

Distributed Inference in the Presence of Eavesdroppers: A Survey

Bhavya Kailkhura*, *Student Member, IEEE*, V. Sriram Siddhardh Nadendla*, *Student Member, IEEE*,
Pramod K. Varshney, *Fellow, IEEE*

Abstract—The distributed inference framework comprises of a group of spatially distributed nodes which acquire observations about a phenomenon of interest. Due to bandwidth and energy constraints, the nodes often quantize their observations into a finite-bit local message before sending it to the fusion center (FC). Based on the local summary statistics transmitted by nodes, the FC makes a global decision about the presence of the phenomenon of interest. The distributed and broadcast nature of such systems makes them quite vulnerable to different types of attacks. This paper addresses the problem of secure communication in the presence of eavesdroppers. In particular, we focus on efficient mitigation schemes to mitigate the impact of eavesdropping. We present an overview of the distributed inference schemes under secrecy constraints and describe the currently available approaches in the context of distributed detection and estimation followed by a discussion on avenues for future research.

Index Terms—Distributed inference, distributed detection, distributed estimation, eavesdroppers, secrecy, confidentiality

I. INTRODUCTION

Distributed inference networks have attracted much recent attention due to a variety of applications in civilian and military domains. These include surveillance, environment monitoring, cognitive radio networks and cyber physical systems. Distributed inference networks employ a group of sensing entities that collaborate to sense and make inferences about a given phenomenon of interest (POI). In the traditional framework of centralized inference networks, nodes transmit raw observations to the FC. These transmissions are not attractive in practice as raw observations require a large bandwidth (or energy) for reliable reception at the FC. Therefore, distributed inference networks have been proposed where the nodes transmit compressed observations which are obtained by processing original observations into a finite and tractable alphabet set.

In this paper, we denote the POI with a variable $\theta \in \Theta$, where Θ is the set of possible states that the phenomenon can take. Consider a distributed network, as shown in Figure 1, which comprises of N sensors and a central entity known as the fusion center (FC), which makes inferences about the POI. We assume that the i^{th} node makes an observation Y_i and compresses it into a symbol v_i using a quantizer γ_i . The compressed symbol v_i is then transmitted to the FC through a channel, which is represented as a function $C_F^i(\cdot)$. We denote the received symbols at the FC as $u_i = C_F^i(v_i)$, corresponding to the i^{th} sensor's transmission. The FC uses the fusion rule

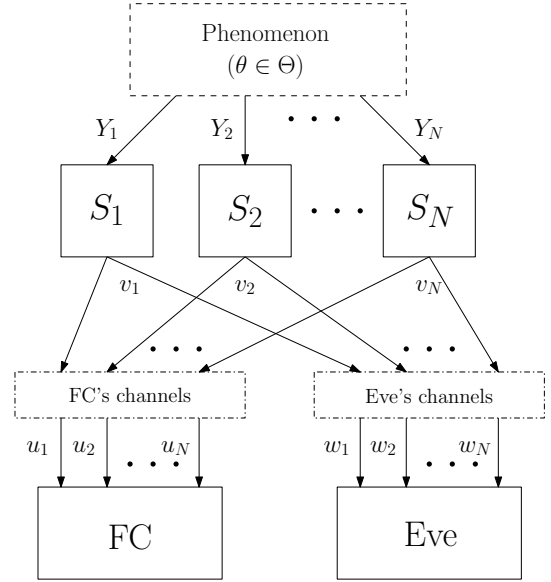


Fig. 1: Distributed Inference Network in the Presence of an Eavesdropper

Γ_{FC} to integrate the symbols $\mathbf{u} = \{u_1, \dots, u_N\}$ into a global inference $\hat{\theta}_{FC} \in \Theta$ about the unknown phenomenon θ .

Although the problem of distributed inference encompasses a broader set of problems, in this paper, we focus our attention on two fundamental problems, namely, *distributed detection* and *distributed estimation*. The fundamental difference in the two problems lies in the definition of the set Θ . In the case of *distributed detection*, $\Theta \in \{0, 1\}$ and in the case of *distributed estimation*, Θ is a continuous set. Practical applications of distributed detection include radar networks where the network may be interested in detecting the presence of an aircraft, or a cognitive radio (CR) network where the secondary users are interested in vacant primary user (PU) channels. On the other hand, examples of distributed estimation include location-estimation and surveillance using spatially distributed sensors.

