Optimal Sequential Parameter Estimation

Yasin Yılmaz
Electrical Engineering Department,
Columbia University
New York, NY 10027
Email: yasin@ee.columbia.edu

George V. Moustakides
Dep. of Electrical & Computer Eng.
University of Patras
26500 Rion, Greece
Email: moustaki@upatras.gr

Xiaodong Wang
Electrical Engineering Department,
Columbia University
New York, NY 10027
Email: wangx@ee.columbia.edu

***Abstract*—"THIS PAPER IS ELIGIBLE FOR THE STUDENT PAPER AWARD."**
We develop optimal centralized sequential estimators under different formulations of the problem. Decentralized sequential estimation is also considered for wireless sensor networks. We propose an asymptotically optimal decentralized scheme based on level-triggered sampling, a non-uniform sampling technique. Performance of the proposed scheme is analyzed.

I. Introduction

Sequential estimation is a classical problem in *sequential analysis*. It was, together with sequential hypothesis testing, first studied by Wald in his seminal work [1]. In sequential estimation, the sample size is not fixed, as opposed to the fixed sample size methods. Instead, it is determined by the collected samples according to a predefined stopping rule [2].

Decentralized parameter estimation is a fundamental signal processing task that can be realized in wireless sensor networks. Due to stringent bandwith and energy requirements imposed by sensors it is typically performed under the constraints of low bandwith usage and low communication rate. In this paper we are interested in sequential decentralized estimators rather than fixed-sample-size ones. There are a few works considering the sequential decentralized estimation in the literature, e.g., [3], [4], in which sensors employ the conventional uniform-in-time samplers to sample and transmit their local observations. On the other hand, similar to [5], in this paper we will consider using level-triggered sampling, a non-uniform sampling strategy, which perfectly fits to transmitting information in decentralized systems as recently shown in [6], [7].

II. Problem Formulation and Background Information

We represent scalars with lower-case letters, vectors with upper-case letters and matrices with upper-case bold letters. Consider the following linear signal model,

$$y_t = H_t^T X + w_t,\ t \in \mathbb{N}, \tag{1}$$

where $y_t \in \mathbb{R}$ is the observed sample, $X \in \mathbb{R}^{n\times 1}$ is the deterministic but unknown vector of parameters to be estimated, $H_t \in \mathbb{R}^{n\times 1}$ is the random vector of scaling coefficients (e.g., channel gain vector in a multiaccess channel) and $w_t \in \mathbb{R}$ is the additive noise. We observe, at each time t, the sample y_t and the coefficient vector H_t. Hence, at each time t, $\{(y_n, H_n)\}_{n=1}^t$ are known. We assume $\{w_t\}$ are i.i.d. with $\mathsf{E}[w_t] = 0$ and $\mathsf{Var}(w_t) = \sigma^2$.

The ordinary least squares (OLS) estimator minimizes the sum of squared errors, i.e.,

$$\hat{X}_t = \arg\min_X \sum_{n=1}^t (y_n - H_n^T X)^2, \tag{2}$$

and given by

$$\hat{X}_t = \left(\sum_{n=1}^t H_n H_n^T\right)^{-1} \sum_{n=1}^t H_n y_n = (\boldsymbol{H}_t^T \boldsymbol{H}_t)^{-1} \boldsymbol{H}_t^T Y_t, \tag{3}$$

where $\boldsymbol{H}_t = [H_1, \ldots, H_t]^T$ and $Y_t = [y_1, \ldots, y_t]^T$.

Under Gaussian noise, $w_t \sim \mathcal{N}(0, \sigma^2)$, the OLS estimator coincides with the minimum variance unbiased estimator (MVUE). That is to say, the OLS estimator achieves the Cramer-Rao lower bound (CRLB), i.e., $\mathsf{Cov}(\hat{X}_t|\boldsymbol{H}_t) =$ CRLB. To compute the CRLB we first write the log-likelihood ratio of the vector Y_t given X and $\boldsymbol{H}_t$ as

$$L_t = \log f(Y_t|X, \boldsymbol{H}_t) = -\sum_{n=1}^t \frac{(y_n - H_n^T X)^2}{2\sigma^2} - \frac{t}{2}\log(2\pi\sigma^2). \tag{4}$$

Then, we have

$$\text{CRLB} = \left(\mathsf{E}\left[-\frac{\partial^2}{\partial X^2} L_t \big| \boldsymbol{H}_t\right]\right)^{-1} = \sigma^2 \boldsymbol{U}_t^{-1}, \tag{5}$$

where $\mathsf{E}\left[-\frac{\partial^2}{\partial X^2} L_t \big| \boldsymbol{H}_t\right]$ is the Fisher information matrix and we defined $\boldsymbol{U}_t \triangleq \boldsymbol{H}_t^T \boldsymbol{H}_t$. Since $\mathsf{E}[Y_t|\boldsymbol{H}_t] = \boldsymbol{H}_t X$ and $\mathsf{Cov}(Y_t|\boldsymbol{H}_t) = \sigma^2 \boldsymbol{I}$, from (3) we have $\mathsf{E}[\hat{X}_t|\boldsymbol{H}_t] = X$ and $\mathsf{Cov}(\hat{X}_t|\boldsymbol{H}_t) = \sigma^2 \boldsymbol{U}_t^{-1}$, thus from (5) $\mathsf{Cov}(\hat{X}_t|\boldsymbol{H}_t) =$ CRLB. Note that the maximum likelihood estimator, maximizing (4), coincides with the OLS estimator in (3).

