array} \right] \quad \mathbf{Q} = \mathbf{G}\mathbf{\Gamma}\mathbf{G}^\top \quad (7)$$

and:

$$\mathbf{G} = \left[\begin{array}{ccc|ccc} 0.5\Delta t^2 & 0 & 0 & & & \\ 0 & 0.5\Delta t^2 & 0 & & & \\ 0 & 0 & 0.5\Delta t^2 & & & \\ \hline \Delta t & 0 & 0 & & & \\ 0 & \Delta t & 0 & & & \\ 0 & 0 & \Delta t & & & \end{array} \right] \quad \mathbf{\Gamma} = \begin{bmatrix} \sigma_e^2 & 0 & 0 \\ 0 & \sigma_n^2 & 0 \\ 0 & 0 & \sigma_u^2 \end{bmatrix} \quad (8)$$

C. Predicted Measurement Calculation $\hat{\mathbf{z}}_k$

The total predicted measurement vector $\hat{\mathbf{z}}_k$ consists of the predicted measurement vectors $\hat{\mathbf{z}}_{k,i}$ of each individual sensor location i . The predicted measurement vectors of EO camera with or without radar consist of range ρ (if radar is included), elevation α , and azimuth γ , and are described as follows:

$$\hat{\mathbf{z}}_{k,i} = \left[\hat{\mathbf{z}}_{k,i,\text{radar}} \quad \hat{\mathbf{z}}_{k,i,\text{EOcam}} \right]^\top \quad (9)$$

$$\hat{\mathbf{z}}_{k,i,\text{radar}} = \left[\hat{\rho}_i \quad \hat{\alpha}_i \quad \hat{\gamma}_i \right]^\top \quad (10)$$

$$\hat{\mathbf{z}}_{k,i,\text{EOcam}} = \left[\hat{\alpha}_i \quad \hat{\gamma}_i \right]^\top \quad (11)$$

since the EO cameras in this scenario do not measure distance information due to the lack of stereo view and depth perception. The individual measured variables are defined as follows:

$$\hat{\rho}_i = \sqrt{\hat{X}_{S_i}^2 + \hat{Y}_{S_i}^2 + \hat{Z}_{S_i}^2} = \sqrt{\hat{E}_i^2 + \hat{N}_i^2 + \hat{U}_i^2} \quad (12)$$

$$\hat{\alpha}_i = \arctan \frac{-\hat{Z}_{S_i}}{\hat{\rho}_{i,\text{hor}}} \quad (13)$$

$$\hat{\gamma}_i = \arctan \frac{\hat{Y}_{S_i}}{\hat{X}_{S_i}} \quad (14)$$

where:

$$\hat{\rho}_{i,\text{hor}} = \sqrt{\hat{X}_{S_i}^2 + \hat{Y}_{S_i}^2} \quad (15)$$

and:

$$\begin{bmatrix} \hat{X}_{S_i} \\ \hat{Y}_{S_i} \\ \hat{Z}_{S_i} \end{bmatrix} = \mathbf{M}_i^\top \cdot \begin{bmatrix} \hat{E}_i \\ \hat{N}_i \\ \hat{U}_i \end{bmatrix}_k \quad (16)$$

where \mathbf{M}_i^\top is the rotation matrix from the i^{th} sensor-based ENU reference frame to the sensor fixed reference frame of the i^{th} sensor, as defined in the Appendix. It rotates the East-North-Up coordinates for each sensor location towards the sensor orientation (viewing direction of camera (and radar if applicable)). These local ENU _{i} coordinates are calculated

from the aforementioned predicted state in WGS-84 coordinates $\left[\phi_t \quad \lambda_t \quad h_t\right]_{\text{pred}}^T$ via the transformation elaborated in the appendix, resulting in the predicted target position coordinates in ECEF, and a subsequent rotation from ECEF to the local sensor's ENU as follows:

$$\begin{bmatrix} \hat{E}_i \\ \hat{N}_i \\ \hat{U}_i \end{bmatrix}_k = \mathbf{T}_i^T \begin{bmatrix} \hat{X} \\ \hat{Y} \\ \hat{Z} \end{bmatrix}_{0_{\text{ECEF}}} \quad (17)$$

where \mathbf{T}_i^T is the rotation matrix from the ECEF reference frame to the local sensor ENU reference frame, as defined in the Appendix.

D. Kalman Filter Gain W

The following steps are performed in order to calculate the gain of the filter:

- 1) Calculate measurement matrix \mathbf{H}
- 2) Calculate innovation covariance matrix \mathbf{S}
- 3) Calculate gain matrix \mathbf{W}

Each step will be elaborated in detail next:

1. Measurement matrix \mathbf{H}

The measurement matrix \mathbf{H} is basically the Jacobian of the predicted measurement vector:

$$\hat{\mathbf{z}}_k = \left[\hat{\mathbf{z}}_{k,1} \quad \dots \quad \hat{\mathbf{z}}_{k,i} \quad \dots \quad \hat{\mathbf{z}}_{k,N} \right]^T$$

with N the number of sensors and $\hat{\mathbf{z}}_{k,i} = \left[\hat{\mathbf{z}}_{k,i_{\text{radar}}} \quad \hat{\mathbf{z}}_{k,i_{\text{EOcam}}} \right]^T = \left[\hat{\mathbf{h}}_{\text{radar}_i}(\hat{\mathbf{x}}_k) \quad \hat{\mathbf{h}}_{\text{EOcam}_i}(\hat{\mathbf{x}}_k) \right]^T$ with respect to the state vector $\hat{\mathbf{x}}_k = \left[\hat{\mathbf{x}}_{k_1} \quad \hat{\mathbf{x}}_{k_2} \right]^T$, with \mathbf{x}_{k_1} being the position components and \mathbf{x}_{k_2} the velocity components, both in the ENU reference frame. The measurement matrix \mathbf{H} basically consists of the blocks $\mathbf{H} = \left[\mathbf{H}_1 \quad \dots \quad \mathbf{H}_i \quad \dots \quad \mathbf{H}_N \right]^T$ with N the number of sensors and decomposing each \mathbf{H}_i in 4 blocks results in the following structure:

$$\mathbf{H}_i(\mathbf{x}_k) = \frac{d\hat{\mathbf{h}}_i(\hat{\mathbf{x}}_k)}{d\hat{\mathbf{x}}_k} = \begin{bmatrix} \frac{d\hat{\mathbf{h}}_{\text{radar}_i}(\hat{\mathbf{x}}_{k_1})}{d\hat{\mathbf{x}}_{k_1}} & \frac{d\hat{\mathbf{h}}_{\text{radar}_i}(\hat{\mathbf{x}}_{k_1})}{d\hat{\mathbf{x}}_{k_2}} = \mathbf{0}_{3 \times 3} \\ \frac{d\hat{\mathbf{h}}_{\text{EOcam}_i}(\hat{\mathbf{x}}_{k_1})}{d\hat{\mathbf{x}}_{k_1}} & \frac{d\hat{\mathbf{h}}_{\text{EOcam}_i}(\hat{\mathbf{x}}_{k_1})}{d\hat{\mathbf{x}}_{k_2}} = \mathbf{0}_{2 \times 3} \end{bmatrix} \quad (18)$$

The measurements depend only on the position components and thus \mathbf{x}_{k_1} . Both nonzero blocks in Eq. (18) are elaborated next:

$$\frac{d\hat{\mathbf{h}}_{\text{radar}_i}(\hat{\mathbf{x}}_{k_1})}{d\hat{\mathbf{x}}_{k_1}} = \begin{bmatrix} \frac{d\hat{\rho}_i}{dE_i} & \frac{d\hat{\rho}_i}{dN_i} & \frac{d\hat{\rho}_i}{dU_i} \\ \frac{d\hat{\alpha}_i}{dE_i} & \frac{d\hat{\alpha}_i}{dN_i} & \frac{d\hat{\alpha}_i}{dU_i} \\ \frac{d\hat{\gamma}_i}{dE_i} & \frac{d\hat{\gamma}_i}{dN_i} & \frac{d\hat{\gamma}_i}{dU_i} \end{bmatrix} = \begin{bmatrix} \frac{d\hat{\rho}_i}{dX_{S_i}} & \frac{d\hat{\rho}_i}{dY_{S_i}} & \frac{d\hat{\rho}_i}{dZ_{S_i}} \\ \frac{d\hat{\alpha}_i}{dX_{S_i}} & \frac{d\hat{\alpha}_i}{dY_{S_i}} & \frac{d\hat{\alpha}_i}{dZ_{S_i}} \\ \frac{d\hat{\gamma}_i}{dX_{S_i}} & \frac{d\hat{\gamma}_i}{dY_{S_i}} & \frac{d\hat{\gamma}_i}{dZ_{S_i}} \end{bmatrix} \cdot \mathbf{M}^T = \begin{bmatrix} \frac{\hat{X}_{S_i}}{\hat{\rho}_0} & \frac{\hat{Y}_{S_i}}{\hat{\rho}_0} & \frac{\hat{Z}_{S_i}}{\hat{\rho}_0} \\ \frac{\hat{X}_{S_i} \cdot \hat{Z}_{S_i}}{\hat{\rho}_{\text{hor}_0} \hat{\rho}_0^2} & \frac{\hat{Y}_{S_i} \cdot \hat{Z}_{S_i}}{\hat{\rho}_{\text{hor}_0} \hat{\rho}_0^2} & \frac{-\hat{\rho}_{i_{\text{hor}}}}{\hat{\rho}_0^2} \\ \frac{-\hat{Y}_{S_i}}{\hat{\rho}_{\text{hor}_0}^2} & \frac{\hat{X}_{S_i}}{\hat{\rho}_{\text{hor}_0}^2} & 0 \end{bmatrix} \cdot \mathbf{M}^T \quad (19)$$