There are many benefits of distributed inference networks, such as bandwidth efficiency, cost effectiveness and improved reliability. However, the distributed and broadcast nature of the communication links makes the network susceptible to a breach in confidentiality. Thus, a breach in confidentiality of distributed inference networks is an important problem, especially when the network is a part of a larger cyber-physical system. In a fundamental sense, there are two motives for any

The authors are with Department of EECS, Syracuse University, Syracuse, NY 13244. (Email: bkailkhu@syr.edu; vnadendl@syr.edu; varshney@syr.edu)

*These authors contributed equally to this work.

eavesdropper (Eve), namely *selfishness* and *maliciousness*, to compromise the confidentiality of a given distributed inference network. For instance, some of the nodes within a CR network may selfishly take advantage of the FC's inferences and may compete against the CR network in using the PU's channels without paying any participation costs to the network moderator. In another example, if the radar decisions are leaked to a malicious aircraft, the adversary aircraft can maliciously adapt its strategy against a given distributed radar network accordingly so as to remain invisible to the radar and in clandestine pursuit of its mission. Therefore, in the recent past, there has been a lot of interest in the research community in addressing confidentiality in distributed inference networks.

To set the notations, we represent the channel between the i^{th} sensor and the Eve as a function $C_E^i(\cdot)$. The symbol corresponding to the i^{th} node received at the Eve is denoted by $w_i = C_E^i(v_i)$ (See Figure 1). In other words, the total information leakage is a function of $\mathbf{w} = \{w_1, \dots, w_N\}$. Similar to the FC, we assume that Eve uses a decision rule Γ_E to integrate the symbols \mathbf{w} into its own global inference $\hat{\theta}_E$. Several metrics have been proposed in the literature to quantify secrecy or the information leakage to the Eve. Some of them include equivocation, Kullback-Leibler (KL) Divergence, Fisher Information (FI) and probability of error. Ideally, we expect to minimize this information leakage to the maximal extent possible. For example, if KL Divergence or conditional FI is the chosen metric, then *perfect secrecy* is achieved only when KL Divergence or conditional FI at the Eve becomes zero.

In this paper, we survey the state-of-the-art approaches proposed to address secrecy in the context of distributed inference networks. We first introduce a taxonomy in Section II where we present a survey on the state-of-the-art on secrecy in distributed inference networks. Then, in Sections III and IV, we specifically focus on distributed detection and estimation frameworks respectively where we present a detailed account on how secrecy is addressed in each of these frameworks. Finally, we present some important open problems while designing a secure distributed inference network in the presence of eavesdroppers in Section V.

II. APPROACHES TO MITIGATE THREATS ON CONFIDENTIALITY

There are fundamentally four approaches to address secrecy in the context of distributed inference networks which we discuss next.

A. Design of Sensor Quantizers and Fusion Rule

In this approach, the network designer takes advantage of the difference in the channels (C_F^i, C_E^i) , for all $i = 1, \dots, N$, while designing sensor quantizers and the fusion rule. We denote by $\gamma = \{\gamma_1, \dots, \gamma_N\}$ the vector of all sensor quantizers in the distributed inference network. We assume that the quantizer γ_i at the i^{th} sensor lies within the set \mathbb{R}_i , for all $i = 1, \dots, N$. Similarly, we denote the set of decision rules at the FC and Eve as \mathbb{R}_{FC} and \mathbb{R}_E , respectively.

Without any loss of generality, we denote the performance metric at the FC and Eve as Ω_{FC} and Ω_E , respectively. Consider a scenario where the network has a tolerable upper bound on the amount of information leaked to the Eve. Mathematically, this can be quantified in terms of a constraint α on the Eve's performance metric Ω_E . Then, one way of finding the distributed inference system design in terms of sensor quantizers and the fusion rule at the FC is stated as follows.

Problem 1. Find (γ, Γ_{FC}) such that Ω_{FC} is maximized while satisfying the constraints:

- 1) $\max_{\Gamma_E \in \mathbb{R}_E} \Omega_E$ lies below a tolerable value α ,
- 2) quantizers satisfy $\gamma_i \in \mathbb{R}_i, \forall i = 1, \dots, N$,
- 3) fusion rule at the FC satisfy $\Gamma_{FC} \in \mathbb{R}_{FC}$.