In general, under a non-Gaussian distribution the OLS estimator is the best linear unbiased estimator (BLUE). In other words, any linear unbiased estimator $\boldsymbol{A}_t Y_t$, where $\mathsf{E}[\boldsymbol{A}_t Y_t|\boldsymbol{H}_t] = X$ for any $\boldsymbol{A}_t \in \mathbb{R}^{n\times t}$ which is a function of $\boldsymbol{H}_t$, has a covariance no smaller than that of the OLS estimator in (3), i.e., $\mathsf{Cov}(\boldsymbol{A}_t Y_t|\boldsymbol{H}_t) \geq \sigma^2 \boldsymbol{U}_t^{-1}$ in the positive semidefinite sense. To see this result we write $\boldsymbol{A}_t = (\boldsymbol{H}_t^T \boldsymbol{H}_t)^{-1} \boldsymbol{H}_t^T + \boldsymbol{B}_t$ for some $\boldsymbol{B}_t \in \mathbb{R}^{n\times t}$, and then $\mathsf{Cov}(\boldsymbol{A}_t Y_t|\boldsymbol{H}_t) = \sigma^2 \boldsymbol{U}_t^{-1} + \sigma^2 \boldsymbol{B}_t \boldsymbol{B}_t^T$, where $\boldsymbol{B}_t \boldsymbol{B}_t^T$ is a positive semidefinite matrix.

The recursive least squares (RLS) algorithm enables us to compute $\hat{X}_t$ in a much simpler way than (3), which requires a matrix inversion at each time t. Using RLS, at each time t, we can update $\hat{X}_t$ as

$$\begin{aligned} &\hat{X}_t = \hat{X}_{t-1} + K_t(y_t - H_t^T \hat{X}_{t-1}) \\ \text{where} \quad &K_t = \frac{\boldsymbol{P}_{t-1} H_t}{1 + H_t^T \boldsymbol{P}_{t-1} H_t} \\ \text{and} \quad &\boldsymbol{P}_t = \boldsymbol{P}_{t-1} - K_t H_t^T \boldsymbol{P}_{t-1}, \end{aligned} \tag{6}$$

$K_t \in \mathbb{R}^{n \times 1}$ being the gain vector and $\boldsymbol{P}_t = \boldsymbol{U}_t^{-1}$. While applying RLS we first initialize $\hat{X}_0 = 0$ and $\boldsymbol{P}_0 = \delta^{-1}\boldsymbol{I}$, where 0 represents a zero vector and δ is a very small number, and then at each time t compute K_t, $\hat{X}_t$ and $\boldsymbol{P}_t$ as in (6), respectively.

III. Optimal Sequential Estimators

In this section we aim to find the optimal pair $(\mathcal{T}, \hat{X}_\mathcal{T})$ of stopping time and estimator. The stopping time for an estimator is selected as the first time it achieves a target accuracy level. We assess the accuracy of an estimator by using either its covariance matrix $\mathsf{Cov}(\hat{X}_t)$ or conditional covariance matrix $\mathsf{Cov}(\hat{X}_t | \boldsymbol{H}_t)$. Specifically, we have the following constrained optimization problems,

$$\min_{\mathcal{T}, \hat{X}_\mathcal{T}} \mathsf{E}[\mathcal{T} | \boldsymbol{H}_\mathcal{T}] \quad \text{such that} \quad f\left(\mathsf{Cov}(\hat{X}_\mathcal{T} | \boldsymbol{H}_\mathcal{T})\right) \leq C, \tag{7}$$

$$\text{and} \quad \min_{\mathcal{T}, \hat{X}_\mathcal{T}} \mathsf{E}[\mathcal{T}] \quad \text{such that} \quad f\left(\mathsf{Cov}(\hat{X}_\mathcal{T})\right) \leq C, \tag{8}$$

under the conditional and unconditional setups, respectively, where $f(\cdot)$ is a function from $\mathbb{R}^{n \times n}$ to $\mathbb{R}$ and $C \in \mathbb{R}$ is the target accuracy level.

Note that the constraint in (7) is stricter than the one in (8) since it requires that $\hat{X}_\mathcal{T}$ satisfies the target accuracy level for each realization of $\boldsymbol{H}_\mathcal{T}$, whereas in (8) it is sufficient that $\hat{X}_\mathcal{T}$ satisfies the target accuracy level on average. In other words, even if for some realizations of $\boldsymbol{H}_\mathcal{T}$ we have $f\left(\mathsf{Cov}(\hat{X}_\mathcal{T} | \boldsymbol{H}_\mathcal{T})\right) > C$, we can still have $f\left(\mathsf{Cov}(\hat{X}_\mathcal{T})\right) \leq C$. The accuracy function f should be a monotone function of the covariance matrices $\mathsf{Cov}(\hat{X}_\mathcal{T} | \boldsymbol{H}_\mathcal{T})$ and $\mathsf{Cov}(\hat{X}_\mathcal{T})$ in order to make fair accuracy assessments. Two popular and easy-to-compute choices are the trace $\mathsf{Tr}(\cdot)$ and the Frobenius norm $\| \cdot \|_F$. We will next deal with (7) and (8) separately.

