Effect of Retransmission and Retrodiction on Estimation and Fusion in Long-Haul Sensor Networks

Qiang Liu, *Member, IEEE,* Xin Wang, *Member, IEEE,*
Nageswara S. V. Rao, *Fellow, IEEE,*
Katharine Brigham, *Student Member, IEEE,* and B. V. K. Vijaya Kumar, *Fellow, IEEE*

***Abstract*—In a long-haul sensor network, sensors are remotely deployed over a large geographical area to perform certain tasks, such as target tracking. In this work, we study the scenario where sensors take measurements of one or more dynamic targets and send state estimates of the targets to a fusion center via satellite links. The severe loss and delay inherent over the satellite channels reduce the number of estimates successfully arriving at the fusion center, thereby limiting the potential fusion gain and resulting in suboptimal accuracy performance of the fused estimates. In addition, the errors in target-sensor data association can also degrade the estimation performance. To mitigate the effect of imperfect communications on state estimation and fusion, we consider retransmission and retrodiction. The system adopts certain retransmission-based transport protocols so that lost messages can be recovered over time. Besides, retrodiction/smoothing techniques are applied so that the chances of incurring excess delay due to retransmission are greatly reduced. We analyze the extent to which retransmission and retrodiction can improve the performance of delay-sensitive target tracking tasks under variable communication loss and delay conditions. Simulation results of a ballistic target tracking application are shown in the end to demonstrate the validity of our analysis.**

***Index Terms*—Long-haul sensor networks, state estimation and fusion, data association, message retransmission, prediction and retrodiction, mean-square-error (MSE) and root-mean-square-error (RMSE) performance.**

I. Introduction

Sensor networks, consisting of sensor nodes that are equipped with sensing, data processing, and communication components, have been deployed for a wide variety of applications, including healthcare, cyber security, environmental monitoring, and national defense, among others [1]. In a long-haul sensor network, sensors are deployed to cover a vast geographical area, which could be a continent or even the entire globe depending on the specific application. We consider a class of such networks in which state estimates (e.g., position and velocity) of certain dynamic targets – such as aircrafts or ballistic missiles [6] – are sent from the remote sensors to a fusion center so that a global estimate can be obtained by fusing the individual estimates. Under certain circumstances, satellite links may be the only type of cost-effective medium for such long-range communications because of the prohibitive cost of extending submarine and terrestrial fiber connections extensively to rough terrains and sparsely populated areas. We in particular focus on such satellite link-based monitoring and tracking applications, wherein measurement data are collected at the remote sensors and state estimates are individually generated. These estimates are then sent via long-haul satellite links to the fusion center, which, upon successful reception of a subset of these estimates, applies a certain fusion rule and obtains the final estimate to be reported.

This work is motivated by the many challenges arising from the imperfect communications over the long-haul satellite links. Because of the long distance (tens of thousands of miles), the signal propagation time is significant. For example, the round-trip time (RTT) for signal propagation with a geostationary earth orbit (GEO) satellite is more than a half second [21]. More importantly, communication over the satellite links is characterized by sporadic high bit-error rates (BERs) and burst losses. Losses either incurred during transmission or resulting from the high BERs could further effectively reduce the number of messages available at the fusion center. It is well known that fusion of estimates from different sensors is a viable means of reducing the estimation error; with high loss rates, however, only a portion of the potential fusion gain could be achieved and the quality of the fused estimate output obtained may be deemed unacceptable by the system operator. In addition, prior to fusion, the estimates generated by the sensors must be associated with the corresponding targets. When multiple targets are in close proximity or in clutter [5], an estimate is very likely to be associated with a wrong target; this is only to be exacerbated by the long-haul target-sensor link. Apparently, all the above-mentioned drawbacks of the satellite links could work against the very purpose of the underlying task – to promptly and accurately report state estimates – and may result in failure to comply with requirements on the worst-case estimation error.

Manuscript received ... ; revised ... .

This work was supported by the Mathematics of Complex, Distributed, Interconnected Systems Program, Office of Advanced Computing Research, U.S. Department of Energy, and the SensorNet Project of Office of Naval Research, and was performed at Stony Brook University with NSF awards CNS 1247924, ECCS 1231800, and ECCS 1408247.

Q. Liu and X. Wang are with the Department of Electrical and Computer Engineering, Stony Brook University, Stony Brook, NY, 11794. E-mail: {qiangliu, xwang}@ece.sunysb.edu.

N. S. V. Rao is with Oak Ridge National Laboratory, Oak Ridge, TN, 37831. Email: raons@ornl.gov.

K. Brigham and B. V. K. Vijaya Kumar are with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, 15213. E-mail: {kbrigham, kumar}@ece.cmu.edu.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identier

A. Related Works

State estimation under imperfect communications has been studied in the literature. State augmentation [23] can handle fixed delay up to several sampling periods. In [7], [9], [22], [24], estimation and fusion performance using Kalman filters (KFs) under variable packet loss rates has been considered. A dynamic selective fusion method based on information gain is proposed in [15] so that fusion is deferred till enough information has arrived at the fusion center. A staggered estimation scheduling scheme is proposed in [14] that aims to explore the temporal domain relationships of adjacent data within an estimation interval to improve the estimation and fusion performance. More recently, [16] has considered an information feedback mechanism where fused estimates are fed back to a subset of sensors in order to improve their information quality, and in turn, the overall fusion performance.

In contrast to these works, we consider the recovery of missing information over time. One way to counteract the effect of the lossy transmission link is to adopt certain transport protocols in which message retransmission is implemented and some lost messages can be recovered after one or multiple rounds of retransmission. Not to be overlooked, however, is another aspect of the system requirement – the delay performance. Owing to often near real-time requirements of the monitoring/tracking tasks, the system often allows for only a small time gap between the time of interest and the time when the estimate should be finally obtained and reported. This often comes as a predefined reporting deadline before which an estimate must be reported by the fusion center. Message retransmission may exacerbate the reporting delay performance by incurring extra time on top of the already relatively large propagation and transmission latency. The fusion center may have to increase its reporting time significantly in order to recover the lost messages, even at the risk of violating the stipulated reporting deadline.

The transmission control protocol (TCP) implemented in wired Internet and wireless local area networks (WLANs) is still garnering research efforts that are too numerous to list. Analysis of TCP-like transport protocols over satellite links can be found in studies such as [2] and [10]. Commonly acknowledged are the difficulties in applying "conventional" TCP protocols to transmission over satellite links, mainly because of the very large propagation delay not encountered in other networks. The specificity of our application also somewhat distinguishes our analysis from the ones geared toward the voice- and video-based broadcasting and data-based Internet access, both of which have continuous data in flight. Also of note is that in our settings, state estimates from the remote sensors are generally intermittently sent over a wideband satellite channel – with the interval possibly ranging from a few times within a second to once every few minutes – and thus congestion is not as much a concern as in conventional TCP applications. Hence, we assume a simplified transport protocol in which retransmission is performed on the message-level basis.

In many state estimation applications, retrodiction, also known as smoothing, serves as the "backward prediction" of an earlier estimate. Depending on the relationship between the length of data used and the time of interest, we can categorize retrodiction roughly into fixed-point, fixed-lag, and fixed-interval retrodiction [23]. Whereas the conventional retrodiction techniques are used mainly for improving estimates that have been obtained, for instance, in the context of out-of-order measurement (OOSM) problems [18], [26], [27], we are primarily interested in how missing estimates can be interpolated from retrodiction[1] and the corresponding improvement in estimation errors following such retrodiction.

Another important aspect of the multi-target tracking task is that prior to fusion, the data association algorithm determines the groups of estimates, so that each group is hypothesized to correspond to the same target. Many types of association and fusion algorithms for tracking and navigation applications have been studied [4]. However, our focus here is not on performance comparison among different algorithms; rather, of interest to us is the performance improvement from retransmission and retrodiction following a given simple set of association and fusion algorithms.

B. Our Contributions

In this work, we provide analytical models to systematically study the impact from retransmission and retrodiction on target-tracking performance under variable loss and delay conditions in a long-haul sensor network. This work is among the very first to link both communication (message retransmission) and computation (prediction, retrodiction) with state estimation performance, accounting for both data fusion and association in target tracking applications. Simulations of a coasting ballistic target tracking example are conducted and results under various conditions are shown in the end to validate our analysis.

The remainder of this paper is organized as follows: We first provide analysis of the delivery rate of a message after retransmission in Sec. II, and derive the arrival time distribution in Sec. III. Following analytical studies on retransmission, we explore the effect of retrodiction on the estimation performance improvement in Sec. IV, where non-cooperative and cooperative types of retrodiction are discussed. Simulation results of a coasting ballistic target tracking application are presented and analyzed in Sec. V before we conclude the paper in Sec. VI.

II. Message Retransmission

In a long-haul sensor network, a remote sensor sends out a message containing the state estimate; upon successful reception of this message, the fusion center sends back an acknowledgment (ACK) message to the sensor. A failed arrival of the ACK message before the expiration of the timeout T_{TO} – due to loss and/or long delay of the message itself or the ACK – will prompt the sensor to retransmit the message. Typically, T_{TO} could be several times the RTT of the connection, and over long-haul connections it could be of the order of

[1]In the meantime, an available estimate is retrodicted by subsequent estimate(s) as well whenever applicable as in conventional retrodiction.

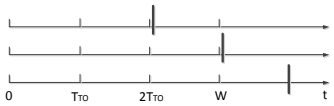

Fig. 1: Timing of message retransmission: the timeout is T_{TO}, retransmission window is W, and different choices of the cutoff time T_{CO} are marked by bold lines

seconds. Setting T_{TO} too long could reduce the maximum number of retransmissions, thereby limiting the potential to recover the lost message; on the other hand, a short T_{TO} may incur many rounds of retransmission (often unnecessarily) when the sensor could have waited a bit longer to receive the ACK. Such retransmission continues till the acknowledgment is received by the sensor, or the retransmission window W expires. This window should ideally contain multiple T_{TO} periods so that under adverse link conditions, it's likely that the message can eventually be recovered after multiple tries.

