DISTRIBUTED KALMAN FILTERING BASED ON SEVERELY QUANTIZED WSN DATA *

Alejandro Ribeiro and Georgios. B. Giannakis

Dept. of ECE, University of Minnesota
200 Union Street, Minneapolis, MN 55455, USA

ABSTRACT

This paper deals with recursive random parameter or state estimation for use in distributed tracking applications implemented with a Wireless Sensor Network (WSN). Bandwidth and energy limitations encountered with WSNs, motivate quantization of individual sensor observations before their digital transmission to the fusion center, where tracking is to be performed. Recent results investigating the intertwining between quantization and batch parameter estimation with WSNs, hint that quantization to a single bit per sensor may lead to a small penalty in state estimation variance. Relying on a dynamical model, we derive a Kalman-like Filter (KF) based on what we term "sign-differential" quantization, and establish that for all cases of practical interest, its asymptotic variance comes surprisingly close to the asymptotic variance of the clairvoyant minimum mean-square error KF state estimator which is based on the original (analog) observations. In a nutshell, this paper establishes the rather unexpected result that tracking with a WSN can simply rely on sensor observations quantized to a single bit.

1. INTRODUCTION

Continuously reducing the size and price of electronic devices has made possible the concept of Wireless Sensor Networks (WSN) in which a large number of small, power-limited, low-cost devices is deployed to perform monitoring, surveillance and control tasks [7]. The unique characteristics of WSNs require rethinking of many standard algorithms when bandwidth and power resources are reduced by orders of magnitude. Such rethinking is advanced in the context of distributed detection, where a WSN is deployed to decide among a number of possible hypotheses; see e.g., [15, 16] and references therein. The related problem of distributed estimation however, has received relatively less attention in the WSN context. Spatial correlations have been exploited to reduce the amount of information exchanged among sensors [2, 3, 4, 5, 8, 11, 12]. Recent works –to which the present paper belongs to– explore the intertwining between quantization and estimation that arises due to the distributed nature of WSNs [1, 9, 10, 13, 14]. If a digital communication system is to be used, individual observations have to be quantized and estimation can only rely on this set of *binary observations*, a problem undoubtedly different from estimation based on the original (analog) observations. The information loss due to quantization certainly increases the achievable variance [10]. Results in [10] hint that when the noise power is at least comparable with the range of the parameter to be estimated, quantization to a single bit per observation can lead to a minimal increase in variance. Recent studies in [13, 14] show that this holds true for a large class of estimation problems pertaining to different signal models and different degrees of knowledge about the additive noise statistics corrupting the observations.

* Work in this paper was prepared through collaborative participation in the Communications and Networks Consortium sponsored by the U. S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0011. The U. S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

Going a step forward, this paper introduces a dynamical model for the parameter to be estimated. This is a challenging generalization since the customary Kalman Filter (KF) that is known to offer the Minimum Mean-Squared (MMSE) estimator cannot be directly applied because of the required quantization step. Considering this fact, our first contribution in this paper is a Kalman-like filter (that we abbreviate as Dymbo filter) to closely approximate the MMSE estimator based on binary observations. We further study the Mean-Squared Error (MSE) achieved by the Dymbo filter, and establish the surprising result that for almost all cases of practical interest the MSE of the Dymbo filter is within a factor of 2 of the Kalman filter's MSE based on the analog observations. To appreciate the importance of this result, recall that the Dymbo filter requires transmission of only a single bit per observation.

The rest of the paper is organized as follows. Section 2 presents ad-hoc and hierarchical WSN topologies and describes the transmission protocol used by the sensors to recursively estimate the state of the dynamical system under consideration. In Section 3, we review the clairvoyant KF, emphasizing the fact that for a KF to be worth implementing the variance of the driving input should be smaller than the variance of the measurement noise. The core of the paper is in Section 4, where the Dymbo filter based on the sign of the difference between the current estimate and the current prediction is introduced and shown to exhibit asymptotic variance very close to that of the KF for most cases of practical interest. No matter how useful the asymptotic variance is, the convergence rate of the Dymbo filter can be slower than the convergence rate of the clairvoyant KF. To alleviate this problem, we study initialization issues of the Dymbo filter in Section 5, where we show that the proposed initialization requires a different message only in the first step thus having a minimal effect in the overall bandwidth requirements. Section 6 demonstrates two implementations of the Dymbo filter, and Section 7 concludes the paper.