A. Conditional Problem

It is known that, in general, with an unconstrained stopping time the sequential CRLB is not attainable under any kind of noise (Gaussian or non-Gaussian) except Bernoulli-distributed-noise [8]. We will next show that, with a stopping time $\mathcal{T}$ that $\{\mathcal{H}_t\}$-adapted, the OLS estimator attains the sequential CRLB, i.e., $\hat{X}_\mathcal{T}$ is the sequential MVUE, under Gaussian noise and it is also the sequential BLUE under non-Gaussian noise. Denote the sigma-algebra and the filtration generated by the coefficient vectors $H_1, \dots, H_t$ with $\mathcal{H}_t$ and $\{\mathcal{H}_t\}$, respectively. Similarly denote the sigma-algebra and the filtration generated by the sample $y_1, \dots, y_t$ with $\mathcal{F}_t$ and $\{\mathcal{F}_t\}$, respectively. Then, we are interested in $\{\mathcal{H}_t\}$-adapted stopping times. Note that an unconstrained stopping time could in general be $\{\mathcal{F}_t \cup \mathcal{H}_t\}$-adapted, for which unfortunately we know that there is no optimal sequential estimator.

Lemma 1. *Having a monotone accuracy function f and an $\{\mathcal{H}_t\}$-adapted stopping time $\mathcal{T}$ we can write, for the constraint in* (7)*,*

$$f\left(\mathsf{Cov}(\hat{X}_\mathcal{T} | \boldsymbol{H}_\mathcal{T})\right) \geq f\left(\sigma^2 \boldsymbol{U}_\mathcal{T}^{-1}\right) \tag{9}$$

for all unbiased estimators under Gaussian noise, and for all linear unbiased estimators under non-Gaussian noise. And the OLS estimator satisfies this inequality with equality.

Proof: In the previous section, the OLS estimator was shown to be MVUE under Gaussian noise and BLUE under non-Gaussian noise. It was also shown that $\mathsf{Cov}(\hat{X}_t | \boldsymbol{H}_t) = \sigma^2 \boldsymbol{U}_t^{-1}$. Hence, we write

$$\begin{aligned} &f\left(\mathsf{Cov}(\hat{X}_\mathcal{T} | \boldsymbol{H}_\mathcal{T})\right) \\ &= f\left(\mathsf{E}\left[\sum_{t=1}^{\infty} (\hat{X}_t - X)(\hat{X}_t - X)^T \ \mathbb{1}_{\{t=\mathcal{T}\}} \Big| \boldsymbol{H}_t\right]\right) \\ &= f\left(\sum_{t=1}^{\infty} \mathsf{E}\left[(\hat{X}_t - X)(\hat{X}_t - X)^T \Big| \boldsymbol{H}_t\right] \ \mathbb{1}_{\{t=\mathcal{T}\}}\right) &(10) \\ &\geq f\left(\sum_{t=1}^{\infty} \sigma^2 \boldsymbol{U}_t^{-1} \ \mathbb{1}_{\{t=\mathcal{T}\}}\right) &(11) \\ &= f\left(\sigma^2 \boldsymbol{U}_\mathcal{T}^{-1}\right), &(12) \end{aligned}$$

for all unbiased estimators under Gaussian noise and for all linear unbiased estimators under non-Gaussian noise. We used the fact that the event $\{\mathcal{T} = t\}$ is $\mathcal{H}_t$-measurable and $\mathsf{E}[(\hat{X}_t - X)(\hat{X}_t - X)^T | \boldsymbol{H}_t] = \mathsf{Cov}(\hat{X}_t | \boldsymbol{H}_t) \geq \sigma^2 \boldsymbol{U}_t^{-1}$ to write (10) and (11), respectively. ■

Since $\mathcal{T}$ is $\{\mathcal{H}_t\}$-adapted, we have $\mathsf{E}[\mathcal{T} | \boldsymbol{H}_\mathcal{T}] = \mathcal{T}$, and thus from (7) we want to find the first time that a member of our class of estimators (i.e., unbiased estimators under Gaussian noise and linear unbiased estimators under non-Gaussian noise) satisfies the constraint $f\left(\mathsf{Cov}(\hat{X}_\mathcal{T} | \boldsymbol{H}_\mathcal{T})\right) \leq C$, and also the estimator that attains this earliest stopping time. From Lemma 1 it is seen that the OLS estimator achieves the earliest stopping time among its competitors. Hence, for the conditional problem the optimal pair of stopping time and estimator is $(\mathcal{T}, \hat{X}_\mathcal{T})$ where $\mathcal{T}$ is given by

$$\mathcal{T} = \min\{t \in \mathbb{N} : f\left(\sigma^2 \boldsymbol{U}_t^{-1}\right) \leq C\}, \tag{13}$$

and from (3), $\hat{X}_\mathcal{T} = \boldsymbol{U}_\mathcal{T}^{-1} \boldsymbol{H}_\mathcal{T}^T Y_\mathcal{T}$, which can be computed recursively as in (6). The recursive computation of $\boldsymbol{U}_t^{-1}$ in the test statistic in (13) is also given in (6). Note that for an accuracy function f such that $f(\sigma^2 \boldsymbol{U}_t^{-1}) = \sigma^2 f(\boldsymbol{U}_t^{-1})$, e.g., $\mathsf{Tr}(\cdot)$ and $\| \cdot \|_F$, we can use the following stopping time,

$$\mathcal{T} = \min\{t \in \mathbb{N} : f\left(\boldsymbol{U}_t^{-1}\right) \leq C'\}, \tag{14}$$

where $C' = C/\sigma^2$ is the relative target accuracy with respect to the noise power. Hence, given C' we do not need to know the noise variance σ^2 to run the test given by (14).