Note that error exponents are asymptotic performance metrics at the FC and Eve that represent exponential decay rates of the error probability of their respective "optimal" detectors. Therefore, if the performance metric chosen is an error-exponent such as KL Divergence (for Neyman-Pearson detection setup) or Chernoff Information (for Bayesian detection setup), Problem 1 becomes independent of the fusion rules Γ_{FC} and Γ_E at both the FC and Eve respectively, and reduces to the design of the sensor quantizers alone.

B. Stochastic Encryption

As an alternative to the first approach where the network is designed within the tolerable bounds on information leakage to the Eve, one can pursue a more active approach where the sensors flip their decisions randomly in order to confuse the Eve. In this case, the FC is assumed to have a better knowledge about the sensors than Eve, since the FC either deterministically knows the flipping sensors, or has knowledge about the flipping probability, about which the Eve is completely ignorant. This introduces a significant difference in the channels (C_F^i, C_E^i) , for all $i = 1, \dots, N$, thus, reducing the information leakage to the Eve.

Let the alphabet set of the compressed symbols v_i at the i^{th} sensor be denoted as \mathcal{A} , where the size of \mathcal{A} is denoted by M . In other words, the i^{th} sensor employs an M -ary quantizer to compress the observation Y_i into one of the M symbols. Let us denote the flipping probability matrices as $\mathcal{P} = \{P_1, \dots, P_N\}$, where P_i denotes the flipping probability matrix at the i^{th} sensor which can be interpreted as pre-shared keys between the nodes and the FC. Note that P_i is a stochastic matrix for any $i = 1, \dots, N$, since all of its row elements sum up to unity. The basic problem in this case can be stated as

Problem 2. Find $\mathcal{P} = \{P_1, \dots, P_N\}$ such that Ω_{FC} is maximized while satisfying the constraints:

- 1) $\max_{\Gamma_E \in \mathbb{R}_E} \Omega_E$ lies below a tolerable value α ,
- 2) P_i is a row-stochastic matrix, for all $i = 1, \dots, N$.

Note that, several variants of this problem can be investigated depending on the amount of knowledge the FC has regarding the stochastic encryption process. For example, one may consider that the FC has complete knowledge about the flipping probability matrices \mathcal{P} , however, does not know

exactly whether the sensor messages are flipped or not. In this case, the FC can improve the secrecy performance at the expense of detection performance. On the other hand, the ideal scenario is the case where the FC acquires exact instantaneous knowledge regarding which sensor messages are flipped. This can be done by spending energy in the mechanism that facilitates communication between the FC and the flipping sensors.

C. Artificial Noise Injection

Another approach, similar to the case of stochastic encryption, is the addition of artificial noise to the sensor transmissions. Note that, both stochastic encryption and the addition of artificial noise to the sensor transmissions are data-falsification schemes that are employed to confuse the Eve.

In this paper, we denote the artificial noise added to the i^{th} sensor's transmissions as η_i . Then, the i^{th} sensor transmits \mathbf{x}_i to the FC and Eve, where $\mathbf{x}_i = v_i + \eta_i$. Let $f_i(\eta_i)$ denotes the distribution of η_i . Also, let $\mathcal{F} = \{f_1(\eta_1), \dots, f_N(\eta_N)\}$ denote the set of artificial noise distributions employed by all the sensors in the network. Then the problem can be stated as follows.

Problem 3. Find $\mathcal{F} = \{f_1(\eta_1), \dots, f_N(\eta_N)\}$ such that Ω_{FC} is maximized while satisfying the constraints:

- 1) $\max_{\Gamma_E \in \mathbb{R}_E} \Omega_E$ lies below a tolerable value α ,
- 2) $f_i(\eta_i)$ is a probability density function of η_i , for all $i = 1, \dots, N$.

D. MIMO Beamforming

In order to ensure minimal performance loss at the FC as a trade-off to attaining the secrecy constraint at the Eve, another alternative approach is to use MIMO beamforming, where the sensor messages are directed towards the FC. In this case, we assume that the sensors are equipped with multiple antennas to transmit their messages to the FC. The beamforming mechanism is designed in such a way that some of the available energy is invested in the beams directed towards the FC, while the nulls towards the Eve.

