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

Adaptive Filtering of Imprecisely Time-stamped Measurements with Application to AIS Networks

Leonardo M. Millefiori*, Paolo Braca*, Karna Bryan*, and Peter Willett†
*NATO STO-CMRE, La Spezia, Italy, Email: {leonardo.millefiori,paolo.braca,karna.bryan}@cmre.nato.int
†University of Connecticut, Storrs CT, Email: willett@engr.uconn.edu

***Abstract*—Driven by real-world issues in maritime surveillance, we consider the problem of estimating the target state from a sequence of observations that can be imprecisely time-stamped. That is, the time between two consecutive observations can be affected by an unknown error or delay.**

We propose an adaptive filtering strategy able to sequentially detect the time delays and correctly estimate the target state. Two decision statistics for the presence of delay are derived, the first is non-parametric while the second is based on the Generalized Likelihood Ratio Test (GLRT). When a delayed measurement is detected, the Maximum Likelihood (ML) estimate of the delay can be used to correct the timestamps of the target observation used in the filter.

The validation of the proposed method is carried out using Monte Carlo computer simulations and analyzing real-world data collected by a global network of Automatic Identification System (AIS) receivers.

***Keywords*—Unknown timestamp delay, filtering algorithms, Kalman filters, maximum likelihood estimation, real-world data**

I. INTRODUCTION

The oceans connect nations globally through an interdependent network of economic, financial, social and political relationships. The statistics are compelling: 70% of the Earth is covered in water; 80% of the world's population lives within 100 miles of the coast; 90% of the world's commerce is seaborne and 75% of that trade passes through a few vulnerable canals and international straits. The maritime environment includes trade routes, choke points, ports, and other infrastructure such as pipelines, oil and natural gas platforms and trans-oceanic telecommunications cables [1]. The maritime security environment is a priority for many nations and international organizations, and ship traffic monitoring represents one of the biggest challenges to law enforcement, search and rescue, environmental protection and resource management.

Cooperative vessel self-reporting systems, including Automatic Identification System (AIS), provide near real time information [2], [3]. The International Maritime Organization's (IMO) International Convention for the Safety of Life at Sea (SOLAS) [2] requires AIS to be fitted aboard international voyaging ships with gross tonnage (GT) of 300 or more, and all passenger ships regardless of size. Each AIS transmitting vessel will report its position depending on factors such as the speed and maneuvering status. In order to make the most efficient use of the bandwidth available, vessels that are anchored or moving slowly transmit less frequently than those that are moving faster or are maneuvering. The update interval ranges from 3 minutes for anchored or moored vessels, to 2 seconds for fast moving or maneuvering vessels, the latter being similar to that of conventional marine radar.

While AIS was originally conceived for collision avoidance, and the main use of the system is for local and real time applications, there are increasing possibilities for the use of AIS beyond this scope. Coastal states are also able to receive, plot and log the data by means of receiving stations along the coast, in the air, or in space. The amount of information reported by AIS providers is impressive, and for reference we report in Fig. 1 the worldwide density of traffic computed using AIS data collected during six months at the NATO Science and Technology Organization Centre for Maritime Research and Experimentation (STO-CMRE).

The NATO STO-CMRE uses AIS data from a variety of sources, including government-to-government near real time collection networks based on coastal receivers, external providers, and single coastal receiving stations. These, together with other experimental data, are used for scientific purposes and the development of algorithms. Among the possible applications, a developmental tool, namely Traffic Route Extraction for Anomaly Detection (TREAD), has been recently developed [4]–[6], aimed to process historical data in order to automatically obtain knowledge about maritime traffic and Pattern of Life (PoL) of ships at sea. AIS information is also typically used as ground truth in order to estimate the performance of coastal radars, see e.g. [7]–[14].

While in theory AIS information provides high-fidelity target kinematic estimates, in practice several kinds of errors are often observed. Usually, a filtering mechanism, such as the Kalman filter, is used to fit the observation sequence to a given dynamic model [15]. In this case, one assumption that is frequently made is that the measurement propagates from the sensing device to the filter occurs without delay. However, in practice time delays can occur between when an observation is taken by the sensor on board the vessel and when it becomes available to the filtering algorithm leading to time delayed measurements. Moreover, if receiver clocks are not synchronized, the unknown time delay may also be negative –without violating causality.

Generally speaking, any sensor network poses the issue of clock synchronization among receivers, but AIS networks are particularly subject to this kind of problem, because a complete timestamp of the transmission time is not present in the positional message broadcast by AIS devices on-board