In a real system, the reporting deadline may limit the potential gain from retransmission as the overall time before reporting can be very short. A cutoff time T_{CO} is defined to mark the end of the waiting at the fusion center. This cutoff on the one hand limits the total number of retransmissions, and on the other limits certain messages from being eventually delivered due to the randomness of the delay. In Fig. 1, the effect of this time cutoff is shown. The window W is set to be $3T_{TO}$ and hence there are a total of two rounds of retransmision (at T_{TO} and $2T_{TO}$). In the first case, T_{CO} is small so that the last round of retransmitted message cannot arrive in time for fusion. While in the last case, setting T_{CO} way beyond the end of the retransmission window is not likely to significantly increase the chance of receiving the message. Therefore, the system should guarantee that the retransmission window – at the sensors side – is commensurate with the cutoff time – at the fusion center, as in the second case in the figure, so that the fusion center could benefit from all rounds of retransmission while not wasting time attempting to recover the pending message after the retransmission has ended.

The message-level loss and delay characteristics are determined by the long-haul link conditions. We assume that each message sent by a sensor is lost during transmission with probability p_L independently of other messages. Normally, the latency that a message experiences before arriving at the fusion center consists of the initial detection and measurement delay, data processing delay by both the sensor and the fusion center, propagation delay, and transmission delay, among others[2]. These are collectively considered as the minimum delay that a message must undergo to reach the fusion center, which is bound mostly by factors such as the distance of the satellite link, the transmission data rate, and length of the message. The extra random delay is often due to link conditions such as weather and terrain. We suppose a pdf $f(t)$ can model the overall delay t that a message experiences to be successfully delivered to the fusion center. One typical example is that of the shifted exponential distribution:

$$f(t) = \frac{1}{\mu_D} \exp^{-\frac{t-T}{\mu_D}}, \text{ for } t \geq T. \tag{1}$$

[2]The queueing delay is also minimal with little/no congestion.

in which T serves as the common link and processing delay, and μ_D is the mean of the random delay beyond T. In a real system, the empirical values of the message delay can be measured over time and thus an approximate function $\tilde{f}$ can be estimated. In the following analysis, however, we still use the generic function $f(t)$ to model the arrival delay.

We are interested in the average probability of a message being successfully delivered by a certain time, that is, by the cutoff time T_{CO}. An estimate is only counted once even if it arrives multiple times due to retransmission. The duplicate messages received by the FC can simply be ignored as they will not contribute further to the fusion performance.

With the time of interest being regarded as time zero in this section, the maximum number of retransmissions before the cutoff time T_{CO} is

$$K_{retx} = \left\lceil \frac{\min\{T_{CO}, W\}}{T_{TO}} \right\rceil - 1. \tag{2}$$

From the definition, $K_{retx}+1$ is the total rounds of transmission, including the original and subsequent retransmissions.

We define $p^k_{del,t}$ as the probability that a message is delivered by time t after k rounds of retransmissions, and

$$T_{retx,k} = T_{CO} - kT_{TO}, \text{ for } k = 0, 1, ..., K_{retx} \tag{3}$$

as the duration of the period $[kT_{TO}, T_{CO}]$ in which the k-th retransmitted message is in flight and could be potentially delivered to the fusion center.

When there is no retransmission within $[0, t]$, the probability of a message being delivered by time t is

$$p^0_{del,t} = (1 - p_L)F(t), \tag{4}$$

in which $F(t) = \int_0^t f(u)\,\mathrm{d}u$ is the cdf of the arrival delay. Its complement, the probability that the original message is unavailable at time t, is denoted as

$$p^0_{loss,t} = 1 - p^0_{del,t} = p_L + (1 - p_L)\overline{F}(t), \tag{5}$$

in which $\overline{F}(t) = 1 - F(t)$ is the tail distribution. With these two probabilities, we can derive the message delivery rate $p^{K_{retx}}_{del,T_{CO}}$.

The original message is delivered by T_{CO} with probability

$$p^0_{del,T_{CO}} = p^0_{del,T_{retx,0}} = (1 - p_L)F(T_{retx,0}). \tag{6}$$

And with the first round of retransmission, the delivery probability totals

$$p^1_{del,T_{CO}} = p^0_{del,T_{CO}} + p^0_{loss,T_{retx,0}} p^0_{del,T_{retx,1}}. \tag{7}$$

In general, for the k-th ($0 < k \leq K_{retx}$) round of message retransmission, we have

$$p^k_{del,T_{CO}} = p^{k-1}_{del,T_{CO}} + p^0_{del,T_{retx,k}} \left(\prod_{i=0}^{k-1} p^0_{loss,T_{retx,i}} \right). \tag{8}$$

In other words, the extra delivery rate from the k-th round is realized when all the previous $k-1$ retransmissions and the original message are not available by T_{CO}. Subsequently, we

can obtain the overall message delivery probability within time $[0, T_{CO}]$ by summing up all such probabilities for $K_{retx} > 0$:

$$\begin{aligned} p^{K_{retx}}_{del,T_{CO}} &= p^0_{del,T_{CO}} + \sum_{k=1}^{K_{retx}} p^0_{del,T_{retx,k}} \left(\prod_{i=0}^{k-1} p^0_{loss,T_{retx,i}} \right) \\ &= (1-p_L)F(T_{CO}) + \\ &(1-p_L) \sum_{k=1}^{K_{retx}} F(T_{retx,k}) \left\{ \prod_{i=0}^{k-1} [1-(1-p_L)F(T_{retx,i})] \right\}. \end{aligned} \tag{9}$$

III. Arrival Time Pattern Distribution

So far we have combined the message-level loss rate and one-way arrival delay distribution to obtain the message delivery probability with a certain deadline requirement as specified by T_{CO}. In practice, the fusion center should be afforded some flexibility in deciding its actual cutoff time so that it does not violate the systemic value T_{CO}. First, sometimes reporting an abnormal change promptly is more crucial than initially providing an accurate estimate because the alert level can be increased immediately facilitating further investigation. As the maximum allowable delay, T_{CO} may nevertheless be too large for such rare but time-critical incidences. On the other hand, owing to the dynamic environment of the field, sometimes the fusion center may decide to reduce its waiting time for the retransmitted messages because of increased computational effort to obtain the final estimate, due to an increased number of sensors or increased state dimensionality when multiple closely positioned targets are in clutter [5].

In this section, we aim to derive the distribution of the arrival time, and more particularly, the cdf of the first instance of arrival before T_{CO}. This provides us a view of the internal structure of the arrival process within $[0, T_{CO}]$, which could be explored for the above flexible scheduling of early cutoff.

A. Arrival Time: One-way Communication Analysis

We treat loss and delay as two independent processes, although a lost message can be regarded as having an arrival delay of infinity. In the previous section, only the one-way delay characterized by the pdf $f(t)$ is considered because we noticed the equivalence of the final delivery probability no matter how acknowledgments are actually received. In deriving the arrival time, however, we need to consider two-way delay as well: loss and latency of ACKs would affect the total number of retransmissions, which in turn decides the distribution of the arrival time. For ease of explication though, we first work on the case in which there are exactly K_{retx} retransmissions – as if no ACKs were ever sent back by the fusion center – and later extend the results to an arbitrary number of retransmissions.

First, we define the cdf of a truncated nonnegative random variable Y_T with the upper truncation point $b > 0$ as[3]

$$F^b_T(y) = \frac{F(y)}{F(b)}, \text{ for all } 0 \le y \le b. \tag{10}$$

[3]Note that the subscript "T" appearing in cdf's and pdf's denote the function describes a truncated random variable.

And the associated pdf is

$$f^b_T(y) = \frac{f(y)}{F(b)}, \text{ for all } 0 \le y \le b. \tag{11}$$

In our study, we are interested in a series of truncated cdfs and pdfs corresponding to different retransmission cycles. More specifically, we consider the k-th round of retransmitted message has a truncated cdf by time $T_{retx,k}$ as

$$F^{T_{retx,k}}_T(t) = \frac{F(t)}{F(T_{retx,k})}, \text{ for all } 0 \le t \le T_{retx,k}. \tag{12}$$

Let D denote the arrival time of the message, and more specifically, $D_k = d_k + kT_{TO}$ the arrival time of the k-th retransmitted message (or the original message when $k = 0$). Apparently, d_k denotes the arrival delay of the k-th retransmission.

We are interested in deriving the distribution of $D_{(1)}$ – the time of the first arrival – before the cutoff time T_{CO}. When there are a maximum number of K_{retx} retransmissions before T_{CO}, we let

$$\begin{aligned} D^{K_{retx}}_{(1),T_{CO}} &= \min_{k=0,\ldots,K_{retx}} D^{K_{retx}}_{k,T_{CO}} \\ &= \min_{k=0,\ldots,K_{retx}} \{d^{K_{retx}}_{k,T_{CO}} + kT_{TO}\} \end{aligned} \tag{13}$$

be the time of the first arrival among all $K_{retx}+1$ messages sent out by the sensor. We note

$$\Pr\{D^{K_{retx}}_{(1),T_{CO}} \le T_{CO}\} = p^{K_{retx}}_{del,T_{CO}}, \tag{14}$$

where the right-hand side of the equation was given in Eq. (9). Our goal is to derive the distribution of $D^{K_{retx}}_{(1),T_{CO}}$, that is, $\Pr\{D^{K_{retx}}_{(1),T_{CO}} \le t\}$ for any $0 < t < T_{CO}$. For ease of presentation, in the remainder of this section, we assume that a certain T_{CO} value has been specified along with the resulting K_{retx} and drop them from the notations unless otherwise specified.