In this paper, we denote the number of antennas at the i^{th} sensor as L_i . Therefore, the i^{th} sensor constructs a vector \mathbf{x}_i based on the symbol v_i and transmits it to the FC and Eve, respectively. Based on the channel gains at the FC and Eve, this \mathbf{x}_i is designed to appear very noisy at the Eve, and simultaneously have significant information about the compressed symbol v_i at the FC. For example, let \mathbf{x}_i be constructed as $\mathbf{x}_i = \mathbf{b}_i v_i$, where \mathbf{b}_i is the beamforming gain vector of the i^{th} sensors' signal. Assuming that both the FC and Eve have only a single antenna, the resulting received symbol at the FC and Eve are given by u_i and w_i respectively. Let n_{FC_i} and n_{E_i} denote the noise at the FC and Eve respectively. Then, $u_i = \mathbf{h}_i^T \mathbf{x}_i + n_{FC_i} = v_i \mathbf{h}_i^T \mathbf{b}_i + n_{FC_i}$ and $w_i = \mathbf{g}_i^T \mathbf{x}_i + n_{E_i} = v_i \mathbf{g}_i^T \mathbf{b}_i + n_{E_i}$. Let the beamforming matrix be denoted as $B = [\mathbf{b}_1 \dots \mathbf{b}_N]$. Since, any practical sensor is energy-constrained, we assume that the total energy available at the i^{th} sensor is denoted by E_i . Then, the design problem can be formally stated as follows.

Problem 4. Find $B = [\mathbf{b}_1 \dots \mathbf{b}_N]$ such that Ω_{FC} is maximized while satisfying the constraints:

- 1) $\max_{\Gamma_E \in \mathbb{R}_E} \Omega_E$ lies below a tolerable value α ,
- 2) \mathbf{b}_i is chosen such that the total transmit energy is within the prescribed limit E_i , for all $i = 1, \dots, N$.

Note that, all of the above approaches can be combined together to design system in a holistic manner and attain a better performance in terms of Ω_{FC} , given a tolerable Eve's constraint α .

III. SECRECY IN DISTRIBUTED DETECTION

In this section, we provide a survey on the state-of-the-art on how secrecy is addressed within the framework of classical and compressive detection networks respectively. In both these frameworks, we organize the survey according to the four different approaches listed in Section II.

A. Classical Distributed Detection

First, we focus our attention to the first approach where the distributed detection network (i.e., sensor quantizers and fusion rule) is optimized while satisfying the secrecy constraints at the Eve. Nadendla *et al.*, made the first attempt in the year 2010 in [1] where they considered an unconstrained differential secrecy problem. Let us denote KL Divergences at the FC and Eve by D_{FC} and D_E , respectively. Now, Problem 1 in their setup reduces to the design of sensor quantizers alone, with $\Omega_{FC} = D_{FC} - D_E$ and $\alpha = \infty$. It was assumed that the channel-state information is completely known at both the FC and the Eve. The authors showed that in the case of eavesdropper with noisier channels, the optimal local detectors are always on the boundaries of the achievable region of sensor's ROC and, therefore, are likelihood-ratio tests (LRTs). Later, the authors also considered Problem 1 with $\Omega_{FC} = D_{FC}$ and $\Omega_E = D_E$, in which case, the structure of an optimal local detector was conjectured to be a LRT-based test based on numerical results.

Marano *et al.* [2], in 2009, considered the problem of designing optimal decision rules for a sensor network where the sensors perform censoring in order to save energy. It was assumed that the eavesdropper does not have access to the sensors' transmitted data but can monitor the transmission activity of the channel and exploit the busy/idle state of the channel for detecting the hypothesis. KL Divergence was used as the performance metric for both the FC and the Eve, and a censoring strategy is developed in order to maximize the divergence of FC while ensuring that the divergence of Eve was zero (perfect secrecy). Although their framework of censoring sensor networks is more general, they assumed that the Eve can only determine whether an individual sensor transmits its decision or not. In reality, Eve can extract more information than just merely determining the presence or absence of transmission, and hence can make a reasonably good decision based on its receptions.

Li *et al.*, in 2014, investigated the problem of Bayesian distributed detection with two nodes in the network in the presence of an eavesdropper in [3], where the Eve has access

to only one of the sensor's transmissions. Here, Ω_{FC} and Ω_E were assumed to be negative expected detection costs at the FC and the Eve respectively. The authors proved that LRT-based tests were optimal at the sensors if the network is designed to minimize the expected detection cost at the FC such that the minimum average cost at the Eve is no greater than a prescribed non-negative value α .