The denominators $\hat{\rho}_0$ and $\hat{\rho}_{\text{hor}_0}$ in Eq. (19) are singularity protected as follows:

$$\hat{\rho}_0 = \max(\hat{\rho}_i, \varepsilon_\rho) \quad (20)$$

$$\hat{\rho}_{\text{hor}_0} = \max(\hat{\rho}_{i_{\text{hor}}}, \varepsilon_{\rho_{\text{hor}}}) \quad (21)$$

where ε_ρ and $\varepsilon_{\rho_{\text{hor}}}$ are fairly small positive values to prevent singularities.

$$\frac{d\hat{\mathbf{h}}_{\text{EOcam}_i}(\hat{\mathbf{x}}_{k_1})}{d\hat{\mathbf{x}}_{k_1}} = \begin{bmatrix} \frac{d\hat{\alpha}_i}{dE_i} & \frac{d\hat{\alpha}_i}{dN_i} & \frac{d\hat{\alpha}_i}{dU_i} \\ \frac{d\hat{\gamma}_i}{dE_i} & \frac{d\hat{\gamma}_i}{dN_i} & \frac{d\hat{\gamma}_i}{dU_i} \end{bmatrix} = \begin{bmatrix} \frac{d\hat{\alpha}_i}{dX_{S_i}} & \frac{d\hat{\alpha}_i}{dY_{S_i}} & \frac{d\hat{\alpha}_i}{dZ_{S_i}} \\ \frac{d\hat{\gamma}_i}{dX_{S_i}} & \frac{d\hat{\gamma}_i}{dY_{S_i}} & \frac{d\hat{\gamma}_i}{dZ_{S_i}} \end{bmatrix} \cdot \mathbf{M}^T = \begin{bmatrix} \frac{\hat{X}_{S_i} \cdot \hat{Z}_{S_i}}{\hat{\rho}_{\text{hor}_0} \hat{\rho}_0^2} & \frac{\hat{Y}_{S_i} \cdot \hat{Z}_{S_i}}{\hat{\rho}_{\text{hor}_0} \hat{\rho}_0^2} & \frac{-\hat{\rho}_{i_{\text{hor}}}}{\hat{\rho}_0^2} \\ \frac{-\hat{Y}_{S_i}}{\hat{\rho}_{\text{hor}_0}^2} & \frac{\hat{X}_{S_i}}{\hat{\rho}_{\text{hor}_0}^2} & 0 \end{bmatrix} \cdot \mathbf{M}^T \quad (22)$$

2. Innovation covariance matrix \mathbf{S}

The innovation covariance matrix is calculated as follows:

$$\mathbf{S} = \mathbf{H}\hat{\mathbf{P}}_k\mathbf{H}^T + \mathbf{R}_{\text{sensors}} \quad (23)$$

where: $\mathbf{R}_{\text{sensors}} = \text{diag} [\mathbf{R}_{\text{sensor}_1} \dots \mathbf{R}_{\text{sensor}_i} \dots \mathbf{R}_{\text{sensor}_N}]$ with $\mathbf{R}_{\text{sensor}_i} = \text{diag} [\sigma_{\rho_{\text{radar}}}^2 \ \sigma_{\alpha_{\text{radar}}}^2 \ \sigma_{\gamma_{\text{radar}}}^2 \ \sigma_{\alpha_{\text{EOcam}}}^2 \ \sigma_{\gamma_{\text{EOcam}}}^2]$. A diagonal matrix of sensor valid flags is defined similarly: $\mathbf{M}_{\text{sensors}_{i_{\text{valid}}}} = \text{diag} [\mu_{\text{radar}} \ \mu_{\text{radar}} \ \mu_{\text{radar}} \ \mu_{\text{EOcam}} \ \mu_{\text{EOcam}}]$. These valid flags μ_{\bullet} are one for valid sensor readings or zero for invalid ones. These flags allow to take into account lower sampling rates for specific sensors, signal dropouts (e.g. object is outside field of view), sensor malfunctions, invalid signals, etc. This setup assumes that there is a monitoring algorithm in place that drives the valid flag values accordingly. This valid flag matrix is needed for calculating the inverse of the innovation covariance matrix:

$$\mathbf{S}_{\text{valid}}^{-1} = \mathbf{M}_{\text{sensors}_{\text{valid}}} \cdot \mathbf{S}^{-1} \cdot \mathbf{M}_{\text{sensors}_{\text{valid}}} \quad (24)$$

3. Gain matrix \mathbf{W}

The Kalman filter gain matrix is calculated with all the aforementioned information as follows:

$$\mathbf{W} = \hat{\mathbf{P}}_k\mathbf{H}^T\mathbf{S}_{\text{valid}}^{-1} \quad (25)$$

$$\mathbf{W}_{\text{valid}} = \mathbf{W}\mathbf{M}_{\text{sensors}_{\text{valid}}} \quad (26)$$

Invalid flags in $\mathbf{M}_{\text{sensors}_{\text{valid}}}$ will result in zero columns in $\mathbf{W}_{\text{valid}}$. As a consequence, that sensor information will not be used for the correction step. This means that the prediction step will have a higher relative importance for the states that are the main drivers of these invalid measurements, except when there are other valid sensor measurements that contain significant information of these states. When valid measurements are sparse, the Kalman Filter will extrapolate a trajectory that will gradually diverge from reality. This growing uncertainty is indicated by the rank and/or condition number of the observability matrix, as well as by the standard deviation of the estimates. More information in Sec. V.B.

E. State Update \mathbf{x}_k

The state update step relies on the total sensor measurements vector covering all sensor locations, which will give only measurement values for the locations where the tracked object is in the field of view. This is illustrated in Fig. 12. The vehicle is flying from East to West in Fig. 12(a). The yellow circles correspond to the time instants for which vertical dashed lines are shown in Fig. 12(b). The corresponding shape of the measurement vector for these specific time samples is shown in Tab. 4. This table shows that the total sensor measurements vector is only partially populated where observations are available. Throughout the flight, the non-empty elements move from the left hand side of the vector to the right hand side.




Fig. 12 Sample measurements at discrete time instants along the track

time	MHP		ANS		BBE		YBA		BBW		SFT		FMG	
	α	γ	α	γ	α	γ	α	γ	α	γ	α	γ	α	γ
$t = 300s$	0	0	0	0	0	0	0	0	0	0	0	0
$t = 325s$	0	0	0	0	0	0	0	0	0	0	0	0
$t = 350s$	0	0	0	0	0	0	0	0	0	0
$t = 375s$	0	0	0	0	0	0	0	0	0	0	0	0
$t = 400s$	0	0	0	0	0	0	0	0	0	0	0	0
$t = 450s$	0	0	0	0	0	0	0	0	0	0

Table 4 Measurement vector values for specific time samples

The innovation is calculated first, which is the difference between the predicted measurements $\hat{\mathbf{z}}_k$ and the actual sensor measurements \mathbf{z}_k^* , including sensor limitations and perturbations:

$$\boldsymbol{\varepsilon} = \mathbf{z}_k - \hat{\mathbf{z}}_k = \mathbf{z}_k - \mathbf{h}(\hat{\mathbf{x}}_k) \quad (27)$$

This innovation $\boldsymbol{\varepsilon}$ is then used to calculate the state update by means of the previously computed Kalman Gain matrix:

$$\mathbf{x}_k = \hat{\mathbf{x}}_k + \mathbf{W}_{\text{valid}} \cdot \boldsymbol{\varepsilon} \quad (28)$$

F. Geodetic Coordinates Calculation

Next step is calculating the absolute position of the target in WGS-84 coordinates. This is achieved by combining the known sensor position and the relative position of the target with respect to the sensor position, which is the first half of the corrected state vector \mathbf{x}_{k_1} . Since both are defined in different ENU reference frames, they are first converted to the ECEF (earth centered earth fixed) reference frame, then added together and converted back to WGS-84.