Notation: We will use $P_x(x)$ to denote the probability density function (pdf), of the random variable x taking the value x. The notation $P_{x|y(n),\ldots,y(0)}(x)$ will stand for the conditional pdf of x given $y(n),\ldots,y(0)$. The notation $P_x = \mathcal{N}(\mu,\sigma^2)$ abbreviates the fact that x is normally distributed with mean μ and variance σ^2; i.e., $P_x(x) = 1/(\sqrt{2\pi}\sigma)\exp[-(x-\mu)^2/\sigma^2]$.

2. NETWORK SETUP AND SYSTEM MODEL

Two different WSN architectures characterized by the presence or absence of a Fusion Center (FC) are considered. When an FC is present, the WSN is termed hierarchical in the sense that sensors act as information gathering devices for the FC that is in charge of processing this information; see Fig. 1. Sensor S_n sends the message $m(n)$ to the FC through a multiple access channel, and the FC

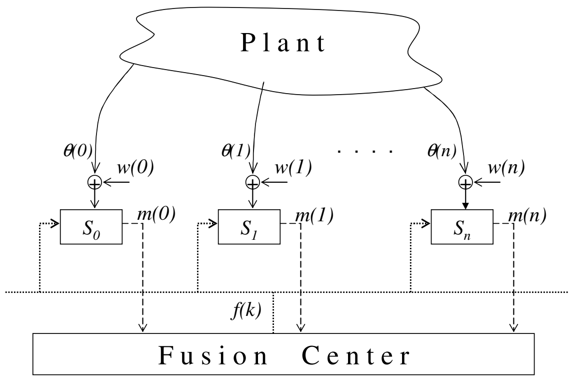

Fig. 1. A hierarchical WSN is used to track the plant parameter $\theta(n)$ that follows an AR-1 model. Based on the observation $x(n)$ a message $m(n)$ is constructed and sent to the FC, which broadcasts it back to all sensors.

broadcasts feedback information $f(k)$ to be used for the recursive parameter or state estimation. In ad-hoc WSNs, the network itself is responsible for processing, and to this end sensors communicate with each other through the shared wireless medium; see Fig. 2. We assume that the message $m(n)$ sent by sensor S_n is received by all other sensors, using a forwarding mechanism the details of which go beyond the scope of the present paper.

An important distinction between ad-hoc and hierarchical architectures pertains to the amount of information available to each sensor. In ad-hoc WSNs, the messages $m(n)$ percolate through all sensors, and information of their original observations is thus available at each sensor. In hierarchical WSNs, on the other hand, the information at each sensor becomes available through the feedback messages $f(k)$. In order to make the two architectures identical from the perspective of the problem at hand, the FC must broadcast back the last received message; i.e., we require $f(k) = m(k)$. It will also become clear later on that as far as communications resources are concerned, this feedback is the optimum strategy for the FC.

With either one of these network models, the WSN is deployed to monitor a dynamically evolving physical phenomenon, $\theta(n)$, that will be referred to as the plant. Although generalizations to matrix-vector models are possible, for simplicity in exposition, we will consider here an AR-1 process for the plant

$$\theta(n) = a\theta(n-1) + u(n), \tag{1}$$

whose driving input $u(n)$ is an independent identically distributed (i.i.d.) random process with $P_u = \mathcal{N}(0, \sigma_u^2)$. For the AR-1 process to be stable, we require $|a| < 1$. Notice that n here denotes both the sensor index as well as the time (or space) evolution of the plant.

At time n, when the parameter takes the value $\theta(n)$ the sensor S_n collects a noisy observation,

$$x(n) = \theta(n) + w(n), \tag{2}$$

where $w(n)$ denotes additive white Gaussian noise with pdf $P_{w(n)} = \mathcal{N}(0, \sigma^2)$.

