18th International Conference on Information Fusion
Washington, DC - July 6-9, 2015

Censoring in Distributed Radar Tracking Systems with Various Feedback Models

Armond Conte and Ruixin Niu
Dept. of Electrical and Computer Engineering
Virginia Commonwealth University
Richmond, VA, 23284, U.S.A
{conteas, rniu}@vcu.edu

Abstract* – In this paper, various methods to control censoring of local state estimates in a distributed multi-sensor radar tracking system are proposed. The data flow architectures that are used include three different feedback methods in order to achieve the goal of adequate state estimation using the fewest number of local state estimates sent to the fusion center as possible. The main novelty introduced in this paper is the use of J-Divergence for censoring the local state estimates in conjunction with the various feedback architectures. A simulation was run on these architectures with three spatially diverse sensors and a simple non-maneuvering target. Results show that distributed architectures that provide feedback from the fusion center to the local censoring processes or to the local state estimators, yield better estimates with more one-way communication savings.

Keywords: Tracking, Information Fusion, Extended Kalman Filtering, Censoring

1 Introduction

Taking advantage of the information gained through the combination of data in multiple sensors has shown dramatic potential in a wide variety of applications. Two common data flow architectures used in data fusion include centralized and distributed structures. In a centralized architecture the measurements from sensors are combined at a fusion center to obtain a state estimate. A distributed architecture has each sensor generate and maintain local state estimates that are sent to a fusion center to be fused to obtain a global state estimate. This therefore requires that the sensors have their own state estimation algorithms. One of the major motivations to use a distributed architecture rather than a centralized one is that the rate at which information is sent to the fusion center does not have to be as high as the measurement sampling rate in order to maintain a reasonable level of communication cost. Another reason why a distributed architecture might be used is the inherent redundancy, since local estimates are maintained at the sensors in addition to the fusion center. Some of the more general work in centralized and distributed architectures for data fusion is provided in [1-2]. A comparison between a centralized and a distributed architecture as well as an overview of the use of each in track association was provided in [3]. The concept of limiting the data rate between the local sensor and the fusion center was explored in [4].

Feedback from the fusion center to the local estimators has been studied in the past. The classical concept is that the prior state estimates at the local estimators are replaced by the global one. Some articles provide information on the use of feedback [1, 2, 5, 6]. According to [2], providing feedback at too low of a rate can cause divergence of the state estimators at local sensors.

Each individual sensor's data rate can be controlled by selectively censoring the information that goes to the fusion center. Intuitively selecting the data that go to the fusion center allows for potentially reduced data rate, while maintaining a reasonable degree of accuracy in the global state estimate. Some recent work that explored various censoring algorithms is provided in [6-8]. Of note, for both linear and nonlinear filtering problems, [6] used the normalized innovation squared (NIS) as a metric for determining the viability of measurements that should be censored along with feedback and a fusion center.

This paper analyzes the use of a distance measure (the J-divergence) between the posterior and prior state probability density functions (pdfs) to censor data at the local sensor level, using a 2-D radar scenario. It is assumed that the local sensors have a significantly lower energy budget than the fusion center, and that the cost of the local sensors receiving data from the fusion center is negligible. Thus, for the purposes of this paper, only the one-way communication savings from the local sensor to the fusion center are taken into account. A main difference between this paper and [6] is that in [6], the raw local sensor measurements are either censored or transmitted to the fusion center, while in this paper it is the local state estimates which are either censored or transmitted to the fusion center.

The remainder of this paper is organized as follows: Section 2 describes general background information including the equations in an extended Kalman filter (EKF) for use in 2-D radar tracking, more in-depth comparisons between a centralized architecture and a distributed architecture as well as equations that are used for fusion. Section 3 describes censoring local state estimates by the use of thresholding, the J-divergence criterion, and various censoring architectures involving feedback in a distributed system. Section 4 gives the outline of parameters for a simulation that was run to show the effects of censoring using various architectures as well as the results of this simulation. Section 5 provides a discussion of the simulation results. Section 6 concludes the paper.

2 Extended Kalman Filter

2.1 Assumptions

The EKF allows the performance of non-linear state estimation using general Kalman filter formulas by linearizing state-transition and measurement equations. The difference in the EKF is that the state transition and measurement equations may not be linear. The solution is linearizing the state transition and measurement equations. A general overview of the theory behind both the Kalman filter and the EKF can be found in [2] and [9].