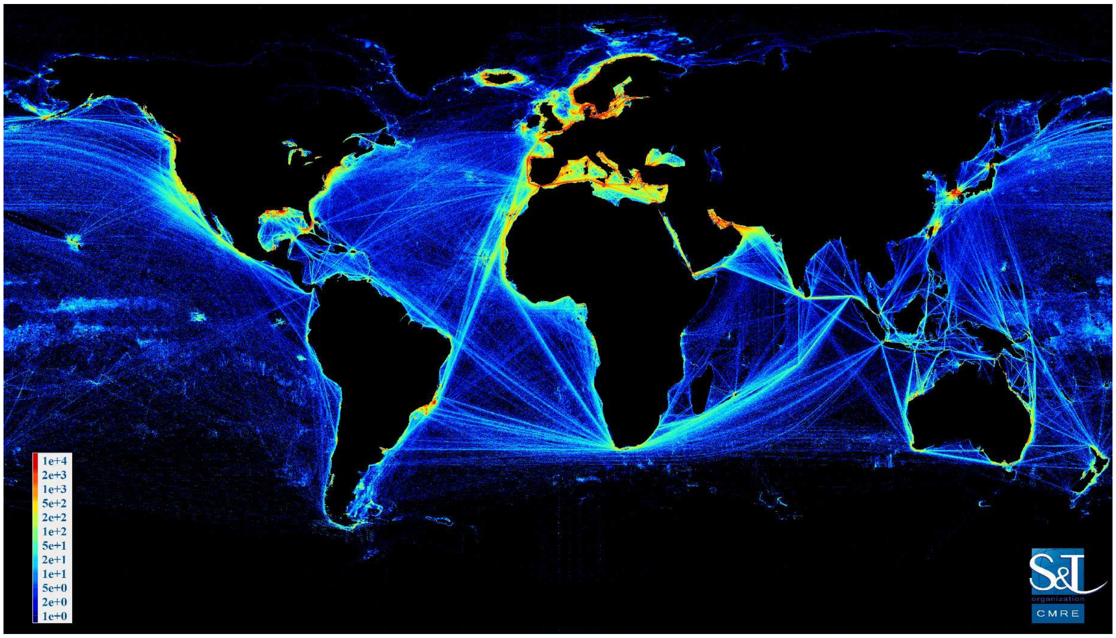

Fig. 1. Density of Automatic Identification System (AIS) messages collected from multiple AIS networks from April to September 2012. Each pixel covers a 4 nmi (one-fifteenth-degree) square on the ground and its color is (logarithmically) proportional to the number of ships whose reported position fall within its footprint.

vessels. When measurements with different time delays are interleaved with one another, this is known as the Out-of-Sequence Measurement (OOSM) problem [16], [17]. If time delays are known, the filter can be extended to account for the delays, and in the recent years a number of works have dealt with this problem. In [16] the exact solution for the OOSM is provided, while in [18] 1-step-lag algorithms are efficiently generalized to handle an arbitrary lag while preserving their main feature of solving the update problem without iterating. The extension to the particle filtering toolbox that enables nonlinear/non-Gaussian filtering with arbitrary lag is proposed in [19]. In [20] the author studied the optimal estimation procedure when the sensor measurements are subject to delay or might even be completely lost.

All of the aforementioned methods assume that, although the measurements are delayed, the –potentially random– amount of delay is known. However, as discussed in [21], situations can arise where the time delay is not known perfectly. To the best of authors' knowledge there are only few works dealing with such a problem. In [21] the problem is addressed using the Covariance Union (CU) technique [22]. In [23] the least squares filtering problem is investigated when observations are affected by stochastic delays, where the delay is random and can amount to at most one sampling time. In [24] authors analyze several existing methods to incorporate possible (often uncertain) measurement delays, typically applied for a variety of chemical processes systems. While in [21], [23], [24] the time delay is modeled as a discrete value multiple of the sensor sampling time, in the case of AIS, the ships' reporting activity is asynchronous by design, because the data originate from a variety of transponders.

In this work we propose a different method to deal with unknown delayed measurements based on the adaptive filtering that is able to sequentially detect the time delays and correctly estimate the target state. Two decision statistics for the presence of time error are derived, the first being non-parametric and the second based on the Generalized Likelihood Ratio Test (GLRT). When a time error is detected the Maximum Likelihood (ML) estimate of the error can be used to correct the timestamps of target observations used in the filter.

The paper is organized as follows. In Sec. II the problem is formulated, in Sec. III an adaptive filter solution is proposed and detailed, in Sec. IV results using synthetic and real-world data are reported. Finally, in Sec. V the final remarks are presented.

II. Problem formulation

Let us consider the following state transition and observational model

$$\boldsymbol{x}_k = \boldsymbol{f}\left(t_k, \boldsymbol{x}_{k-1}, \boldsymbol{v}_k\right), \tag{1}$$

$$\boldsymbol{y}_k = \boldsymbol{h}\left(\boldsymbol{x}_k, \boldsymbol{n}_k\right), \tag{2}$$

$$\tau_k = t_k + \delta_k, \tag{3}$$

where $\boldsymbol{f}$ and $\boldsymbol{h}$ are respectively the state transition and measurement function, t_k is the time between two consecutive target states, $\boldsymbol{v}_k$ and $\boldsymbol{n}_k$ are respectively the process noise and the measurement noise, both often assumed as a sequence of independent and identically distributed (i.i.d.) random variables.

In this work we consider the possibility that the acquisition time of a measurement $\boldsymbol{y}_k$ at step k, indicated with τ_k, is affected by an error. We model this error by using an *unknown* time delay, denoted by δ_k. Consequently, at step k we observe the measurement $\boldsymbol{y}_k$ and its acquisition time τ_k.