For digital transmission, $x(n)$ has to be quantized to yield the message $m(n) = q[x(n)]$. Thus, our WSN-based tracking problem can be posed as the problem of estimating $\theta(n)$ given the messages $m(0), \ldots, m(n)$ sent respectively by sensors $S_0, \ldots, S_n$. Adopting the MMSE as our optimality criterion, the estimator we are after is given by the conditional expectation

$$\hat{\theta}(n) = \mathrm{E}[\theta(n)|m(n), \ldots, m(0)]. \tag{3}$$

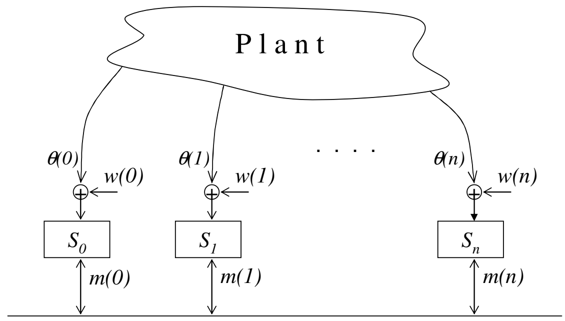

Fig. 2. Ad-hoc WSN architecture deployed to track $\theta(n)$. The messages $m(n)$ percolate through all sensors in the network.

Our first objective is to find the estimator in (3) for properly designed messages $m(n)$. Clearly, the MSE of the obtained estimates will depend on how the messages are designed. To this end, we will exploit the implicit information that previous estimates carry about the current observation, seeking quantization functions of the form:

$$m(n) = q[x(n), \hat{\theta}(n-1)]. \tag{4}$$

That is, we look for quantization functions that depend on the value of the observation and the previous parameter (or state) estimate.

In summary, given the plant-observation model (1)-(2) we want to design the quantizer function in (4) and find the MMSE optimum estimator as in (3). Before tackling this problem, let us review the clairvoyant KF from a WSN angle.

3. THE KF BENCHMARK

We begin by considering the case in which $m(n) = x(n)$. Note that this scheme is not realizable, since with $x(n)$ being analog, its *digital* transmission requires infinite bandwidth. However, the estimator based on $x(n)$ is a proper clairvoyant estimator to benchmark the loss associated with the quantized messages to be introduced in the next section.

When $m(n) = x(n)$, we deal with a classical estimation problem. It is known that the MMSE estimator and corresponding variance, given by

$$\begin{aligned} \hat{\theta}_{KF}(n|n) &:= \mathrm{E}[\theta(n)|x(n)\ldots x(0)], \\ M_{KF}(n|n) &:= \mathrm{var}[\theta(n|n)|x(n)\ldots x(0)], \end{aligned} \tag{5}$$

can be computed recursively by means of a KF [6, Chap.13].

The KF is derived by considering a prediction phase in which $\theta(n)$ is predicted from past observations $\{x(k)\}_{k=0}^{n-1}$, followed by a correction phase in which the current observation $x(n)$ is incorporated to correct the prediction.

Specifically, consider the estimate $\hat{\theta}_{KF}(n-1|n-1)$ and its variance $M_{KF}(n-1|n-1)$ at the $(n-1)^{st}$ step, and assume that the conditional distribution of $\theta(n-1)$ is Gaussian:

$$P_{\theta(n-1)|x(n-1),\ldots,x(0)} = \mathcal{N}[\hat{\theta}_{KF}(n-1|n-1), M_{KF}(n-1|n-1)]). \tag{6}$$

Under this assumption, it follows immediately that all the information from past observations is contained in $\hat{\theta}_{KF}(n-1|n-1)$ and $M_{KF}(n-1|n-1)$. By using the plant model (1), the 1-step prediction, $\hat{\theta}_{KF}(n|n-1) := \mathrm{E}[\theta|x(n-1)\ldots x(0)]$, is given by

$$\hat{\theta}_{KF}(n|n-1) = a\hat{\theta}_{KF}(n-1|n-1), \tag{7}$$

$$M_{KF}(n|n-1) = a^2 M_{KF}(n-1|n-1) + \sigma_u^2, \tag{8}$$

where $M_{KF}(n|n-1) := \text{var}[\theta(n|n-1)|x(n-1)\dots x(0)]$.