The sensor position is converted straight from WGS-84 (i.e. latitude, longitude and altitude) to ECEF via the transformation elaborated in the appendix, resulting in the position coordinates $[X \ Y \ Z]_{\text{ECEF}}^T$. The relative position of the target with respect to the sensor position in ECEF coordinates is calculated as follows:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix}_{(S \rightarrow T)_{\text{ECEF}}} = \mathbf{T} \begin{bmatrix} E \\ N \\ U \end{bmatrix}_k = \mathbf{T} \mathbf{x}_{k_1} \quad (29)$$

where \mathbf{T} is the rotation matrix from the ENU reference frame to the ECEF reference frame, as defined in the Appendix.

The absolute position of the target in ECEF coordinates is then calculated as follows:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix}_{T_{\text{ECEF}}} = \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}_{\text{ECEF}} + \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}_{(S \rightarrow T)_{\text{ECEF}}} \quad (30)$$

As the last step, this absolute target position is converted back to WGS-84 (i.e. latitude, longitude and altitude) via the reverse transformation elaborated in the Appendix.

G. Velocity Update

The velocity components of the target, which are the second half of the corrected state vector \mathbf{x}_{k_2} are defined with respect to the sensor ENU reference frame, these are transformed to the target ENU reference frame as follows:

$$\begin{bmatrix} E \\ N \\ U \end{bmatrix}_{k_T} = \mathbf{L}^T \begin{bmatrix} E \\ N \\ U \end{bmatrix}_k = \mathbf{L}^T \mathbf{x}_{k_2} \quad (31)$$

*In the numerical calculations, an 'empty' measurement scalar in \mathbf{z}_k is usually represented by a 'NaN'. In order to prevent numerical issues these have to be replaced by zeros together with their respective valid flag being set to zero in $\mathbf{W}_{\text{valid}}$ so that they are ignored.

H. State Error Covariance Update \mathbf{P}_k

By first defining:

$$\mathbf{W}_H = \mathbf{I}_{6 \times 6} - \mathbf{W}_{\text{valid}}\mathbf{H} \quad (32)$$

where $\mathbf{I}_{6 \times 6}$ is a 6-by-6 identity matrix, the state error covariance matrix is then updated in the reference point reference frame:

$$\mathbf{P}_k = \mathbf{W}_H \hat{\mathbf{P}}_k \mathbf{W}_H^T + \mathbf{W}_{\text{valid}} \mathbf{R}_{\text{sensors}} \mathbf{W}_{\text{valid}}^T \quad (33)$$

which is finally converted to the target reference frame via the aforementioned rotation matrix \mathbf{L} :

$$\mathbf{P}_{k_T} = \begin{bmatrix} \mathbf{L} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{L} \end{bmatrix} \mathbf{P}_k \quad (34)$$

V. Results and Discussion

Four different sets of results are included here. First, the nominal results for a flight across the bay are presented. These are followed by a failure case of the same scenario where all the ground based optical sensors are impacted by reduced visibility. The third set focuses on the landing scenario. Finally, an in-depth study is performed to quantify the impact of the sensor latency on the accuracy of the estimates.

A. Nominal Results for Flight Across the Bay

Fig. 13 shows a map view comparison of the estimated track with the true path, along with the locations of the cameras and their viewing angles. A three dimensional comparison of the estimated track and the true path in East-North-Up coordinates is given in Fig. 14. Both figures confirm a good fit of the estimate with the true path.




Fig. 13 Map view comparison of estimate and true path

Fig. 15 shows the fit errors component-wise in two different reference frames, namely east-north-up and World Geodetic System. It shows the importance of the reference frame, as the WGS figure shows minimal errors in three dimensions, but the impact on the ENU reference frame is significant as that fit error is several meters per dimension. The initial spike around $t=275s$ is caused by the convergence of the estimate. Around $t=460s$ is the time instant when the vehicle leaves the field of view of the Salesforce tower camera, leaving only the camera and radar at the vertiport FMG for tracking. This results in a significant increase in noise, since the single 2D camera cannot fully compensate for the 3D noise from the radar measurements. The runaway around $t=525s$ occurs when the vehicle leaves the field of view of the sensors at the vertiport FMG.

Fig. 16 shows the standard deviations for the estimated position and velocity components in the ENU reference frame. From top to bottom one can observe the standard deviations for the predicted position, the corrected position and the corrected velocity estimate, with each graph depicting the three components east north and up. Besides showing similar trends for the aforementioned time instants $t=460s$ and $t=525s$, another significant change is observed around




Fig. 14 Three dimensional comparison of estimate and true path in East-North-Up coordinates

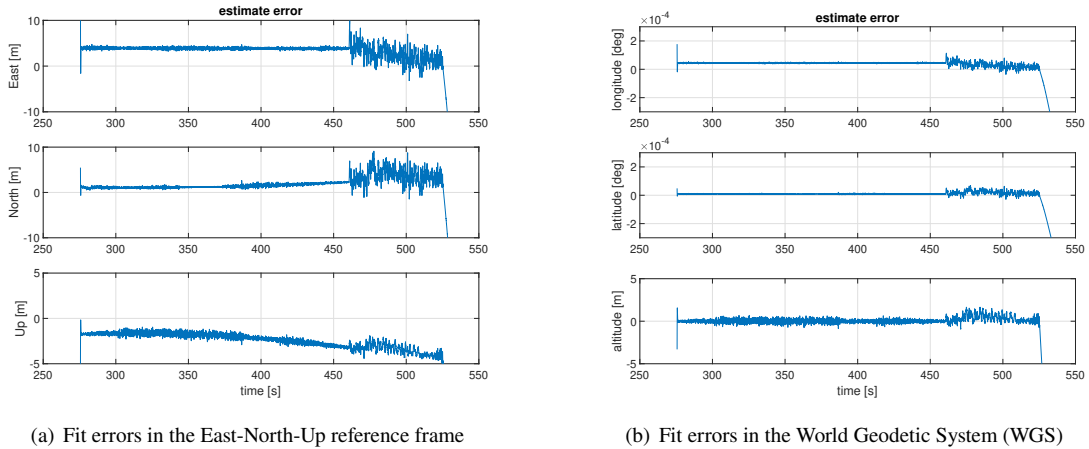


Fig. 15 Fit errors in two different reference frames

$t=415s$. This is the time instant when the distributed sensing tracker transitions from 3 non-collocated observers (2 cameras and 1 combined camera - radar) to only 2 observers (1 camera and 1 combined camera - radar). It is noteworthy that the tracker manages to reduce the standard deviation gradually over time after the latter transition. This does not happen with the former events.

As a validation check of the object tracker results, the expected sensor measurements are reconstructed based on the estimated positions and these are compared with the actual sensor measurements. This is done in Fig. 17 for the radar and Fig. 18 for the camera. Fig. 17(a) shows the reconstructed radar measurements, based on the estimated positions. In Fig. 17(b), these reconstructed radar measurements are used to compare the actual radar disturbances (left column) with the fit errors (right column). The radar disturbances are the differences between radar observation and ground truth. These disturbances are caused by the sensor properties such as noise, resolution, latency, etc. The fitting error is the difference between the Kalman Filter estimate and the ground truth. The difference between sensor disturbances and fitting error is the influence of the Kalman Filter. First of all, radar is only available around the vertiports, which is reflected by the measurements only being available in the beginning and at the end. The latency, which is modeled at 0.1s in this study, has the largest effect on the range measurements, which is primarily due to the orientation of the radar with respect to the path. For the radar at the starting point, the measurement is closer by than the real position. For the radar at the arrival site, the measurement is farther away than the real position. Thanks to the availability of other




Fig. 16 Estimate standard deviation in the ENU reference frame

camera measurements, the tracker is capable of filtering out much of the radar induced noise. However, the tracker cannot compensate for the latency, and after $t=460$ s only the camera at the vertiport is left. A collocated 2D camera and 3D radar result that the camera can compensate for the radar induced noise in azimuth and elevation, but not in the range due to the lack of depth since this is not a stereo camera.

