Sequential Estimation of Distributed Parameters in Networks

Saurabh Sihag, Javad Heydari and Ali Tajer
Electrical, Computer, and Systems Engineering Department
Rensselaer Polytechnic Institute

***Abstract*—The problem of estimating a set of unknown parameters in a multi-agent network is considered. Each agent can make noisy observations from a subset of the unknown parameters, and different agents can potentially observe common parameters. The objective of each agent is to estimate its observed unknown parameters. This paper focuses on sequentially estimating the parameters such that, in the quickest fashion, all the agents form reliable estimates for their designated parameters. A proper estimation cost function is adopted in order to signify the fidelity of the estimates to the ground truth, and to ensure consistency in the estimates of different agents. By imposing practical constraints on the number of data points that the network affords to process, the sequential strategy dynamically decides about the minimum number of measurements required to form reliable estimates, the agents from which these measurements should be collected, and the optimal estimators for each agent. Specifically, a sequential strategy is proposed, which consists of the stopping rule of the sampling process, a data-adaptive control policy for selecting the agents over time, and a set of estimators, combination of which admits asymptotic optimality.**

I. INTRODUCTION

Due to the technological advances in information generation, sensing, and communication, many application domains have evolved toward large-scale networks of interconnected agents. Network of measurement devices in a power grid, array of cameras in a surveillance system, and communities in online social networks are a few examples. The data generated by such networks is often processed for various inferential purposes. Due to their large scale as well as the varying degrees of connectivity and interaction among the agents, the amount of information provided by various agents for a desired inference goal is volatile. For instance, in a power grid, during a fault event the data from the locality of the event is more informative about the type and location of the fault, or in a large-scale surveillance system cameras closer to an external object provide more informative data. Therefore, it is of interest to quickly focus the sensing resources to the areas in the network that are most informative about the inference goal of interest.

This paper considers the problem of *sequentially* estimating a set of unknown parameters in a multi-agent network. It assumes that each agent can observe only a subset of the unknown parameters, and different agents can potentially observe common parameters. The ultimate goal is i) to provide each agent with reliable estimates of the unknown parameters it observes, and ii) the estimates of the same parameter obtained by different agents are consistent. Forming a centralized optimal estimate by using all the measurements collected in the network is computationally prohibitive. On the other hand, forming local estimates at each agent solely based on their own local observations, while being computationally efficient, is agnostic to the structure of the network, and subsequently, inherently suboptimal [1]. Therefore, in order to strike a balance between these two extremes, in this paper we assume that agents are capable of forming estimates based on their own data and sharing their information among themselves. In this sense, the setting of this paper is related to distributed estimation, which is often studied under diffusion or consensus setting. In consensus based estimation algorithms, the agents exchange information among themselves to converge to a state estimate for the systems [2] and [3], while in the diffusion approach the agents interact with each other on a local level and diffuse information across the network to estimate the unknown parameters in a distributed manner [4] and [5]. The distinction of this paper from these two approaches is that there exists no constraint on the communication of agents, but rather the structure of the estimation cost determines how the information should be communicated among the agents.

Estimating unknown parameters observed by a network arises in a number of application domains such as biochemistry [6], systems biology [7] and [8], social science [9], and power grid [10]. The existing literature on parameter estimation often focuses on settings in which the estimation strategy and data-acquisition processes are decoupled, and their focus is placed on forming the most reliable estimates based on a given set of measurements [11]–[13]. Driven by controlling communication, sensing, and decision delay costs, this paper considers a sequential sampling strategy, which is specified by the required number of observations, as well as the order in which they are collected. When the order is pre-specified, determining the optimal sampling strategy reduces to minimizing the number of measurements by dynamically deciding whether to take more measurements, or to terminate the process and form the estimates [14]–[16]. However, incorporating dynamic decisions about the order of sampling, especially for inference over networks, is less-investigated. Forming such dynamic decisions that pertain to data acquisition for inference naturally arises in a broad range of applications such as sensor management [17], inspection and classification [18], medical diagnosis [19], cognitive science [20], generalized binary search [21], and channel coding with feedback [22]. One directly applicable approach to treat such coupled sampling and decision-making problem is controlled sensing, in which besides deciding when to