The target state is denoted by $\boldsymbol{x}_k = \left[x_k^{(1)}, \dot{x}_k^{(1)}, x_k^{(2)}, \dot{x}_k^{(2)}\right]^T$, where $x_k^{(1)}$ and $x_k^{(2)}$ are the positional coordinates and $\dot{x}_k^{(1)}$ and $\dot{x}_k^{(2)}$ are the velocities in the two dimensions. A common formulation for (1) is the near constant velocity (NCV) model

$$\boldsymbol{x}_k = \boldsymbol{F}(t_k)\boldsymbol{x}_{k-1} + \boldsymbol{A}(t_k)\boldsymbol{v}_k, \tag{4}$$

where

$$\boldsymbol{F}(t) = \begin{bmatrix} 1 & t & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & t \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{5}$$

$$\boldsymbol{A}(t) = \begin{bmatrix} t^2/2 & 0 \\ t & 0 \\ 0 & t^2/2 \\ 0 & t \end{bmatrix}, \tag{6}$$

where $\boldsymbol{v}_k$ is the two-dimensional acceleration noise vector modeled as an i.i.d. random process with Gaussian distribution $\mathcal{N}\left(\boldsymbol{v}; \boldsymbol{0}, \sigma_v^2 \boldsymbol{I}_2\right)$.

We also consider a linear measurement function, so we have

$$\boldsymbol{y}_k = \boldsymbol{H}\boldsymbol{x}_k + \boldsymbol{n}_k, \tag{7}$$

where $\boldsymbol{H}$ is the measurement matrix and $\boldsymbol{n}_k$ is distributed as $\mathcal{N}(\boldsymbol{n}; \boldsymbol{0}, \boldsymbol{R})$ with a covariance matrix $\boldsymbol{R}$.

The aim of our work is to estimate $\boldsymbol{x}_k$ based on the data $\{\boldsymbol{y}_i, \tau_i\}_{i=1}^k$. Note that the system of state transition (4) and measurement equation (7) is no longer linear because of the fact that t_k is unknown.

Assuming that a delay event ($\delta_k \neq 0$) can be present or not present, we can model the problem as a hypothesis testing problem (see the background on [25], [26])

$$\tau_k = \begin{cases} t_k, & \text{if } \mathcal{H}_k = \bar{\mathcal{D}}, \\ t_k + \delta_k, & \text{if } \mathcal{H}_k = \mathcal{D}, \end{cases} \tag{8}$$

where the current hypothesis $\mathcal{H}_k \in \left\{\bar{\mathcal{D}}, \mathcal{D}\right\}$, and $\bar{\mathcal{D}}$ denotes the *simple* hypothesis with absence of delay while $\mathcal{D}$ denotes the *composite* hypothesis with presence of delay where δ_k can be modeled as an unknown parameter.

III. Adaptive filter for unknown delayed measurements

Based on a suitable decision statistic T_k at step k for the hypothesis test (8), the delay δ_k can be

- detected,
- estimated, and
- used to correct the state estimate.

Let us assume that the estimate $\boldsymbol{x}_{k-1|k-1}$ at step $k-1$ (the common notation $\boldsymbol{x}_{m|n}$ indicates the estimate of $\boldsymbol{x}_m$ based on observations up to time n) is Gaussian with a covariance matrix $\boldsymbol{P}_{k-1|k-1}$. Then, under the hypothesis that there is a delay δ_k and by using the linear-Gaussian assumptions made in the previous section, the prediction is Gaussian

$$\boldsymbol{x}_{k|k-1}(\delta_k) = \boldsymbol{F}(\tau_k - \delta_k)\, \boldsymbol{x}_{k-1|k-1}, \tag{9}$$

with a covariance given by

$$\begin{aligned} \boldsymbol{P}_{k|k-1}(\delta_k) &= \boldsymbol{F}(\tau_k - \delta_k)\, \boldsymbol{P}_{k-1|k-1} \boldsymbol{F}(\tau_k - \delta_k)^T \\ &\quad + \sigma_v^2 \boldsymbol{A}(\tau_k - \delta_k)\, \boldsymbol{A}^T(\tau_k - \delta_k). \end{aligned} \tag{10}$$

Under the hypothesis that there is a delay δ_k, the innovation $\boldsymbol{\nu}_k(\delta_k) = \boldsymbol{y}_k - \boldsymbol{H}\boldsymbol{x}_{k|k-1}(\delta_k)$ is a zero mean Gaussian with covariance given by

$$\boldsymbol{S}_k(\delta_k) = \boldsymbol{H}\boldsymbol{P}_{k|k-1}(\delta_k)\, \boldsymbol{H}^T + \boldsymbol{R}. \tag{11}$$