Note that $\boldsymbol{U}_t$, being the summation of covariance matrices up to time t, is a non-decreasing positive semidefinite matrix, and thus, from the monotonicity of f, the test statistic $f\left(\sigma^2 \boldsymbol{U}_t^{-1}\right)$ is a non-increasing scalar function of time. Specifically, for accuracy functions $\mathsf{Tr}(\cdot)$ and $\|\cdot\|_F$ we can show that if the minimum eigenvalue of $\boldsymbol{U}_t$ tends to infinity as $t \to \infty$, then the stopping time is finite, i.e., $\mathcal{T} < \infty$.

For the special case of scalar parameter estimation, we do not need a function f to assess the accuracy of the estimator since instead of a covariance matrix we now have a variance $\frac{\sigma^2}{u_t}$, where $u_t = \sum_{n=1}^t h_n^2$ and h_t is the scaling coefficient in (1). Hence, from (14) the stopping time in the scalar case is given by

$$\mathcal{T} = \min\left\{t \in \mathbb{N} : u_t \geq \frac{1}{C'}\right\}, \tag{15}$$

where $\frac{u_t}{\sigma^2}$ is the Fisher information at time t. This result is in accordance with [9, Eq. (3)].

B. *Unconditional Problem*

In this case we assume $\{H_t\}$ is i.i.d.. From the constrained optimization problem in (8), using a Lagrange multiplier λ we obtain the following unconstrained optimization problem,

$$\min_{\mathcal{T}, \hat{X}_\mathcal{T}} \mathsf{E}[\mathcal{T}] + \lambda f\left(\mathsf{Cov}(\hat{X}_\mathcal{T})\right). \tag{16}$$

We are again interested in $\{\mathcal{H}_t\}$-adapted stopping times to use the optimality property of the OLS estimator in the sequential sense. For the sake of simplicity assume a linear accuracy function f so that $f(\mathsf{E}[\cdot]) = \mathsf{E}[f(\cdot)]$, e.g., the trace function $\mathsf{Tr}(\cdot)$, which is also monotone. Then, our constraint function turns out to be the sum of the individual variances, i.e., $\mathsf{Tr}\left(\mathsf{Cov}(\hat{X}_\mathcal{T})\right) = \sum_{i=1}^n \mathsf{Var}(\hat{x}_\mathcal{T}^i)$. Since $\mathsf{Tr}\left(\mathsf{Cov}(\hat{X}_\mathcal{T})\right) = \mathsf{Tr}\left(\mathsf{E}\left[\mathsf{Cov}(\hat{X}_\mathcal{T}|\boldsymbol{H}_\mathcal{T})\right]\right) = \mathsf{E}\left[\mathsf{Tr}\left(\mathsf{Cov}(\hat{X}_\mathcal{T}|\boldsymbol{H}_\mathcal{T})\right)\right]$, we rewrite (16) as

$$\min_{\mathcal{T}, \hat{X}_\mathcal{T}} \mathsf{E}\left[\mathcal{T} + \lambda \mathsf{Tr}\left(\mathsf{Cov}(\hat{X}_\mathcal{T}|\boldsymbol{H}_\mathcal{T})\right)\right], \tag{17}$$

where expectation is with respect to $\boldsymbol{H}_\mathcal{T}$.

From Lemma 1, we have $\mathsf{Tr}\left(\mathsf{Cov}(\hat{X}_\mathcal{T}|\boldsymbol{H}_\mathcal{T})\right) \geq \mathsf{Tr}\left(\sigma^2 \boldsymbol{U}_\mathcal{T}^{-1}\right)$ where $\sigma^2 \boldsymbol{U}_t^{-1}$ is the covariance matrix of the OLS estimator at time t. Note that $\boldsymbol{U}_t/\sigma^2$ is the Fisher information matrix at time t [cf. (5)]. Using the OLS estimator we minimize the objective function in (17), hence $\hat{X}_\mathcal{T} = \boldsymbol{U}_\mathcal{T}^{-1} \boldsymbol{H}_\mathcal{T}^T Y_\mathcal{T}$ [cf. (6) for recursive computation] is the optimal estimator also in the unconditional problem.