This research was supported in part by the U. S. National Science Foundation under the CAREER Award ECCS-1554482 and grant DMS-1737976.

terminate the sampling process, the sequence of agents whose data are collected over time should also be specified. The study in [23] considers the setting in which the parameters observed by different agents are unique, and provide asymptotically optimal decision rules. In [24] the results are generalized to the setting in which observations depend on common unknown parameters. The problem considered in this paper is closely related to [24] where, at each time instant and based on the collected information up to that time, one of a finite number of agents are selected. However, the presence of multiple estimators, each one of which requiring to make reliable estimates of the unknown parameters, makes the problem of this paper fundamentally different from [24]. Furthermore, the agent selection rules in this paper are fundamentally co-dependent, while in [24] they are independent.

Controlled inference has also been studied in domains other than sequential estimation [25]–[29]. The studies in [25]–[28] focus on sequential hypothesis testing, with controlled actions, while [29] adopts a controlled sensing based approach to graph classification. In another direction, the problem of sequential joint detection and estimation without any controlled action is studied in [30] and [31]. The major distinction of this paper from [30] and [31] is that, besides having a detection action upon stopping, in these studies the observation process is fixed and the only dynamic of the sampling process is the stopping time. Furthermore, there exist two sequences to observe and both of them are being observed at each time, while in this paper multiple sequences are available and only one of them is observed at each time instance.

II. Model and Formulation

A. Data Model

Consider a network of K agents indexed by $\mathcal{U} \triangleq \{1,\dots,K\}$, collectively making noisy observations from a set of unknown and random parameters denoted by $\mathcal{X} \triangleq \{X_i : i \in \mathcal{M}\}$, where $\mathcal{M} \triangleq \{1,\dots,m\}$. These parameters are assumed to be statistically independent. Due to the large scale of the network, each agent can observe only a subset of the unknown parameters $\mathcal{X}$. We denote the set of parameters observed by agent $i \in \mathcal{M}$ by $\mathcal{X}_i \subseteq \mathcal{X}$. Furthermore, for any $i \in \mathcal{M}$ the set of the agents that can observe X_i is denoted by $\mathcal{S}_i$. We consider a fully sequential sampling process, in which only one agent is selected at-a-time to make a noisy observation. By defining $u(t) \in \mathcal{U}$ as the index of the agent that makes the measurement at time $t \in \mathbb{N}$, the measurement of agent $i \in \mathcal{U}$ at time t is given by

$$Y_t^i = \begin{cases} g_i(\mathcal{X}_i) + N_t^i & \text{if } u(t) = i \\ \emptyset & \text{if } u(t) \neq i \end{cases} , \tag{1}$$

where the convention $\emptyset$ denotes lack of a measurement, g_i is a known function that captures the measurement model of agent i, and N_t^i accounts for the measurement noise, which for each agent i is independent and identically distributed over time. Accordingly, we define $\mathcal{Y}_t^i \triangleq \{Y_1^i,\dots,Y_t^i\}$ as the ordered set of measurements collected by agent $i \in \mathcal{U}$ up to time t. The prior probability distribution function (pdf) of the scalar parameter X_i for $i \in \mathcal{M}$ is denoted by π_i, and when $Y_t^i \neq \emptyset$, we denote the pdf of Y_t^i by f_i.

The information collected sequentially up to time t generates the filtration $\{\mathcal{F}_t : t \in \mathbb{N}\}$ where

$$\mathcal{F}_t \triangleq \sigma\big(Y_1^{u(1)}, Y_2^{u(2)}, \dots, Y_t^{u(t)}\big) . \tag{2}$$

In the sequentially sampling process, at any time t and based on the information accumulated up to that time, i.e., $\mathcal{F}_t$, the sensing mechanism takes one of the following two intertwined actions.