As commonly applied in the context of adaptive filtering, e.g. in the case of input estimation or target maneuver detection, see [27], a decision statistic is given by the innovation. In our case, under $\bar{\mathcal{D}}$ (no delay), the innovation is zero mean Gaussian with covariance $\boldsymbol{S}_k(0)$, consequently the normalized innovation squared $\boldsymbol{\nu}_k(0)^T \boldsymbol{S}_k(0)^{-1} \boldsymbol{\nu}_k(0)$ is distributed as chi-square with n_y, dimension of the measurement, degree of freedom. A delay manifests itself as a "large" innovation (as for the case of a maneuvering target [27]), and a simple detection procedure for such an occurrence can be based on the normalized innovation squared

$$T_k = \boldsymbol{\nu}_k(0)^T \boldsymbol{S}_k(0)^{-1} \boldsymbol{\nu}_k(0) \begin{cases} < \gamma & \text{decide } \bar{\mathcal{D}}, \\ \geq \gamma & \text{decide } \mathcal{D}, \end{cases} \tag{12}$$

where γ is the threshold which determines the probability of false alarm $P_{FA}(\gamma) = \mathbb{P}\left[\text{decide } \mathcal{D} \,\middle|\, \bar{\mathcal{D}}\right]$ and the probability of detection $P_D(\gamma) = \mathbb{P}\left[\text{decide } \mathcal{D} \,|\, \mathcal{D}\right]$. In the case of a decision for $\mathcal{D}$, the delay can be estimated by using a Maximum Likelihood (ML) estimator of δ_k given that the distribution of the innovation is a zero mean Gaussian with covariance $\boldsymbol{S}_k(\delta_k)$, see (11). The ML estimator is then given by

$$\begin{aligned} \delta_k^{\text{ML}} &= \arg\max_{\delta \in \Delta} \left\{\mathcal{N}\left(\boldsymbol{\nu}_k(\delta_k); \boldsymbol{0}, \boldsymbol{S}_k(\delta_k)\right)\right\} \\ &= \arg\min_{\delta \in \Delta} \left\{\log\left|\boldsymbol{S}_k(\delta_k)\right| + \boldsymbol{\nu}_k(\delta_k)^T \boldsymbol{S}_k(\delta_k)^{-1} \boldsymbol{\nu}_k(\delta_k)\right\} \end{aligned} \tag{13}$$

where Δ is the interval of admissible delays and $|\cdot|$ is the determinant operator. Another detection strategy is based on the Generalized Likelihood Ratio Test (GLRT) [25], [26]

$$T_k = \frac{\mathcal{N}\left(\boldsymbol{\nu}_k\left(\delta_k^{\text{ML}}\right); \boldsymbol{0}, \boldsymbol{S}_k\left(\delta_k^{\text{ML}}\right)\right)}{\mathcal{N}\left(\boldsymbol{\nu}_k(0); \boldsymbol{0}, \boldsymbol{S}_k(0)\right)} \begin{cases} < \gamma & \text{decide } \bar{\mathcal{D}}, \\ \geq \gamma & \text{decide } \mathcal{D}, \end{cases} \tag{14}$$

where γ rules $P_{FA}(\gamma)$ and $P_D(\gamma)$.

In Fig. 2 there is a comparison between the nonparametric test (12) based on the normalized innovation squared and the GLRT (14). As expected, the GLRT exhibits a gain in terms of performance with respect to the nonparametric test. However, an advantage of the latter is that the threshold can be controlled analytically from the fact that $P_{FA}(\gamma)$ is a tail of a chi square cumulative distribution. Conversely, in the case of the GLRT, the threshold can only be determined numerically using a Monte Carlo simulation.

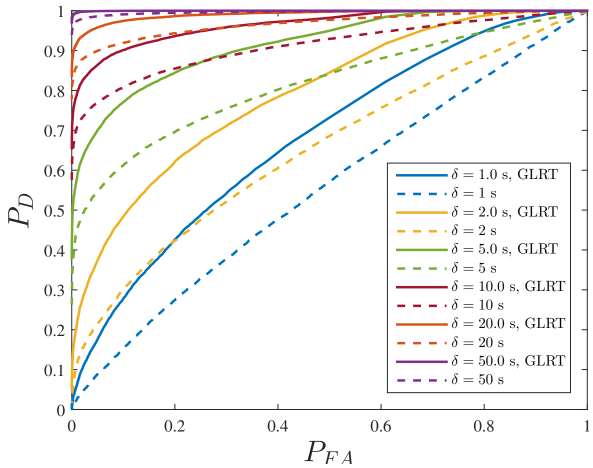

Fig. 2. Receiver operating characteristic (ROC) curve generated using a Monte Carlo simulation with $2 \cdot 10^4$ iterations. Parameters: $\boldsymbol{R} = \operatorname{diag}\left(50^2, 1, 50^2, 1\right)$, $\sigma_v^2 = 10^{-2}$.

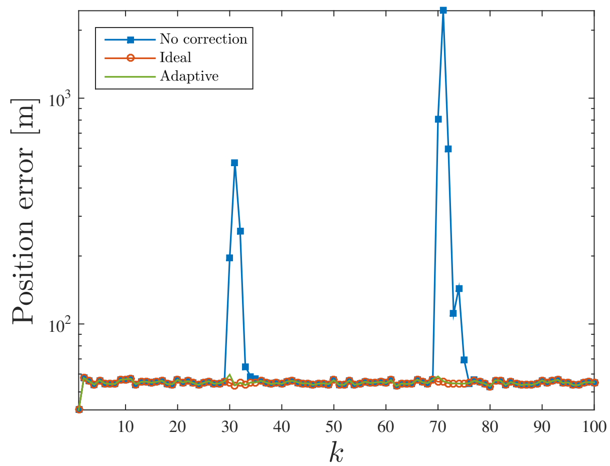

Fig. 3. Mean position errors using Monte Carlo simulation with 10^3 iterations. Comparison among the standard filter (no correction), the ideal filter (perfect correction) and the proposed adaptive filter. Parameters: $\boldsymbol{R} = \operatorname{diag}\left(50^2, .25, 50^2, .25\right)$, $\sigma_v^2 = 10^{-3}$.