One may first be tempted to directly apply the result of the distribution of the minimum of n random variables, which is a special case of the order statistics [19]. Despite the seemingly similar relationship, the problem at hand is more complicated. First, from Eq. (13), we need to find the minimum of D_k for $k = 0, 1, 2, ..., K_{retx}$, which are from different distributions for different k. Although results for independently and non-identically distributed random variables have been studied in the literature [3], the operations involve substantial use of matrix manipulation, and one must enumerate all $2^{K_{retx}+1}$ possible arrival patterns since each would yield a distinct result for the distribution of the minimum. To circumvent the issue, we follow another approach by finding the probability that the k-th retransmitted message (and the original message when $k = 0$) is the earliest to arrive, denoted as $\Pr\{\mathbb{I}_{D_{(1)}} = k\}$, in which $\mathbb{I}$ is the indicator for the earliest arriving message. As stated above, we have

$$\Pr\{\mathbb{I}_{D_{(1)}} = k\} = \Pr\{D_k \le D_j\} \text{ for all } j \ne k. \tag{15}$$

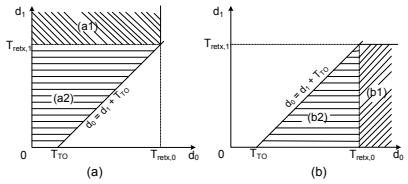

Fig. 2: Distributions of d_0 and d_1

1) $K_{retx} = 0$: For the simplest case where there is no retransmission before the cutoff time T_{CO}, that is, if $K_{retx} = 0$, the cdf of the first instance of arrival is simply the truncated cdf as shown in Eq. (12) for $k = 0$. To show this, we first have (a) the delivery probability in Eq. (9) is $(1 - p_L)F(T_{retx,0})$; and (b) the associated probability that the arrival time – which happens to be the delay of the original message as well in this case – is no greater than t is $(1 - p_L)F(t)$. And the cdf can thus be obtained by having (b) divided by (a):

$$F_{D^0_{(1)}}(t) = \Pr\{D^0_{(1)} \le t\} = \frac{(1-p_L)F(t)}{(1-p_L)F(T_{retx,0})} = F_T^{T_{retx,0}}(t). \tag{16}$$

2) $K_{retx} = 1$: When $K_{retx} = 1$, the sensor can retransmit the message at most once before the cutoff time. There are basically two scenarios for the first instance of arrival, namely:

1) the original message arrives first before T_{CO};
2) the retransmitted message arrives first before T_{CO}.

The first case can be further subdivided into

a1) the retransmitted message is not available by T_{CO};
a2) the retransmitted message is also delivered by T_{CO}, but its random arrival delay d_1 must satisfy $d_0 \le d_1 + T_{TO}$.

Likewise, for the second scenario, we have

b1) the original message is not available by T_{CO};
b2) the original message is also delivered by T_{CO} with its random arrival delay d_0 satisfying $d_1 + T_{TO} < d_0$.

All different scenarios are illustrated in Fig. 2. Their probabilities are calculated as

$$\begin{aligned}\Pr\{\text{a1}\} &= p^0_{loss,T_{retx,1}} p^0_{del,T_{retx,0}} \\ &= (1-(1-p_L)F(T_{retx,1}))(1-p_L)F(T_{retx,0}).\end{aligned} \tag{17}$$

$$\begin{aligned}\Pr\{\text{a2}\} &= (1-p_L)^2 \Pr\{d_0 - T_{TO} \le d_1 < T_{retx,1}\} \\ &= (1-p_L)^2 \int_0^{T_{retx,1}} \left(\int_0^{d_1+T_{TO}} f(d_0)\,\mathrm{d}d_0 \right) f(d_1)\,\mathrm{d}d_1 \\ &= (1-p_L)^2 \int_0^{T_{retx,1}} F(t+T_{TO}) f(t)\,\mathrm{d}t.\end{aligned} \tag{18}$$

Similarly, we have

$$\begin{aligned}\Pr\{\text{b1}\} &= p^0_{loss,T_{retx,0}} p^0_{del,T_{retx,1}} \\ &= (1-(1-p_L)F(T_{retx,0}))(1-p_L)F(T_{retx,1}).\end{aligned} \tag{19}$$

$$\begin{aligned}\Pr\{\text{b2}\} &= (1-p_L)^2 \Pr\{d_1 + T_{TO} < d_0 < T_{retx,0}\} \\ &= (1-p_L)^2 \int_{T_{TO}}^{T_{retx,0}} \left(\int_0^{d_0-T_{TO}} f(d_1)\,\mathrm{d}d_1 \right) f(d_0)\,\mathrm{d}d_0 \\ &= (1-p_L)^2 \int_0^{T_{retx,1}} F(t) f(t+T_{TO})\,\mathrm{d}t.\end{aligned} \tag{20}$$

Note that the sum of Eqs. (18) and (20) is $p^0_{del,T_{retx,0}} p^0_{del,T_{retx,1}} = (1-p_L)F(T_{retx,0}) \cdot (1-p_L)F(T_{retx,1})$, the probability that both the original and the retransmitted messages are available at the cutoff time.

The resulting cdf of $D^1_{(1)}$ is hence

$$F_{D^1_{(1)}}(t) = \Pr\{D^1_{(1)} \le t\} = \frac{\Pr\{\mathbb{I}_{D^1_{(1)}} = 0\} F_T^{T_{retx,0}}(t) + \Pr\{\mathbb{I}_{D^1_{(1)}} = 1\} F_T^{T_{retx,1}}(t - T_{TO})}{p^1_{del,T_{CO}}}, \tag{21}$$

in which $\Pr\{\mathbb{I}_{D^1_{(1)}} = 0\}$ and $\Pr\{\mathbb{I}_{D^1_{(1)}} = 1\}$ are the sums of Eqs. (17) and (18), and Eqs. (19) and (20), respectively. From Eq. (9), we have the total delivery probability $p^1_{del,T_{CO}}$ for "normalizing" the probabilities to obtain the cdf.

3) $K_{retx} > 1$: When K_{retx} is any number greater than one, thanks to the independence of different retransmissions, we can carry out the above pairwise comparison for any arbitrary pair of arrivals. In fact, if we generalize the above results, we have for any pair i and j $(i < j)$ of effective retransmissions

$$\begin{aligned}\Pr\{D_i \le D_j\} = (1-p_L)^2 &\int_0^{T_{retx,j}} F\left(t + (j-i)T_{TO}\right) f(t)\,\mathrm{d}t \\ &+ \left(1 - (1-p_L)F(T_{retx,j})\right)(1-p_L)F(T_{retx,i}),\end{aligned} \tag{22}$$

and

$$\begin{aligned}\Pr\{D_i > D_j\} = (1-p_L)^2 &\int_0^{T_{retx,j}} F(t) f(t + (j-i)T_{TO})\,\mathrm{d}t \\ &+ \left(1 - (1-p_L)F(T_{retx,i})\right)(1-p_L)F(T_{retx,j}).\end{aligned} \tag{23}$$

Repeating for all possible pairs, we have the probability of D_k being the minimum, that is, the k-th retransmitted message is received first, as

$$\begin{aligned}\Pr\{\mathbb{I}_{D^{K_{retx}}_{(1)}} = k\} = \prod_{\substack{j=0 \\ j\ne k}}^{K_{retx}} \Pr\{D_k \le D_j\} = (1-p_L)\cdot \\ \prod_{\substack{j=0 \\ j\ne k}}^{K_{retx}} \Bigg\{ (1-p_L) \int_0^{T_{retx,\max\{k,j\}}} F\left(t + \max\{0,(j-k)T_{TO}\}\right) \cdot \\ f\left(t + \max\{0,(k-j)T_{TO}\}\right)\,\mathrm{d}t \\ + \left(1 - (1-p_L)F(T_{retx,j})\right) F(T_{retx,k}) \Bigg\}.\end{aligned} \tag{24}$$

This leads to the overall distribution of the $D^{K_{retx}}_{(1)}$:

$$\begin{aligned}F_{D^{K_{retx}}_{(1)}}(t) &= \Pr\{D^{K_{retx}}_{(1)} \le t\} \\ &= \frac{\sum_{k=0}^{K_{retx}} \Pr\{\mathbb{I}_{D^{K_{retx}}_{(1)}} = k\} F_T^{T_{retx,k}}(t - kT_{TO})}{p^{K_{retx}}_{del,T_{CO}}}.\end{aligned} \tag{25}$$

B. Arrival Time: Two-way Communication Analysis

Eq. (25) describes the distribution of the earliest arrival time under the condition that all K_{retx} rounds of retransmissions are sent out after the original message. In reality, though, the total number of retransmissions can be anywhere from 0 to K_{retx}. In this subsection, we consider the two-way communication that determines the number of the actual retransmissions, which in turn affects the overall distribution of the earliest arrival time before the cutoff.

For satellite systems with conventional bent pipe type of transponders [21], one uplink (sensor $\rightarrow$ satellite) and downlink (satellite $\rightarrow$ FC) pair is used for the forward link, and the reverse link similarly consists of the uplink (FC $\rightarrow$ satellite) and downlink (satellite $\rightarrow$ sensor) pair. Depending on specific channel allocation schemes (e.g., TDMA- or FDMA-based), that is, whether the forward and reverse channels are assigned the same frequency band, the delay distribution of the ACK could vary from that of the messages[4]. Regardless, we have the pdf of the sum of two random delay values being expressed as the convolution of their respective pdfs:

$$h(t) = f(t) \star g(t) = \int_{t=0}^{\infty} f(u)g(t-u)\,\mathrm{d}u, \tag{26}$$

in which f and g are the distributions of the forward and reverse links, respectively. Meanwhile, if the ACK message is lost over the reverse link with a probability $p_{L,ACK}$, the overall probability that the ACK message can be eventually delivered is $(1-p_L)(1-p_{L,ACK})$, and its complement

$$p_{L,T} = 1-(1-p_L)(1-p_{L,ACK}) \tag{27}$$

is the "total" loss rate of the "super-message" that includes both the estimate message and ACK. With this loss rate and $h(t)$ function, we can derive a general form of the arrival time.

1) Probability of Having k_1 ($0 \le k_1 \le K_{retx}$) Retransmissions: First, we have the trivial case in which $K_{retx} = 0$ as $T_{CO} \le T_{TO}$, then with probability one, there is no retransmission. Next we focus on the cases where $K_{retx} \ge 1$.