A$_1$) *Observation:* Due to lack of sufficient information, forming the final estimates is deferred and one more measurement is taken from one of the agents. Under this action, the agent to be selected is specified.

A$_2$) *Estimation:* Sensing process is terminated and each agent forms final estimates of its associated unknown parameters. Under this action, the stopping time and the final estimation rules upon stopping are specified.

To formalize the *observation* action, we define a data-adaptive control policy denoted by μ. Specifically, at time t, the control policy μ leverages all the past measurements, denoted by $\widetilde{\mathcal{Y}}^{t-1} \triangleq \{\mathcal{Y}_{t-1}^1,\dots,\mathcal{Y}_{t-1}^K\}$, and all the past agent selection decisions, denoted by $\mathcal{U}^{t-1} \triangleq \{u(1),\dots u(t-1)\}$, to determine $u(t)$, i.e.,

$$\mu : \mathcal{U}^{t-1} \times \widetilde{\mathcal{Y}}^{t-1} \rightarrow \mathcal{U} . \tag{3}$$

In order to formalize the *estimation* action, we define $T \in \mathbb{N}$ as the stochastic stopping time of the process, at which the information gathering process is terminated and the estimates are formed. Each agent $i \in \mathcal{U}$ needs to form an estimate of the set $\mathcal{X}_i$, the set of unknown parameters that it can observe through Y_t^i. For this purpose, we define

$$\delta_i^j : \mathcal{U}^t \times \widetilde{\mathcal{Y}}^t \rightarrow \mathbb{R} \tag{4}$$

as the estimator of agent $i \in \mathcal{U}$ for unknown parameter $X_j \in \mathcal{X}_i$, and denote the estimate of X_j by agent i at time t by $\hat{X}_j^i(t)$. Finally, we define $\Delta \triangleq \{\delta_i^j : i \in \mathcal{U} \ , \ j \in \mathcal{M}\}$ as the set of estimation rules in the network.

B. Problem Formulation

We aim to design a data acquisition policy μ, the set of estimators Δ, and the stopping time T. Designing the optimal strategy for achieving reliable estimates involves striking a balance between two opposing figures of merit pertinent to the quality of the estimates on the one hand, and the aggregate number of measurements made, on the other hand. The quality of the process is captured by the reliability of the estimates at individual agents. Also, we assume that each agent desires to have its estimates remain consistent with and in close proximity of those of the other agents that observe the same unknown parameters. To accommodate this, the cost function for quantifying the estimation quality of the network at any time t is defined according to

$$J(t,\mu,\Delta) \triangleq \sum_{i\in\mathcal{U}}\sum_{j\in\mathcal{I}_i}\mathbb{E}[(X_j-\hat{X}^i_j(t))^2 \mid \mathcal{Y}^i_t,\,\mathcal{U}^t] + \frac{1}{2}\sum_{j\in\mathcal{M}}\sum_{i\in\mathcal{S}_j}\sum_{k\in\mathcal{S}_j}|\hat{X}^i_j(t)-\hat{X}^k_j(t)|^2\ , \quad (5)$$

where $\mathcal{I}_i$ denotes the set of the indices of the unknown parameters that are observed by agent i, i.e., $\mathcal{I}_i \triangleq \{j : X_j \in \mathcal{X}_i\}$ In this estimation cost function, the term $\mathbb{E}[(X_j-\hat{X}^i_j(t))^2|\mathcal{Y}^i_t,\,\mathcal{U}^t]$ ensures that each agent $i\in\mathcal{U}$ forms an estimate of each $X_j\in\mathcal{X}_i$ with high fidelity, and $|\hat{X}^i_j(t)-\hat{X}^k_j(t)|^2$ guarantees that the estimates of X_j made by all the agents observing X_j stay consistent. By assuming that the cost of each measurement is a constant $c>0$, we integrate both estimation and the sampling costs into a unified cost function. Specifically, at time t, we denote the unified cost function by

$$\bar{J}(t,\mu,\Delta,c) \triangleq J(t,\mu,\Delta)+tc\ , \quad (6)$$

where c controls the balance between the quality of the estimates and the sensing cost.