An important consequence of the linearity of (1), is that $\theta(n)$ given past observations is also normally distributed:

$$P_{\theta(n)|x(n-1),\dots,x(0)} = \mathcal{N}[\hat{\theta}_{KF}(n|n-1), M_{KF}(n|n-1)]. \quad (9)$$

Thus, the problem of estimating $\theta(n)$ after observing $x(n)$ is a simple Gaussian prior – Gaussian noise problem whose solution is known to be given by:

$$\hat{\theta}_{KF}(n|n) = \hat{\theta}_{KF}(n|n-1) + \frac{M_{KF}(n|n-1)}{M_{KF}(n|n-1)+\sigma^2}\tilde{x}(n), \quad (10)$$

where $\tilde{x}(n) := x(n) - \hat{\theta}_{KF}(n|n-1)$. The estimation variance also follows readily as

$$M_{KF}(n|n) = M_{KF}(n|n-1) - \frac{M^2_{KF}(n|n-1)}{M_{KF}(n|n-1)+\sigma^2}. \quad (11)$$

The last property of the Gaussian prior – Gaussian noise problem is that the posterior probability is Gaussian $P_{\theta(n)|x(n),\dots,x(0)} = \mathcal{N}[\hat{\theta}_{KF}(n|n), M_{KF}(n|n)]$. This justifies the previous assumption about the distribution $P_{\theta(n-1)|x(n-1),\dots,x(0)}$ and completes the KF derivation.

A particularly important performance indicator for KFs is the asymptotic (steady state) variance defined as

$$M_{KF_\infty} := \lim_{n\to\infty} M_{KF}(n|n), \quad (12)$$

that can be easily obtained by taking $M_{KF}(n|n) = M_{KF}(n-1|n-1)$, which in the limit obeys the so called Ricatti equation [c.f. (8) and (11)]:

$$a^2\left(\frac{M_{KF_\infty}}{\sigma^2}\right)^2 + \left(\frac{\sigma_u^2}{\sigma^2} + (1-a^2)\right)\left(\frac{M_{k_\infty}}{\sigma^2}\right) - \frac{\sigma_u^2}{\sigma^2}. \quad (13)$$

The asymptotic MSE given by the positive solution of (13), is seen to depend on $\gamma := \sigma_u^2/\sigma^2$, and is proportional to σ^2. Quantity M_{KF_∞}/σ^2 can be interpreted as the variance reduction achieved by the KF with respect to the crude observations $x(n)$.

This reduction is plotted in Fig. 3 as a function of γ, for three representative values of a. It is apparent that for $\gamma > 1$, the variance reduction achieved by the KF is small $(M_{KF_\infty}/\sigma^2) \approx 0.5$; an expected behavior, since for large σ_u subsequent states are essentially independent. Apart from intuition, the important point is that KF is practically effective only when $\gamma < 1$. Having brought this point to the reader's attention, we are set up to prove that in a WSN setup 1-bit messages $m(n)$ lead to a Kalman-like filter whose asymptotic variance is almost identical to the variance of the clairvoyant KF reviewed in this section.

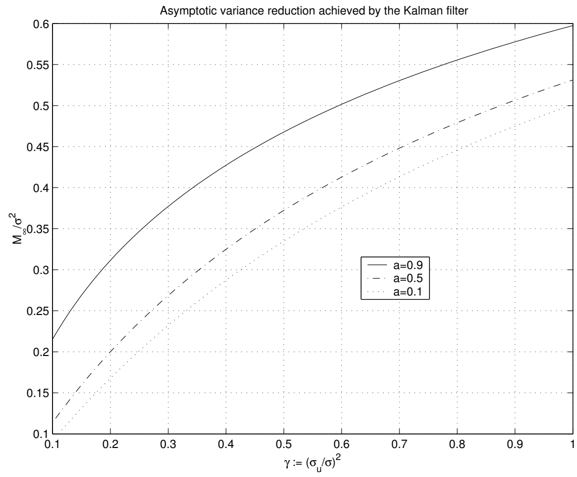

Fig. 3. Variance reduction achieved by a KF versus $\gamma := \sigma_u^2/\sigma$. The smaller σ_u^2 is relative to σ^2, the more effective KFs are.