Having exactly k_1 ($0 \le k_1 < K_{retx}$) retransmissions means that the earliest reception of the ACK message by the sensor occurs in the interval $[k_1T_{TO}, (k_1+1)T_{TO})$. In other words, the first instance of the ACK arrival at the sensor

$$D_{T,(1)}^{k_1} = \min_{k=0,\ldots,k_1}\{D_{T,k}\} \tag{28}$$

must satisfy

$$k_1T_{TO} \le D_{T,(1)}^{k_1} < (k_1+1)T_{TO}, \tag{29}$$

in which $D_{T,(1)}^{k_1}$ is similarly defined as in Eq. (13), with the subscript T specifying that this is the arrival time accounting for the total delay from both forward and reverse links. Meanwhile, we define the delivery rate for the "super-message" described earlier as $p_{T,del,t}^{K_{retx}}$ with the maximum number of retransmissions K_{retx}. Then we have

$$\begin{aligned}&\Pr\{\text{There are exactly } k_1 \text{ retransmissions, } 0 \le k_1 < K_{retx}\}\\ &= p_{T,del,(k_1+1)T_{TO}}^{K_{retx}} - p_{T,del,k_1T_{TO}}^{K_{retx}}.\end{aligned} \tag{30}$$

[4]Also the initial delay could be quite different too, owing to the usually much smaller size of the ACK messages.

The delivery probabilities can be similarly calculated as in Eq. (9), with p_L being replaced by $p_{L,T}$ and $F(t)$ by $H(t) = \int_0^t h(u)\,\mathrm{d}u$, respectively.

On the other hand, having K_{retx} retransmissions means that none of the ACKs have been received by $K_{retx}T_{TO}$, and we have

$$\begin{aligned}&\Pr\{\text{There are exactly } K_{retx} \text{ retransmissions}\}\\ &= \prod_{k=0}^{K_{retx}-1} \Pr\{D_{T,k}^{K_{retx}} > K_{retx}T_{TO}\}\\ &= \prod_{k=0}^{K_{retx}-1} \left[1-(1-p_{L,T})H((K_{retx}-k)T_{TO})\right].\end{aligned} \tag{31}$$

2) Distribution of the Arrival Time for a Given T_{CO}: With K_{retx} in Eqs. (9) and (24) being replaced by any k_1 ($0 \le k_1 \le K_{retx}$), we can easily find the delivery rate $p_{del,T_{CO}}^{k_1}$ after k_1 rounds of retransmissions and the probability $\Pr\{\mathbb{I}_{D_{(1)}^{k_1}} = k\}$ that the k-th retransmission marks the earliest arrival among all k_1+1 sent out messages. And then we have the cdf with exactly k_1 retransmissions as

$$\begin{aligned}F_{D_{(1)}^{k_1}}(t) &= \Pr\{D_{(1)}^{k_1} \le t\}\\ &= \frac{\sum_{k=0}^{k_1}\Pr\{\mathbb{I}_{D_{(1)}^{k_1}} = k\}F_T^{T_{retx,k}}(t-kT_{TO})}{p_{del,T_{CO}}^{k_1}}.\end{aligned} \tag{32}$$

Finally, we can combine Eqs. (30), (31), and (32) to obtain the distribution of the arrival time for any given T_{CO}:

$$F_{D_{(1)}}(t) = \sum_{k_1=0}^{K_{retx}} F_{D_{(1)}^{k_1}}(t)\Pr\{\text{There are } k_1 \text{ retransmissions}\}. \tag{33}$$

IV. State Estimation and Fusion with Retransmission and Retrodiction

We have seen that retransmission can effectively increase the message delivery rate, which in turn is expected to improve the estimation quality. However, at times, we wish to further expedite this recovery process so that the final estimate can be reported earlier; besides, the system may impose rather stringent requirements on the estimation errors so that given the same allocated time for retransmission, we want more accurate estimates from the output of the fusion center. This section addresses these concerns by means of utilizing estimate retrodiction.

A. Estimate Retrodiction

Estimation of a target state at a particular time based on measurements collected beyond that time is generally called retrodiction or smoothing. Retrodiction improves the accuracy of the estimates, thanks to the use of more information, at the cost of extra delay. Nevertheless, the inherent link delay in a long-haul network occurring before the final reporting entails that the fusion center can exploit the opportunities for potential retrodiction to improve the accuracy of the fused estimate. Moreover, the randomness of the arrival delay of different

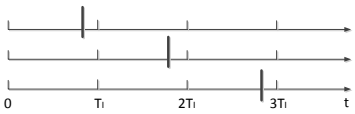

Fig. 3: Timing of estimate retrodiction: the estimation interval is T_I and different choices of the cutoff time T_{CO} are marked by bold lines

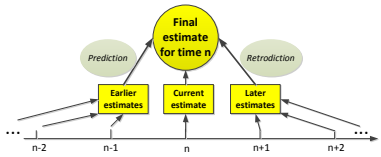

Fig. 4: Prediction and retrodiction: the estimate at time n is to be obtained

messages also facilitates the fusion center to opportunistically interpolate the missing messages from the available ones for subsequent time periods.

Consider a discrete-time system in which estimates are generated regularly at a certain fixed period T_I. An important note here is that a subsequent estimate must be available in order for retrodiction to happen. As such, the retransmission window W and the cutoff time T_{CO} must contain at least one estimation interval T_I (and also the initial latency of the latest estimate to arrive at the fusion center) so that the next estimate could potentially arrive therein. This is shown in Fig. 3. In the first case, no retrodiction can be possibly performed because the cutoff time comes before the end of the same estimation interval; while in the second and third case, a maximum of one and two, respectively, rounds of retrodiction can be possibly carried out. In many actual systems, the near real-time reporting requirement dictates that the estimation intervals are chosen fairly small. As a result, the next estimate can be generated and sent to the fusion center in time for retrodiction before the reporting deadline.

Let $\mathbf{x}_n$ denote the true target state at time $t = nT_I$ and $\hat{\mathbf{x}}_n$ its estimate. One note is that since we have both continuous time – when addressing latency – and the discrete time indices for retrodiction analysis, all time subsequently labeled as "n" – the time of interest to us – should be understood as the continuous time nT_I and the retransmission parameters T_{TO} and W and the cutoff time T_{CO} are all defined relative to this time instant.

Fig. 4 demonstrates the effect of prediction and retrodiction on the quality of the state estimate. For each sensor, for example, we have the following types of estimates when retrodiction of *up to one step* is performed:

1) $\hat{\mathbf{x}}_n$, shorthand for $\hat{\mathbf{x}}_{n|n}$, the "default" estimate sent from the sensor;
2) $\hat{\mathbf{x}}_{n^-}$, the predicted estimate;
3) $\hat{\mathbf{x}}_{n^-|n+1}$, the predicted & retrodicted estimate; and
4) $\hat{\mathbf{x}}_{n^+|n+1}$, the retrodicted estimate.

In the cases a) and d), the estimate $\hat{\mathbf{x}}_n$ is received by the fusion center; whereas in both other cases, this estimate is missing and hence prediction of one or multiple steps is at first necessary. As is well-known in filtering theory, estimates derived from prediction alone generally have higher errors when system uncertainty exists; and the errors will increase with the number of prediction steps that have accrued. For example, $\hat{\mathbf{x}}_{n|n-2}$ is a worse estimate than $\hat{\mathbf{x}}_{n|n-1}$ in terms of accuracy. On the other hand, the presence of the subsequent estimate $\hat{\mathbf{x}}_{n+1}$ (in the last two cases) helps improve the quality of the estimate of time n. Of course, the improvement is on top of the predicted estimate in case c) but on the already received $\hat{\mathbf{x}}_n$ in case d). Of concern here is whether an interpolated estimate from retrodiction – such as that in case c) – can adequately substitute the default estimate in case a); and what is the probabilistic performance of obtaining these different types of estimates so that the system requirement on estimation errors can be met.

The quantitative performance of retrodiction also depends on the specific algorithm. In this work, we apply Rauch-Tung-Striebel (RTS) retrodiction [23] because not only the algorithm is easy to implement, with relatively low computational cost, but the algorithm involves only the state estimates and their covariances – rather than the raw measurement data – which is well suited in our settings when the fusion center needs to run the algorithm.

Based on the criteria that whether the sensors actively participate in retrodiction during message retransmission, we categorize our schemes into non-cooperative and cooperative retrodiction. In the former case, message retransmission is carried out in exactly the same way as before; it is up to the fusion center to opportunistically apply retrodiction whenever applicable. In contrast, cooperative retrodiction means that the sensors themselves, upon request, send out the retrodicted estimates during retransmission so that the fusion center can directly fuse such retrodicted values if successfully delivered. In what follows, we derive the delivery rates of different types of estimates during retransmission for both types of retrodiction and consider their impact on the final error performance.

B. Non-Cooperative Retrodiction

The fusion center can opportunistically apply retrodiction whenever applicable; this can of course be combined with message retransmission introduced earlier. In what follows, we derive the delivery rates of different types of estimates following retransmission and retrodiction. Although here only one-step retrodiction is shown, analysis for retrodiction of two or more steps can be similarly obtained, albeit in a more exhaustive manner, as the number of possible scenarios grows exponentially with the number of steps[5].

We have the following probabilities at the cutoff time $nT_I + T_{CO}$:

- p_{n^-}, the probability that $\hat{\mathbf{x}}_{n^-}$ is reported (neither $\hat{\mathbf{x}}_n$ or $\hat{\mathbf{x}}_{n+1}$ is delivered by the cutoff);
- $p_{n|n}$, the probability that $\hat{\mathbf{x}}_{n|n}$ is reported ($\hat{\mathbf{x}}_{n+1}$ is not delivered yet);
- $p_{n^-|n+1}$, the probability that $\hat{\mathbf{x}}_{n^-|n+1}$ is reported ($\hat{\mathbf{x}}_{n+1}$ has been delivered but not $\hat{\mathbf{x}}_n$);
- $p_{n^+|n+1}$, the probability that $\hat{\mathbf{x}}_{n^+|n+1}$ is reported (both $\hat{\mathbf{x}}_n$ and $\hat{\mathbf{x}}_{n+1}$ have been delivered).

[5]Another caveat is that message-level loss and delay may worsen as significantly more data are sent simultaneously with increasing retrodiction steps.

