



**Autonomous Vehicle Simulation (AVS) Laboratory,
University of Colorado**

Basilisk Technical Memorandum

Document ID: Basilisk-test_sunlineEKF

EKF MODULE AND TEST

Prepared by	T. Teil
-------------	---------

Status: Initial document
Scope/Contents
This module implements and tests a Extended Kalman Filter in order to estimate the sunline direction.

Rev:	Change Description	By
Draft	Initial Revision	T. Teil

Contents

1	Introduction	1
2	Filter Set-up, initialization, and I/O	1
2.1	Dynamics	1
2.2	User initialization	2
2.3	Inputs and Outputs	2
3	Filter Algorithm	2
4	Test Design	4
4.1	sunline_individual_test	4
4.2	StatePropStatic	5
4.3	StatePropVariable	6
4.4	Full Filter tests	6

1 Introduction

The Extended Kalman filter (EKF) in the AVS Basilisk simulation is a sequential filter implemented to give the best estimate of the desired states. In this method we estimate the sun heading as well as it's rate of change in the body frame. The EKF reads in the message written by the coarse sun sensor, and writes a message containing the sun estimate.

This document summarizes the content of the module, how to use it, and the test that was implemented for it. More information on the filter derivation can be found in Reference [?].

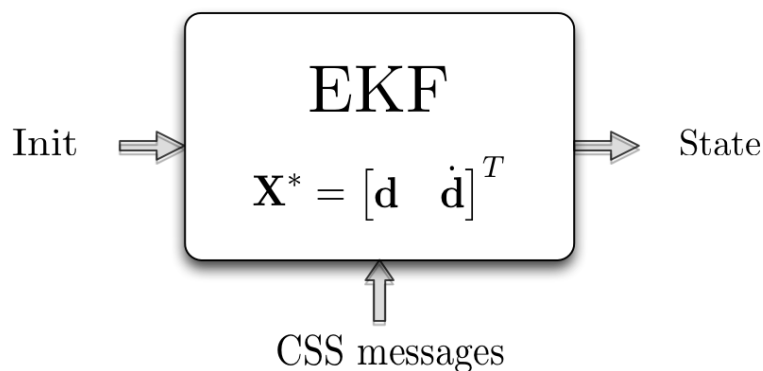


Fig. 1: Filter module

2 Filter Set-up, initialization, and I/O

2.1 Dynamics

The states that are estimated in this filter are the sunline vector, and it's rate of change $\mathbf{X}^* = [\mathbf{d} \quad \dot{\mathbf{d}}]^T$. The star superscript represents that this is the reference state.

The dynamics are given in equation 1. Given the nature of the filter, there is an unobservable state component: the rotation about the \mathbf{d} axis. In order to remedy this, we project the states along this axis and subtract them, in order to measure only observable state components.

$$\mathbf{F}(\mathbf{X}) = \begin{bmatrix} \mathbf{F}_1(\mathbf{d}) \\ \mathbf{F}_2(\dot{\mathbf{d}}) \end{bmatrix} = \begin{bmatrix} \dot{\mathbf{d}} - \left((\mathbf{d} \cdot \dot{\mathbf{d}}) \frac{\mathbf{d}}{\|\mathbf{d}\|^2} \right) \\ -\frac{1}{\Delta t} \left((\mathbf{d} \cdot \dot{\mathbf{d}}) \frac{\mathbf{d}}{\|\mathbf{d}\|^2} \right) \end{bmatrix} \quad (1)$$

This leads us to the computation of the dynamics matrix $A = \left[\frac{\partial \mathbf{F}(\mathbf{X}, t_i)}{\partial \mathbf{X}} \right]^*$. The partials are given in equation 2, and were verified in Mathematica.

$$A = \begin{bmatrix} \frac{\partial \mathbf{F}_1(\mathbf{X}, t_i)}{\partial \mathbf{d}} & \frac{\partial \mathbf{F}_1(\mathbf{X}, t_i)}{\partial \dot{\mathbf{d}}} \\ \frac{\partial \mathbf{F}_2(\mathbf{X}, t_i)}{\partial \mathbf{d}} & \frac{\partial \mathbf{F}_2(\mathbf{X}, t_i)}{\partial \dot{\mathbf{d}}} \end{bmatrix} \quad (2)$$