4. KF WITH SEVERELY QUANTIZED DATA

As noted in Section 1, quantization is integral to the distributed nature of WSNs. This section considers estimation based on severely quantized data, and for that matter we define the binary observations

$$\begin{aligned} b(n) &= \text{sign}[x(n) - \hat{\theta}(n-1|n-1)] \\ &:= \begin{cases} +1, & \text{if } x(n) \geq \hat{\theta}(n-1|n-1) \\ -1, & \text{if } x(n) < \hat{\theta}(n-1|n-1) \end{cases}, \end{aligned} \quad (14)$$

and consider messages $m(n) = b(n)$. This implies that at step n, sensor S_n compares its observation $x(n)$ with the previous estimate and transmits the sign of the difference. Note that this setup requires a single bit per sensor –hence the name binary observation– and quantization amounts to taking the sign of $x(n) - \hat{\theta}(n-1|n-1)$ –hence the name "sign-differential" quantization.

The first goal of this section is to find an approximation to the MMSE estimator based on the binary observations $b(n)$ and the corresponding MSE:

$$\begin{aligned} \hat{\theta}(n|n) &:= \text{E}[\theta(n)|b(n)\dots b(0)], \\ M(n|n) &:= \text{var}[\theta(n|n)|b(n)\dots b(0)]. \end{aligned} \quad (15)$$

A second, more ambitious goal, is to establish that the asymptotic MSE

$$M_\infty := \lim_{n\to\infty} M(n|n), \quad (16)$$

is very close to the M_{KF_∞} obtained from (13) for $\gamma < 1$. This will establish that except for a transient behavior, quantization to a single bit per observation leads to an estimator whose variance is almost equal to the variance of the estimator based on the original (analog) observations for practically *all* cases of interest.

4.1. Dymbo filter

As a first step, let us consider the problem of MMSE estimation of a Gaussian random variable when a single binary observation is given. For this case, the integration required by the MMSE estimator is computable in closed form, as we prove in the following proposition.

Proposition 1 *Consider a normally distributed random variable θ, with $P_\theta = \mathcal{N}(\mu_\theta, \sigma_\theta^2)$; an observation of this random variable $x = \theta + w$, with $P_w = \mathcal{N}(0,\sigma)$; and define the binary observation $b = \text{sign}(x - \mu_\theta)$. Then,*

(a) the MMSE estimator of θ given b, and the corresponding MSE are given by:

$$\hat{\theta} := \text{E}[\theta|b] = \mu_\theta + \sqrt{\frac{2}{\pi}}\frac{\sigma_\theta^2}{\sqrt{\sigma_\theta^2+\sigma^2}}b, \quad (17)$$

$$\text{var}[\theta|b] = \sigma_\theta^2 - \frac{2}{\pi}\frac{\sigma_\theta^4}{\sigma_\theta^2+\sigma^2}. \quad (18)$$

(b) when $\sigma \to \infty$, the distribution of θ given b converges uniformly to a Gaussian distribution:

$$\lim_{\sigma\to\infty} P_{\theta|b} \stackrel{u}{=} \mathcal{N}[\hat{\theta}, \text{var}(\theta)|b]. \quad (19)$$

Proposition 1 can be used to construct an approximation to the MMSE estimator in (15). Towards this objective, let us assume as we did in Section 3 that at the $(n-1)^{st}$ step the estimate $\hat{\theta}(n-1|n-1)$ and its MSE $M(n-1|n-1)$ are known, and the conditional distribution is Gaussian; i.e., $P_{\theta(n-1)|b(n-1),\ldots,b(0)} = \mathcal{N}[\hat{\theta}(n-1|n-1), M(n-1|n-1)]$. If this is the case, it follows from the signal model (1) that the 1-step prediction $\hat{\theta}(n|n-1) := E[\theta(n)|b(n-1),\ldots,b(0)]$ is given by

$$\hat{\theta}(n|n-1) = a\hat{\theta}(n-1|n-1), \quad (20)$$
$$M(n|n-1) = a^2 M(n-1|n-1) + \sigma_u^2, \quad (21)$$

where $M(n|n-1) := \mathrm{var}[\theta(n|n-1)|b(n-1)\ldots b(0)]$.