2.2 2-D Radar Application

In 2-D Radar with a near-constant-velocity model the state vector is defined as (1) with its elements in Cartesian coordinates.

$$\mathbf{x} = \begin{bmatrix} x \\ \dot{x} \\ y \\ \dot{y} \end{bmatrix} \tag{1}$$

The measurement vector is defined in range and bearing as shown in (2), with elements in polar coordinates.

$$\mathbf{z} = \begin{bmatrix} R \\ \theta \end{bmatrix} \tag{2}$$

Using a linear white noise acceleration motion model [9], the state-transition matrix $\mathbf{F}$ is defined in (4), with ΔT as the difference in time between state estimate updates. $\mathbf{v}$ is the process noise with zero-mean and a covariance of $\mathbf{Q}$ as defined in (5), where $\tilde{q}$ is the power spectral density of the continuous process noise before its discretization over time. The state evolves linearly according to (3).

$$\mathbf{x}_{k+1k} = \mathbf{F}\mathbf{x}_{k|k} + \mathbf{v}_{k,i} \tag{3}$$

$$\mathbf{F} = \begin{bmatrix} 1 & \Delta T & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & \Delta T \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{4}$$

$$\mathbf{Q} = \begin{bmatrix} \frac{(\Delta T)^3}{3} & \frac{(\Delta T)^2}{2} & 0 & 0 \\ \frac{(\Delta T)^2}{2} & \Delta T & 0 & 0 \\ 0 & 0 & \frac{(\Delta T)^3}{3} & \frac{(\Delta T)^2}{2} \\ 0 & 0 & \frac{(\Delta T)^2}{2} & \Delta T \end{bmatrix} \tilde{q} \tag{5}$$

The measurements in 2-D Radar are nonlinear because of the translation between Cartesian and polar coordinate systems. Assuming that the position of the target at time k is given as (x_k, y_k) and the position of Radar i is given as (x_i, y_i). Then we define $x_{k,i}$ as in (5) and $y_{k,i}$ as in (6).

$$x_{k,i} = x_k - x_i \tag{5}$$

$$y_{k,i} = y_k - y_i \tag{6}$$

Similarly the measurement equation is defined in (7) as

$$\mathbf{z}_{k,i} = \begin{bmatrix} \sqrt{{x_{k,i}}^2 + y_{k,i}^2} \\ \mathrm{atan}\left(\frac{y_{k,i}}{x_{k,i}}\right) \end{bmatrix} + \mathbf{w}_{k,i} \tag{7}$$

where $\mathbf{w}_{k,i}$ is the measurement noise with a covariance defined as $\mathbf{R}$ in (8)

$$\mathbf{R} = \begin{bmatrix} \sigma_R^2 & 0 \\ 0 & \sigma_\theta^2 \end{bmatrix} \tag{8}$$

The measurement matrix in 2-D radar, denoted as $\mathbf{H}$, is given as (9) and is required for use in the normal Kalman filtering measurement prediction covariance and filter gain. This is approximated as the Jacobian of the coordinate conversions in (7).

$$\mathbf{H}_{k,i} = \begin{bmatrix} \frac{x_{k,i}}{\sqrt{x_{k,i}^2 + y_{k,i}^2}} & 0 & \frac{y_{k,i}}{\sqrt{x_{k,i}^2 + y_{k,i}^2}} & 0 \\ -\frac{y_{k,i}}{x_{k,i}^2 + y_{k,i}^2} & 0 & \frac{x_{k,i}}{x_{k,i}^2 + y_{k,i}^2} & 0 \end{bmatrix} \tag{9}$$

3 Fusion Architectures

3.1 Centralized Architecture

A centralized architecture is one where the individual sensor measurements are taken and combined at the fusion center. A diagram of this architecture is given in Fig. 1. It is expected that an increase in the total number of sensors that are providing measurements will yield a decrease in the total root mean squared error (RMSE) of the state estimate. In a centralized architecture, we can evaluate the utility of a measurement in two ways: 1) before the measurement is available, by using either information theoretic measures, such as the mutual information [12-13], or estimation theoretic measures, such as the conditional posterior Cramer-Rao lower bound (PCRLB) [14-15] to predict the utility of the measurement; 2) or after the measurement is available, by using the Kullback-Leibler (K-L) divergence or J-divergence between the prior and the posterior distributions, as proposed later in this paper.