Li *et al.*, also investigated the detection problem under the Neyman-Pearson setup for the same network as in [4]. The sensor quantizers and the fusion rule were designed to maximize the FC's probability of detection (Ω_{FC}) in the presence of constraints on false-alarm probabilities at the FC and Eve, along with the probability of detection at the Eve (Ω_E). Note that the false-alarm constraints at both the FC and the Eve are captured by the feasibility sets \mathbb{R}_{FC} and \mathbb{R}_E respectively. Here, the authors proved that the optimal local quantizer is a deterministic LRT, while the fusion rule may still be a randomization between two or more LRTs. Later, in 2014, Nadendla *et al.* investigated a more general framework in [5] with N sensors. Here, they proved the conjecture stated in [1] in the context of binary symmetric channels between the sensors, FC and Eve. An algorithm was also presented to find optimal thresholds for the likelihood-ratio quantizers when the sensor observations are corrupted by additive Gaussian noise. Figure 2 depicts the behavior of FC's performance in terms of both receiver operating characteristics (ROC) and KL Divergence at the FC as a function of tolerable limits on Eve's KL Divergence. Note that, the optimal quantizer is always on the intersection of the ROC and the Eve's constraint curve. The authors also showed that the network with non-identical sensors and channels can be designed by solving N sequential problems, where the order of this sequence is dictated by the quality of the corresponding sensor's channel.

Next, we survey the literature that addresses the second mitigation approach where a stochastic cipher is employed to confuse the Eve regarding the true phenomenon. In [6], Soosahabi *et al.* employ J-divergence as the performance metric for both the FC and Eve and design a network that guarantees perfect secrecy. This is achieved by fixing $\alpha = 0$ in Problem 3. Probabilistic ciphers were also studied in [7] where the performance metric chosen was the error probability in the case of both FC and Eve. Note that both [6], [7] assume the existence of an underlying key-exchange mechanism that is secure from Eve. Alternatively, channel-aware stochastic ciphers use seeds that are obtained by exploiting randomness in the channel-gains between the node and the FC. For example, Jeon *et al.*, in [8], proposed a type-based multiple access (TBMA) protocol for a distributed detection network with a multiple access channel (MAC). Here, some of the nodes in the network are selected to deliberately transmit interfering signals so as to minimize degradation in the FC's detection performance, while simultaneously preventing Eve from identifying the sensors generating interference. Note that the above scheme requires full channel-state information at the sensors, and therefore, may be impractical in some scenarios. In order to alleviate this problem, efforts such as [9] have been made in the literature, where Jeon *et al.* designed a secure transmission strategy for the local nodes in a parallel

distributed detection network, where the FC first broadcasts known symbols and two thresholds to let the nodes measure their channel condition. Depending on the received symbols, the nodes are divided into three groups, non-flipping, flipping, and dormant groups. The non-flipping set of sensors quantize the sensed data and transmit them to the FC, while the flipping sensors transmit flipped decisions in order to confuse the Eve. The sensors within the dormant set sleep, in order to conserve energy and have an energy-efficient sensor network with longer lifetime.


Finally, there have been efforts to design a hybrid mitigation approach that combines the effects of both the first and the second approaches. In this regard, in [10], Nadendla considered the problem of Bayesian distributed detection in the presence of an eavesdropper, where the nodes use identical threshold quantizers to make their binary decisions and encrypt them before transmission using a simple probabilistic cipher. Cipher parameters and threshold were optimized jointly so as to ensure an acceptable probability of error at the FC while maximizing the probability of error at the Eve.

B. Collaborative Compressed Detection



In scenarios where the POI is a high dimensional signal vector, the Collaborative Compressed Detection (CCD) framework has been proposed. In contrast to conventional detection framework, in CCD, the detection problem is solved directly in compressive measurement domain. More specifically, the CCD framework comprises of a group of spatially distributed nodes which acquire observations regarding the high dimensional ($K \times 1$) signal vector to be detected. Nodes compress their observations using a $M \times K$ low dimensional ($M \ll K$) random projection operator ϕ . Each node i sends an un-quantized (or quantized) version of compressed observation vector Y_i to the Fusion Center (FC) where a global decision is made.