Moreover, the corresponding distribution is also Gaussian; i.e., $P_{\theta(n)|b(n-1),\ldots,b(0)} = \mathcal{N}[\hat{\theta}(n|n-1), M(n|n-1)]$. Thus, Proposition 1 applies and allows one to conclude that

$$\hat{\theta}(n|n) = \hat{\theta}(n|n-1) + \frac{(\sqrt{2/\pi})\ M(n|n-1)}{\sqrt{M(n|n-1)+\sigma^2}} b(n), \quad (22)$$
$$M(n|n) = M(n|n-1) - \frac{(2/\pi)\ M^2(n|n-1)}{M(n|n-1)+\sigma^2}. \quad (23)$$

To complete our argument, we should establish that the posterior distribution $P_{\theta(n)|b(n),\ldots,b(0)}$ is also Gaussian, which unfortunately is not the case. Notwithstanding, invoking Proposition 1-(b) this is asymptotically true as $\sigma \to \infty$. This implies that the filter given by (20)-(23) offers asymptotically an approximation to the MMSE for small values of γ. This filter estimates the state of a DYnamical Model based on Binary Observations; hence the abbreviation Dymbo.

It is fair to comment that Proposition 1 is not as revealing with regards to error propagation and eventually error accumulation which could lead to considerable deviation from the actual MMSE. Using extensive simulations we have verified that errors do not propagate, and that $P_{\theta(n)|b(n),\ldots,b(0)}$ is very well approximated by a Gaussian distribution.

4.2. Asymptotic MSE

The striking similarities between (11) and (23) prompt one to compare the corresponding MSEs. Towards this objective, the asymptotic MSE defined by (16) can be obtained by setting $M(n|n) = M(n-1|n-1)$ to obtain the slightly more cumbersome Ricatti equation [c.f. (21) and (23)]:

$$\frac{M_\infty}{\sigma^2} = a^2\frac{M_\infty}{\sigma^2} + \gamma - \frac{(2/\pi)\ [a^2(M_\infty/\sigma^2)+\gamma]^2}{a^2(M_\infty/\sigma^2)+\gamma}, \quad (24)$$

where as before $\gamma = \sigma_u^2/\sigma^2$ and M_∞/σ^2 is the asymptotic variance reduction achieved by the Dymbo filter.

While it is possible to plot M_∞/σ^2 versus γ to obtain curves similar to those of Fig. 3, it is more informative to plot the asymptotic loss associated with the Dymbo filter

$$\ell := \frac{M_{KF\infty}}{M_\infty}. \quad (25)$$

A contour plot of ℓ as a function of a and γ is depicted in Fig. 4, revealing the surprising result that $\ell < 2$ when $\gamma < 1$, for all values of a. This result combined with the comments at the end of Section 3 establish the even more surprising conclusion that *when a KF is worth implementing, it can be implemented with the transmission of a single bit per observation.*

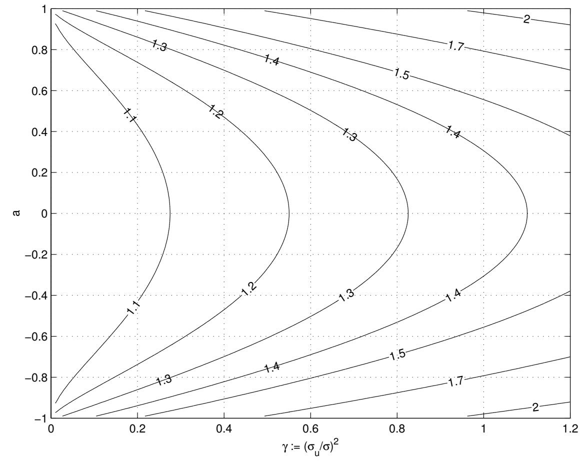

Fig. 4. The asymptotic MSE increase associated with the Dymbo filter, $\ell := (M_{KF\infty}/M_\infty)$. Surprisingly, the asymptotic loss is very small for the region of interest ($\gamma < 1$). Note that the largest penalty is paid for systems close to instability $|a| \approx 1$.