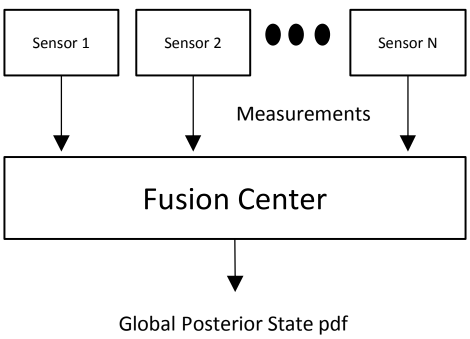

Figure 1: Centralized Fusion Architecture

3.2 Distributed Architecture

In a distributed architecture, each sensor preforms an individual state estimate and provides this to a fusion center, which calculates a global state estimate. A general distributed architecture is shown in Fig. 2. In a distributed architecture, the states do not need to be sent to the fusion center each time a new measurement is obtained. This has the potential to reduce the communication cost in comparison to the centralized architecture, but could also increase the error in the global state estimates if important updates are not sent to the fusion center.

For both linear and nonlinear filtering problems, the optimal formulas governing the combination of local state posterior pdfs at the fusion center for both cases that come with and without feedback have been provided in [1,4]. When an EKF is used as the state estimator, the linear versions of the fusion equations can be used. This is done using the information filter for the purpose of reducing computational complexity. Assuming that $\mathbf{P}(k\,|\,k)$ and $\mathbf{P}(k\,|\,k-1)$ are respectively the posterior and prior covariances, $\hat{\mathbf{x}}(k\,|\,k)$, and $\hat{\mathbf{x}}(k\,|\,k-1)$ are respectively the posterior and prior state estimates, the presence of the subscript i denotes the local covariance or state estimate for sensor i, and N is the total number of sensors being fused, the fusion equations are thereby given as (10) and (11) for cases where the fusion center does not feed its global state estimate back to local sensors. For cases with feedback from the fusion center to the sensors, these are given as (12) and (13).

$$\mathbf{P}(k|k)^{-1} = \mathbf{P}^{-1}(k|k-1) + \sum_{i=1}^{N}\left[\mathbf{P}_i^{-1}(k|k) - \mathbf{P}_i^{-1}(k|k-1)\right] \tag{10}$$

$$\begin{aligned}\mathbf{P}^{-1}(k\,|\,k)\hat{x}(k\,|\,k) &= \mathbf{P}^{-1}(k\,|\,k-1)\hat{\mathbf{x}}(k\,|\,k-1)\\ &+ \sum_{i=1}^{N}\left[\mathbf{P}_i^{-1}(k\,|\,k)\hat{\mathbf{x}}_i(k\,|\,k) - \mathbf{P}_i^{-1}(k\,|\,k-1)\hat{\mathbf{x}}_i(k\,|\,k-1)\right]\end{aligned} \tag{11}$$

$$\mathbf{P}(k|k)^{-1} = \sum_{i=1}^{N}\left[\mathbf{P}_i(k|k)^{-1}\right] - (N-1)\mathbf{P}^{-1}(k|k-1) \tag{12}$$

$$\begin{aligned}\mathbf{P}^{-1}(k\,|\,k)\hat{x}(k\,|\,k) &= \sum_{i=1}^{N}\left[\mathbf{P}_i^{-1}(k\,|\,k)\hat{\mathbf{x}}_i(k\,|\,k)\right]\\ &- (N-1)\mathbf{P}^{-1}(k\,|\,k-1)\hat{\mathbf{x}}(k\,|\,k-1)\end{aligned} \tag{13}$$

It is obvious, but should be noted that $\mathbf{P}(k\,|\,k)$ should be multiplied by the results of (11) and (13) in order to obtain the state estimate.