$$= \begin{bmatrix} -\left(\frac{\mathbf{d}\mathbf{d}^T}{\|\mathbf{d}\|^2} + (\mathbf{d} \cdot \dot{\mathbf{d}}) \frac{\|\mathbf{d}\|^2 \mathbf{I} - 2\mathbf{d}\mathbf{d}^T}{\|\mathbf{d}\|^4} \right) & \mathbf{I} - \frac{\mathbf{d}\mathbf{d}^T}{\|\mathbf{d}\|^2} \\ -\frac{1}{\Delta t} \left(\frac{\mathbf{d}\mathbf{d}^T}{\|\mathbf{d}\|^2} + (\mathbf{d} \cdot \dot{\mathbf{d}}) \frac{\|\mathbf{d}\|^2 \mathbf{I} - 2\mathbf{d}\mathbf{d}^T}{\|\mathbf{d}\|^4} \right) & -\frac{1}{\Delta t} \frac{\mathbf{d}\mathbf{d}^T}{\|\mathbf{d}\|^2} \end{bmatrix} \quad (3)$$

The measurement model is given in equation 4, and the H matrix defined as $H = \left[\frac{\partial \mathbf{G}(\mathbf{X}, t_i)}{\partial \mathbf{X}} \right]^*$ is given in equation 5.

In this filter, the only measurements used are from the coarse sun sensor. For the i^{th} sensor, the measurement is simply given by the dot product of the sunline heading and the normal to the sensor. This yields easy partial derivatives for the H matrix, which is a matrix formed of the rows of transposed normal vectors (only for those which received a measurement). Hence the H matrix has a changing size depending on the amount of measurements.

$$\mathbf{G}_i(\mathbf{X}) = \mathbf{n}_i \cdot \mathbf{d} \quad (4)$$

$$\mathbf{H}(\mathbf{X}) = \begin{bmatrix} \mathbf{n}_1^T \\ \vdots \\ \mathbf{n}_i^T \end{bmatrix} \quad (5)$$

2.2 User initialization

In order for the filter to run, the user must set a few parameters:

- The angle threshold under which the coarse sun sensors do not read the measurement:
`FilterContainer.sensorUseThresh = 0.`
- The process noise value, for instance:
`FilterContainer.qProcVal = 0.001`
- The measurement noise value, for instance:
`FilterContainer.qObsVal = 0.001`
- The threshold in the covariance norm leading to the switch from the EKF update to the linear Kalman Filter update (discussed more closely in the Measurement update part):
`FilterContainer.ekfSwitch = 5`

- The initial covariance:
`Filter.covar =`
`[0.4, 0., 0., 0., 0., 0.,`
`0.,0.4, 0., 0., 0., 0.,`
`0., 0., 0.4, 0., 0., 0.,`
`0., 0., 0., 0.004, 0., 0.,`
`0., 0., 0., 0., 0.004,`
`0., 0., 0., 0., 0.,, 0.004]`
- The initial state :
`Filter.state = [1., 0., 1., 0., 0.1, 0.]`

2.3 Inputs and Outputs

The EKF reads in the measurements from the coarse sun sensors. These are under the form of a list of cosine values. Knowing the normals to each of the sensors, we can therefore use them to estimate sun heading.

3 Filter Algorithm

Once the filter has been properly setup in the python code, it can go through it's algorithm:

Initialization

First the filter is initialized. This can be done at any time during a simulation in order to reset the filter.

- Time is set to t_0
- The state \mathbf{X}^* is set to the initial state \mathbf{X}_0^*
- The state error \mathbf{x} is set to it's initial value \mathbf{x}_0
- The covariance P is set to the initial state P_0

Time Update

At some time t_i , if the update filter method is called, a time update will first be executed.

- The state is propagated using the dynamics \mathbf{F} with initial conditions $\mathbf{X}^*(t_{i-1})$
- Compute the dynamics matrix $A(t) = \left[\frac{\partial \mathbf{F}(\mathbf{X}, t)}{\partial \mathbf{X}} \right]^*$ which is evaluated on the reference trajectory
- Integrate the STM, $\dot{\Phi}(t, t_{i-1}) = A(t)\Phi(t, t_{i-1})$ with initial conditions $\Phi(t_{i-1}, t_{i-1}) = I$

This gives us $\mathbf{X}^*(t_i)$ and $\Phi(t_i, t_{i-1})$.

Observation read in

If no measurement is read in at time t_i :

- $\mathbf{X}^*(t_i)$ previously computed becomes the most recent reference state
- $\mathbf{x}_i = \bar{\mathbf{x}}_i = \Phi(t_i, t_{i-1})\mathbf{x}_{i-1}$ is the new state error
- $P_i = \bar{P}_i = \Phi(t_i, t_{i-1})P_{i-1}\Phi^T(t_i, t_{i-1})$ becomes the updated covariance

If a measurement is read in, the algorithm computes the observation, the observation state matrix, and the Kalman Gain.