First, we focus our attention on the first approach where nodes do not quantize their observations and the FC receives compressed observation vectors, $\mathbf{Y} = [Y_1, \dots, Y_N]$. Kailkhura *et al.* in [11] considered the problem of collaborative signal vector detection using un-quantized compressive measurements under a physical layer secrecy constraint $\Omega_E \leq \alpha$. To counter Eve, the authors proposed to use β fraction of cooperative nodes that assist the FC by injecting artificial noise (adding or subtracting a constant vector D_i from their observation vector Y_i) in the system to confuse the eavesdroppers. The authors employed deflection coefficient, d_i , as the performance metric for both the FC and the Eve, thus, $\Omega_{FC} = d_{FC}$ and $\Omega_E = d_E$. The problem of determining optimal system parameters (i.e., compression ratio c and noise injection parameters (β, D_i)) which maximize d_{FC} , while ensuring perfect secrecy at the eavesdropper (information of the eavesdropper is exactly zero, i.e., $\alpha = 0$) was also considered.

Kailkhura *et al.* in [12] extended the CCD framework to the case where compressive measurements were quantized to one-bit using LRT. The performance metric was assumed to be the probability of error P_E . They proposed to use B



(a) Sensor's ROC in the presence of Eve

(b) D_{FC} as a function of local false-alarm probabilityFig. 2: Sensor performance in the presence of a constraint, $D_E \leq \alpha$, where $\rho_e = 0.1$ [5].Fig. 3: Minimum Probability of Error $\min_K P_E$ as a function of compression ratio c where local sensor threshold $\lambda = 1$, $\beta = 0.2$, SNR= 10dB and $N = 10$ [12].

out of N cooperating trustworthy nodes that assist the FC by providing flipped decisions (stochastic enciphering with $P_i = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ for all $i = 1, \dots, B$) to the Eve to achieve perfect secrecy. The authors considered the problem of designing optimal system parameters (fusion rule, compression ratio c and fraction of data falsifying nodes $\beta = B/N$) such that P_E at the FC is minimized while ensuring perfect secrecy. In Figure 3 the minimum probability of error (for equal prior case), both at the FC and at the eavesdropper, is plotted as a function of compression ratio c . It can be seen from Figure 3 that the detection performance, both at the FC and at the eavesdropper, is a monotonically increasing function of the compression ratio, i. e., detection performance is better with less compression. This suggests that compression improves security performance at the expense of detection performance.

IV. SECRECY IN DISTRIBUTED ESTIMATION

In this section, we survey the state-of-the-art on how breaches in confidentiality are mitigated in distributed estimation networks. Although little work has been published that addresses secrecy in the context of distributed estimation when compared to the richer literature on secrecy in distributed detection, we again focus on each mitigation technique presented in Section II.

First, we survey the first approach in the context of distributed estimation networks where the sensor quantizers and the fusion rule are designed to guarantee the tolerable limits on Eve's performance. For example, Guo *et al.*, in [13], considered the problem of estimating a single point Gaussian source in the presence of Eve, where the sensor observations are transmitted using an amplify-and-forward technique over a slow-fading orthogonal MAC. Two different scenarios have been addressed within this framework: one, where there are multiple nodes, with each node having a single transmit antenna, and another scenario where a single node has multiple antennas. Through appropriate power allocation at the sensors, the network is designed to achieve the minimum mean squared error (MSE) regarding the POI in each of the above mentioned scenarios while guaranteeing MSE at the Eve to be greater than a threshold α . As shown in Figure 4, the authors plot the distortion (MSE) performance at the FC with respect to the security threshold $\alpha = D_{min}$, with the transmission power budget being set to 30mW, for a one-antenna case and a three-antenna case respectively. For comparison, the system performance is depicted under four settings, namely partial CSI, full CSI, full CSI with perfect secrecy, and partial CSI with artificial noise. First, due to the channel knowledge of both the FC and the Eve, it is not surprising to see that the performance of full CSI scenario is superior to the performance of partial CSI, and the gap keeps increasing as we increase the secrecy threshold. Another important observation is the small gap between the MSE in the perfect secrecy setting and the MSE in the setting with artificial noise. A similar