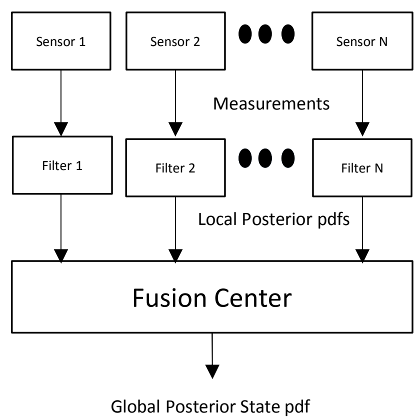

Figure 2: Distributed Fusion Architecture

3.3 Censoring Local State Estimates

Censoring state estimate updates to the fusion center can be done when the goal is to reduce communication cost. Non-Gaussian updates may require a more elaborate determination of the threshold. One method to determine this threshold is based on the K-L divergence. This method can be thought of as a measurement of the "distance" between two probability distributions. This is defined in [10, 16] as (14)

$$D(p(x)\,\|\,q(x)) = \sum_x \left[p(x)\log\left(\frac{p(x)}{q(x)}\right)\right] \tag{14}$$

For the KL-Divergence of two random Gaussian distributions (p(x) and q(x)), this was adapted from [11] as (15). Note that in this equation $\mathbf{P}$ and $\hat{\mathbf{x}}$ are respectively given as the covariance of and state estimate of the subscripted distributions and d is the dimension of $\hat{\mathbf{x}}$.

$$D(p(x)\,\|\,q(x)) = \frac{1}{2}\,\mathrm{tr}\left(\mathbf{P}_{q(x)}{}^{-1}\mathbf{P}_{p(x)}\right) - \ln\frac{\left|\mathbf{P}_{p(x)}\right|}{\left|\mathbf{P}_{q(x)}\right|} - d + \frac{1}{2}\left(\hat{\mathbf{x}}_{q(x)} - \hat{\mathbf{x}}_{p(x)}\right)^T \mathbf{P}^{-1}_{p(x)}\left(\hat{\mathbf{x}}_{q(x)} - \hat{\mathbf{x}}_{p(x)}\right) \tag{15}$$

This is not a formal distance since the value will differ depending on the direction. In order to make the K-L divergence symmetric the J-divergence is introduced. A variation of the J-divergence is defined in (16) as referenced in [10].

$$J(p(x), q(x)) = \left(\frac{1}{2}D(p(x)\,\|\,q(x))^2 + \frac{1}{2}D(q(x)\,\|\,p(x))^2\right)^{\frac{1}{2}} \tag{16}$$

For a distributed sensor system, the inputs to the K-L divergence calculation are the prior and posterior pdfs. There are 3 different methods that are analyzed in this paper, which are described in the next subsections. The differences between each of them are outlined in Table I below.

Table I: Feedback Model Inputs

MODEL #	Threshold Input	Sensor Filter Input
I	Sensor prior pdf	Sensor prior pdf
II	Global prior pdf	Sensor prior pdf
III	Global prior pdf	Global prior pdf

3.4 MODEL I: No Feedback

In the Model I the data flow is straightforward. Each local state estimator compares the prior state pdf and the posterior state pdf. If there is significant amount of information in the new measurement, then the distance between the prior and posterior state pdfs exceeds a given threshold, and the estimate is allowed to be transmitted to and fused at the fusion center.

There is no control at the fusion center regarding censoring when using this method. Therefore whether or not the fusion center diverges is dependent on the rate at which updates are provided to the fusion center.

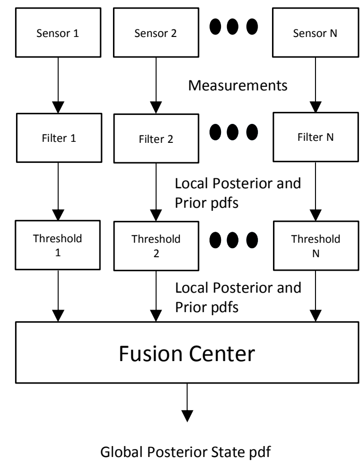

Figure 3: Model I - No Feedback

3.5 MODEL II: Threshold Feedback

Model II feeds the global state estimate back to the threshold calculation and uses it as prior information. The idea is that the decision upon whether or not the information is passed to the fusion center lies in the hands of the fusion center itself rather than the individual sensors. If the global estimate begins to diverge separately from that of the local estimates, the distance between the global and local pdfs should also grow.