III. Sequential Decision Rules

In this section, we provide a stopping rule, a design for the estimators Δ, and a control policy μ, and establish that their combination admits certain optimality properties. To proceed, we define $I_i(\mathcal{X}_i)$ as the Fisher information matrix corresponding to the unknown parameters observed by agent i, i.e.,

$$I_i(\mathcal{X}_i) \triangleq -\mathbb{E}\left[\nabla^2_{\mathcal{X}_i}\log f_i(Y^i \mid \mathcal{X}_i)\right]\ , \quad (7)$$

where the expectation is with respect to Y^i for given $\mathcal{X}_i$. We also define $\mathcal{X}^i_{t,\mathsf{mle}}$ as the maximum likelihood estimate of $\mathcal{X}^i$ formed by agent $i\in\mathcal{U}$ at time t, i.e.,

$$\mathcal{X}^i_{t,\mathsf{mle}} \triangleq \arg\max_{\mathcal{X}_i} f_i(\mathcal{Y}^i_t \mid \mathcal{X}_i,\mathcal{U}^t)\ . \quad (8)$$

A. Estimator Design

In this subsection, we propose a set of estimators, and establish their optimality in Section IV. Specifically, at time t and for any $j\in\mathcal{M}$, the estimators of X_j by all agents $k\in\mathcal{S}_j$ are defined as the solution to the following linear system of $|\mathcal{S}_j|$ equations:

$$(1+|\mathcal{S}_j|)\hat{X}^k_j(t) = \mathbb{E}[X_j \mid \mathcal{Y}^k_t,\mathcal{U}^t] + \sum_{l\in\mathcal{S}_j}\hat{X}^l_j(t)\ . \quad (9)$$

The estimators $\hat{X}^k_j(t)$ which solve the linear system (9) are of the form

$$\hat{X}^k_j(t) = \sum_{l\in\mathcal{S}_j} a^l_{kj}\mathbb{E}[X_j \mid \mathcal{Y}^l_t,\mathcal{U}^t]\ , \quad (10)$$

where constants $\{a^l_{kj}\}$ satisfy

$$\sum_{l\in\mathcal{S}_j} a^l_{kj} = 1\ . \quad (11)$$

This indicates that the estimators $\hat{X}^k_j(t)$ are weighted average of the local posterior means of the parameter X_j from all the agents that observe X_j.

Remark 1. *The estimators designed according to* (10) *minimize the estimation cost $J(t,\mu,\Delta)$ at any time t and for any control policy μ. Even though minimizing $J(t,\mu,\Delta)$ is not directly related to the objective pursued in this paper, it serves as evidence that these estimators are the optimal choices in the non-sequential setting.*

Throughout the rest of the paper, Δ^* represents the set of estimators designed according to (9). Define $\boldsymbol{p}\triangleq[p_1,\dots,p_K]^T$, where p_i is the probability of agent $i\in\mathcal{U}$ being chosen to collect the new measurement.

Lemma 1. *For the given estimation cost function $J(t,\mu,\Delta)$ defined in* (5) *and the estimators determined by solving* (9)*, we have*

$$\mathbb{P}(t\inf_{\mu} J(t,\mu,\Delta^*)>0)=1\ , \quad (12)$$

and $\forall\epsilon>0$,

$$\lim_{t\to\infty}\mathbb{P}\big(t\inf_{\mu} J(t,\mu,\Delta^*)\geq Q(\mathcal{X})-\epsilon\big) \xrightarrow{\mathsf{P}(\mathcal{X})} 1\ , \quad (13)$$

where we have defined

$$Q(\mathcal{X}) \triangleq \inf_{\boldsymbol{p}\in\phi^K}\sum_{i\in\mathcal{U}}Q_i(\mathcal{X},\boldsymbol{p}) + \frac{1}{2}\sum_{i\in\mathcal{U}}\sum_{j\in\mathcal{X}_i}Q_{ij}(\mathcal{X},\boldsymbol{p})\ , \quad (14)$$