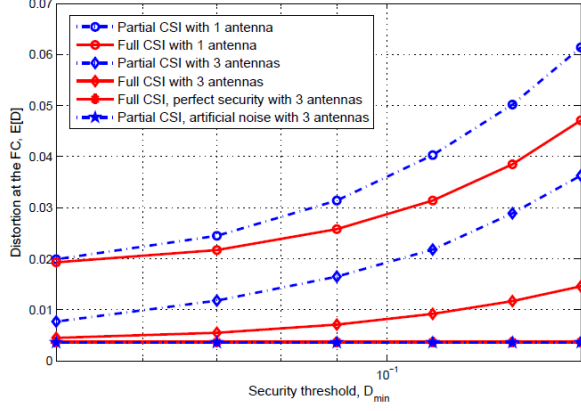


Fig. 4: Performance comparison between full CSI, partial CSI and artificial noise in a multiple-antenna system [13].

performance was also obtained for the multiple nodes network, where each node has only one transmit antenna.

Next, we survey how stochastic encryption is used to achieve the secrecy guarantees within the framework of distributed estimation. Aysal *et al.*, in [14] considered the problem of distributed estimation of a deterministic signal in the presence of an Eve, where each node collects a noisy observation, performs binary quantization, and transmits the 1-bit decision to the FC. The authors assume that both the FC and Eve pursue maximum-likelihood estimation in the presence of a stochastic cipher, for which bias, variance, and MSE were derived in closed form. In the context of symmetric ciphers where $P_i = \begin{pmatrix} 0 & p \\ p & 0 \end{pmatrix}$ for all $i = 1, \dots, N$, the behavior of Eve's bias and MSE and FC's CRLB are characterized in Figure 5. Note that, as $p \rightarrow 0$: 1) the Eve's bias increases; 2) the Eve's MSE increases; and 3) the CRLB decreases. On the other hand, as p tends to unity: 1) the Eve's bias decreases, 2) the Eve's MSE decreases, and 3) the CRLB decreases. In other words, choosing a smaller p is better as it results in a significant amount of bias and MSE at the Eve, with a marginal increase in the estimation variance at the FC. In the case where $P_i = \begin{pmatrix} 0 & p_0 \\ p_1 & 0 \end{pmatrix}$ for all $i = 1, \dots, N$ with $p_0 \neq p_1$, the effect of varying p_0 and p_1 on the FC's CRLB, Eve's bias and Eve's MSE are summarized in Figure 6. In their numerical results, the authors also demonstrated that asymmetric ciphers (i.e., ciphers with asymmetric flipping probability matrices) produce greater bias and MSE than the symmetric ciphers.

V. SUMMARY AND OPEN PROBLEMS

Despite the increasing attention on the problem of secure distributed inference in the presence of eavesdroppers, research in this area is still at an early stage. So far, four different approaches have been proposed to mitigate breaches in confidentiality in the context of distributed inference networks. But, all of these four approaches rely on an important underlying assumption that Eve's channels C_E^i , for all $i = 1, \dots, N$, are




Fig. 5: Effect of varying p on the FC's CRLB (variance of the optimal ML estimator) and the bias and MSE of the Eve [14].

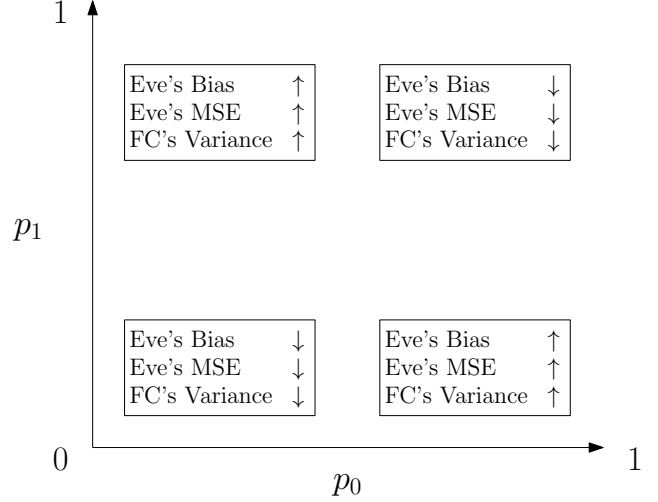


Fig. 6: Effect of varying p_0 and p_1 on the FC's CRLB, Eve's bias and Eve's MSE [14].

completely known at the FC and vice-versa, which may not be true in practice. In fact, there have been no works in the context of inference networks on how one can acquire the information about a passive Eve's channel. This is a hard problem to solve because there is no feedback from the Eve to any of the nodes in the network regarding its presence or activity. An alternative to this roadblock is to assume that Eve's channel belongs to a set \mathcal{C} , and, investigate the best and the worst case performance at the Eve over a class \mathcal{C} . Information regarding this set \mathcal{C} can be obtained from the scene where the network is deployed.