Comparing this method to Model I, the J-Divergence threshold should be tuned differently. Since the fusion center maintains its state estimates separately from the sensors, there may be larger differences rather than simply using the local state estimates alone. If the threshold is set too low, then few, if any, updates will remain unused. Therefore a higher threshold for the J-Divergence metric must be used in Model II.

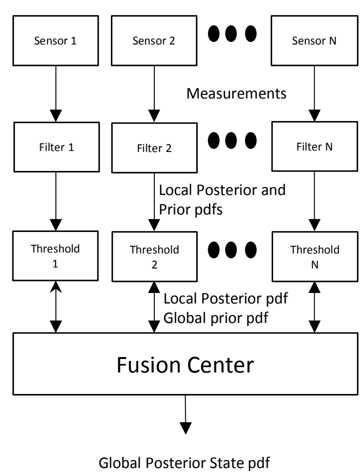

Figure 4: Model II - Feedback to Local Thresholds

3.6 MODEL III: Global Feedback

The third method feeds the information from the fusion center back to the sensor-level estimation algorithms as the prior. Inherently, the prior information in the thresholds is the same as that which was given to the sensor. In this method even more control is provided to the fusion center.

In determining an adequate censoring threshold, it should be considered that the local filters already use the global state estimate. This means that the updates are likely to be closer to each other than in the thresholding performed in Model II. Although this is problem dependent, comparing Model III to Model I for a steady state case should yield a similar number of unused updates.

One difference between Model III and Model I/II is in the use of equations (12) and (13) instead of (10) and (11) at the fusion center. It should be noted that between iterations the number of sensors providing inputs may change dynamically and this needs to be recorded as an input to (12) and (13).

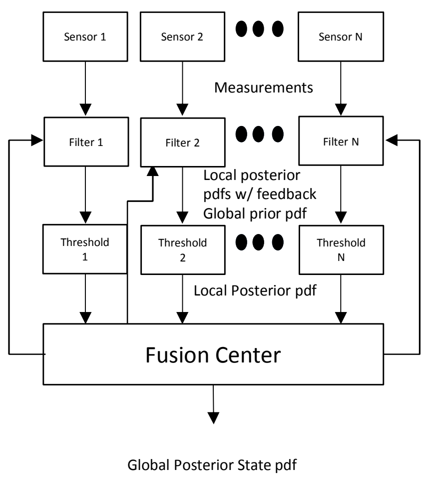

Figure 5: Model III - Feedback to Local State Estimators

4 Simulation

4.1 Parameters

The parameters for the sensors involve their 2-D positions (Pos.) and measurement accuracies (Acc.) in range and bearing. There were 3 sensors in total with accuracies are the same for each sensor and are 100m in range, 1 deg in bearing respectively. Measurements on the target were taken at a rate of 100Hz and it was assumed that all were received and properly ordered by the fusion center. It is assumed that the fusion center has a significantly larger energy budget than the individual sensors and so, the fusion center also performs feedback at a rate of 100Hz, for Models II and III. 50 seconds of simulation time data were taken. For the purpose of analysis, the data were recorded over a set of 100 Monte-Carlo runs. The value of $\tilde{q}$ in (5) was set to 1.

Throughout each run, the sensors maintained constant positions, which are listed in Table II.

Table II: Sensor Positions

Sensor #	Pos. (x,y) (km)
1	(0,10)
2	(10,25)
3	(25,10)

The target follows a near constant velocity model with a set start position and velocity. For the purposes of this simulation this is given as follows.

Table III: Initial Target State

Tar. Param	Value
x – start	5 km
y – start	20 km
x – velocity	300 m/s
y - velocity	-100 m/s

The motivation behind the choice of initial target state and the sensor locations is such that the target moves through the center of the set of sensors. In this way the observability of each sensor will have a significant change as the target traverses the surveillance area. The choice of scaling was dependent upon the velocities, which were selected such that they were just below Mach 1 (~340 m/s). In this scenario, the speed is 316.23 m/s. A plot of the target trajectory along with the sensor locations is given in Fig. 6.