The analysis in Sec. II can be readily applied here, thanks to the independence of the transmission from different time intervals. Similar to Eq. (2), we have the maximum number of retransmissions during $t \in [nT_I + T_I, nT_I + T_{CO}]$ given as

$$K_{retx,retr_1} = \left\lceil \frac{\min\{T_{CO} - T_I, W\}}{T_{TO}} \right\rceil - 1, \tag{34}$$

in which the subscript $retr_1$ denotes that there is a maximum of one-step retrodiction. And the duration in Eq. (3) can also be defined likewise for the above time interval:

$$T_{retx,retr_1,k} = T_{CO} - T_I - kT_{TO}, \quad \text{for } k = 0, 1, ..., K_{retx,retr_1}. \tag{35}$$

To calculate the above probabilities, we need to consider the probability that $\hat{\mathbf{x}}_{n+1}$ is delivered at the cutoff time. This probability, denoted as $p_{del,T_{CO}-T_I}^{K_{retx,retr_1}}$, follows the very same form as in Eq. (9), but with newly defined Eqs. (34) and (35) substituting the corresponding terms. Then we have

$$p_{n^-} = (1 - p_{del,T_{CO}}^{K_{retx}})(1 - p_{del,T_{CO}-T_I}^{K_{retx,retr_1}}) \tag{36}$$

$$p_{n|n} = p_{del,T_{CO}}^{K_{retx}}(1 - p_{del,T_{CO}-T_I}^{K_{retx,retr_1}}) \tag{37}$$

$$p_{n^-|n+1} = (1 - p_{del,T_{CO}}^{K_{retx}})p_{del,T_{CO}-T_I}^{K_{retx,retr_1}} \tag{38}$$

$$p_{n^+|n+1} = p_{del,T_{CO}}^{K_{retx}} p_{del,T_{CO}-T_I}^{K_{retx,retr_1}} \tag{39}$$

C. Cooperative Retrodiction

In this subsection, we provide similar analysis as above – that is, with a maximum of one step of retrodiction being performed – and focus on comparisons between cooperative and non-cooperative retrodiction techniques. In particular, two possible implementations of the cooperative retrodiction are studied, one in which we only consider one-way communications – as we have done so far – and the other requiring two-way analysis that addresses the delivery of the ACK messages as well.

1) Condition for One-Step Cooperative Retrodiction: In non-cooperative retrodiction, message retransmission is scheduled in the manner as described in Sec. II and a sensor is oblivious to the retrodiction process happening at the fusion center. In contrast, cooperative retrodiction requires the retrodicted estimates to be sent out during retransmission. In order for a sensor to actually send out its one-step retrodicted estimates, the retransmission window W should not have expired at the end of the estimation interval T_I; in fact, there should be at least one round of retransmission initiated by the sensor after T_I when the retrodicted estimate can be sent out by the sensor. Hence, compared to non-cooperative retrodiction, tighter conditions are in place for cooperative retrodiction.

2) Cooperative Retrodiction: One-way Communications without ACK: During the time period $[nT_I + T_I, nT_I + T_{CO}]$, instead of the original estimate, the sensor sends out the retrodicted estimate $\hat{\mathbf{x}}_{n^+|n+1}$ directly. The total number of retransmission rounds for the original message during $[nT_I, nT_I + T_I]$ is reduced to

$$K_{retx,coop} = \left\lceil \frac{T_I}{T_{TO}} \right\rceil - 1, \tag{40}$$

while both the new state estimate $\hat{\mathbf{x}}_{n+1|n+1}$ and retrodicted estimate $\hat{\mathbf{x}}_{n^+|n+1}$ are sent after T_I. If $T_I = lT_{TO}$, $l = 1, 2, 3, ...$, both estimates will undergo

$$K_{retx,coop,retr_1} = K_{retx,retr_1} = \left\lceil \frac{T_{CO} - T_I}{T_{TO}} \right\rceil - 1, \tag{41}$$

rounds of retransmission, in which the subscripts "$coop, retr_1$" and "$retr_1$" denote cooperative and non-cooperative retrodiction of up to one step respectively. Similar to our earlier analysis, we can obtain the delivery probabilities of the original estimate $\hat{\mathbf{x}}_{n|n}$ as $p_{del,T_{CO}}^{K_{retx,coop}}$, of the subsequent estimate $\hat{\mathbf{x}}_{n+1|n+1}$ as $p_{del,T_{CO}-T_I}^{K_{retx,retr_1}}$, and of the retrodicted estimate by the sensor $\hat{\mathbf{x}}_{n^+|n+1}$ as $p_{del,T_{CO}-T_I}^{K_{retx,coop,retr_1}}$. With an increased size of the state space, the probabilities of obtaining different types of estimates at the cutoff time can now be computed as

$$p_{n^-} = (1 - p_{del,T_{CO}-T_I}^{K_{retx,coop,retr_1}})(1 - p_{del,T_{CO}}^{K_{retx,coop}})(1 - p_{del,T_{CO}-T_I}^{K_{retx,retr_1}}) \tag{42}$$

$$p_{n|n} = (1 - p_{del,T_{CO}-T_I}^{K_{retx,coop,retr_1}})p_{del,T_{CO}}^{K_{retx,coop}}(1 - p_{del,T_{CO}-T_I}^{K_{retx,retr_1}}) \tag{43}$$

$$p_{n^-|n+1} = (1 - p_{del,T_{CO}-T_I}^{K_{retx,coop,retr_1}})(1 - p_{del,T_{CO}}^{K_{retx,coop}})p_{del,T_{CO}-T_I}^{K_{retx,retr_1}} \tag{44}$$

$$p_{n^+|n+1} = p_{del,T_{CO}-T_I}^{K_{retx,coop,retr_1}} + (1 - p_{del,T_{CO}-T_I}^{K_{retx,coop,retr_1}})p_{del,T_{CO}}^{K_{retx,coop}} p_{del,T_{CO}-T_I}^{K_{retx,retr_1}} \tag{45}$$

Note in Eq. (45) that the estimate $\hat{\mathbf{x}}_{n^+|n+1}$ can be obtained either directly from the sensor, or indirectly in the manner we discussed in the non-cooperative retrodiction case.

3) Cooperative Retrodiction: Two-way Communications with ACK: The analysis above is along the same line as that in Sec. II, where only one-way communication is considered. This can also be seen as the extreme case where no ACK is ever sent back by the fusion center, since the sensors always have their one-step retrodicted estimates sent out after one estimation interval T_I. In reality, though, the ACK might have been successfully received by the sensor within T_I, thereby obviating the need for further retransmission. Next, we carry out two-way communication analysis to account for such scenarios.

Recall from Eq. (27) that $p_{L,T}$ is the loss rate of the "super-message" that includes both the estimate message and ACK. With this loss rate and $h(t)$ function in Eq. (26), we have the probability that the ACK is delivered by time t and hence no more retransmission occurs afterward:

$$\overline{p_{retx}}(t) = (1 - p_{L,T})H(t), \text{ for } t \in [0, \min\{T_{CO}, W\}], \tag{46}$$

in which $H(t) = \int_0^t h(u)f(u)\,\mathrm{d}u$ is the cdf of the two-way communications delay.

After the ACK has been successfully received within one estimation interval, the retrodicted estimate is no longer to be sent out, thereby reducing the chance that the best estimate $\hat{\mathbf{x}}_{n^+|n+1}$ is available at the fusion center. Subsequently, with probability $1 - \overline{p_{retx}}(T_I)$, Eqs. (42)–(45) hold true; on the other hand, with probability $\overline{p_{retx}}(T_I)$, only two types of estimates

are possible to be used by the cutoff time – since $\hat{\mathbf{x}}_{n|n}$ has been received successfully – with probabilities $1-p_{del,T_{CO}-T_I}^{K_{retx,retr_1}}$ for $\hat{\mathbf{x}}_{n|n}$ and $p_{del,T_{CO}-T_I}^{K_{retx,retr_1}}$ for $\hat{\mathbf{x}}_{n^+|n+1}$. Using the law of total probability, we can easily incorporate them to calculate the overall probabilities of obtaining each type of estimate.

D. Impact of Retransmission and Retrodiction on Estimation Error

It's generally difficult to derive the exact error performance analytically. Note that in Eqs. (36) and (38), the time at which the last received estimate was generated is not specified; that is, the number of prediction steps leading up to $\hat{\mathbf{x}}_{n^-}$ is unknown. Even with known error profiles for $\hat{\mathbf{x}}_{n|n}$ and $\hat{\mathbf{x}}_{n^+|n+1}$, an exact evaluation of estimation error for $\hat{\mathbf{x}}_{n^-}$ and $\hat{\mathbf{x}}_{n^-|n+1}$ requires knowledge of the errors with an arbitrary number of prediction steps, which is generally unrealistic. If the off-line error performance can be obtained for the above original, predicted, and retrodicted estimates, an approximate estimation error can be calculated as the probabilistic sum of errors for different types of estimates. Next, via simulations of a ballistic target tracking application, we explore the impact of both retransmission and retrodiction on actual estimation performance.

V. Performance Studies of Ballistic Target Tracking with Retransmission and Retrodiction

We carry out system-level tracking simulations of coasting ballistic targets. First, we focus on the one-target case where fusion is performed by the fusion center to combine the state estimates of a single target sent by the sensors. Then, the multi-target tracking scenario is investigated, where a probability association model is used to capture the performance of data association prior to the estimate fusion. We aim to explore the benefits and limitations of applying retransmission and retrodiction under variable loss and delay statistics.

A. Effect of Retransmission Schedules on Tracking Performance

1) Target Model: The trajectory of a ballistic target, from launch to impact, is divided into three basic phases: boost, coast, and reentry. In this work we focus on the coast phase, which is an exo-atmospheric, free-flight motion, continuing until the Earths atmosphere is reached again. The states of a coasting ballistic target – whose motion is governed predominantly by gravity – are generated using the following state-space model [13]:

$$\dot{\mathbf{x}} \triangleq \begin{bmatrix}\dot{\mathbf{p}}\\ \dot{\mathbf{v}}\end{bmatrix} = \mathbf{f}\left(\begin{bmatrix}\mathbf{p}\\ \mathbf{v}\end{bmatrix}\right) \triangleq \begin{bmatrix}\mathbf{v}\\ \mathbf{a}_G(\mathbf{p})\end{bmatrix}. \tag{47}$$

The target state vector $\mathbf{x} = \begin{bmatrix}\mathbf{p}^T & \mathbf{v}^T\end{bmatrix}^T$, where $\mathbf{p} = \begin{bmatrix}x & y & z\end{bmatrix}^T$ and $\mathbf{v} \triangleq \dot{\mathbf{p}} = \begin{bmatrix}\dot{x} & \dot{y} & \dot{z}\end{bmatrix}^T$ are the target position and velocity vectors, respectively. $\mathbf{a}_G(\mathbf{p})$ is the gravitational acceleration under the spherical Earth model [13]:

$$\mathbf{a}_G(\mathbf{p}) = -\frac{\mu}{p^2}\mathbf{u}_p = -\frac{\mu}{p^3}\mathbf{p}, \tag{48}$$

where $\mathbf{p}$ is the vector from the Earth's center to the target, $p \triangleq \|\mathbf{p}\|$ is its length, $\mathbf{u}_p \triangleq \mathbf{p}/p$ is the unit vector in the direction of $\mathbf{p}$, and $\mu = 3.986012 \times 10^5$ km^3/s^2 is the Earth's gravitational constant.