Also, the designers may extend the aforementioned four fundamentally different approaches into several hybrid approaches by considering two or more of these approaches together to create a more sophisticated and improved system in terms of FC's performance for a given tolerable constraint on Eve. Although there have been a few attempts in this direction, one can still envision many such hybrid mechanisms where the designer may accumulate the benefits of each of these approaches. Of course, there is always a need for any new approach which is fundamentally different from any of the four approaches listed in this paper.

REFERENCES

- [1] V. Nadendla, H. Chen, and P. Varshney, "Secure distributed detection in the presence of eavesdroppers," in *Signals, Systems and Computers (ASILOMAR), 2010 Conference Record of the Forty Fourth Asilomar Conference on*, Nov 2010, pp. 1437–1441.

- [2] S. Marano, V. Matta, and P. Willett, "Distributed Detection With Censoring Sensors Under Physical Layer Secrecy," *Signal Processing, IEEE Transactions on*, vol. 57, no. 5, pp. 1976–1986, May 2009.
- [3] Z. Li, T. Oechtering, and K. Kittichokechai, "Parallel Distributed Bayesian Detection with Privacy Constraints," in *Communications Workshops (ICC), 2014 IEEE International Conference on*, June 2014, pp. 2178–2183.
- [4] Z. Li, T. Oechtering, and J. Jalden, "Parallel Distributed Neyman-Pearson Detection with Privacy Constraints," in *Communications Workshops (ICC), 2014 IEEE International Conference on*, June 2014, pp. 765–770.
- [5] V. S. S. Nadendla and P. K. Varshney, "Design of binary quantizers for distributed detection under secrecy constraints," October 2014. [Online]. Available: <http://arxiv.org/abs/1410.8100>
- [6] R. Soosahabi, M. Naraghi-Pour, D. Perkins, and M. Bayoumi, "Optimal Probabilistic Encryption for Secure Detection in Wireless Sensor Networks," *Information Forensics and Security, IEEE Transactions on*, vol. 9, no. 3, pp. 375–385, March 2014.
- [7] R. Soosahabi and M. Naraghi-Pour, "Scalable PHY-Layer Security for Distributed Detection in Wireless Sensor Networks," *Information Forensics and Security, IEEE Transactions on*, vol. 7, no. 4, pp. 1118–1126, Aug 2012.
- [8] H. Jeon, D. Hwang, J. Choi, H. Lee, and J. Ha, "Secure Type-Based Multiple Access," *Information Forensics and Security, IEEE Transactions on*, vol. 6, no. 3, pp. 763–774, Sept 2011.
- [9] H. Jeon, J. Choi, S. McLaughlin, and J. Ha, "Channel aware encryption and decision fusion for wireless sensor networks," *Information Forensics and Security, IEEE Transactions on*, vol. 8, no. 4, pp. 619–625, April 2013.
- [10] V. S. S. Nadendla, "Secure distributed detection in wireless sensor networks via encryption of sensor decisions," Master's thesis, Louisiana State University, 2009.
- [11] B. Kailkhura, T. Wimalajeewa, and P. K. Varshney, "On Physical Layer Secrecy of Collaborative Compressive Detection," in *48th Annual Asilomar Conference on Signals, Systems, and Computers*, 2014.
- [12] B. Kailkhura, T. Wimalajeewa, L. Shen, and P. K. Varshney, "Distributed Compressive Detection with Perfect Secrecy," in *2nd International Workshop on compressive Sensing in Cyber-Physical Systems*, 2014.
- [13] X. Guo, A. S. Leong, and S. Dey, "Estimation in wireless sensor networks with security constraints," 2014. [Online]. Available: http://people.eng.unimelb.edu.au/asleong/secure_estimation_journal.pdf
- [14] T. Aysal and K. Barner, "Sensor Data Cryptography in Wireless Sensor Networks," *Information Forensics and Security, IEEE Transactions on*, vol. 3, no. 2, pp. 273–289, June 2008.