

Distributed Inference with M-ary Quantized Data in the Presence of Byzantine Attacks

V. Sriram Siddhardh (Sid) Nadendla, *Student Member, IEEE*, Yunghsiang S. Han, *Fellow, IEEE*,
and Pramod K. Varshney, *Fellow, IEEE*

Abstract—The problem of distributed inference with M-ary quantized data at the sensors is investigated in the presence of Byzantine attacks. We assume that the Byzantine nodes attack the inference network by modifying the symbol corresponding to the quantized data to one of the other symbols in the quantization alphabet-set and transmitting falsified symbol to the fusion center (FC). In this paper, we find the optimal Byzantine attack that blinds any distributed inference network. As the quantization alphabet size increases, a tremendous improvement in the security performance of the distributed inference network is observed. In addition to the perfect channel case, in Appendix A, we also analyze the optimal Byzantine attack when the channel between the nodes and the FC is noisy and is modelled as a discrete M-ary channel. We also investigate the optimal attack within the restricted space of highly-symmetric attack strategies, that maximally degrades the performance of the inference network in the presence of resource-constrained Byzantine attacks. A reputation-based scheme for identifying malicious nodes is also presented as the network's strategy to mitigate the impact of Byzantine threats on the inference performance of the distributed sensor network. We also provide asymptotic analysis to find the optimal reputation-based scheme as a function of the fraction of compromised nodes in the network.

Index Terms—Distributed Inference, Network-Security, Sensor Networks, Byzantine Attacks, Kullback-Leibler Divergence, Fisher Information.

I. INTRODUCTION

Distributed inference in sensor networks has been widely studied by several scholars in the past three decades (See [1]–[11] and references therein). The distributed inference framework comprises of a group of spatially distributed sensors which acquire observations about a phenomenon of interest (POI) and send processed data to a fusion center (FC) where a global inference is made. Due to resource-constraints in sensor networks, this data is processed at the sensors in such a way that the observations are mapped to symbols from an alphabet set of size M , prior to transmission to the FC. When $M = 2$, we employ binary quantization to generate processed data. When $M > 2$, we send an M-ary symbol that is assumed to be generated via fine quantization. A sensor decision rule is assumed to be characterized by a set of quantization thresholds. In this paper, we use the phrases ‘mapped to one of the M-ary symbols’ and ‘quantized to an M-ary symbol’ interchangeably. A lot of work in the past has

focussed on the binary quantization case, i.e., $M = 2$. In this paper, we consider the case of more general M , $M = 2$ being a special case. The framework of distributed inference networks has been extensively studied for different types of inference problems such as detection (e.g., [1], [3], [5]–[8], [12]), estimation (e.g., [3], [9], [10]), and tracking (e.g., [3], [11]) in the presence of both ideal and non-ideal channels. In this paper, we focus our attention on two distributed inference problems, namely *detection* and *estimation* in the framework of distributed inference, where sensors quantize their data to M-ary symbols.

Although the area of sensor networks has been a very active field of research in the past, security problems in sensor networks have gained attention only in the last decade [13]–[15]. As the security threats have evolved more specifically directed towards inference networks, attempts have been made at the system-level to either prevent or mitigate these threats from deteriorating the network performance. While there are many types of security threats, in this paper, we address the problem of one such attack, called the Byzantine attack, in the context of distributed inference networks (see a recent survey [16] by Vempaty *et al.*). Byzantine attacks (proposed by Lamport *et al.* in [17]) in general, are arbitrary and may refer to many types of malicious behavior. In this paper, we focus only on the data-falsification aspect of the Byzantine attack wherein one or more compromised nodes of the network send false information to the FC in order to deteriorate the inference performance of the network. A well known example of this attack is the *man-in-the-middle* attack [18] where, on one hand, the attacker collects data from the sensors whose authentication process is compromised by the attacker emulating as the FC, while, on the other hand, the attacker sends false information to the FC using the compromised sensors' identity. In summary, if the i^{th} sensor's authentication is compromised, the attacker remains invisible to the network, accepts the true decision u_i from the i^{th} sensor and sends v_i to the FC in order to deteriorate the inference performance.

Marano *et al.*, in [19], analyzed the Byzantine attack on a network of sensors carrying out the task of distributed detection, where the attacker is assumed to have complete knowledge about the hypotheses. This represents the extreme case of Byzantine nodes having an extra power of knowing the true hypothesis. In their model, they assumed that the sensors quantized their respective observations into M-ary symbols, which are later fused at the FC. The Byzantine nodes pick symbols using an optimal probability distribution that are conditioned on the true hypotheses, and transmit

V. Sriram Siddhardh (Sid) Nadendla and Pramod K. Varshney are with the Department of Electrical Engineering and Computer Science, Syracuse University, Syracuse, NY 13201, USA. E-mail: {vnadendl, varshney}@syr.edu.