Now, to find the optimal stopping time we need to solve the following optimization problem,

$$\min_{\mathcal{T}} \mathsf{E}\left[\mathcal{T} + \lambda \mathsf{Tr}\left(\sigma^2 \boldsymbol{U}_\mathcal{T}^{-1}\right)\right], \tag{18}$$

which can be solved by using *optimal stopping theory*. Writing (18) in the following alternative form

$$\min_{\mathcal{T}} \mathsf{E}\left[\sum_{t=0}^{\mathcal{T}-1} 1 + \lambda \mathsf{Tr}\left(\sigma^2 \boldsymbol{U}_\mathcal{T}^{-1}\right)\right], \tag{19}$$

we see that the term $\sum_{t=0}^{\mathcal{T}-1} 1$ accounts for the cost of not stopping until time $\mathcal{T}$ and the term $\lambda \mathsf{Tr}\left(\sigma^2 \boldsymbol{U}_\mathcal{T}^{-1}\right)$ represents the cost of stopping at time $\mathcal{T}$. Note that $\boldsymbol{U}_t = \boldsymbol{U}_{t-1} + H_t H_t^T$ and given $\boldsymbol{U}_{t-1}$ the current state $\boldsymbol{U}_t$ is (conditionally) independent of all previous states, hence $\{\boldsymbol{U}_t\}$ is a Markov process. That is, optimal stopping time for a Markov process is sought in (19). From [11] the solution is given by

$$V(\boldsymbol{U}) = \min\{\lambda \mathsf{Tr}\left(\sigma^2 \boldsymbol{U}^{-1}\right), 1 + \mathsf{E}[V(\boldsymbol{U} + H_1 H_1^T)|\boldsymbol{U}]\}, \tag{20}$$

where expectation is with respect to H_1 and V is the optimal cost function. The optimal cost function is found by iterating a sequence of functions $\{V_n\}$ where $V(\boldsymbol{U}) = \lim_{n\to\infty} V_n(\boldsymbol{U})$ and

$$\begin{aligned} &V_n(\boldsymbol{U}) \\ &= \min\left\{\lambda \mathsf{Tr}\left(\sigma^2 \boldsymbol{U}^{-1}\right), 1 + \mathsf{E}[V_{n-1}(\boldsymbol{U} + H_1 H_1^T)|\boldsymbol{U}]\right\}. \end{aligned} \tag{21}$$

In optimal stopping theory, the original complex optimization problem in (18) is divided into simpler subproblems given by (20). At each time t we are faced with a subproblem consisting of a stopping cost $F(\boldsymbol{U}_t) = \lambda \mathsf{Tr}\left(\sigma^2 \boldsymbol{U}_t^{-1}\right)$ and an expected sampling cost $G(\boldsymbol{U}_t) = 1 + \mathsf{E}[V(\boldsymbol{U}_{t+1})|\boldsymbol{U}_t]$ to proceed to time $t+1$. The optimal cost function $V(\boldsymbol{U}_t)$, selecting the action with minimum cost (i.e., either continue or stop), determines the optimal policy to follow at each time t. Specifically, the optimal policy, as we will show later in this section, chooses to continue as long as $V(\boldsymbol{U}_t) = G(\boldsymbol{U}_t)$ and stops the first time $V(\boldsymbol{U}_t) = F(\boldsymbol{U}_t)$. We need to analyze the structure of $V(\boldsymbol{U}_t)$, i.e., the cost functions $F(\boldsymbol{U}_t)$ and $G(\boldsymbol{U}_t)$, to show such a behavior for the optimal policy and find the optimal stopping time $\mathcal{T}$.

Note that V, being a function of the symmetric matrix $\boldsymbol{U} \in \mathbb{R}^{n\times n}$, is a function of $\frac{n^2+n}{2}$ variables $\{u_{ij} : i \leq j\}$ where $\boldsymbol{U} = [u_{ij}]$. Analyzing a multi-dimensional optimal cost function proves intractable, hence we will analyze the special case of scalar parameter estimation. Some numerical results for the two-dimensional vector case, which demonstrate how intractable the higher dimensional problems are, can be found in [10].

For the scalar case, from (20) we have the following one-dimensional optimal cost function,

$$V(u) = \min\left\{\frac{\lambda\sigma^2}{u}, 1 + \mathsf{E}[V(u + h_1^2)]\right\}, \tag{22}$$

where expectation is with respect to h_1 and h_1 is a scalar coefficient, scaling the parameter x to be estimated [cf. (1)]. Write V as a function of $w \triangleq 1/u$,

$$V(w) = \min\left\{\lambda\sigma^2 w, 1 + \mathsf{E}\left[V\left(\frac{w}{1 + w h_1^2}\right)\right]\right\}, \tag{23}$$

where as before expectation is with respect to h_1. We need to analyze the cost functions $F(w) = \lambda\sigma^2 w$ and $G(w) = 1 + \mathsf{E}\left[V\left(\frac{w}{1+wh_1^2}\right)\right]$. The former is a line, whereas the latter is in general a nonlinear function of w. We have the following theorem regarding the structure of $V(w)$ and $G(w)$. Its proof

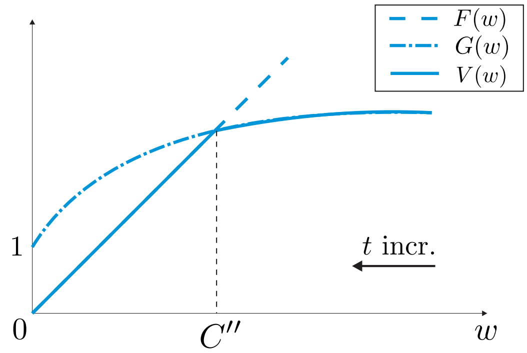

Fig. 1. The structures of the optimal cost function $V(w)$ and the cost functions $F(w)$ and $G(w)$.

is presented in [10, Appendix].