and

$$Q_i(\mathcal{X},\boldsymbol{p}) \triangleq \sum_{j\in\mathcal{X}_i}\sum_{k\in\mathcal{S}_j}\frac{(a^k_{ij})^2}{p_k}[I_k(\mathcal{X}_k)^{-1}]_{jj}\ , \quad (15)$$

and

$$Q_{ij}(\mathcal{X},\boldsymbol{p}) \triangleq \sum_{k\in\mathcal{S}_j}\sum_{l\in\mathcal{S}_j}\frac{(a^l_{ij}-a^l_{kj})^2}{p_l}[I_l(\mathcal{X}_l)^{-1}]_{jj}\ , \quad (16)$$

where $[I_k(\mathcal{X}_k)^{-1}]_{jj}$ is the diagonal entry of the matrix $I_k(\mathcal{X}_k)^{-1}$ corresponding to X_j, and ϕ^K is a the set of probability mass functions (pmfs) over a discrete random variable of size K. $Q(\mathcal{X})$ is a positive random variable and the convergence in (13) *is under* $\mathsf{P}(\mathcal{X})$*, which is the probability distribution of measurements when the true value of the parameters is the set $\mathcal{X}$.*

Lemma 1 is instrumental to establishing the weak asymptotic point-wise optimality of the control policy and stopping rule presented in this paper.

B. Stopping Rule and Control Policy

We now provide the design of the control policy μ^* and the stopping rule T^*.

1) **Control Policy:** At time t, if we decide to collect a new measurement at time $(t+1)$, we form the maximum likelihood estimates $\mathcal{X}^i_{t,\mathsf{mle}}, \forall i\in\mathcal{U}$ based on the measurements collected up to time t. Then at time

$(t+1)$, we randomly choose the agent to collect the measurement from according to the probability distribution $\boldsymbol{p}_t$, where the probabilities $\boldsymbol{p}_t = [p_t(1), \ldots, p_t(K)]^T$ are determined as a solution to

$$\boldsymbol{p}_t = \underset{\boldsymbol{p} \in \phi^K}{\text{argmin}} \sum_{i \in \mathcal{U}} \sum_{j \in \mathcal{X}_i} \sum_{k \in \mathcal{S}_j} \frac{(a_{ji}^k)^2}{p_k} [I_k(\mathcal{X}_{t,\mathsf{mle}}^k)^{-1}]_{jj} + \frac{1}{2} \sum_{i \in \mathcal{U}} \sum_{j \in \mathcal{X}_i} \sum_{k,l \in \mathcal{S}_j} \frac{(a_{ji}^l - a_{jk}^l)^2}{p_l} [I_l(\mathcal{X}_{t,\mathsf{mle}}^l)^{-1}]_{jj} \ . \tag{17}$$

2) **Stopping Time:** At time t, we decide to stop collecting measurements if it is the first time instant at which the estimation cost $J(t, \mu^*, \Delta^*)$ falls below the total sampling cost $(t+1)c$, i.e.,

$$T^* = \inf\{t : J(t, \mu^*, \Delta^*) \leq (t+1)c\} \ . \tag{18}$$

IV. Performance Analysis

The distribution of measurements collected in the network depends on the sampling strategy. Hence, the performance guarantees provided in this paper are based on the notion of weakly asymptotic pointwise optimality defined in [24]. To establish the weakly asymptotic pointwise optimality properties of the designed rules, some regularity conditions are required to be satisfied for the pdfs f_i to ensure the asymptotic efficiency and consistency of maximum likelihood estimators and Bayesian estimators, as well as their normality. The complete list of these conditions can be found in [24] and [32]. We first provide the definition of weak asymptotic point-wise optimality.