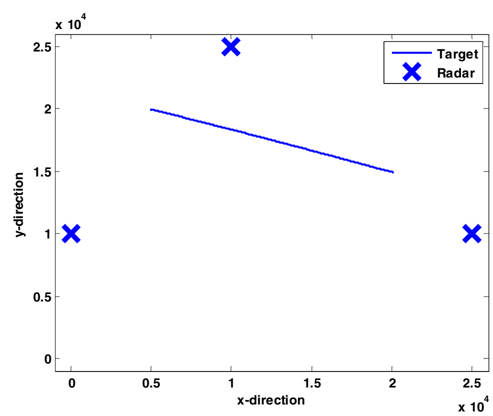

Figure 6: Target Trajectory and Sensor Locations

The sensor state and covariance are initialized using two-point differencing for the debiased converted measurements Kalman Filter as described in [2]. The fusion center is initialized by copying the state estimate of a randomly selected sensor.

The selection of the threshold used in the J-Divergence calculation was performed using trial and error. Over several runs, the number of unused updates for various J-divergence values in each model was noted and the threshold was set based on this. Table IV in the results section provides an overview of this. More theoretical analysis of the relationship between the threshold on the J-divergence and a desired global estimation accuracy will be conducted in our future work. For Models I and III, the J-divergence threshold was set to run between 0 and 0.05 with a step size of 0.01. Model II was set to run between 0 and 10 with a step size of 2.

4.2 Results

Table IV: J-Divergence threshold with Average Unused Updates (One-Way Communication Savings)

J-Divergence I& III/II	Percentage of Unused Updates		
	Model I	Model II	Model III
0/0	0	0	0
0.01/2	65.05	6.83	64.83
0.02/4	89.55	78.13	80.82
0.03/6	93.54	89.75	86.39
0.04/8	95.2	94.84	89.32
0.05/10	96.17	96.14	91.15