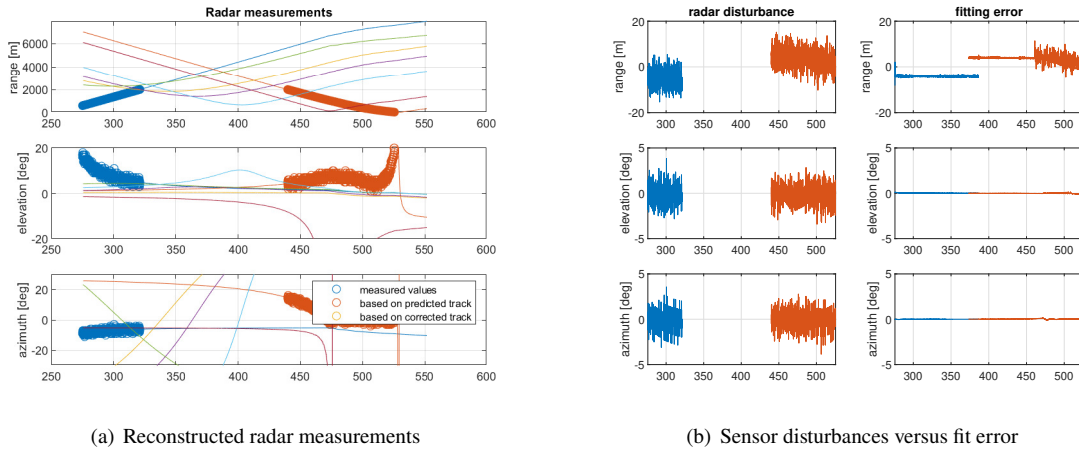


Fig. 17 Radar properties

Fig. 18 shows similar results for the camera measurements. The color coding for the different cameras is consistent with the labeling in Fig. 9. The difference in the scale of the Y-axis shows a much smaller magnitude of camera noise compared to radar noise. Significant camera disturbances are visible for azimuth as well as elevation in the left column of Fig. 18(b). These disturbances are not symmetric around the zero line due to the latency, which is the main driver for these disturbances. Note that the absolute magnitude of the disturbance is proportional to the slope of the corresponding graph in Fig. 18(a), which makes it speed and thus time dependent. This is especially apparent in the azimuth plot. Again, comparing camera disturbance with fit error shows the beneficial effect of the Kalman Filter. The noisy behavior is reduced whenever the geometry and the number of sensors allows, but this setup cannot compensate for the latency.

Next step is analyzing the observability matrix $\mathcal{O} = \begin{bmatrix} \mathbf{H} & \mathbf{H}\mathbf{F} & \mathbf{H}\mathbf{F}^2 & \dots & \mathbf{H}\mathbf{F}^{n-1} \end{bmatrix}^T$, where process matrix \mathbf{F} is defined in Eq. 7 and measurement matrix \mathbf{H} is discussed in section IV.D. n is the order of the system, which is 6 in this case. Fig. 19(a) shows that the observability matrix has full rank, indicating full reconstructibility of the state vector. Fig. 19(b) shows the evolution of the condition number of the observability matrix $\gamma = \frac{\sigma_{\max}}{\sigma_{\min}}$ over time. This qualitative assessment also confirms full reconstructibility of the state vector, but also shows in which timespans the tracking

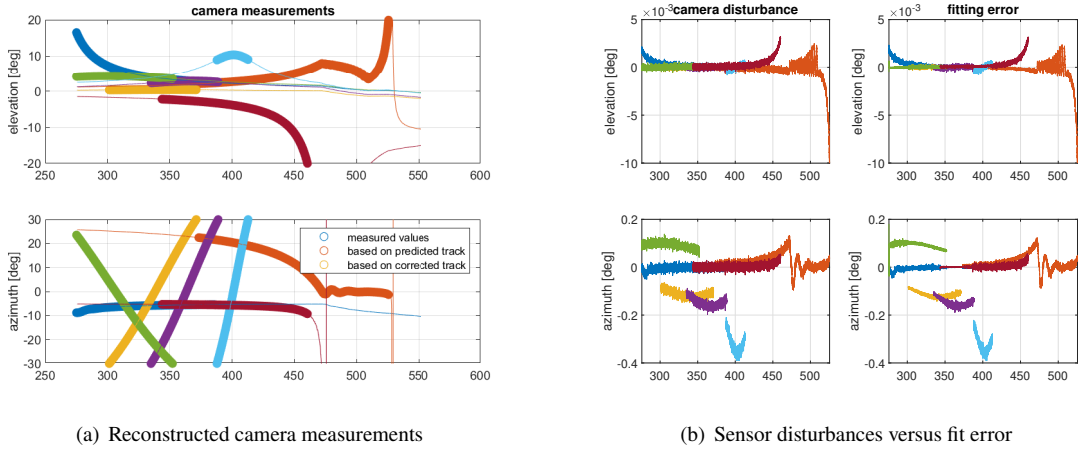


Fig. 18 Camera properties

problem becomes more ill-posed. This is the case up to around $t=325s$, the time instant when the vehicle leaves the noisy radar cone of the takeoff site. Also noteworthy are the aforementioned time instants $t=415s$, when the vehicle transitions from observed by 3 cameras to only 2 cameras and $t=460$, when the vehicle leaves the field of view of the Salesforce tower camera. However, the most significant increase in the condition number occurs around $t=440s$, which is when the vehicle enters the noisy radar cone of the landing site.




Fig. 19 Analysis of observability matrix

Finally, the evolution of the global 3D fit error in the ENU reference frame is studied over time. Fig. 20 shows this value, which is basically the Euclidian norm of the 3 individual fit error components shown in Fig. 15(a). For most of the flight, this fit error is around 5m. However, this value increases significantly, and becomes more noisy, as soon as there is only one collocated camera left to compensate for the noisy radar measurements, which happens around $t=460s$. Also, the general (slightly) decreasing and subsequently increasing trend of the fit error over time correlates well with the number of cameras that have the vehicle in their field of view, as was shown in Fig. 5. More independent non-collocated camera observations result in a lower fit error. But the relative difference is most significant for smaller camera numbers, especially less than 3.




Fig. 20 Evolution of global 3D fit error over time

B. Failure case results for flight across the bay

For the aforementioned flight across the bay scenario, a failure case is considered where all the ground based cameras are impacted by reduced visibility, for example caused by unfavorable weather conditions. As a result, the maximum detection range of the cameras is reduced from 4,000m to 2,000m. The corresponding coverage map and camera measurements are shown in Fig. 21. The coverage map as shown in Fig. 21(a) shows that specific parts of the trajectory, encircled by red ellipses, are not simultaneously observed by two different cameras. This results in a loss of sufficient information for successful triangulation. Note that radar measurements are available at takeoff and landing site to complement the cameras there. Fig. 21(b) shows the camera measurements and confirms that there are four distinct time intervals in which only one camera measurement is available instead of the minimum required number of two.

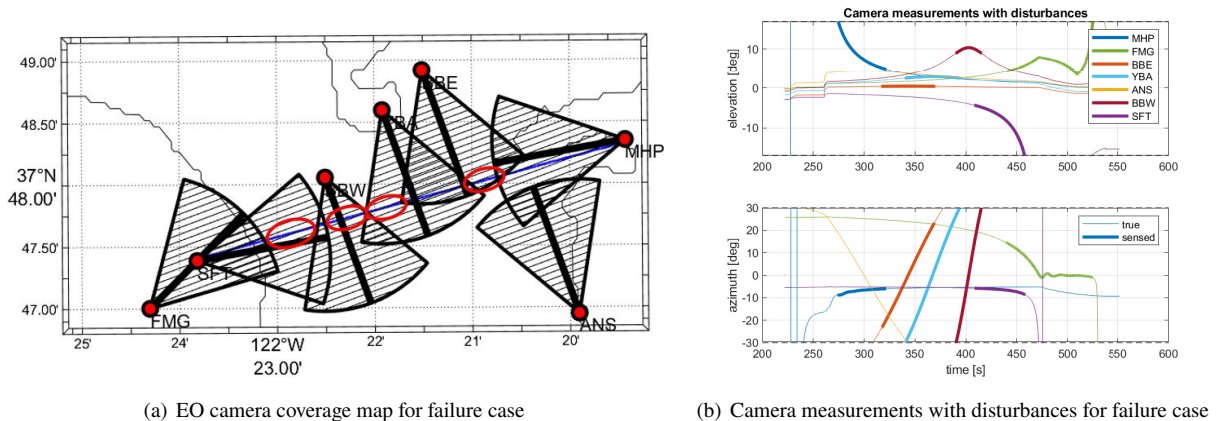


Fig. 21 Camera coverage map and measurements for all the distributed sensors in the failure case

The map in Fig. 22 shows the estimated path and compares it with the true path for the failure case. Comparing this figure with Fig. 21 shows that the estimated path drifts away from the true path exactly when there is not enough information available for triangulation. This drift happens fairly quickly, but the correction is instantaneous and along the line of sight as soon as an extra camera observation becomes available so that triangulation is possible again.