Definition 1. *The sequential decision rules characterized by a control policy μ, a stopping rule T, and a set of estimators Δ, are said to be weak asymptotically point-wise optimal if for any $\epsilon > 0$, and any alternative control policy μ', and stopping time T' we have*

$$\lim_{c \to 0} \mathbb{P} \left(\frac{\bar{J}(T, \mu, \Delta, c)}{\bar{J}(T', \mu', \Delta, c)} \leq 1 + \epsilon \right) \to 1 \ . \tag{19}$$

The sufficient conditions for weak asymptotically point-wise optimality of sequential decision rules are established in [24, Theorem 3.1].

Theorem 1. *The control policy μ^* specified in Section III-B satisfies*

$$tJ(t, \mu^*, \Delta^*) \xrightarrow{\mathsf{P}(\mathcal{X})} Q(\mathcal{X}) \ . \tag{20}$$

The results in (12), (13), and (20) are sufficient to conclude that the combination of control policy μ^*, stopping time T^* and the estimators Δ^* satisfy the condition for weak asymptotic point-wise optimality defined in Definition 1.

V. Numerical Evaluation

In this section, we evaluate the performance of the proposed sequential sampling strategy on a network consisting of 4 agents. The set of unknown parameters is given by $\mathcal{X} = \{X_1, X_2, X_3, X_4, X_5, X_6\}$, and we set $\mathcal{X}_1 = \{X_1, X_2\}$, $\mathcal{X}_2 = \{X_2, X_3, X_4\}$, $\mathcal{X}_3 = \{X_3, X_5\}$, and $\mathcal{X}_4 = \{X_4, X_6\}$. The prior pdf of X_j is set to be $\mathcal{N}(0, \sigma_j^2)$, where $\sigma_1^2 = 1$, $\sigma_2^2 = 2$, $\sigma_3^2 = 1.5$, $\sigma_4^2 = 3$, $\sigma_5^2 = 2$, and $\sigma_6^2 = 4$. The measurement at agent i is assumed to be a vector of the same size as $\mathcal{X}_i$ with the additive noise vectors being i.i.d. and Gaussian for each measurement of each individual agent. By assuming different variance values for measurement noise of different agents, we compare the performance of the proposed rules with both a fixed sample-size setting and a sequential setting with a random selection rule.

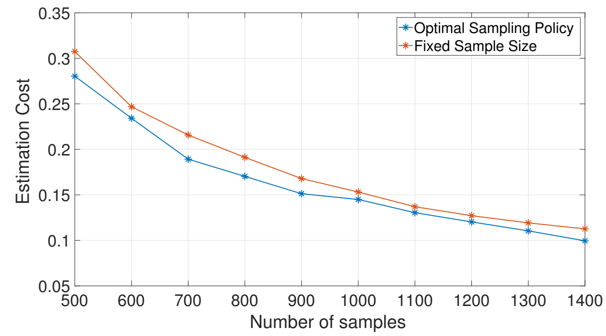

Fig. 1. The estimation cost vs the average number of measurements

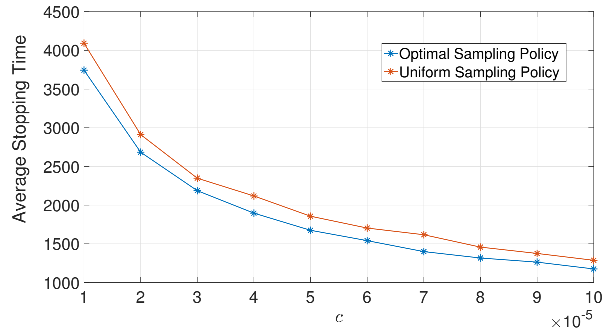

Fig. 2. The average number of measurements vs the sampling cost c

In order to highlight the advantage of sequential sampling, Fig. 1 compares the estimation cost of the proposed strategy with that of a fixed sample-size setting. It is observed that our strategy outperforms the fixed sample setting uniformly for any sample size. To showcase the gain obtained by using the proposed control policy, Fig. 2 compares two sequential setting, one of which uses the proposed control policy for agent selection while the other one uses a random control policy based on a uniform distribution. It is observed that the proposed control policy requires fewer number of measurements for any cost per sample c.