A. Filtering based on the estimated delay

Given the decision at time k the estimated delay $\hat{\delta}_k$ is zero if $\bar{\mathcal{D}}$ is declared true, otherwise is δ_k^{ML}, equivalently we have

$$\hat{\delta}_k = I_{\{T_k \geq \gamma\}} \delta_k^{\text{ML}}, \tag{15}$$

where $I_{\mathcal{A}}$ is the indicator function of $\mathcal{A}$.

If the estimated inter-measurement interval $\tau_k - \hat{\delta}_k > 0$, then the update of the adaptive filter follows the Kalman Filter (KF) update by using

$$\boldsymbol{x}_{k|k}\left(\hat{\delta}_k\right) = \boldsymbol{x}_{k|k-1}\left(\hat{\delta}_k\right) + \boldsymbol{K}_k\left(\hat{\delta}_k\right)\boldsymbol{\nu}_k\left(\hat{\delta}_k\right), \tag{16}$$

$$\boldsymbol{K}_k\left(\hat{\delta}_k\right) = \boldsymbol{P}_{k|k-1}\left(\hat{\delta}_k\right)\boldsymbol{H}^T\boldsymbol{S}_k\left(\hat{\delta}_k\right)^{-1}, \tag{17}$$

$$\boldsymbol{P}_{k|k}\left(\hat{\delta}_k\right) = \left(I - \boldsymbol{K}_k\left(\hat{\delta}_k\right)\boldsymbol{H}\right)\boldsymbol{P}_{k|k-1}, \tag{18}$$

where $\boldsymbol{x}_{k|k-1}\left(\hat{\delta}_k\right)$ and $\boldsymbol{P}_{k|k}\left(\hat{\delta}_k\right)$ are given in (9) and (10) respectively. Alternatively, if the estimated time interval $\tau_k - \hat{\delta}_k < 0$ then the measurement $\boldsymbol{y}_k$ has to be considered as an OOSM, i.e. $\boldsymbol{y}_k$ occurred before than $\boldsymbol{y}_{k-1}$. In this case it is possible to use any OOSM filtering procedure, *e.g.* see [16], [17]. In the following we adopt the exact solution to the OOSM problem.

IV. Experimental Results

In this section we report experimental results using computer simulated trajectories and real-world AIS data.

A. Computer experiment

The main purpose of this simulation is to evaluate the accuracy gain of the GLRT described in Sec. III over a conventional KF, as well as its feasibility for a realistic target tracking scenario.

A Monte Carlo simulation of the position error has been carried out. Three cases have been analyzed:

- ideal case (no delay)
- two unknown delays that are not compensated, and
- two unknown delays faced by the proposed adaptive using the GLRT as detector.

The Monte Carlo simulation has been run with $N_{\text{MC}} = 10^3$ iterations to ensure statistically significant results. At each iteration, a synthetic trajectory of $K = 10^2$ data points with a sampling time of $T_s = 3 \cdot 10^2$ seconds is generated. The same trajectory is then used to feed our adaptive filter and two conventional KFs, the first one that knows precisely the time of measurements, while the second observes delayed (by an unknown amount) measurements.

Fig. 3 illustrates how the error varies with the time index k. Two unknown delays have been introduced in the synthetic measurements, one at $k = 30$ and another one at $k = 70$. The delay –which is known only to the *ideal* filter– amounts to 10^3 seconds at $k = 30$ and -10^3 seconds at $k = 70$.

It has been shown that the accuracy of our adaptive filter is always as good as or better than the one that would be obtained by using a conventional filter. In particular, when a delayed measurement arrives, the error of the adaptive filter is two orders of magnitude lower than the conventional filter, showing also a faster convergence rate to the steady state error of the ideal case. Furthermore, the adaptive filter behaves exactly as the ideal filter outside the region of delayed measurements, leading to the same accuracy.

B. Real-world example

Global satellite and terrestrial AIS data from multiple receiving stations have been used to validate the approach presented in Sec. III. A dataset of received messages that may be affected by timing errors from mixed terrestrial and satellite sensors has been identified and extracted from the STO-CMRE historical data archive, and this has been used to validate the approach proposed in Sec. III. The dataset spans a period of approximately 20 hours, during which time positions reported