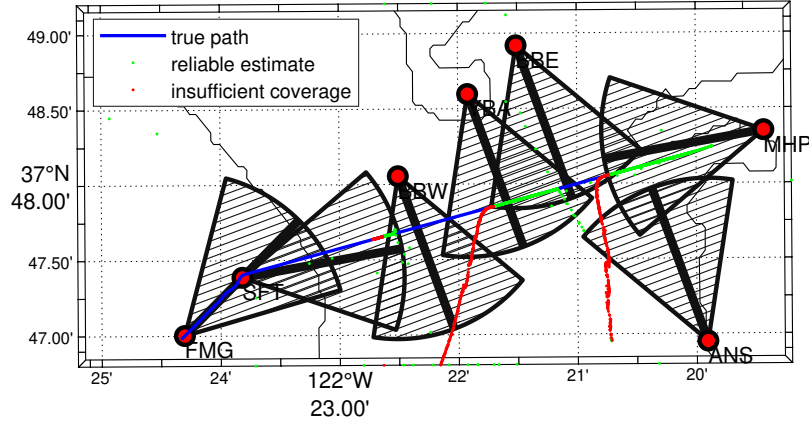


Fig. 22 Map view comparison of estimate and true path for the failure case

Analysis of the observability matrix for the failure case, shown in Fig. 23, predicts when the inferior track estimate can be expected. The rank of the observability matrix, shown in Fig. 23(a), as well as its condition number depicted in Fig. 23(b), both indicate lack of information for the track estimate in the four time intervals that were already pinpointed in the camera measurement time histories in Fig. 21(b), where a second camera observation is missing.




Fig. 23 Analysis of observability matrix for the failure case

C. Results for landing scenario

The landing scenario was introduced in Section II.B and this section discusses some select results, which are presented in a similar way as in section V.A. The map view in Fig. 24 compares the estimate with the true path for the landing scenario, and shows fairly accurate results, despite this scale not being optimal for a detailed accuracy analysis.

Fig. 25(a) shows the three components of the fit error between estimate and true path in the East-North-Up reference frame, which allows for a more detailed accuracy analysis. The error is more noisy in each axis as long as there are less than three camera observations available, which is up to $t=507s$. The close proximity and unfavorable orientation of both cameras at FMG and SFC with respect to the track in this phase of flight (almost collocated) also impacts the estimate accuracy. The addition of the non-collocated observation information from SFT as a third camera from $t=507$ onwards has a significant impact on the accuracy of the estimate. The magnitude of the standard deviations in Fig. 25(b) confirms these observations. The additional incremental improvement of the standard deviation around $t=526s$ happens when a fourth camera observation is added into the mix.

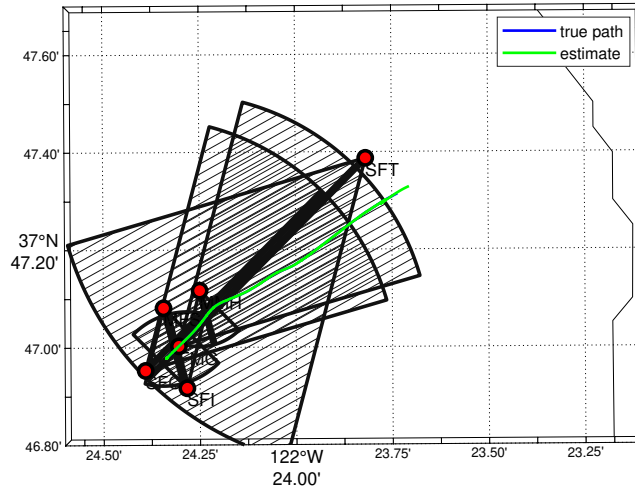
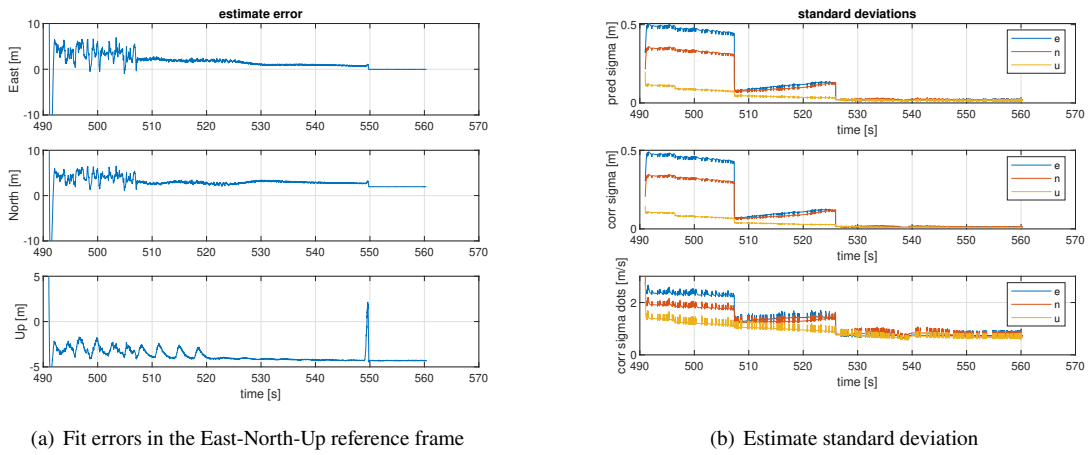


Fig. 24 Map view comparison of estimate and true path for landing scenario



(a) Fit errors in the East-North-Up reference frame

(b) Estimate standard deviation

Fig. 25 Fit errors and estimate standard deviation for the landing scenario

Fig. 26 compares sensor disturbances versus fit error to analyze the effect of the Kalman Filter on the sensor disturbances. Fig. 26(a) gives this information for the radar and Fig. 26(b) for all the cameras. Comparing left column with right column in each figure shows that the filter reduces the noise levels significantly. Just like in the nominal scenario, the latency is the dominant driver for the sensor disturbances and the fit error, since it is not compensated for by the filter.

The low values for the condition number of the observability matrix, shown in Fig. 27, show that this problem is well-posed throughout the entire scenario. The slight jumps in the values correspond to the changes in the number of cameras that have the tracked vehicle in sight, as shown in fig. 8.

Finally, the evolution of the global 3D fit error in the ENU reference frame is studied over time. Fig. 28 shows this value, which is basically the Euclidian norm of the 3 individual fit error components shown in Fig. 25(a). For most of the later part of the landing, where more than 4 cameras are tracking the vehicle, this fit error is steady and around 5m. However, this value increases significantly, and becomes more noisy, as soon as there are fewer camera observations available, which happens in the initial phase of the landing. The general decreasing trend of the fit error over time correlates well with the number of cameras that have the vehicle in their field of view, as was shown in Fig. 8. More

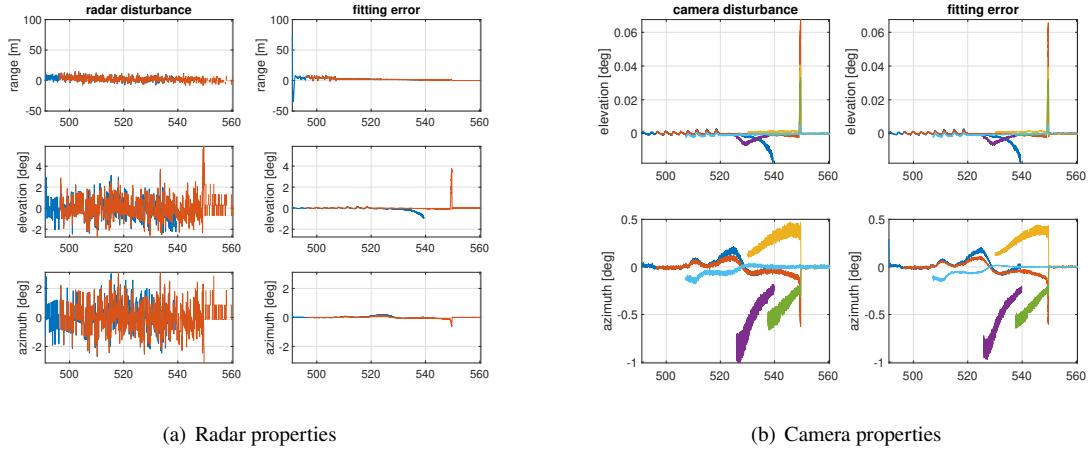


Fig. 26 Sensor disturbances versus fit error for the landing scenario




Fig. 27 Analysis of observability matrix condition number for the landing scenario

independent non-collocated camera observations result in a lower fit error. But the relative difference is most significant for smaller camera numbers, especially less than 3. This confirms earlier observations in section V.A for the nominal flight across the bay. The final phase of flight close to touchdown shows that the fit error is around 5m, with touchdown at $t=550$ s which causes the negative spike due to the combined effect of latency and sudden stop in vertical descent versus constant velocity assumption in the Kalman Filter. The time span between $t=550$ s - 560 s is with the vehicle on the landing platform. An achieved tracking accuracy around 5m, with sufficient cameras around the landing zone, is necessary and sufficient, given the size of the FATO (final approach and takeoff area, not less than 6m) and TLOF (touch down and lift off area, not less than 15m) on a typical helipad [15].