VI. Conclusion

In this paper, we have investigated the problem of parameter estimation from noisy observations collected sequentially in a multi-agent network, where each agent can observe only a subset of the unknown parameters. The goal is to form sufficiently reliable estimates at all agents with the minimum number of measurements and a constraint on the number of measurements that can be collected at any time instant. For the design of estimators, we have adopted a cost function that reflects the fidelity of the estimates formed at the agents and the consistency of the estimates of the parameters commonly observed by multiple agents. We have provided the design for data-adaptive sequential decision rules consisting of the design of optimal estimators at all agents, a control policy for selection of agents to collect measurements from, and a stopping rule to determine whether to stop or continue collecting more measurements. We have established the weak asymptotic point-wise optimality properties for these decision rules.

REFERENCES

[1] G. Fellouris, "Asymptotically optimal parameter estimation under communication constraints," *The Annals of Statistics*, vol. 40, no. 4, pp. 2239–2265, Aug. 2012.

[2] S. Kar and J. M. F. Moura, "Distributed consensus algorithms in sensor networks with imperfect communication: Link failures and channel noise," *IEEE Transactions on Signal Processing*, vol. 57, no. 1, pp. 355–369, Jan. 2009.

[3] S. Kar, J. M. F. Moura, and K. Ramanan, "Distributed parameter estimation in sensor networks: Nonlinear observation models and imperfect communication," *IEEE Transactions on Information Theory*, vol. 58, no. 6, pp. 3575–3605, Jun. 2012.

[4] S. Y. Tu and A. H. Sayed, "Diffusion strategies outperform consensus strategies for distributed estimation over adaptive networks," *IEEE Transactions on Signal Processing*, vol. 60, no. 12, pp. 6217–6234, Dec. 2012.

[5] F. S. Cattivelli and A. H. Sayed, "Diffusion LMS strategies for distributed estimation," *IEEE Transactions on Signal Processing*, vol. 58, no. 3, pp. 1035–1048, Mar. 2010.

[6] P. Meyer, T. Cokelaer, D. Chandran, K. H. Kim, P.-R. Loh, G. Tucker, M. Lipson, B. Berger, C. Kreutz, A. Raue, B. Steiert, J. Timmer, E. Bilal, H. M. Sauro, G. Stolovitzky, and J. Saez-Rodriguez, "Network topology and parameter estimation: from experimental design methods to gene regulatory network kinetics using a community based approach," *BMC Systems Biology*, vol. 8, no. 1, p. 13, Feb. 2014.

[7] J. Sun, J. M. Garibaldi, and C. Hodgman, "Parameter estimation using metaheuristics in systems biology: A comprehensive review," *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol. 9, no. 1, pp. 185–202, Jan. 2012.

[8] H. Eydgahi, W. W. Chen, J. L. Muhlich, D. Vitkup, J. N. Tsitsiklis, and P. K. Sorger, "Properties of cell death models calibrated and compared using Bayesian approaches," *Molecular Systems Biology*, vol. 9, no. 1, 2013.

[9] B. Golub and M. O. Jackson, "Naive learning in social networks and the wisdom of crowds," *American Economic Journal: Microeconomics*, vol. 2, no. 1, Feb. 2010.

[10] L. Xie, D. H. Choi, S. Kar, and H. V. Poor, "Fully distributed state estimation for wide-area monitoring systems," *IEEE Transactions on Smart Grid*, vol. 3, no. 3, pp. 1154–1169, Sep. 2012.

[11] S. Shahrampour and A. Jadbabaie, "Exponentially fast parameter estimation in networks using distributed dual averaging," in *Proc. IEEE Conference on Decision and Control*, Dec. 2013, pp. 6196–6201.

[12] A. Tajer, "Diversified parameter estimation in complex networks," in *Proc. IEEE Global Conference on Signal and Information Processing*, Dec. 2014, pp. 905–908.

[13] K. R. Rad and A. Tahbaz-Salehi, "Distributed parameter estimation in networks," in *Proc. IEEE Conference on Decision and Control*, Dec. 2010, pp. 5050–5055.