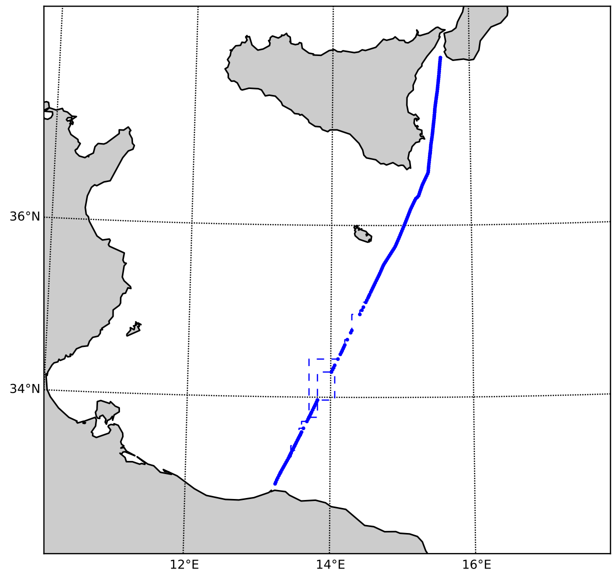

Fig. 4. Step plot of a trajectory of a container ship going from Tripoli to South Italy. Each point is connected to the time-subsequent one by a step. Steps underlying data points represent OOSMs.

by ships all over the world have been considered and, among all the trajectories, two have been used to demonstrate the effectiveness of our approach with real data. In this section we present the results achieved with the proposed model against two trajectories affected by timing problems, in the detection, estimation and possible correction of the errors.

In this example we considered the trajectory of a container ship going from Tripoli to South Italy. The ship's trajectory, as shown in Fig. 4, is mostly rectilinear. With an average observed time interval between subsequent positional messages of about 250 seconds, the ship's AIS reporting activity can be considered regular.

Fig. 4 shows a *step-plot* of the trajectory: given two consecutive data points, a step is drawn that connects them. At this zoom level, the step is only noticeable for subsequent data points that are very far from each other: visible steps represents two time-subsequent data points that are very far in space from each other. Therefore, noticeable steps might be symptomatic of errors in the time of measurements, in the sense that they highlight coverage gaps or OOSMs.

A detail of the trajectory is shown in Fig. 5a, where the measurements have been converted to Cartesian coordinates. Observations in this figure are also labeled with an ordinal number that represents the order of arrival, formally $t_i > t_j \Leftrightarrow i > j, \forall i, j \in \mathbb{N}_0$.

Even by eye it is not difficult to spot that the measurements from t_{17} to t_{19} are OOSM with respect to preceding and following data samples. In fact, if the time of the measurements from t_{17} to t_{19} were not wrong, it would have meant that a 166-by-28 meters, 15,000 GT container ship would have been capable of dramatically changing its heading of 180° twice in

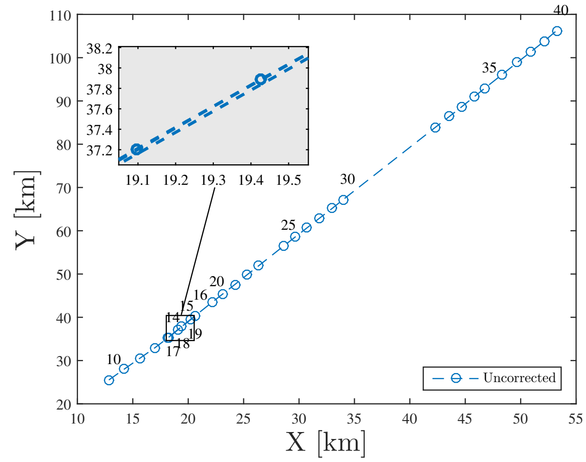

(a) Without correction

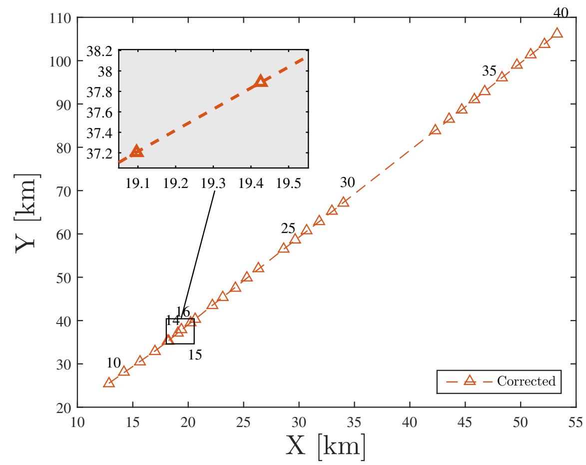

(b) With correction

Fig. 5. Detail of the ship's trajectory. Data points are labeled with an ordinal number that indicates the order of arrival, i.e. $t_i > t_j \Leftrightarrow i > j$.

a few minutes, recovering also all the velocity lost during the maneuver.

Filtering the trajectory using the approach described in Section III leaves us with the situation depicted in Fig. 5b, where again the labels indicate the progressive number of the observation. The proposed GLRT has been able to detect time errors, to estimate them, and ultimately to adjust the filter estimate. OOSMs were also detected and corrected. The parameters used for the filter are: $\boldsymbol{R} = \text{diag}\left(50^2, 1, 50^2, 1\right)$, $\sigma_v^2 = 10^{-3}$.