Theorem 1. *The optimal cost V and the expected sampling cost G, given in* (23)*, are non-decreasing, concave and bounded functions of w.*

The cost functions $F(w)$ and $G(w)$ are continuous functions as F is linear and G is concave. From (23) we have $V(0) = \min\{0, 1 + V(0)\} = 0$, hence $G(0) = 1 + V(0) = 1$. Then, using Theorem 1 we show $F(w)$ and $G(w)$ in Fig. 1. The optimal cost function $V(w)$, being the minimum of F and G [cf. (23)], is also shown in Fig. 1, justifying Theorem 1. Note that as t increases w tends from infinity to zero. Hence, we continue until the stopping cost $F(w_t)$ is lower than the expected sampling cost $G(w_t)$, i.e., until $w_t \leq C''$. In other words, the stopping time in the scalar case of the unconditional problem is given by

$$\mathcal{T} = \min\left\{t \in \mathbb{N} : u_t \geq \frac{1}{C''}\right\}, \tag{24}$$

similar to the scalar case of the conditional problem [cf. (15)]. Note that the threshold C'' is determined by the Lagrange multiplier λ, which is selected so that $\mathsf{E}\left[\frac{\sigma^2}{u_\mathcal{T}}\right] = C$, i.e., the variance of the estimator exactly hits the target accuracy level C, [cf. (16)]. Accordingly, we have $C'' \geq C/\sigma^2 = C'$ since the upper bound $\sigma^2 C''$ on the conditional variance $\sigma^2 w_\mathcal{T}$ [cf. (24)] is also an upper bound for the variance $\mathsf{E}[\sigma^2 w_\mathcal{T}] = C$. This result implies that the stopping time of the unconditional problem will be smaller than that of the conditional problem.

IV. Decentralized Implementation

In this section, we will develop asymptotically optimal decentralized sequential estimators for the scalar case of the conditional problem. Consider the problem of estimating a non-random parameter, $x \in \mathbb{R}$, at a central unit, i.e., the fusion center (FC), via noisy observations collected at K distributed nodes, i.e., sensors. Let y_t^k, $t \in \mathbb{N}$, $k = 1, \ldots, K$, denote the discrete-time noisy sample observed by the k-th sensor at time t, given by

$$y_t^k = x h_t^k + w_t^k, \tag{25}$$

where x is the constant parameter to be estimated, $h_t^k \in \mathbb{R}$ is the random channel gain and observed by the k-th sensor, and $w_t^k \sim \mathcal{N}(0, \sigma_k^2)$ is the Gaussian noise assumed to be independent and identically distributed (i.i.d.) across time and independent but not necessarily identically distributed across sensors. Accordingly, given h_t^k we have $y_t^k \sim \mathcal{N}(x h_t^k, \sigma_k^2)$, i.e., y_t^k is conditionally Gaussian. Note that random h_t^k corresponds to the fading channels. If sensors transmit their observations in whole by using infinite number of bits, then the FC will have access to all local observations $\{y_t^k\}_{t,k}$ [1], which corresponds to the conventional *centralized* estimation problem, as discussed in the previous section. However, in practice, due to power and bandwith constraints, sensors typically sample their observations and transmit only a few bits per sample to the FC. In such *decentralized* setup, the FC can only obtain a summary of local observations based on which it performs estimation.

In the scalar case of the conditional problem, optimal stopping time is given in (15) and from (3) we write the optimal estimator as $\hat{x}_\mathcal{T} = \frac{v_\mathcal{T}}{u_\mathcal{T}}$ where $u_t = \sum_{k=1}^K u_t^k = \sum_{k=1}^K \sum_{n=1}^t (h_n^k)^2$ and $v_t = \sum_{k=1}^K v_t^k = \sum_{k=1}^K \sum_{n=1}^t h_n^k y_n^k$ due to the independence among sensors. Each sensor k computes (updates) its local processes u_t^k and v_t^k after observing y_t^k and h_t^k at time t. However, the FC, which determines the stopping time and computes the estimator, has no access to the local processes. Hence, sensors should report both $\{u_t^k\}_k$ and $\{v_t^k\}_k$ to the FC.

Imitating the optimal centralized scheme we propose a decentralized scheme $(\tilde{\mathcal{T}}, \tilde{x}_{\tilde{\mathcal{T}}})$ based on level-triggered sampling. We propose that each sensor k, via level-triggered sampling, informs the FC whenever considerable change occurs in its local processes u_t^k and v_t^k. The level-triggered sampling is a simple form of event-triggered sampling, in which sampling (communication) times $\{t_{n,u}^k\}_n$ and $\{t_{n,v}^k\}_n$ [2] are not deterministic, but rather dynamically determined by the random processes u_t^k and v_t^k, respectively, i.e.,

$$t_{n,u}^k \triangleq \min\{t > t_{n-1,u}^k : u_t^k - u_{t_{n-1,u}^k}^k \geq e_k\}, \tag{26}$$

$$t_{n,v}^k \triangleq \min\{t > t_{n-1,v}^k : v_t^k - v_{t_{n-1,v}^k}^k \notin (-d_k, d_k)\}, \tag{27}$$

where $n \in \mathbb{N}$, $t_{0,u}^k = 0$, $t_{0,v}^k = 0$. The threshold parameters d_k and e_k are constants known by both sensor k and the FC. Note that in (26) we use a single threshold different from (27) since $u_t^k = \sum_{n=1}^t (h_n^k)^2$ is a nondecreasing process.