The algorithm for state propagation can be found in [25] and the initial target state is [11]: $\begin{bmatrix}113.75 & 3960 & 5150 & 0.988 & 3.33 & -6.01\end{bmatrix}^T$, in which the position and velocity values are in the units of km and km/s respectively.

2) Sensor and Noise Profiles: A total of $N = 5$ sensors are deployed for reporting their state estimates of the dynamic target defined above. A state estimate $\hat{\mathbf{x}}_i(k)$ from sensor i at time k is generated by adding random Gaussian noise to the true states[6] as in [20]:

$$\hat{\mathbf{x}}_i(k) = \mathbf{x}(k) + \mathbf{n}_i(k), \tag{49}$$

where $\mathbf{x}(k)$ is the true target state for a target at time k, and $\mathbf{n}_i(k) \sim \mathcal{N}(\underline{0}, \mathbf{\Sigma})$. $\mathbf{\Sigma}$ is a diagonal matrix:

$$\mathbf{\Sigma} = diag\left(\begin{bmatrix}\sigma_x^2 & \sigma_y^2 & \sigma_z^2 & \sigma_{\dot{x}}^2 & \sigma_{\dot{y}}^2 & \sigma_{\dot{z}}^2\end{bmatrix}\right), \tag{50}$$

where σ_x^2, σ_y^2, and σ_z^2 are the position error variances, and $\sigma_{\dot{x}}^2$, $\sigma_{\dot{y}}^2$, and $\sigma_{\dot{z}}^2$ are the velocity error variances. The following state estimation errors are set commonly for all sensors: $\sigma_x^2 = \sigma_y^2 = \sigma_z^2 = 1$ and $\sigma_{\dot{x}}^2 = \sigma_{\dot{y}}^2 = \sigma_{\dot{z}}^2 = 10^{-4}$.

3) Fusion Rule: We apply the linear fuser defined as follows:

$$\mathbf{P}_F = \left(\sum_{i=1}^{L}\mathbf{P}_i^{-1}\right)^{-1}, \text{ and } \hat{\mathbf{x}}_F = \mathbf{P}_F\sum_{i=1}^{L}\mathbf{P}_i^{-1}\hat{\mathbf{x}}_i, \tag{51}$$

where $\hat{\mathbf{x}}_F$ is the fused estimate and $\mathbf{P}_F$ is its error covariance matrix. $\mathbf{P}_i$ and $\hat{\mathbf{x}}_i$ are similarly defined for sensor i. A total of L ($L \leq N$) state estimates are combined at the fusion center. This simple fuser is a special form of the track-to-track fuser [5] where the process noise is zero.

4) Communication Link Statistics: The default forward link loss rate is set to be $p_L = 0.5$, compared to that of the reverse link $p_{L,ACK} = 0.1$. The arrival delay of both directions satisfies the shifted exponential distribution defined in Eq. (1), with $\mu_{D,F} = 0.3$ s and $\mu_{D,R} = 0.2$ s for the forward and reverse links respectively and a common $T = 0.5$ s. The default T_{TO} and W are set to be 1.5 s and 4.5 s respectively, which are both integer multiples of the measured average RTT at 1.5 s.

5) MSE analysis: From Eq. (51), the mean-square-error (MSE) of the position and velocity estimates of the linearly fused estimate satisfies

$$|\hat{\mathbf{p}}_F - \mathbf{p}|^2 = |\hat{\mathbf{p}}_i - \mathbf{p}|^2/L, \text{ and } |\hat{\mathbf{v}}_F - \mathbf{v}|^2 = |\hat{\mathbf{v}}_i - \mathbf{v}|^2/L, \tag{52}$$

if the sensors have the same estimation error profiles (that is, the same MSEs). Suppose the system imposes its maximum tolerable errors of position and velocity estimates as $MSE_{max,\hat{\mathbf{p}}}$ and $MSE_{max,\hat{\mathbf{v}}}$ respectively, with independent

[6]This is a simplified model for scenarios in which the dependency of estimation errors on the underlying states are negligible.

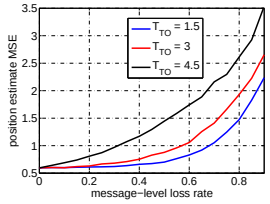

Fig. 5: Position MSE (km^2) versus message-level loss rates with $T_{CO} = 5$ s and different T_{TO}

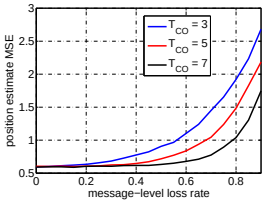

Fig. 6: Position MSE (km^2) versus message-level loss rates with $T_{TO} = 1.5$ s and different T_{CO}

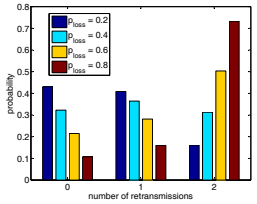

Fig. 7: probability of different number of retransmissions with different loss rates

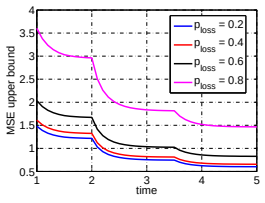

Fig. 8: Upper bound of position MSE (km^2) with different message-level loss rates

transmissions from up to N sensors, the minimum delivery probability should satisfy

$$p_{del,min} = \frac{L_{min}}{N} = \frac{1}{N}\max\{\frac{|\hat{\mathbf{p}}_i - \mathbf{p}|^2}{MSE_{max,\hat{\mathbf{p}}}}, \frac{|\hat{\mathbf{v}}_i - \mathbf{v}|^2}{MSE_{max,\hat{\mathbf{v}}}}\}, \tag{53}$$

in which L_{min} is the minimum number of estimates that should be delivered. The actual delivery probability from message retransmission as expressed in Eq. (9) must be checked against this minimum delivery rate to ensure the MSE requirements are met.

The distribution of the arrival time, as in Eq. (33), can be used to derive an upper bound of the estimation error. For a given cutoff time T_{CO}, when the fusion center decides to finalize the estimate at an earlier time t, according to Eq. (51), the expected contribution of the position estimate from the sensor i is lower bounded by $\frac{\alpha}{|\hat{\mathbf{p}}_i - \hat{\mathbf{p}}|^2}$, in which $\alpha = F_{D_{(1)}^{T_{CO}}}(t)$. The other part, weighted by $1-\alpha$, depends on how the fuser chooses to substitute the missing estimate, such as using one- or multi-step prediction. If we consider the error performance of the fused estimate, the upper bound – as in the worst-case scenario in which a missing estimate results in zero contribution – is

$$|\hat{\mathbf{p}}_F(t) - \mathbf{p}|^2_{max} = \frac{|\mathbf{p}_F(T_{CO}) - \mathbf{p}|^2}{F_{D_{(1)}^{T_{CO}}}(t)}, \text{ for } 0 < t \le T_{CO}, \tag{54}$$

in which $\mathbf{p}_F(T_{CO})$ is the fused position estimate at time T_{CO}. A similar result can be found for the velocity estimate.

6) Timeout T_{TO} and Cutoff T_{CO}: Figs. 5 and 6 demonstrate the effect of different retransmission timeout and fusion cutoff. There is no retransmission when $W = T_{TO} = 4.5$ s. When T_{TO} is set to be 1.5 s, however, two rounds of retransmissions can effectively reduce the MSE of the estimate. For example, when the loss rate is 60%, the error is reduced by more than 50%. Likewise, for a given T_{TO}, when the reporting deadline requirement is tightened, as reflected by decreasing cutoff time T_{CO}, the estimation MSE will increase given the same loss rate. On the other hand, when read horizontally, the plots indicate that to meet the same MSE requirement, with increasing loss rates, T_{TO} should be reduced and/or T_{CO} should be increased. A good rule of thumb to determine the MSE performance is the ratio T_{CO}/T_{TO}, although this rule may at times fail due to the periodicity of the retransmission process, especially when the values being compared are close.

7) Retransmission Performance: In Fig. 7, we plotted the proportion of different numbers of retransmissions with respect to various loss rates. Note that the last group (labeled as "2") does not indicate the message will be delivered within this round, but rather this is the last try as $K_{retx} = 2$. As expected, an increased message-level loss rate requires more rounds of retransmissions so that the message can be recovered with the same probability over a longer period of time.

8) Upper Bound of the MSEs From Early Cutoff: Finally, in Fig. 8, the upper bounds of the MSEs resulting from earlier-than-scheduled cutoff are shown. The singular points in these plots indicate the arrival of a new round of retransmission. The concavity of the bounds (excluding the singular points) implies that the "best" time for early fusion is roughly in the middle of each round (accounting for the link delay), where the deepest descent in this round has occurred. In addition, as time inches closer to the cutoff time T_{CO}, the bound also approaches the actual MSE obtained when the cutoff time is set at that point, and is a good indicator of the actual MSE performance.

B. Tracking of One Coasting Target

1) Target Model: The target model remains the same as that introduced in the previous subsection. We note the generated trajectory there – which happens to be the very last few seconds before the target enters the re-entry stage starting at an altitude of about 100 km – is very short. Hence, we have backtracked a segment of the earlier trajectory so we can focus on the position state estimates during the last 30 seconds prior to the re-entry stage.