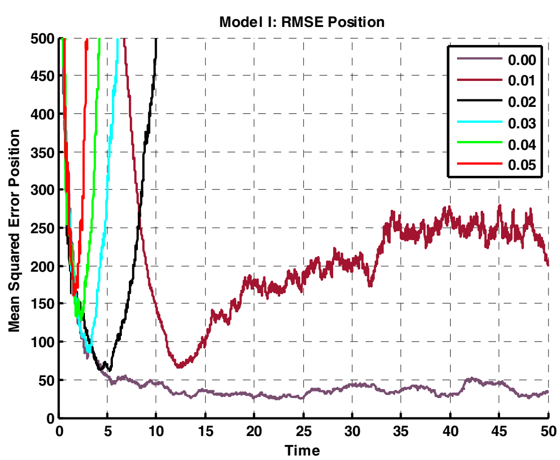

Figure 7: Model I - Positional RMSE vs. Time

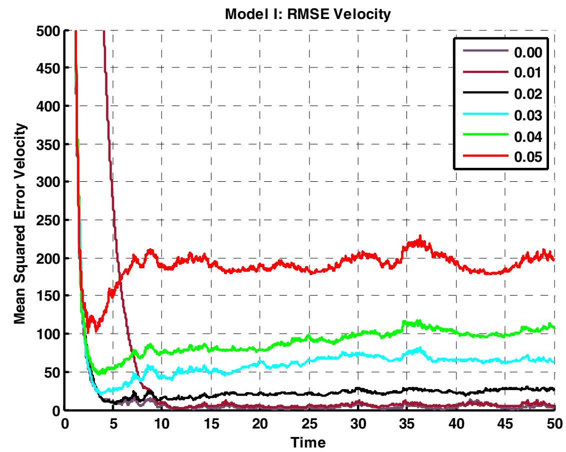

Figure 8: Model I - Velocity RMSE vs. Time

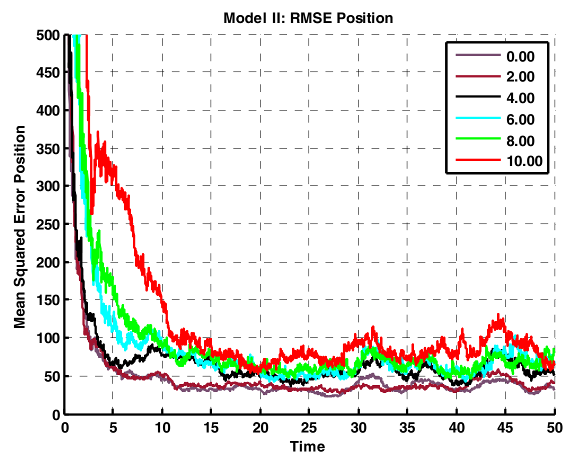

Figure 9: Model II - Positional RMSE vs. Time

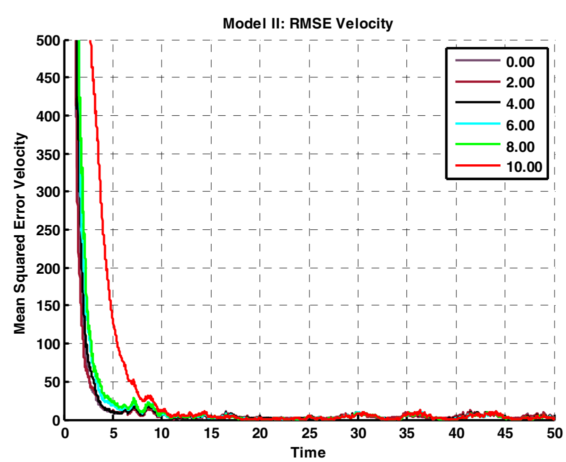

Figure 10: Model II - Velocity RMSE vs. Time

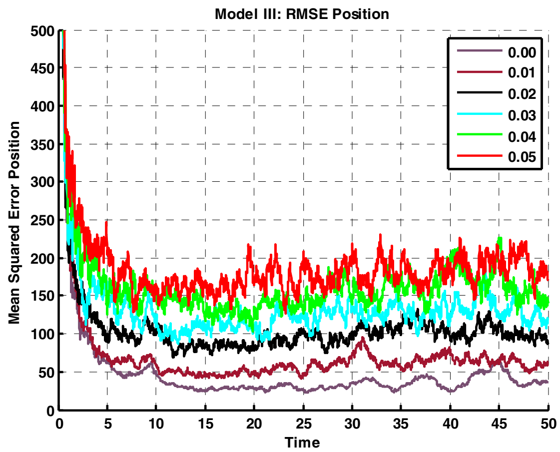

Figure 11: Model III - Positional RMSE vs. Time

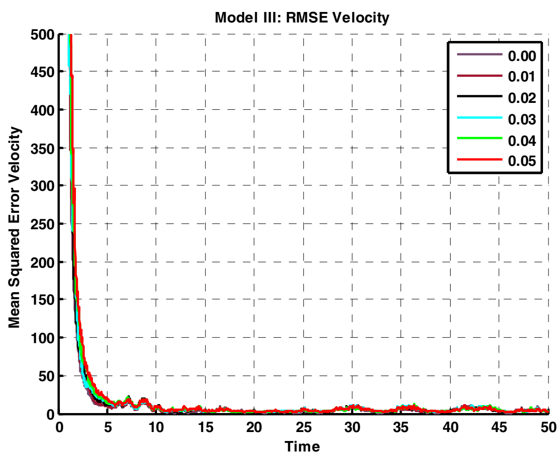

Figure 12: Model III - Velocity RMSE vs. Time

5 Discussion

It can be inferred from Table IV and Figs. 7-12 that Model II and Model III are able to better stabilize the RMSE with an increase in the number of unused updates in comparison to Model I. This is likely because of the control enabled by the global estimator. When updates are unused the covariance is expected to increase each time a prediction is made. Allowing feedback from the fusion center to the threshold implies that after a certain covariance is reached, the global estimator will attempt to reduce this value. By only allowing control of the threshold at the sensor level, as in Model I, the covariance at the global estimator may drastically increase when the sensors do not determine their information to be useful enough to perform an update (consequently yielding a J-divergence above the threshold). The effect shown in the Model I figures show an extreme increase in MSE with a similar number of unused updates to Model II and Model III.

In contrasting the results of Model II to Model III, it appears that the Model III results are more stable than Model II, but it is unclear which one is better. Fig. 9 and Fig. 10 (Model II) show fluctuations in the RMSE at the end of the run, but generally appear to yield similar RMSE results to Figs. 11 and 12 (Model III). This could be due to the increased stability provided by feeding back the fusion center state estimates to the local state estimators rather than to just the threshold calculations. Based on Table IV, Model II is able to maintain this degree of estimation accuracy with an average higher number of unused updates, and hence more communication savings.