- The observation (\mathbf{Y}_i) is compared to the observation model, giving the innovation: $\mathbf{y}_i = \mathbf{Y}_i - G(\mathbf{X}_i^*, t_i)$
- Compute the observation matrix along the reference trajectory: $\tilde{H}_i = \left[\frac{\partial G(\mathbf{X}, t_i)}{\partial \mathbf{X}} \right]^*$
- Compute the Kalman Gain $K_i = \bar{P}_i \tilde{H}_i^T \left(\tilde{H}_i \bar{P}_i \tilde{H}_i^T + R_i \right)^{-1}$

Measurement Update

Depending on the covariance, the filter can either update as a classic, linear Kalman Filter, or as the Extended Kalman filter. This is done in order to assure robust and fast filter convergence. Indeed in a scenario with a very large initial covariance, the EKF's change in reference trajectory could delay or inhibit the convergence. In order to remedy this, a few linear updates are performed if the maximum value in the covariance is greater than a user-set threshold.

Linear update:

- The state error is updated using the time updated value: $\mathbf{x}_i = \bar{\mathbf{x}}_i + K_i \left[\mathbf{y}_i - \tilde{H}_i \bar{\mathbf{x}}_i \right]$
- The covariance is updated using the Joseph form of the covariance update equation: $P_i = \left(I - K_i \tilde{H}_i \right) \bar{P}_i \left(I - K_i \tilde{H}_i \right)^T + K_i R_i K_i^T$
- The reference state stays the same, and it's propagated value $\mathbf{X}^*(t_i)$ becomes $\mathbf{X}^*(t_{i-1})$

EKF update:

- The state error is updated using the innovation and the Kalman Gain: $\mathbf{x}_i = K_i \mathbf{y}_i$
- The reference state is changed by the state error: $\mathbf{X}^*(t_i) = \mathbf{X}^*(t_i) + \mathbf{x}_i$
- The covariance is updated using the Joseph form of the covariance update equation: $P_i = \left(I - K_i \tilde{H}_i \right) \bar{P}_i \left(I - K_i \tilde{H}_i \right)^T + K_i R_i K_i^T$
- The new reference state is now used $\mathbf{X}^*(t_i)$ becomes $\mathbf{X}^*(t_{i-1})$

4 Test Design

The unit test for the sunlineEKF module is located in:

```
fswAlgorithms/attDetermination/sunlineEKF/_UnitTest/test_SunlineEKF.py
```

As well as another python file containing plotting functions:

```
fswAlgorithms/attDetermination/sunlineEKF/_UnitTest/SunlineEKF_test_utilities.py
```

The test is split up into 4 subtests, the last one is parametrized in order to test different scenarios. The first test checks up all of the individual filter methods and tests them individually. The second test verifies that in the case where the state is zeroed out from the start of the simulation, it remains at zero. The third test verifies the behavior of the time update in a general case. The final test is a full filter test.

4.1 sunline_individual_test

In each of these individual tests, random inputs are fed to the methods and their values are computed in parallel in python. These two values are then compared to assure that the correct computations are taking place.

- Dynamics Matrix: This method computes the dynamics matrix A . Tolerance to absolute error $\epsilon = 10^{-10}$.
Passed
- State and STM propagation: This method propagates the state using the F function as well as the STM using $\dot{\Phi} = A\Phi$. Tolerance to absolute error $\epsilon = 10^{-10}$.
Passed
- H and y propagation: This method computes the H matrix, and compares the measurements to the expected measurements given the state. Tolerance to absolute error $\epsilon = 10^{-10}$.
Passed
- Kalman gain: This method computes the K matrix. Tolerance to absolute error $\epsilon = 10^{-10}$.
Passed
- EKF update: This method performs the measurement update in the case of an EKF. Tolerance to absolute error $\epsilon = 10^{-10}$.
Passed
- Linear Update: This method performs the measurement update in the linear case. Tolerance to absolute error $\epsilon = 10^{-10}$.
Passed

4.2 StatePropStatic

This test runs the filter with no measurements. It initializes with a zeroed state, and assures that at the end of the simulation all values are still at zero. Plotted results are seen in Figure 2.

Tolerance to absolute error: $\epsilon = 10^{-10}$
Passed

4.3 StatePropVariable

This test also takes no measurements in, but gives a random state with rate of change. It then tests that the states and covariance are as expected throughout the time of simulation. Plotted results are seen in Figure 3. We indeed see that the state and covariance for the test and the code overlap perfectly.