2) Sensor Profiles: In this subsection, a more complex sensor measurement model is used that can describe state-dependent measurement errors. A total of $N = 3$ sensors are deployed for reporting their state estimates with the common estimation interval set as $T_I = 2$ s. The measurements ($\mathbf{z}$) of the range (r), elevation (E), and azimuth (A) of the target are generated using the following measurement model [11]:

$$\mathbf{z} = \mathbf{h}(\mathbf{x}) + \mathbf{v}, \tag{55}$$

where the target state $\mathbf{x}$ is in Cartesian coordinates, but the measurement $\mathbf{z}$ and additive noise $\mathbf{v}$ are in the sensor spherical coordinates. If $[x \quad y \quad z]^T$ is the true position of the target, then the measurement[7] is given as

$$\mathbf{z} = \begin{bmatrix} r \\ E \\ A \end{bmatrix} + \mathbf{v} = \begin{bmatrix} \sqrt{x^2+y^2+z^2} \\ \tan^{-1}\left(\frac{z}{\sqrt{x^2+y^2}}\right) \\ \tan^{-1}\left(\frac{x}{y}\right) \end{bmatrix} + \mathbf{v}, \tag{56}$$

[7]Note that this measurement has been normalized to any sensor's own known location.

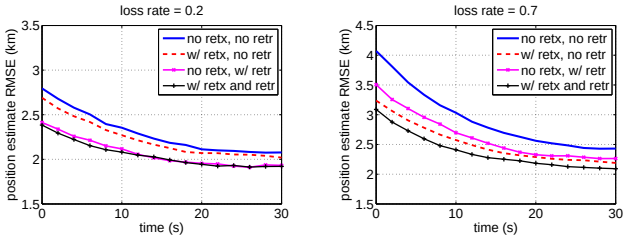

(a) (b)

Fig. 10: Position estimate RMSE over time, with loss rates of (a) 0.2; and (b) 0.7

$$\mathbf{v} \sim \mathcal{N}(0, \mathbf{R}), \quad \mathbf{R} = \begin{bmatrix} \sigma_r^2 & 0 & 0 \\ 0 & \sigma_E^2 & 0 \\ 0 & 0 & \frac{\sigma_E^2}{\cos^2(E)} \end{bmatrix}. \tag{57}$$

A simplified radar model is used to define the range and elevation error standard deviations. The σ_r and σ_E values of a ballistic target/satellite tracking phased array radar, the Cobra Dane, are 15 ft and 0.05° respectively, according to [8]; and these values are assumed for all the sensors. The sensors apply the recursive best linear unbiased estimator (BLUE) proposed in [28] which improves upon the measurement-conversion approach [12]; that is, the output has the minimum MSE among all linear unbiased filters in Cartesian coordinates. As in [28], the filter is initialized with an effectively large state error covariance and a highly inaccurate state estimate.

3) Fusion Rule: We apply a linear fuser similar to that defined in (51). With imperfect communications, the number of received state estimates for any time of interest is often less than the desired value N. The fusion center can apply two approaches: First, as in Eq. (51), it uses the actual number of available estimates L ($L \leq N$) and ignores the remaining, if any, unavailable tracks. Alternatively, it can substitute the predicted (or retrodicted) estimates in place of the original, if the latter is unavailable, and use all N tracks for fusion. While the first approach is easier to implement, its estimation error is often worse; in other words, a predicted estimate still more or less provides error reduction for the fused estimate when compared to the case where the track is simply dropped. Therefore, in what follows, we pursue the second method of fusing the estimates in which all the tracks are used.

4) Communication Link Statistics: The following communication link statistics are used for the remainder of this section. Two link loss rates of 0.2 and 0.7, are studied, representing respectively the low and high loss scenarios. The arrival delay satisfies the shifted exponential distribution defined in Eq. (1), with $\mu_D = 0.3$ s and the initial latency $T = 0.5$ s. With the cutoff time set as $T_{CO} = 3$ s – the same for the retransmission window W – we explore two options with the retransmission timeout period T_{TO} set as 1.5 s (up to one round of retransmission) and 3 s (no retransmission) respectively.

5) Position Estimate RMSE Performance: We run Monte Carlo simulations to test the position estimate root-mean-square-errors[8] (RMSEs) during the last 30 seconds of the coasting phase, and the results are shown in Fig. 10. Whereas the performance improvement of cooperative retrodiction over the non-cooperative counterpart has been shown in a numerical example in [17], the improvement is usually not as significant. We focus here on the case – often in practical implementation – where the sensors do not participate in cooperation. The estimation errors become noticeably higher with a higher message-level loss rate; but regardless of the loss rate, applying retransmission and retrodiction can effectively reduce the estimation errors. Given the communication parameters, according to Eq. (9), the total delivery probability of the original message at time cutoff T_{CO} is approximately $(1 - p_L^2)$ where p_L denotes the message-level loss rate. As such, the delivery rates for $p_L = 0.2$ and 0.7 at T_{CO} are approximately 96% and 50% respectively. The increased delivery rate effectively prevents the fusion center from using the predicted estimates – statistically worse than the original – thereby improving the estimation performance compared to the case without retransmission.

During this 30-second period[9], the sensors produce state estimates whose quality progressively improves over time – the favorable condition for applying retrodiction. Even with just one-step retrodiction, its effect on improving the accuracy can already be seen from the plots, where the RMSE is reduced generally by over 10%. Interestingly, retransmission seems to play a more prominent role than retrodiction when the loss rate is high – note the two curves in the middle appear "flipped" when the loss changes from 0.2 to 0.7 – since the chance of applying retrodiction is rather small without retransmission under higher loss rates.

C. Tracking of Multiple Coasting Targets

In this subsection, we consider tracking of multiple ballistic targets coasting in close proximity to one another. The initial states for $M = 4$ different targets (in which Target 1 is the one we used in the previous subsection) are shown in Table I. The trajectories are generated in the same manner as before.

TABLE I: Initial target states

Target	Position x, y, z (km)	Velocity $\dot{x}$, $\dot{y}$, $\dot{z}$ (km/s)
1	76.1, 3829.4, 5373.1	0.995, 3.54, -5.73
2	77.0, 3827.1, 5368.6	0.985, 3.48, -5.75
3	75.0, 3828.2, 5370.0	0.970, 3.52, -5.74
4	77.8, 3831.3, 5375.8	0.998, 3.51, -5.72

1) Association Properties: An important task prior to fusion is that every sensor estimate (track) must be associated with a certain target. Rather than focus on a particular type of data association algorithm, we explore the effect of available time on data association and on the final estimation performance by means of the probabilistic model [20] as follows.

The state estimate of Target i is assigned to Target j with probability a_{ij}:

$$a_{ij} = \begin{cases} 1 - (M-1)p_A, & i = j \\ p_A, & i \neq j \end{cases}, \quad i, j = 1, ..., M, M \geq 2,$$

[8]The RMSEs are used more often in the tracking literature. We used MSEs earlier in Sec. V.A. because the upper bounds were used to approximate fusion performance.

[9]The filters have been initialized prior to the "zero" time in the figures.

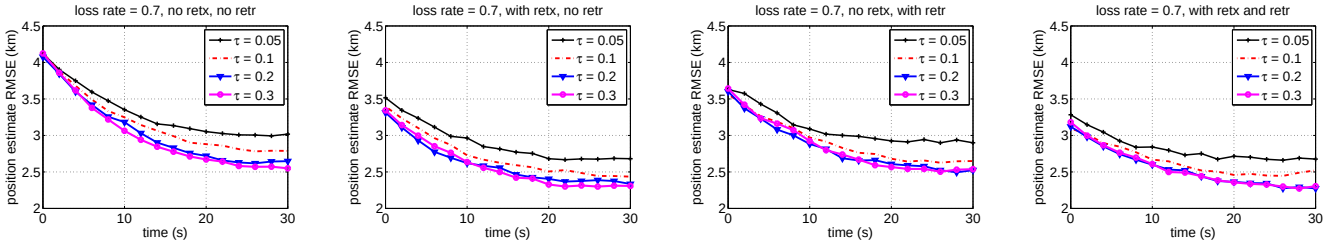

Fig. 9: Multi-target position estimate RMSE over time with variable τ and the loss rate as 0.7: (a) no retransmission or retrodiction; (b) with retransmission, but no retrodiction; (c) with retrodiction, but no retransmission; (d) with both retransmission and retrodiction

(58)

where M is the number of targets, and p_A is the probability that we assign a state estimate to any incorrect target[10]. We model p_A here to be the following function of the number of targets, M, and time τ:

$$p_A = \frac{1}{M(1+\beta\tau)}, \quad p_A \leq \frac{1}{M}, \tag{59}$$

where β is a scaling parameter. This function has the following properties: (a) as the number of targets M increases, the total probability of misassignment, $(M-1)p_A$, increases; (b) as τ increases, p_A decreases; (c) if $\tau = 0$, there is an equal probability of assigning a state estimate to any target; and (d) as $\tau \to \infty$, $p_A \to 0$. However, of note is that p_A is unlikely to be zero given an unlimited execution time for the correlator; more likely there may be some lower bound $p_{A,min}$ on the probability p_A, such that if $\tau \to \infty$, $p_A \to p_{A,min}$. This model is suited for the scenario in our study because the targets remain in close proximity and thus the distance between any two does not undergo dramatic change; as such, no target lies far apart from the rest from the view of a long-haul sensor. A fixed amount of time τ is allocated right before T_{CO} for computation, and the scaling factor β depends on the specific algorithms used. Here we set $\beta = 80$.

2) Position Estimate RMSE Performance: In Fig. 9, the position estimate RMSEs for Target 1 under different retransmission-retrodiction combinations are plotted for a message loss rate of $p_L = 0.7$. First, by comparing the results from each category with that in Fig. 10(b), we observe the estimation errors have increased due to mis-assigned tracks, although with the same allocated time τ, the improvement after retransmission and/or retrodiction is on the level as before. When the allocated time for computation is short, say $\tau = 0.05$, the degradation in error performance is much higher. As τ gradually goes up, the errors are seen to decrease, thanks to the improved performance of the correct association probabilities.