5. INITIALIZATION

As with KFs, Dymbo filters require a Gaussian prior distribution $P_{\theta(-1)} = \mathcal{N}(\mu_\theta, \sigma_\theta)$ to initialize the recursions using $\hat{\theta}(-1|-1) = \mu_\theta$, and $M(-1|-1) = \sigma_\theta^2$.

While we just proved that after reaching steady state there is no practical difference between Kalman and Dymbo filters, a quick comparison of (11) and (23) shows that the rate of convergence of the latter is slower, a particularly important problem when $M(-1|-1) \gg \sigma^2$. To alleviate this problem, we can allow transmission of a different message in the first step

$$m(0) = x_\tau(0) := \tau\mathrm{round}(x(n)/\tau)\ . \quad (26)$$

Specifically, we quantize the first observation to a grid of resolution (size) τ. If we select

$$\hat{\theta}(0) := m(0), \quad (27)$$

then it is straightforward to show that the MSE of this first estimate is bounded as specified in the following proposition.

Proposition 2 *The MSE of the first step estimator,* (27) *satisfies*

$$\mathrm{E}[(\hat{\theta}(0)-\theta(0))^2] \leq \sigma^2\left(1+\frac{\tau}{\sigma}+\frac{\tau^2}{4\sigma^2}\right). \quad (28)$$

Proof: See Appendix B.

By selecting $\tau \ll \sigma$, we can bring the variance of $\hat{\theta}(0)$ close to the noise variance σ^2, thus alleviating the slow convergence problem of the Dymbo filter. The effect on the overall bandwidth requirement is, of course, negligible.

6. SIMULATIONS

Including the initialization step, the messages transmitted by the Dymbo filter are given by

$$\begin{array}{lcll} m(n) & = & x_\tau(0) & n = 0, \\ m(n) & = & \mathrm{sign}[x(n) - \hat{\theta}(n-1|n-1)] & n \geq 1. \end{array} \quad (29)$$

The first step estimate $\hat{\theta}(0)$ is given by (27), and subsequent estimates are obtained by iterative application of (20)-(23).

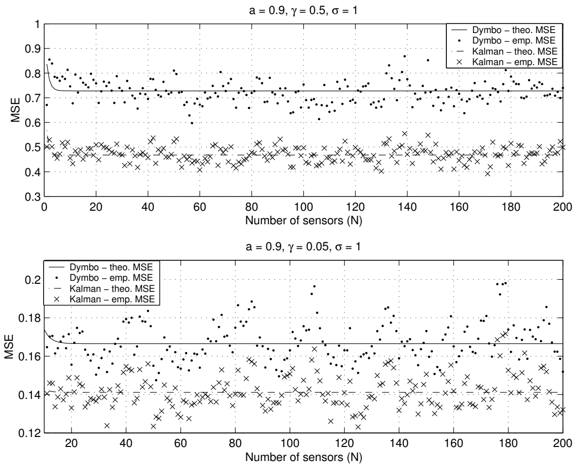

Fig. 5. Two implementations of the Dymbo filter compared with the corresponding KF. There is practically no difference in the behavior of both filters.

For the simulations we took $\sigma = 1$, and considered $a = 0.9$ since from Fig. 4 larger values of a lead to a larger loss of the Dymbo filter when compared to the KF. The initial quantization step was set to $\tau = \sigma$. Simulation results are shown in Fig. 5 for $\gamma = 0.5$, and $\gamma = 0.05$ confirming that in both cases the asymptotic MSE as given by (16) is a good prediction of the experimental results. This very fact also testifies that the approximate MMSE estimator defined by (20)-(23) is, indeed, a good approximation to the MMSE.

It can also be seen that the convergence of the Dymbo filter with proper initialization is very fast and virtually indistinguishable from the Kalman filter.