Fig. 6 provides information on the inner functioning of the proposed approach. The chart at the middle of the figure illustrates how the LLR varies with time. Shown in the boxes at the top and at the bottom of the figure are, in order, the plots of the estimated time error and target speed. The improvement over the conventional filter (which ignores the possibility of timing errors) accruing from the correction of the filter estimate is especially clear from this last chart. Finally,

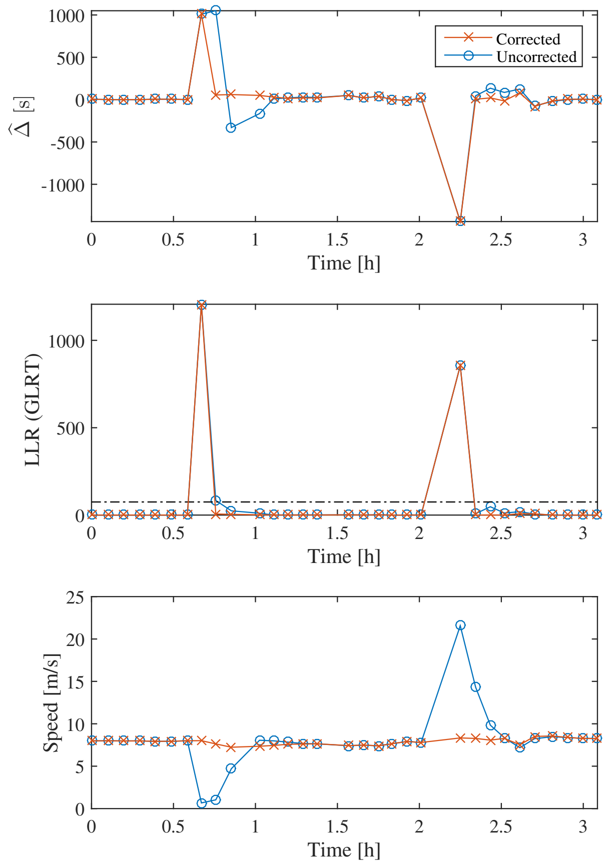

Fig. 6. From the top to the bottom, in order: estimated time delay of measurements, LLR and target speed, over time.

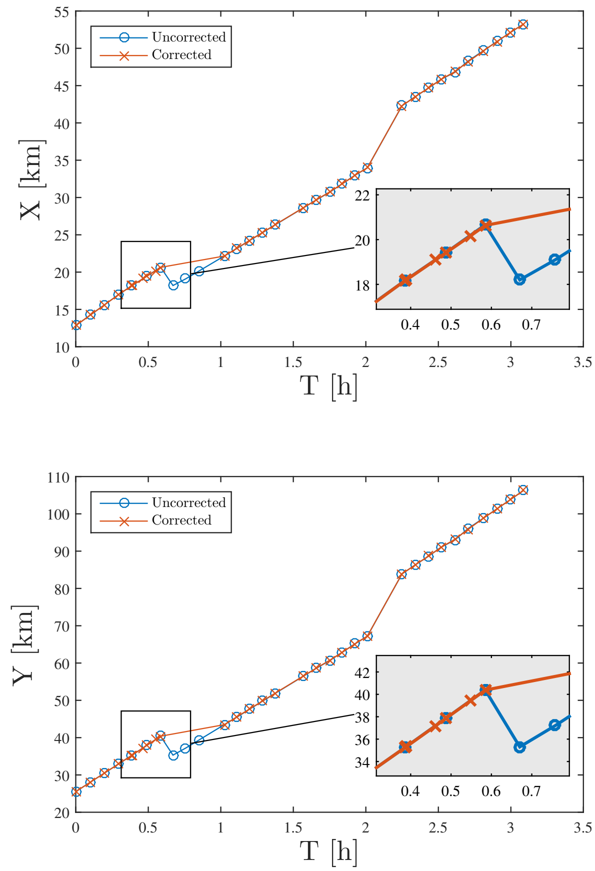

Fig. 7. Cartesian components of the filtered target position vector over time, in both corrected (red) and uncorrected (blue) versions.

Fig. 7 shows the effect of the filter on the components of the target positions over time.

V. Conclusion

This paper deals with the problem of estimating the target state from a sequence of observations that can be imprecisely time-stamped. This is a typical situation that arises in the scenario when a vessel is observed from a network of AIS sensor receivers.

An adaptive filtering strategy able to detect possible time delays and to correctly estimate the target state is proposed. Two decision statistics based on the KF innovation are derived, the first is non-parametric while the second is based on the GLRT. When a delayed measurement is detected, the Maximum Likelihood (ML) estimate of the delay is used to correct the timestamps of the target observation.

The proposed method is validated using Monte Carlo computer simulations and real-world data collected by a global network of AIS receivers.

Acknowledgment

L. M. Millefiori, P. Braca and K. Bryan would like to thank the NATO Allied Command Transformation (ACT) for supporting the CMRE project Data Knowledge Operational Effectiveness (DKOE).