On the other hand, we observe that as τ is further increased, for example from 0.2 to 0.3, except in Case (a), where no retransmission or retrodiction is implemented, there is no clearly discernible performance improvement in accuracy. Actually, by this time, the probability for wrong associations has become very low. But as τ, the allocated time within T_{CO}, increases, the remaining time for communications is shortened. As such, when τ becomes higher, the slight improvement in computation can hardly offset reduced communication opportunities, which could either be the retransmitted messages (in (b)) or the next estimate itself used for retrodiction (in (c)) or both (in (d)). This demonstrates the trade-offs between computation and communication when the total allocated time T_{CO} is fixed.

3) Discussions: In this subsection, we have assumed the communication parameters are given for a certain long-haul network. For a given deadline T_{CO}, however, if an existing schedule cannot meet the required or desired estimation accuracy levels, the fusion center can schedule more frequent retransmission and/or estimation intervals (at the cost of higher communication overhead for the sensors), to counteract the increasing demand for computation due to more sensors being deployed and/or more targets being monitored.

VI. Conclusion

Many estimation- and fusion-based civilian and military applications, notably target tracking and monitoring, are subject to constraints over the communication link, in the computation-constrained sensors and devices, and as a result of various environmental factors. It is thus very challenging to meet both the accuracy and timeliness requirements mandated by most systems. In this paper, we studied target tracking in a long-haul sensor network where communication loss and delay are severe enough to degrade the estimation accuracy performance significantly. We derived the improved delivery probabilities of estimate messages and the distribution of the first arrivals. We also analyzed the probabilities of obtaining different types of estimates by the fusion center when retransmission and retrodiction techniques are applied. Simulation results for tracking of coasting ballistic targets – with the effect of track association and fusion being accounted for – demonstrate the effectiveness of our retransmission- and retrodiction-based mechanisms and the extent to which they can be applied so that the system requirements on estimation errors can be potentially met. Extensions of this work may include adaptive determination of different retransmission and retrodiction schedules in the context of more complex target-measurement models and highly dynamic network environment.

[10]For the case where $M = 1$, data association is not required since there is only one target to which the state estimate can belong (i.e., for $M = 1$, $p_A = 0$).

REFERENCES

[1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. A survey on sensor networks. *Communications Magazine, IEEE*, 40(8):102–114, Aug. 2002.

[2] M. Allman, C. Hayes, H. Kruse, and S. Ostermann. TCP performance over satellite links. In *Proceedings of Fifth International Conference on Telecommunications Systems*, Nashville, TN, Mar. 1997.

[3] R. B. Bapat and M. I. Beg. Order statistics for nonidentically distributed variables and permanents. *Sankhy: The Indian Journal of Statistics, Series A*, 51(1):79–93, Feb. 1989.

[4] Y. Bar-Shalom, T. Kirubarajan, and X. R. Li. *Estimation with Applications to Tracking and Navigation*. John Wiley & Sons, Inc., New York, NY, 2002.

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

[6] W. Boord and J. B. Hoffman. *Air and Missile Defense Systems Engineering*. CRC Press, 2014.

[7] A. Chiuso and L. Schenato. Information fusion strategies and performance bounds in packet-drop networks. *Automatica*, 47:1304–1316, Jul. 2011.

[8] E. Filer and J. Hartt. Cobra dane wideband pulse compression system. In *Proc. IEEE EASCON*, pages 26–29, 1976.

[9] V. Gupta, B. Hassibi, and R. M. Murray. On sensor fusion in the presence of packet-dropping communication channels. In *Proceedings of 44th IEEE Conference on Decision and Control 2005 (CDC 05)*, pages 3547–3552, Seville, Spain, Dec. 2005.

[10] T. R. Henderson and R. H. Katz. Transport protocols for internet-compatible satellite networks. *IEEE Journal on Selected Areas in Communications*, 17:326–344, 1999.

[11] T. H. Kerr. Streamlining measurement iteration for EKF target tracking. *IEEE Transactions on Aerospace and Electronic Systems,*, 27(2):408–421, Mar. 1991.

[12] X. R. Li and V. P. Jilkov. A survey of maneuvering target tracking, part iii: Measurement models. In *Proceedings of the 2001 SPIE Conference on Signal and Data Processing of Small Targets*, volume 4473, pages 423–446, San Diego, CA, Jul–Aug. 2001.

[13] X. R. Li and V. P. Jilkov. Survey of maneuvering target tracking. part ii: Motion models of ballistic and space targets. *Aerospace and Electronic Systems, IEEE Transactions on*, 46(1):96–119, Jan. 2010.

[14] Q. Liu, X. Wang, and N. S. V. Rao. Staggered scheduling of estimation and fusion in long-haul sensor networks. In *Proc. Information Fusion (FUSION), 2013 16th International Conference on*, pages 1699–1706, Istanbul, Turkey, Jul. 2013.

[15] Q. Liu, X. Wang, and N. S. V. Rao. Fusion of state estimates over long-haul sensor networks with random loss and delay. *IEEE/ACM Transactions on Networking*, 2014.

[16] Q. Liu, X. Wang, and N. S. V. Rao. Information feedback for estimation and fusion in long-haul sensor networks. In *Proc. Information Fusion (FUSION), 2014 17th International Conference on*, Salamanca, Spain, Jul. 2014.

[17] Q. Liu, X. Wang, N. S. V. Rao, K. Brigham, and B. V. K. Vijaya Kumar. Fusion performance in long-haul sensor networks with message retransmission and retrodiction. In *Proc. Mobile Ad-Hoc and Sensor Systems (MASS), 2012 9th International Conference on*, pages 407–415, Las Vegas, NV, Oct. 2012.

[18] M. Mallick and K. Zhang. Optimal multiple-lag out-of-sequence measurement algorithm based on generalized smoothing framework. In *Proc. SPIE, Signal and Data Processing of Small Targets*, San Diego, CA, Apr. 2005.

[19] A. Papoulis and S. U. Pillai. *Probability, random variables, and stochastic processes*. McGraw-Hill series in electrical and computer engineering. McGraw-Hill, 2002.

[20] N. S. V. Rao, K. Brigham, B. V. K. Vijaya Kumar, Q. Liu, and X. Wang. Effects of computing and communications on state fusion over long-haul sensor networks. In *Proc. Information Fusion (FUSION), 2012 15th International Conference on*, pages 1570–1577, Singapore, Singapore, Jul. 2012.

[21] D. Roddy. *Satellite Communications*. McGraw-Hill, 2006.

[22] E. I. Silva and M. A. Solis. An alternative look at the constant-gain kalman filter for state estimation over erasure channels. *Automatic Control, IEEE Transactions on*, 58(12):3259–3265, Dec. 2013.

[23] D. Simon. *Optimal state estimation: Kalman, H [infinity] and nonlinear approaches*. Wiley-Interscience, 2006.

[24] S. Sun, L. Xie, W. Xiao, and Y. C. Soh. Optimal linear estimation for systems with multiple packet dropouts. *Automatica*, 44:1333–1342, May 2008.

[25] M. Yeddanapudi, Y. Bar-Shalom, K. R. Pattipati, and S. Deb. Ballistic missile track initiation from satellite observations. *IEEE Transactions on Aerospace and Electronic Systems*, 31(3):1054–1071, Jul. 1995.

[26] K. Zhang, X. R. Li, and Y. Zhu. Optimal update with out-of-sequence measurements. *Signal Processing, IEEE Transactions on*, 53(6):1992–2004, Jun. 2005.

[27] S. Zhang and Y. Bar-Shalom. Optimal update with multiple out-of-sequence measurements with arbitrary arriving order. *Aerospace and Electronic Systems, IEEE Transactions on*, 48(4):3116–3132, Oct. 2012.

[28] Z. Zhao, X. R. Li, and V. P. Jilkov. Best linear unbiased filtering with nonlinear measurements for target tracking. *IEEE Transactions on Aerospace and Electronic Systems*, 40(4):1324–1336, Oct. 2004.

Qiang Liu received his Ph.D. degree in Electrical and Computer engineering from Stony Brook University. His research interests include networked data/information fusion, signal and target detection/estimation, statistical signal processing, wireless sensor networks, as well as spectrum sensing and allocation for cognitive radios.

Xin Wang is currently an Associate Professor in the Department of Electrical and Computer Engineering at Stony Brook University, Stony Brook, NY. Before joining Stony Brook, she was a Member of Technical Staff in the area of mobile and wireless networking at Bell Labs Research, Lucent Technologies, NJ, and an Assistant Professor in the Department of Computer Science and Engineering of SUNY at Buffalo, Buffalo, NY. Her research interests include algorithm and protocol design in wireless networks and communications, mobile and distributed computing, as well as networked sensing and detection. She has served in executive committee and technical committee of numerous conferences and funding review panels, and is the referee for many technical journals. She serves as an associate editor of IEEE Transactions on Mobile Computing since 2013. Dr. Wang achieved the NSF career award in 2005, and ONR challenge award in 2010.

Nageswara S. V. Rao is currently is a Corporate Fellow in Computer Science and Mathematics Division, Oak Ridge National Laboratory, where he joined in 1993. He was on assignment at Missile Defense Agency as the founding Technical Director, C2BMC Knowledge Center during 2008-2010. He published more than 350 technical conference and journal papers in the areas of sensor networks, information fusion, high-performance networking and fault diagnosis. His research projects have been funded by multiple federal agencies including National Science Foundation, Office of Naval Research, Department of Energy, Department of Defense, Domestic Nuclear Detection Office, and Defense Advanced Research Projects Agency. He received 2005 IEEE Technical Achievement Award for his contributions to information fusion area.

Katharine Brigham received a B.E. in Electrical Engineering and Mathematics from Vanderbilt University, and an M. Eng. in Systems Engineering from Cornell University. She is currently working towards her Ph.D. degree with the Department of Electrical and Computer Engineering at Carnegie Mellon University in Pittsburgh, PA. Her research interests include machine learning/pattern recognition, signal and image processing, and computer vision.

B. V. K. Vijaya Kumar is currently a Professor of Electrical and Computer Engineering at Carnegie Mellon University, Pittsburgh, PA, where he is also the Associate Dean for Graduate and Faculty Affairs with the College of Engineering. He is the author or a coauthor of 15 book chapters, over 500 technical papers, and the book *Correlation Pattern Recognition*. His research interests include pattern recognition, biometrics and coding, and signal processing for data storage systems.