D. Impact of sensor latency on accuracy of estimate

Based on the results for the nominal flight across the bay in section V.A, a sensitivity analysis was performed to investigate the impact of the sensor latency on the accuracy of the estimate, as shown in Fig. 29. Based on the speed of the vehicle in the scenario, which is around 30 m/s, and the chosen value for the sensor latency, a certain theoretical

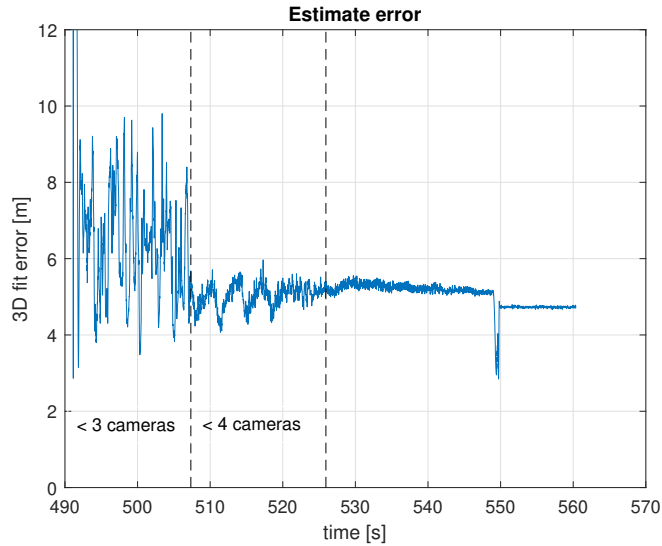


Fig. 28 Evolution of global 3D fit error over time for the landing scenario

value for the fit error is depicted, which is a linear function of the latency. However, the addition of noise, resolution, etc. cause further fit errors in addition to this theoretical value, and these are shown in the figure. Also this trend is a linear function of the latency. The purpose of this graph is to derive a value of a certain maximum allowable latency such that the resulting estimate accuracy is within a certain required threshold. For example, if one needs a maximum in flight fit error of not more than 10 m (among others for traffic separation purposes) for typical speeds used in this scenario, then this graph shows that the maximum allowable total sensor latency cannot exceed 280ms. This is a requirement that applies to the combined pipeline of sensors latency, signal transportation delay, and sensor fusion computational load. One can expect significantly higher accuracy requirements in the terminal area around the touchdown point of the vertiport, but this flight phase will also involve significantly lower vehicle speeds, which will at least partially offset the requirements for the maximum allowable sensor latency in the vertiport touchdown area.

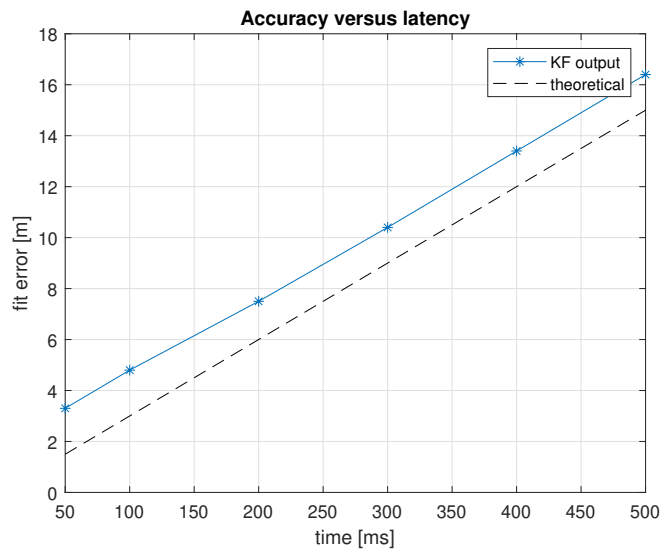


Fig. 29 Estimate accuracy as a function of sensor latency for a vehicle speed of 30 m/s

VI. Conclusions and Recommendations

In this research, a baseline distributed ground sensor based Kalman Filter is developed as an object tracker in the context of Urban/Advanced Air Mobility, where a capability is needed such that eVTOL vehicles as well as non-cooperative objects (drones, birds) can be tracked from the ground based on ground based sensors only. The sensors considered in this baseline approach are strategically placed non-collocated EO cameras and radars at the vertiports. These cameras are placed at geographically likely locations and not geometrically optimized, however their viewing directions are strategically chosen in order to achieve full camera coverage of the flightpath. Both cameras and radars have their specific advantages and drawbacks. Radar measurements are notably more noisy than camera observations. However, camera measurements provide azimuth and elevation only, thus lacking range information. Therefore multiple cameras are needed en-route to permit triangulation. Other sensor modalities than EO cameras and radars may offer drastically different benefits (and new disadvantages). This sensor fusion strategy is made adaptive, so that lower sampling rates of specific sensors are taken into account, as well as loss of signal (for example because the object is outside the field of view or the useful range of a particular camera, or because of a sensor malfunction resulting in a sensor signal dropout), and invalid data. This capability is achieved by incorporating valid flags in the weighting matrix of the Kalman filter. Initial results shown in this research confirm satisfactory performance as long as enough cameras are available to permit triangulation. The accuracy of the estimate depends on many factors, from which the latency and the number of cameras observing the target are the most important ones. It was observed that less than four and especially less than three cameras observing the target simultaneously deteriorates the accuracy of the estimate significantly. The impact of the latency also depends on the vehicle speed, among other effects. Finally, adding even only a single camera helps to remove radar noise partially, although not along the line of sight of the camera because the camera provides only two dimensional information (azimuth and elevation).

Subsequent research in this project will focus on the estimate to track association, which is an important aspect for tracking multiple vehicles. Replacing the synthetic camera measurements by YOLO measurements based on the X-plane renderings, as was done in [8], is an important next step to analyse the feasibility of the approach in a controlled environment. An important challenge here is finding the centroid of the detected ellipsoid, especially at closer ranges, as elaborated in [8]. The research team is also looking into robust strategies to deal with the aforementioned latency effects. As a next step, flight tests with real camera and radar hardware are absolutely necessary to understand further the practical implications of this approach.

Appendix: Reference Frame Transformations for Kalman Filter

Throughout this paper, five different reference frames are used: geodetic (WGS-84), Earth-Centered Earth-Fixed (ECEF), East-North-Up (ENU), body fixed and sensor spherical reference frames. They are all illustrated in Fig. 30.




Fig. 30 Reference frames used by Kalman Filter

The geodetic reference frame expresses the spherical coordinates as longitude λ_s with respect to the Greenwich meridian, latitude ϕ_s with respect to the equator and altitude h_s above the average sea-level. The ECEF reference frame is Cartesian with the origin in the Earth center and its coordinates are as follows: the Z_{ECEF} axis extends through true North, X_{ECEF} points through the earth surface where the equator and Greenwich meridian intersect, and Y_{ECEF} is finally perpendicular on the two previous axes. The ENU Cartesian reference frame is sensor centered with the axes pointing East–North–Up. As illustrated in Fig. 30, the ‘up’ direction is location dependent, as well as the others. As a consequence, the ENU frames for different sensors and target are all differently oriented. This difference increases for increasing distances between all of them. This must be reflected in the calculations. Next is the body fixed reference frame, which is Cartesian and centered at the sensor’s center of gravity. The X_B axis points forward, Y_B axis points to the right and the Z_B axis is downwards. Finally, the sensor centered spherical reference frame expresses the coordinates as range ρ (distance from center of gravity), elevation α (vertical angle with respect to the $X_B - Y_B$ -plane) and azimuth γ (horizontal angle with respect to the forward axis). The sensor measurements from radar and camera are expressed in this spherical reference frame. Note that both latter sensor fixed reference frames are sensor attitude dependent, i.e. the viewing direction.