References

[1] "NATO Alliance Maritime Strategy," *Official texts*, Mar. 2011.
[2] *Safety of Life at Sea (SOLAS) Convention*. Chapter V, Regulation 19.
[3] C. of the European Communities, "Document com(2008) 310 final - 2005/0239 cod, brussels 11- 06-2008."
[4] M. Vespe, I. Visentini, K. Bryan, and P. Braca, "Unsupervised learning of maritime traffic patterns for anomaly detection," in *Data Fusion Target Tracking Conference (DF TT 2012): Algorithms Applications, 9th IET*, May 2012.
[5] G. Pallotta, M. Vespe, and K. Bryan, "Vessel pattern knowledge discovery from AIS data: A framework for anomaly detection and route prediction," *Entropy*, vol. 15, no. 6, pp. 2218–2245, Jun. 2013.
[6] G. Pallotta, S. Horn, P. Braca, and K. Bryan, "Context-enhanced vessel prediction based on Ornstein-Uhlenbeck processes using historical AIS traffic patterns: Real-world experimental results," in *Proc. of the* 17^{th} *Intern. Conf. on Inform. Fusion (FUSION)*, Salamanca, Jul. 2014.
[7] P. Braca, R. Grasso, M. Vespe, S. Maresca, and J. Horstmann, "Application of the JPDA-UKF to HFSW radars for maritime situational awareness," in *Proc. of the* 15^{th} *Intern. Conf. on Inform. Fusion (FUSION)*, Singapore, Jul. 2012.
[8] P. Braca, M. Vespe, S. Maresca, and J. Horstmann, "A novel approach to high frequency radar ship tracking exploiting aspect diversity," in *Proc. of the IEEE Intern. Geo. and Rem. Sens. Symp. (IGARSS)*, Munich, Germany, Jul. 2012.

[9] S. Maresca, P. Braca, and J. Horstmann, "Data fusion performance of HFSWR systems for ship traffic monitoring," in *Proc. of the* 16^{th} *Intern. Conf. on Inform. Fusion (FUSION)*, Istanbul, Jul. 2013.

[10] L. Bruno, P. Braca, J. Horstmann, and M. Vespe, "Experimental evaluation of the range-Doppler coupling on HF surface wave radars," *IEEE Geosci. Remote Sens. Lett.*, vol. 10, no. 4, pp. 850–854, Jul. 2013.

[11] S. Maresca, P. Braca, J. Horstmann, and R. Grasso, "Maritime surveillance using multiple high-frequency surface-wave radars," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 8, pp. 5056–5071, Aug. 2014.

[12] G. Vivone, P. Braca, and J. Horstmann, "Knowledge-based multi-target tracking via UKF-JPDA variable structure IMM estimator in HF surface wave radar systems," *IEEE Trans. Geosci. Remote Sens.*, 2015.

[13] G. Papa, P. Braca, S. Horn, S. Marano, V. Matta, and P. Willett, "Adaptive bayesian tracking under time-varying target detection capability of sensor networks with measurement origin uncertainty," *IEEE Trans. Signal Process.*, submitted.

[14] ——, "Adaptive bayesian tracking with unknown time-varying sensor network performance," in *40th IEEE Int. Conf. on Acoust., Speech and Signal Process. (ICASSP)*, Brisbane, 2015.

[15] G. Papa, S. Horn, P. Braca, K. Bryan, and G. Romano, "Estimating sensor performance and target population size with multiple sensors," in *Proc. of the* 15^{th} *Intern. Conf. on Inform. Fusion (FUSION)*, jul 2012, pp. 2102–2109.

[16] Y. Bar-Shalom, "Update with out-of-sequence measurements in tracking: exact solution," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 38, no. 3, pp. 769–777, Jul 2002.

[17] S. Challa, R. J. Evans, and X. Wang, "A Bayesian solution and its approximations to out-of-sequence measurement problems," *Information Fusion*, vol. 4, pp. 185–199, 2003.

[18] Y. Bar-Shalom, H. Chen, and M. Mallick, "One-step solution for the multistep out-of-sequence-measurement problem in tracking," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 40, no. 1, pp. 27–37, Jan 2004.

[19] M. Orton and A. Marrs, "Particle filters for tracking with out-of-sequence measurements," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, no. 2, pp. 693–702, April 2005.

[20] L. Schenato, "Optimal estimation in networked control systems subject to random delay and packet drop," *IEEE Trans. Autom. Control*, vol. 53, no. 5, pp. 1311–1317, June 2008.

[21] S. J. Julier and J. K. Uhlmann, "Fusion of time delayed measurements with uncertain time delays," in *Proc. of the American Control Conference (ACC05)*, 2005, pp. 4028–4033.

[22] J. K. Uhlmann, "Covariance consistency methods for fault-tolerant distributed data fusion," *Information Fusion*, vol. 4, no. 3, pp. 201–215, 2003.

[23] A. Hermoso-Carazo and J. Linares-Prez, "Extended and unscented filtering algorithms using one-step randomly delayed observations," *Applied Mathematics and Computation*, vol. 190, no. 2, pp. 1375–1393, 2007.

[24] "Incorporating delayed and infrequent measurements in extended Kalman filter based nonlinear state estimation," *Journal of Process Control*, vol. 21, no. 1, pp. 119–129, 2011.

[25] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume II: Detection Theory*. Englewood Cliffs, New Jersey: Prentice Hall, 1998.

[26] E. Lehmann, *Testing Statistical Hypotheses*. Springer, 2005.

[27] Y. Bar-Shalom, X. Rong Li, and K. T., *Estimation with Applications to Tracking and Navigation*. New York: John Wiley & Sons, 2001.