Tolerance to absolute error: $\epsilon = 10^{-10}$
Passed

4.4 Full Filter tests

This test the filter working from start to finish. No measurements are taken in for the first 20 time steps. Then a heading is given through the CSS message. Halfway through the simulation, measurements stop, and 20 time steps later a different heading is read. The filter must be robust and detect this change. This test is parametrized for different test lengths, different initial conditions, different measured headings, and with or without measurement noise. All these are successful.

Tolerance to absolute error without measurement noise: $\epsilon = 10^{-10}$

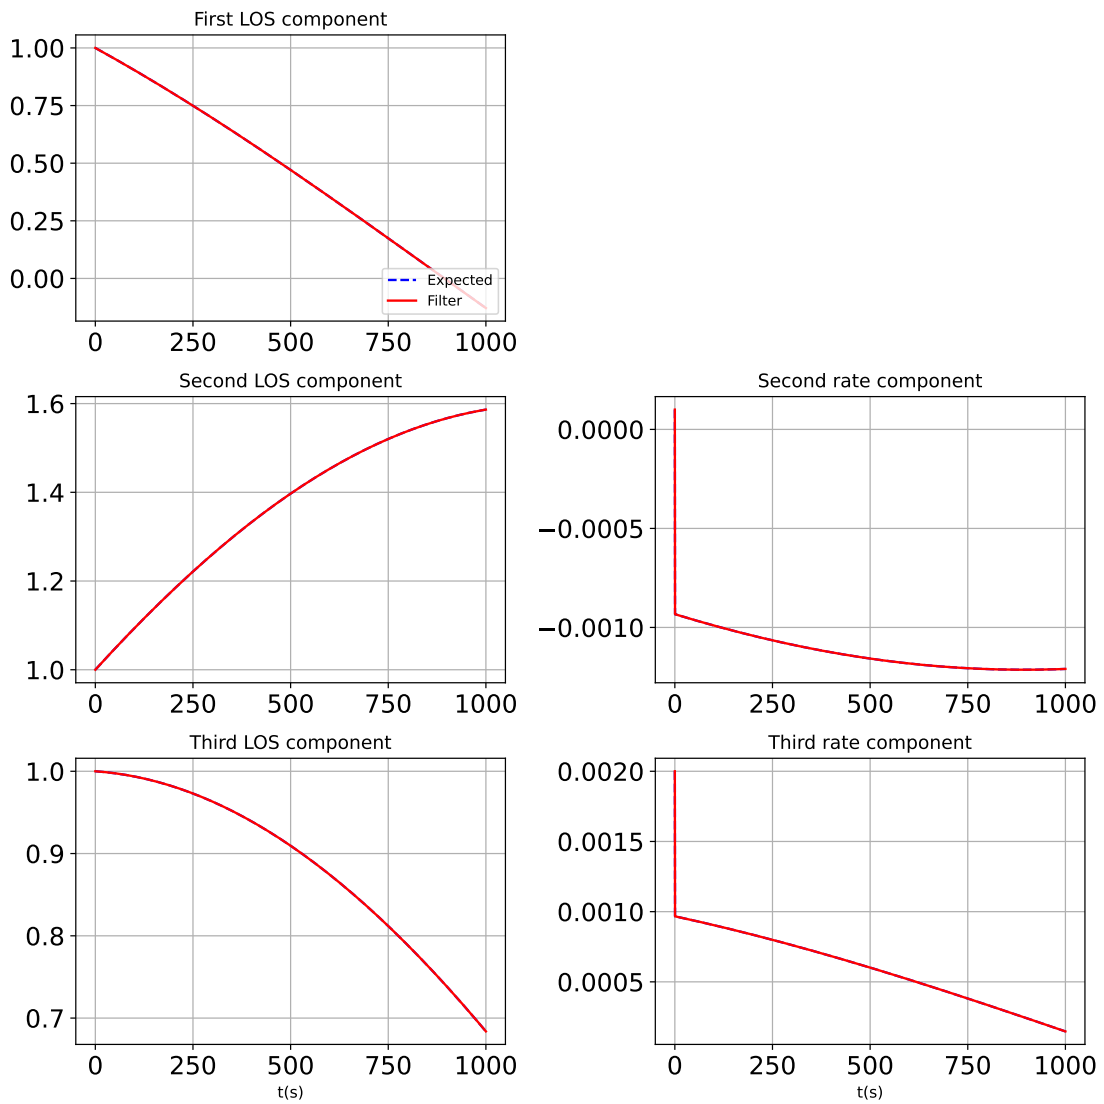


Fig. 2: States vs true states in static case

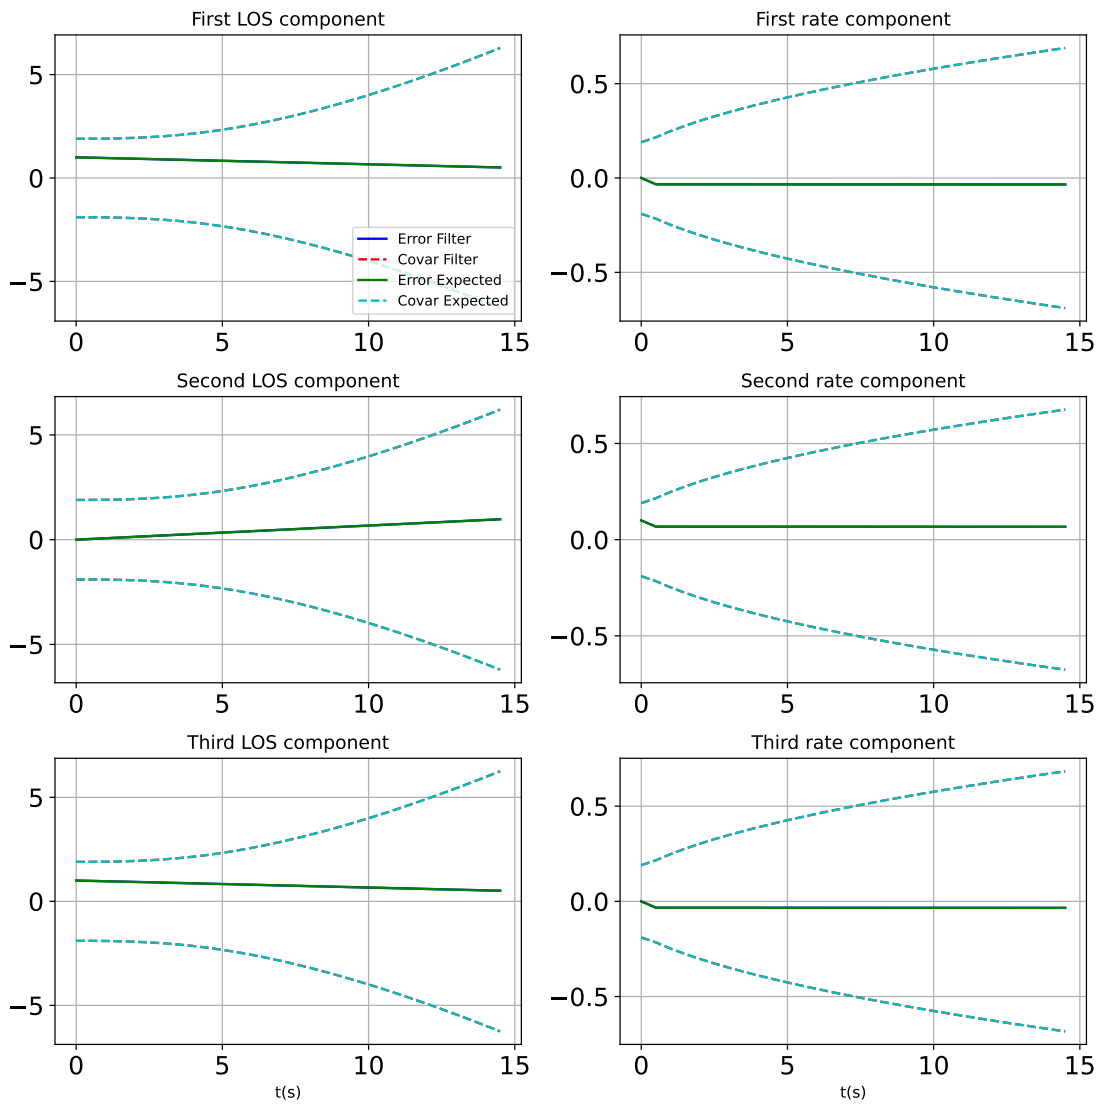


Fig. 3: State error and covariance vs expected Values

Tolerance to absolute error with measurement noise: $\epsilon = 10^{-2}$

Passed

Plotted results are seen in Figures 4, 5, and 6. Figure 4 shows the state error and covariance over the run. We see the covariance initially grow, then come down quickly as measurements are used. It grows once again as the measurements stop before bringing the state error back to zero with a change in sun heading.

Figure 5 shows the evolution of the state vector compared to the true values. The parts where there is a slight delay is due to the fact that no observations are read in.

Figure 6 shows the post fit residuals for the filter, with the 3σ measurement noise values. We see that the observations are read in well and that the residuals are brought back down to noise. We do observe a slight bias in the noise. This could be due to the equations of motion, and is not concerning.




Fig. 4: State error and covariance




Fig. 5: States tracking target values




Fig. 6: Post Fit Residuals