The transformations between the different aforementioned reference frames are illustrated in Fig. 31. The transformations between body-fixed, ENU, ECEF as well as between different ENU’s of different sensors and target at different locations are pure angular rotations by means of the respective rotation matrices $\mathbf{M}(\phi, \theta, \psi)$ over the three Euler angles between body fixed and ENU, $\mathbf{T}(\lambda_s, \phi_s)$ over longitude and latitude between ENU and ECEF and $\mathbf{L}(\lambda_s, \phi_s, \lambda_t, \phi_t)$ over longitude and latitude of both locations between different ENU’s at different locations. The mappings between body fixed and sensor spherical as well as between geodetic and ECEF are more complex nonlinear transformations. Each of them is elaborated next.

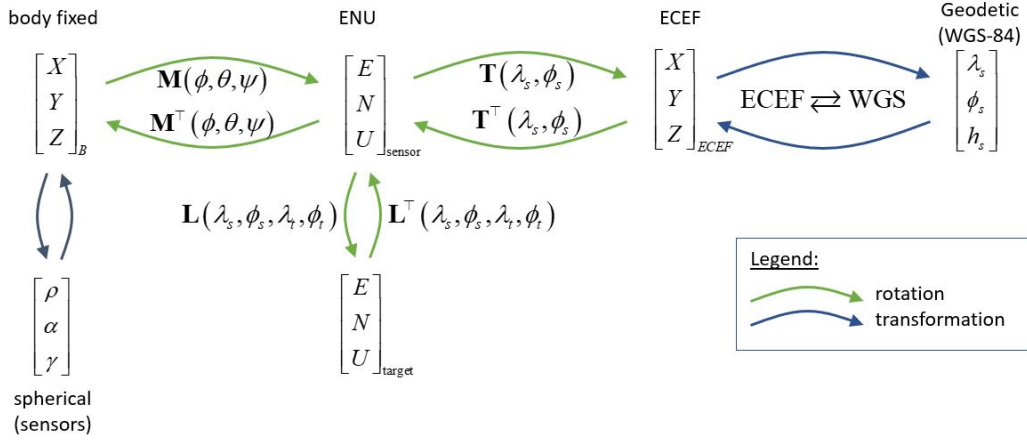


Fig. 31 Transformations between the different reference frames

Transformation from body-fixed to spherical and backwards

The sensor spherical coordinates are defined as a function of the body-fixed coordinates as follows:

$$\rho = \sqrt{X_B^2 + Y_B^2 + Z_B^2} \quad (35)$$

$$\alpha = \arctan \frac{-Z_B}{\rho_{\text{hor}}} \quad (36)$$

$$\gamma = \arctan \frac{Y_B}{X_B} \quad (37)$$

$$(38)$$

where:

$$\rho_{\text{hor}} = \sqrt{X_B^2 + Y_B^2} \quad (39)$$

And the other way around:

$$X_B = \rho \cos \alpha \cos \gamma \quad (40)$$

$$Y_B = \rho \cos \alpha \sin \gamma \quad (41)$$

$$Z_B = \rho \sin \alpha \quad (42)$$

Rotation from body-fixed to ENU and backwards

From body-fixed to ENU is a sequential rotation over the three Euler angles roll ϕ , pitch attitude θ and yaw ψ [16]:

$$\mathbf{M}(\phi, \theta, \psi) = \begin{bmatrix} \sin \psi \cos \theta & \cos \phi \cos \psi + \sin \phi \sin \psi \sin \theta & -\sin \phi \cos \psi + \cos \phi \sin \psi \sin \theta \\ \cos \psi \cos \theta & -\cos \phi \sin \psi + \sin \phi \cos \psi \sin \theta & \sin \phi \sin \psi + \cos \phi \cos \psi \sin \theta \\ \sin \theta & -\sin \phi \cos \theta & -\cos \phi \cos \theta \end{bmatrix} \quad (43)$$

The transpose is used for the backwards direction.

Rotation from ENU to ECEF and backwards

From the vehicle centered ENU to the earth centered ECEF is a sequential rotation over the longitude λ_s and latitude ϕ_s of the current position[16]:

$$\mathbf{T}(\lambda_s, \phi_s) = \begin{bmatrix} -\sin \lambda_s & -\cos \lambda_s \sin \phi_s & \cos \lambda_s \cos \phi_s \\ \cos \lambda_s & -\sin \lambda_s \sin \phi_s & \sin \lambda_s \cos \phi_s \\ 0 & \cos \phi_s & \sin \phi_s \end{bmatrix} \quad (44)$$

The transpose is used for the other way around.

Rotation between ENU's

A rotation between ENU's is basically a rotation from the first ENU to ECEF and subsequently from ECEF back to the second ENU:

$$\mathbf{L}(\lambda_s, \phi_s, \lambda_t, \phi_t) = \mathbf{T}^T(\lambda_t, \phi_t) \cdot \mathbf{T}(\lambda_s, \phi_s) \quad (45)$$

Substituting Eq. (44) results in:

$$\begin{aligned} \mathbf{L}(\lambda_s, \phi_s, \lambda_t, \phi_t) &= \\ &= \mathbf{T}^T(\lambda_t, \phi_t) \cdot \mathbf{T}(\lambda_s, \phi_s) \\ &= \begin{bmatrix} -\sin \lambda_t & -\cos \lambda_t \sin \phi_t & \cos \lambda_t \cos \phi_t \\ \cos \lambda_t & -\sin \lambda_t \sin \phi_t & \sin \lambda_t \cos \phi_t \\ 0 & \cos \phi_t & \sin \phi_t \end{bmatrix}^T \begin{bmatrix} -\sin \lambda_s & -\cos \lambda_s \sin \phi_s & \cos \lambda_s \cos \phi_s \\ \cos \lambda_s & -\sin \lambda_s \sin \phi_s & \sin \lambda_s \cos \phi_s \\ 0 & \cos \phi_s & \sin \phi_s \end{bmatrix} \\ &= \begin{bmatrix} \cos(\lambda_t - \lambda_s) & \sin(\lambda_t - \lambda_s) \sin \phi_s & -\sin(\lambda_t - \lambda_s) \cos \phi_s \\ -\sin(\lambda_t - \lambda_s) \sin \phi_t & \cos \phi_t \cos \phi_s + \sin \phi_t \sin \phi_s \cos(\lambda_t - \lambda_s) & \cos \phi_t \sin \phi_s - \sin \phi_t \cos \phi_s \cos(\lambda_t - \lambda_s) \\ \sin(\lambda_t - \lambda_s) \cos \phi_t & \sin \phi_t \cos \phi_s - \cos \phi_t \sin \phi_s \cos(\lambda_t - \lambda_s) & \sin \phi_t \sin \phi_s + \cos \phi_t \cos \phi_s \cos(\lambda_t - \lambda_s) \end{bmatrix} \end{aligned} \quad (46)$$

Transformation from ECEF to geodetic and backwards

ECEF \Rightarrow *geodetic*

The target position in WGS-84 geodetic coordinates is obtained following the 15-step process presented below (extracted from RTCA DO-317A)[14]:

$$r_{\text{ecef}} = \sqrt{X_t^2 + Y_t^2} \quad (47)$$

$$E^2 = a^2 - b^2 \quad (48)$$

$$F = 54b^2Z_t^2 \quad (49)$$

$$G = r_{\text{ecef}}^2 + (1 - \varepsilon^2)Z_t^2 - \varepsilon^2E^2 \quad (50)$$

$$C = \frac{\varepsilon^4 F r_{\text{ecef}}^2}{G^3} \quad (51)$$

$$S = \sqrt[3]{1 + C + \sqrt{C^2 + 2C}} \quad (52)$$

$$P = \frac{F}{3\left(S + \frac{1}{S} + 1\right)^2 G^2} \quad (53)$$

$$Q = \sqrt{1 + 2\varepsilon^4 P} \quad (54)$$

$$r_o = \frac{-P\varepsilon^2 r_{\text{ecef}}}{1 + Q} + \sqrt{\frac{a^2}{2} \left(1 + \frac{1}{Q}\right) - \frac{P(1 - \varepsilon^2)Z_t^2}{Q(1 + Q)} - \frac{Pr_{\text{ecef}}^2}{2}} \quad (55)$$

$$U = \sqrt{(r_{\text{ecef}} - \varepsilon^2 r_o)^2 + Z_t^2} \quad (56)$$

$$V = \sqrt{(r_{\text{ecef}} - \varepsilon^2 r_o)^2 + (1 - \varepsilon^2)Z_t^2} \quad (57)$$

$$Z_o = \frac{b^2 Z_t}{aV} \quad (58)$$

$$\phi_t = \tan^{-1} \left(\frac{Z_t + \varepsilon'^2 Z_o}{r_{\text{ecef}}} \right) \quad (59)$$

$$\lambda_t = \tan^{-1} \left(\frac{Y_t}{X_t} \right) \quad (60)$$

$$h_t = U \left(1 - \frac{b^2}{aV} \right) \quad (61)$$

where:

$a = 6.378137 \cdot 10^6 m$, semimajor axis of the Earth

$b = 6.3567523142 \cdot 10^6 m$, semiminor axis of the Earth

$\varepsilon = 0.081819190842622$, first eccentricity of the Earth

$\varepsilon' = 0.082094437935831$, second eccentricity of the Earth

ϕ_t = target latitude [rad]

λ_t = target longitude [rad]

h_t = target geodetic altitude [m]

geodetic \Rightarrow *ECEF*

The ownship's Earth-Centered, Earth-Fixed (ECEF) position is calculated by[14]:

$$\begin{bmatrix} X_s \\ Y_s \\ Z_s \end{bmatrix} = \begin{bmatrix} (N_s + h_s) \cos \phi_s \cos \lambda_s \\ (N_s + h_s) \cos \phi_s \sin \lambda_s \\ (N_s (1 - \varepsilon^2) + h_s) \sin \phi_s \end{bmatrix} \quad (62)$$

where:

ϕ_s = ownship geodetic latitude [rad]

λ_s = ownship geodetic longitude [rad]
 h_s = ownship geodetic altitude [m]
 $N_s = \frac{a}{\sqrt{1 - \varepsilon^2 \sin^2 \phi_s}}$
 $a = 6.378137 \cdot 10^6 m$, semimajor axis of the Earth
 $\varepsilon = 0.081819190842622$, first eccentricity of the Earth

If the altitude of the ownship is measured as height above Mean Sea Level (MSL) from a barometric sensor, it should be converted to WGS-84 Height Above Ellipsoid (HAE) using the following steps:

$$\xi(\phi_s) = g_{0_{\text{eq}}} \frac{1 + k \sin^2 \phi_s}{\sqrt{1 - \varepsilon^2 \sin^2 \phi_s}} \quad (63)$$

$$R(\phi_s) = \frac{a}{1 + f + m_r - 2f \sin^2 \phi_s} \quad (64)$$

$$h_S = \frac{R(\phi_s)^h}{\frac{R(\phi_s)\xi(\phi_s)}{\xi(\frac{\pi}{4})} - h} \quad (65)$$

where:

ϕ_s = ownship geodetic latitude [rad]
 $\varepsilon = 0.081819190842622$, first eccentricity of the Earth
 $g_{0_{\text{eq}}} = 9.7803253359 m/s^2$, theoretical gravity at the equator
 $g_{0_{\text{pl}}} = 9.8321849378 m/s^2$, theoretical gravity at the poles
 $k = \frac{b \cdot g_{0_{\text{pl}}}}{a \cdot g_{0_{\text{eq}}}} - 1$
 $\xi(\phi_s)$ = normal gravity on the surface of an ellipsoid at ϕ_s [m/s^2]
 $a = 6.378137 \cdot 10^6 m$, semimajor axis of the Earth
 $b = 6.3567523142 \cdot 10^6 m$, semiminor axis of the Earth
 $f = \frac{a-b}{a}$, flattening of the Earth
 $m_r = 0.003449787$, gravity ratio
 $R(\phi_s)$ = radius of the Earth at ϕ_s [m]
 h = measured ownship height above MSL
 h_s = ownship geodetic altitude [m]

Acknowledgments

This work was performed under the NASA Aeronautics Research Mission Directorate (ARMD), Transformative Tools and Technologies (TTT) Project. The authors would like to thank the project management members Kelley Hashemi, Vanessa Aubuchon and Lisa Le Vie, as well as NASA Armstrong colleague Nelson Brown, for their invaluable contributions to this project.

References

- [1] Aubuchon, V. V., Hashemi, K. E., Shively, R. J., and Wishart, J. M., "Multi-Vehicle (m:N) Operations in the NAS - NASA's Research Plans," *AIAA AVIATION 2022 Forum*, American Institute of Aeronautics and Astronautics, 2022. doi:10.2514/6.2022-3758.
- [2] Chong, C. Y., Chang, K. C., and Mori, S., "Distributed Tracking in Distributed Sensor Networks." *Proceedings of the American Control Conference*, 1986, pp. 1863–1868. doi:10.23919/acc.1986.4789229.
- [3] Qi, H., Iyengar, S. S., and Chakrabarty, K., "Distributed sensor networks - A review of recent research," *Journal of the Franklin Institute*, Vol. 338, 2001, pp. 655–668. doi:10.1016/S0016-0032(01)00026-6.
- [4] Horling, B., Vincent, R., Mailler, R., Shen, J., Becker, R., Rawlins, K., and Lesser, V., "Distributed sensor network for real time tracking," *Proceedings of the International Conference on Autonomous Agents*, 2001, pp. 417–424. doi:10.1145/375735.376399.

- [5] Chamberland, J.-F., and Veeravalli, V., “Decentralized detection in sensor networks,” *IEEE Transactions on Signal Processing*, Vol. 51, 2003, pp. 407–416. doi:10.1109/TSP.2002.806982, URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1166675>.
- [6] La, H. M., and Sheng, W., “Distributed sensor fusion for scalar field mapping using mobile sensor networks,” *IEEE Transactions on Cybernetics*, Vol. 43, 2013, pp. 766–778. doi:10.1109/TSMCB.2012.2215919.
- [7] Chatterjee, A., Flores, H., Sen, S., Hasan, K. S., and Mani, A., “Distributed location detection algorithms using IoT for commercial aviation,” *Proceedings - 2017 3rd IEEE International Conference on Research in Computational Intelligence and Communication Networks, ICRCICN 2017*, Vol. 2017-Decem, 2017, pp. 126–131. doi:10.1109/ICRCICN.2017.8234493.
- [8] Lombaerts, T., Shish, K., Keller, G., Stepanyan, V., Cramer, N., and Ippolito, C., “Adaptive Multi-Sensor Fusion Based Object Tracking for Autonomous Urban Air Mobility Operations,” *AIAA SciTech Forum*, 2022. doi:10.2514/6.2022-0362.
- [9] Lombaerts, T., Kaneshige, J., and Feary, M., “Control Concepts for Simplified Vehicle Operations of a Quadrotor eVTOL Vehicle,” *AIAA Aviation Conference*, American Institute of Aeronautics and Astronautics, American Institute of Aeronautics and Astronautics, 2020. doi:10.2514/6.2020-3189.
- [10] Silva, C., Johnson, W. R., Solis, E., Patterson, M. D., and Antcliff, K. R., “VTOL Urban Air Mobility Concept Vehicles for Technology Development,” *2018 Aviation Technology, Integration, and Operations Conference*, American Institute of Aeronautics and Astronautics, 2018. doi:10.2514/6.2018-3847.
- [11] Kannan, K., Baculi, J., Lombaerts, T., Kawamura, E., Gorospe, G., Holforty, W., Ippolito, C., Stepanyan, V., Dolph, C., and Brown, N., “A Simulation Architecture for Air Traffic Over Urban Environments Supporting Autonomy Research in Advanced Air Mobility,” *AIAA SciTech 2023 Forum*, American Institute for Aeronautics and Astronautics, 2023.
- [12] Kawamura, E., Dolph, C., Kannan, K., Lombaerts, T., and Ippolito, C. A., “Simulated Vision-based Approach and Landing System Advanced Air Mobility,” *AIAA SciTech 2023 Forum*, 2023.
- [13] Shish, K. H., Cramer, N. B., Gorospe, G., Lombaerts, T., Stepanyan, V., and Kannan, K., “Survey of Capabilities and Gaps in External Perception Sensors for Autonomous Urban Air Mobility Applications,” *AIAA Scitech 2021 Forum*, American Institute of Aeronautics and Astronautics, 2021. doi:10.2514/6.2021-1114.
- [14] *Minimum Operational Performance Standards for Detect and Avoid Systems*, DO-365, RTCA Inc., 2019, techreport Appendix F: Example Functional Description For Airborne Surveillance Data Processor (ASDP) Subsystem. URL <https://standards.globalspec.com/std/14281907/rtca-do-365>.
- [15] “Advisory Circular: Heliport Design,” techreport 150/5390-2C, FAA, Apr. 2012. URL https://www.faa.gov/documentLibrary/media/Advisory_Circular/150_5390_2c.pdf.
- [16] Mohinder, S., Grewal, P., Lawrence, R., Weill, A. P., and Andrews, “Appendix C: Coordinate Transformations,” *Global Positioning Systems, Inertial Navigation, and Integration*, John Wiley & Sons, Inc., 2006, pp. 456–501. doi:10.1002/9780470099728.app3.