6 Conclusion

Censoring approaches based on J-divergence between the local posterior and prior state pdfs have been proposed for distributed tracking architectures with different feedback mechanisms, with the goal of reducing communication costs while maintaining a reasonably good tracking performance. These were tested on a multiple 2-D radar simulation using a near-constant velocity target model. Results showed that providing feedbacks from the global estimator to the thresholds or to the local estimators, has the potential to yield better results with greater communication savings.

7 References

[1] C. Y. Chong, S. Mori, and K. C. Chang, Distributed Multitarget Multisensor Tracking, in Y. Bar-shalom, ed., Multitarget-Multisensor Tracking: Advanced Applications, Artech House, Chapter 8, 1990.

[2] Y. Bar-Shalom, P.K. Willett, and X. Tian, Tracking and Data Fusion: A Handbook of Algorithms, YBS Publishing, 2011.

[3] H. Chen, T. Kirubarajan, and Y. Bar-Shalom, “Performance limits of track-to-track fusion versus centralized estimation: theory and application,” IEEE Trans. Aerosp. Electron. Syst., vol. 39, no. 2, pp. 386–400, Apr. 2003.

[4] F. Govaers and W. Koch, “An exact solution to track-to-track-fusion at arbitrary communication rates,” IEEE Trans. Aerosp. Electron. Syst., vol. 48, no. 3, pp. 2718–2729, Jul. 2012.

[5] Y. M. Zhu, Z. S. You, J. Zhou, K. S. Zhang, and X. R. Li, “The optimality for the distributed Kalman filtering fusion with feedback,” Automatica, vol. 37, pp. 1489–1493, Sep. 2001.

[6] Y. Zheng, R. Niu, and P. K. Varshney, “Sequential Bayesian Estimation with Censored Data for Multi-Sensor Systems”. IEEE Trans. Signal Process., vol. 62, no. 10, pp. 2626-2641, May 2014.

[7] C. Rago, P. Willett, and Y. Bar-Shalom “Censoring Sensors: A Low-Communication-Rate Scheme for Distributed Detection”. IEEE Trans. Aerospace and Electronic Systems., vol. 32, no. 2, pp. 554-568, Apr 1996.

[8] R. Jiang and B. Chen, “Fusion of censored decisions in wireless sensor networks,” IEEE Trans. Wireless Commun., vol. 4, pp. 2668–2673 Nov. 2005.

[9] Y. Bar-Shalom, X. Rong Li, and T. Kirubarajan, Estimation with Applications to Tracking and Navigation, John Wiley & Sons, Inc., 2001.

[10] J. C. Principe, Information Theoretic Learning: Renyi’s Entropy and Kernel Perspectives, Springer Science+Business Media, LLC., 2010.

[11] S. J. Roberts and W. D. Penny, “Variational Bayes for generalized autoregressive models,” IEEE Trans. Signal Process. vol. 50, pp. 2245–2257, Sep. 2002.

[12] J. Williams, J. Fisher, and A. Willsky, “Approximate dynamic programming for communication-constrained sensor network management,” IEEE Trans. Signal Process. vol. 55, no. 8, pp. 4300–4311, Aug. 2007.

[13] G. M. Hoffmann and C. J. Tomlin, “Mobile sensor network control using mutual information methods and particle filters,” IEEE Trans. Autom. Control, vol. 55, no. 1, pp. 32–47, Jan. 2010.

[14] L. Zuo, R. Niu, and P.K. Varshney, “Conditional Posterior Cramer-Rao Lower Bounds for Nonlinear Sequential Bayesian Estimation,” IEEE Transactions on Signal Processing, Vol. 59, No. 1, pp. 1-14, January 2011.

[15] E. Masazade, R. Niu, and P.K. Varshney, “Dynamic Bit Allocation for Object Tracking in Wireless Sensor Networks,” IEEE Transactions on Signal Processing, Vol. 60, No. 10, pp. 5048-5063, October 2012.

[16] T. M. Cover and J. A. Thomas, Elements of Information Theory, 2nd Edition, Wiley-Interscience, 2006.