At each sampling time $t_{n,v}^k$, sensor k transmits r_v bits, $b_{n,1}^k b_{n,2}^k \ldots b_{n,r_v}^k$, to the FC. The first bit, $b_{n,1}^k$, indicates the threshold crossed (either d_k or $-d_k$) by the incremental process $\delta_n^k \triangleq v_{t_{n,v}^k}^k - v_{t_{n-1,v}^k}^k$, i.e.,

$$b_{n,1}^k = \text{sign}(\delta_n^k). \tag{28}$$

The remaining $(r_v - 1)$ bits, $b_{n,2}^k \ldots b_{n,r_v}^k$, are used to quantize the over(under)shoot $q_n^k \triangleq |\delta_n^k| - d_k$ into $\tilde{q}_n^k$. We assume that the parameter to be estimated is bounded, i.e., $|x| < \mathcal{X}$, and so does the term $h_t^k y_t^k$, i.e., $|h_t^k y_t^k| < \phi < \infty$, $\forall k, t$. At each sampling time $t_{n,v}^k$, the overshoot value q_n^k cannot exceed

[1] The subscripts t and k in the set notation denote $t \in \mathbb{N}$ and $k = 1, \ldots, K$, respectively.

[2] The subscript n in the set notation denotes $n \in \mathbb{N}$.

the magnitude of the last sample in the incremental process $\delta_n^k = \sum_{n=t_{n-1,v}^k+1}^{t_{n,v}^k} h_n^k y_n^k$, i.e., $0 \le q_n^k < \phi$. Hence, the interval $[0,\phi)$ is uniformly divided into 2^{r_v-1} subintervals. The FC, upon receiving the bits $b_{n,1}^k b_{n,2}^k \dots b_{n,r_v}^k$ from the sensor k at time $t_{n,v}^k$, recovers the quantized value of δ_n^k by computing

$$\tilde{\delta}_n^k \triangleq b_{n,1}^k (d_k + \tilde{q}_n^k). \tag{29}$$

Then, it sequentially sums up $\{\tilde{\delta}_n^k\}_{n,k}$, at the sampling (communication) times $\{t_{n,v}^k\}_{n,k}$ to obtain an approximation $\tilde{v}_t$ to v_t, i.e.,

$$\tilde{v}_t \triangleq \sum_{k=1}^{K} \sum_{n=1}^{N_t^k} \tilde{\delta}_n^k = \sum_{k=1}^{K} \tilde{v}_t^k, \tag{30}$$

where N_t^k is the number of messages that the FC receives from the sensor k about v_t^k up to time t. During the times the FC receives no message, i.e., $t \notin \{t_{n,v}^k\}_{n,k}$, $\tilde{v}_t$ is kept constant.

At each sampling time $t_{n,u}^k$, sensor k transmits r_u bits to the FC, all of which are used to quantize the overshoot $p_n^k \triangleq \eta_n^k - e_k$ into $\tilde{p}_n^k$, where we defined the incremental process $\eta_n^k \triangleq u_{t_{n,u}^k}^k - u_{t_{n-1,u}^k}^k$. In this case, we do not need to allocate a sign bit. Assume $(h_t^k)^2 < \theta < \infty$, $\forall k, t$, hence we have $0 \le p_n^k < \theta$, and the interval $[0,\theta)$ is uniformly partitioned into 2^{r_u} subintervals. In other words, each sensor k determines the index of p_n^k, and then transmits it to the FC using r_u bits. The FC, upon receiving the r_u bits at time $t_{n,u}^k$, similar to (29) computes

$$\tilde{\eta}_n^k \triangleq e_k + \tilde{p}_n^k. \tag{31}$$

Then, similar to (30) it also computes

$$\tilde{u}_t \triangleq \sum_{k=1}^{K} \sum_{n=1}^{M_t^k} \tilde{\eta}_n^k = \sum_{k=1}^{K} \tilde{u}_t^k, \tag{32}$$

where M_t^k is the number of messages that the FC receives from sensor k about u_t^k up to time t.

The scheme is terminated at the stopping time, $\tilde{\mathcal{T}}$ [cf. (15)], given by

$$\tilde{\mathcal{T}} = \min\left\{ t \in \mathbb{N} : \tilde{u}_t \ge \frac{1}{C'} \right\}, \tag{33}$$

and the following estimator

$$\tilde{x}_{\tilde{\mathcal{T}}} = \frac{\tilde{v}_{\tilde{\mathcal{T}}}}{\tilde{u}_{\tilde{\mathcal{T}}}}, \tag{34}$$

is computed. The following theorem presents the sufficient conditions under which the proposed estimator, based on level-triggered sampling, is asymptotically unbiased, consistent and asymptotically optimal. Its proof can be found in [9, Appendices J & K]. In the theorem, $\xrightarrow{p}$ and $\xrightarrow{d}$ denote convergence in probability and convergence in distribution, respectively, and $u_{\tilde{\mathcal{T}}}/\sigma^2$ is the Fisher information at the stopping time.