7. CONCLUSIONS

In order to enable distributed Kalman filtering for tracking applications with sensors transmitting severely quantized observations, we introduced the Dymbo filter which is based on the sign of the difference between the current observation and the prediction based on past observations. Even if single-bit quantization is most severe, we established that for virtually all cases of practical interest transmission of a single bit per observation (the sign of the difference) leads to a filter whose steady-state estimation variance is within a factor of 2 relative to the corresponding Kalman filter variance. The convergence rate of the Dymbo filter was observed to be a potential drawback, but we showed that by modifying the first (and only the first) estimation step it is possible to ameliorate it. The impact of this initialization in the overall bandwidth requirement is negligible.

In a nutshell, the conclusion of this paper is that *when a KF is worth implementing it can be implemented with the transmission of a single bit per observation*, which is very attractive for WSN-based tracking applications[1].

[1] The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U. S. Government.

8. REFERENCES

[1] M. Abdallah and H. Papadopoulos, "Sequential signal encoding and estimation for distributed sensor networks," in *Proc. of the ICASSP*, vol. 4, pp. 2577–2580, Salt Lake City, Utah, May 2001.

[2] B. Beferull-Lozano, R. L. Konsbruck, and M. Vetterli, "Rate-Distortion problem for physics based distributed sensing," in *Proc. of the ICASSP*, vol. 3, pp. 913–916, Montreal, Canada, May 2004.

[3] D. Blatt and A. Hero, "Distributed maximum likelihood estimation for sensor networks," in *Proc. of the ICASSP*, vol. 3, pp. 929–932, Montreal, Canada, May 2004.

[4] E. Ertin, R. Moses, and L. Potter, "Network parameter estimation with detection failures," in *Proc. of the ICASSP*, vol. 2, pp. 273–276, Montreal, Canada, May 2004.

[5] J. Gubner, "Distributed Estimation and Quantization," *IEEE Transactions on Information Theory*, vol. 39, pp. 1456–1459, 1993.

[6] S. M. Kay, *Fundamentals of Statistical Signal Processing - Estimation Theory*. Prentice Hall, 1993.

[7] S. Kumar, F. Zao, and D. Shepherd, eds., *Special issue on collaborative information processing*, vol. 19 of *IEEE Signal Processing Magazine*, March 2002.

[8] W. Lam and A. Reibman, "Quantizer design for decentralized systems with communication constraints," *IEEE Transactions on Communications*, vol. 41, pp. 1602–1605, Aug. 1993.

[9] Z.-Q. Luo, "Universal Decentralized Estimation in a Bandwidth Constrained Sensor Network," *IEEE Transactions on Information Theory*, November 2003 (submitted). available at http://www.ece.umn.edu/users/luozq/recent_work.html.

[10] H. Papadopoulos, G. Wornell, and A. Oppenheim, "Sequential signal encoding from noisy measurements using quantizers with dynamic bias control," *IEEE Transactions on Information Theory*, vol. 47, pp. 978–1002, 2001.

[11] S. S. Pradhan, J. Kusuma, and K. Ramchandran, "Distributed compression in a dense microsensor network," *IEEE Signal Processing Magazine*, vol. 19, pp. 51–60, March 2002.

[12] M. G. Rabbat and R. D. Nowak, "Decentralized source localization and tracking," in *Proc. of the ICASSP*, vol. 3, pp. 921–924, Montreal, Canada, May 2004.

[13] A. Ribeiro and G. B. Giannakis, "Bandwidth-Constrained Distributed Estimation for Wireless Sensor Networks, Part II: Unknown pdf," *IEEE Transactions on Signal Processing*, submitted August 2004. available at http://www.ece.umn.edu/users/aribeiro/Research/Pubs.html.

[14] A. Ribeiro and G. B. Giannakis, "Bandwidth-Constrained Distributed Estimation for Wireless Sensor Networks, Part I: Gaussian Case," *IEEE Transactions on Signal Processing*, submitted July 2004. available at http://www.ece.umn.edu/users/aribeiro/Research/Pubs.html.

[15] Y. Sung, L. Tong, and A. Swami, "Asymptotic locally optimal detector for large-scale sensor networks under the Poisson regime," in *Proc. of the ICASSP*, vol. 2, pp. 1077–1080, Montreal, Canada, May 2004.

[16] P. K. Varshney, *Distributed Detection and Data Fusion*. Springer-Verlag, 1997.