[14] P. J. Bickel and J. A. Yahav, "Asymptotically pointwise optimal procedures in sequential analysis," in *Proc. Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, 1967.

[15] G. V. Moustakides, T. Yaacoub, and Y. Mei, "Sequential estimation based on conditional cost," in *Proc. IEEE International Symposium on Information Theory*, June 2017, pp. 436–440.

[16] P. Grambsch, "Sequential sampling based on the observed fisher information to guarantee the accuracy of the maximum likelihood estimator," *The Annals of Statistics*, vol. 11, no. 1, pp. 68–77, Mar. 1983.

[17] A. O. Hero and D. Cochran, "Sensor management: Past, present, and future," *IEEE Sensors Journal*, vol. 11, no. 12, pp. 3064–3075, Dec. 2011.

[18] G. A. Hollinger, B. Englot, F. S. Hover, U. Mitra, and G. S. Sukhatme, "Active planning for underwater inspection and the benefit of adaptivity," *International Journal of Robotics Research*, vol. 32, no. 1, pp. 3–18, Jan. 2013.

[19] S. M. Berry, B. P. Carlin, J. J. Lee, and P. Müller, *Bayesian adaptive methods for clinical trials*, ser. Chapman & Hall/CRC Biostatistics Series. CRC Press, Boca Raton, FL, 2011, vol. 38, with a foreword by David J. Spiegelhalter.

[20] P. Shenoy and A. J. Yu, "Rational decision-making in inhibitory control," *Frontiers in Human Neuroscience*, vol. 48, no. 5, pp. 224–236, May 2011.

[21] R. D. Nowak, "The geometry of generalized binary search," *IEEE Transactions on Information Theory*, vol. 57, no. 12, pp. 7893–7906, Dec. 2011.

[22] M. V. Burnashev, "Data transmission over a discrete channel with feedback. random transmission time," *Problemy Peredachi Informatsii*, vol. 12, no. 4, pp. 10–30, 1976.

[23] V. J. Yohai, "Asymptotically optimal bayes sequential design of experiments for estimation," *The Annals of Statistics*, vol. 1, no. 5, pp. 822–837, Sep. 1973.

[24] G. Atia and S. Aeron, "Asymptotic optimality results for controlled sequential estimation," in *Proc. 51st Annual Allerton Conference Communication, Control, and Computing*, Oct. 2013.

[25] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli., "Controlled sensing for multihypothesis testing," in *IEEE Transactions on Automatic Control*, vol. 58, no. 10, Oct. 2013, pp. 2451–2464.

[26] G. K. Atia and V. V. Veeravalli, "Controlled sensing for sequential multihypothesis testing," in *Proc. International Symposium on Information Theory*, Jul. 2012, pp. 2196–2200.

[27] S. Nitinawarat and V. V. Veeravalli, "Controlled sensing for sequential multihypothesis testing with controlled Markovian observations and non-uniform control cost." in *Sequential Analysis*, vol. 34, no. 1, Jan. 2015, pp. 1–24.

[28] J. Heydari, A. Tajer, and H. V. Poor, "Quickest detection of Markov networks," *arXiv:1711.04268v2*, Dec. 2017.

[29] J. G. Ligo, G. K. Atia, and V. V. Veeravalli, "A controlled sensing approach to graph classification," in *IEEE Transactions on Signal Processing*, vol. 62, no. 24, Oct. 2014, pp. 6468–6480.

[30] Y. Yilmaz, S. Li, and X. Wang, "Sequential joint detection and estimation: Optimum tests and applications," *arXiv:1411.1440v1*, 2014.

[31] Y. Yilmaz, G. V. Moustakides, and X. Wang, "Sequential joint detection and estimation," *Theory of Probability & Its Applications*, vol. 59, no. 3, pp. 452–465, 2015.

[32] A. DasGupta, *Asymptotic Theory of Statistics and Probability*. Springer, New York, 2008.