Theorem 2. *Consider the decentralized sequential estimator given in* (33) *and* (34)*. As $C' \to 0$, the estimator is asymptotically unbiased, i.e., $\mathsf{E}[\tilde{x}_{\tilde{\mathcal{T}}} - x | \boldsymbol{H}_{\tilde{\mathcal{T}}}] \to 0$, and consistent, that is, $\tilde{x}_{\tilde{\mathcal{T}}} \xrightarrow{p} x$, if $d_k \to \infty$, and $e_k \to \infty$ at slower rates than $1/C'$, i.e., $d_k = o(1/C')$ and $e_k = o(1/C')$, $\forall k$. Moreover, as $C' \to 0$, it is asymptotically optimal, i.e., $\sqrt{u_{\tilde{\mathcal{T}}}/\sigma^2}(\tilde{x}_{\tilde{\mathcal{T}}} - x) \xrightarrow{d} \mathcal{N}(0,1)$, if $d_k = o(\sqrt{1/C'})$, $r_v = \omega(\log(\sqrt{1/C'}/d_k))$, $e_k = o(\sqrt{1/C'})$, and $r_u = \omega(\log(\sqrt{1/C'}/e_k))$, $\forall k$.*

From Theorem 2 we see that the proposed decentralized estimator asymptotically unbiased and consistent with appropriate thresholds and constant number of bits r_u and r_v. Note that it is desired to have the thresholds $d_k \to \infty$ and $e_k \to \infty$ as fast as possible to attain asymptotically low communication rates. It is also desired to have the number of bits r_u and r_v as small as possible to lower the bandwith consumption. For asymptotic optimality, even though the rates of r_u and r_u are lower bounded by $\log(\sqrt{1/C'}/e_k)$ and $\log(\sqrt{1/C'}/d_k)$, respectively, in practice they can be very slow, i.e., close to zero, if d_k and e_k tend to infinity as fast as possible, i.e., close to $\sqrt{1/C'}$, as desired in practice.

In [9], the proposed estimator, based on level-triggered sampling, was shown, both analytically and numerically, to outperform the decentralized sequential estimator based on the conventional uniform-in-time sampling. In the uniform-sampling-based-estimator, the sampling times $\{t_{n,u}^k\}_n$ and $\{t_{n,v}^k\}_n$ are, as opposed to (26) and (27), deterministic with periods T_u and T_v, respectively, and the incremental processes δ_n^k and η_n^k are quantized using r_u and r_v bits, respectively.

V. Conclusion

Optimal sequential estimators were derived for conditional and unconditional formulations of the problem. A decentralized sequential estimator based on level-triggered sampling was proposed. Sufficient conditions for asymptotic optimality, asymptotically unbiasedness and consistency were presented.

References

[1] A. Wald, *Sequential Analysis*, Wiley, New York, NY, 1947.

[2] M. Ghosh, N. Mukhopadhyay, and P.K. Sen, *Sequential Estimation*, Wiley, New York, NY, 1997.

[3] H.C. Papadopoulos, G.W. Wornell, and A.V. Oppenheim, "Sequential Signal Encoding from Noisy Measurements Using Quantizers with Dynamic Bias Control," *IEEE Trans. Inform. Theory*, vol. 47, no. 3, pp. 978-1002, Mar. 2001.

[4] T. Zhao, and A. Nehorai, "Distributed Sequential Bayesian Estimation of a Diffusive Source in Wireless Sensor Networks," *IEEE Trans. Sig. Proc.*, vol. 55, no. 4, pp. 1511-1524, Apr. 2007.

[5] G. Fellouris, "Asymptotically optimal parameter estimation under communication constraints," *Ann. Statist.*, to be published. *http://arxiv.org/abs/1102.0369*, Jul. 2012.

[6] G. Fellouris, and G.V. Moustakides, "Decentralized sequential hypothesis testing using asynchronous communication," *IEEE Trans. Inform. Theory*, vol. 57, no. 1, pp. 534-548, Jan. 2011.

[7] Y. Yilmaz, G.V. Moustakides, and X. Wang, "Cooperative sequential spectrum sensing based on level-triggered sampling," *IEEE Trans. Sig. Proc.*, vol. 60, no. 9, pp. 4509-4524, Sep. 2012.

[8] B.K. Ghosh, "On the attainment of the Cramer-Rao bound in the sequential case," *Sequential Analysis*, vol. 6, no. 3, pp. 267–288, 1987.

[9] Y. Yilmaz, and X. Wang, "Sequential Decentralized Parameter Estimation under Randomly Observed Fisher Information," *http://arxiv.org/abs/1211.3720*, Nov. 2012.

[10] Y. Yilmaz, G.V. Moustakides, and X. Wang, "Optimal Sequential Vector Estimation," *http://arxiv.org/abs/1301.5701*, Jan. 2013.

[11] A.N. Shiryaev, *Optimal Stopping Rules*, Springer, New York, 2008.