Yunghsiang S. Han is with the Department of Electrical Engineering, National Taiwan University of Science and Technology, Taipei, Taiwan. E-mail: yshan@mail.ntust.edu.tw.

them to the FC in order to maximally degrade the detection performance. Rawat *et al.*, in [20], also considered the problem of distributed detection in the presence of Byzantine attacks with binary quantizers at the sensors in their analysis. Unlike the authors in [19], Rawat *et al.* did not assume complete knowledge of the true hypotheses at the Byzantine attacker. Instead, they assumed that the Byzantine nodes derive the knowledge about the true hypotheses from their own sensing observations. In other words, a Byzantine node potentially flips the local decision made at the node. It does not modify the thresholds at the sensor quantizers. Rawat *et al.* also analyzed the performance of the network in the presence of independent and collaborative Byzantine attacks and modeled the problem as a zero-sum game between the sensor network and the Byzantine attacker. In addition to the analysis of distributed detection in the presence of Byzantine attacks, a reputation-based scheme was proposed by Rawat *et al.* in [20] for identifying the Byzantine nodes by accumulating the deviations between each sensor's decision and the FC's decision over a time window of duration T . If the accumulated number of deviations is greater than a prescribed threshold for a given node, then the FC tags it as a Byzantine node. In order to mitigate the attack, the FC removes nodes which are tagged Byzantine node from the fusion rule. Another mitigation scheme was proposed by Vempaty *et al.* [21], where each sensor's behavior is learnt over time and compared to the known behavior of the honest nodes. Any significant deviation in the learnt behavior from the expected honest behavior is labelled Byzantine node. Having learnt their parameters, the authors also proposed the use of this information to adapt their fusion rule so as to maximize the performance of the FC. In contrast to the parallel topology in sensor networks, Kailkhura *et al.* in [22] investigated the problem of Byzantine attacks on distributed detection in a hierarchical sensor network. They presented the optimal Byzantine strategy when the sensors communicate their decisions to the FC in multiple hops of a balanced tree. They assumed that the cost of compromising sensors at different levels of the tree varies, and found the optimal Byzantine strategy that minimizes the cost of attacking a given hierarchical network.

Soltanmohammadi *et al.* in [23] investigated the problem of distributed detection in the presence of different types of Byzantine nodes. Each Byzantine node type corresponds to a different operating point, and, therefore, the authors considered the problem of identifying different Byzantine nodes, along with their operating points. The problem of maximum-likelihood (ML) estimation of the operating points was formulated and solved using the expectation-maximization (EM) algorithm. Once the Byzantine node operating points are estimated, this information was utilized at the FC to mitigate the malicious activity in the network, and also to improve global detection performance.

Distributed target localization in the presence of Byzantine attacks was addressed by Vempaty *et al.* in [24], where the sensors quantize their observations into binary decisions, which are transmitted to the FC. Similar to Rawat *et al.*'s approach in [20], the authors in [24] investigated the problem of distributed target localization from both the network's

and Byzantine attacker's perspectives, first by identifying the optimal Byzantine attack and second, mitigating the impact of the attack with the use of non-identical quantizers at the sensors.

In this paper, we extend the framework of Byzantine attacks when Byzantine nodes do not have complete knowledge about the true state of the phenomenon-of-interest (POI), and when the sensors generate M-ary symbols instead of binary symbols. We also assume that the Byzantine attacker is ignorant about the quantization thresholds used at the sensors to generate the M-ary symbols.¹ Under these assumptions, we address two inference problems: binary hypotheses-testing and parameter estimation.

The main contributions of the paper are three-fold. First, we define a Byzantine attack model for a sensor network with individual sensors quantizing their observations into one of the M-ary symbols, when the attacker does not have complete knowledge about the true state of the POI and thresholds employed by the sensors. We model the attack strategy as a flipping probability matrix, where $(i, j)^{th}$ entry represents the probability with which the i^{th} symbol is flipped into the j^{th} symbol. Second, we show that quantization into M-ary symbols at the sensors, as opposed to binary quantization, improves both inference as well as security performance simultaneously. As a function of the number of Byzantine nodes in the network, we derive the optimal flipping matrix. Finally, we extend the mitigation scheme presented by Rawat *et al.* in [20] to the more general case where sensors generate M-ary symbols. We present simulation results to illustrate the performance of the reputation-based scheme for the identification of Byzantine nodes in the network.

The remainder of the paper is organized as follows. In Section II, we describe our system model and present the Byzantine attack model for the case where sensors generate M-ary symbols when the attacker has no knowledge about the true state of the phenomenon of interest and quantization thresholds employed by the sensors. Next, we determine the most powerful attack strategy that the Byzantine nodes can adopt in Section III. In the case of resource-constrained Byzantine attacks, where the attacker cannot compromise enough number of nodes in the network to *blind* it (to be defined in Section II), we find the optimal Byzantine attack for a fixed fraction of Byzantine nodes in the network in the context of distributed detection and estimation in Sections IV and V respectively. From the network's perspective, we present a mitigation scheme in Section VI that identifies the Byzantine nodes using reputation-tags. Finally, we present our concluding remarks in Section VII.

II. SYSTEM MODEL

Consider an inference (sensor) network with N sensors, where α fraction of the nodes in the network are assumed to be compromised (Refer to Figure 1a). These compromised sensors transmit false data to the fusion center (FC) in order to deteriorate the inference performance of the network. We assume that the network is designed to infer about a particular

¹The well-known attacker-in-the-middle is one such example.

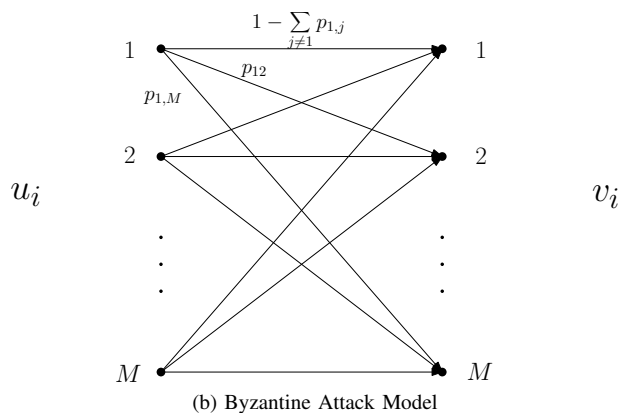
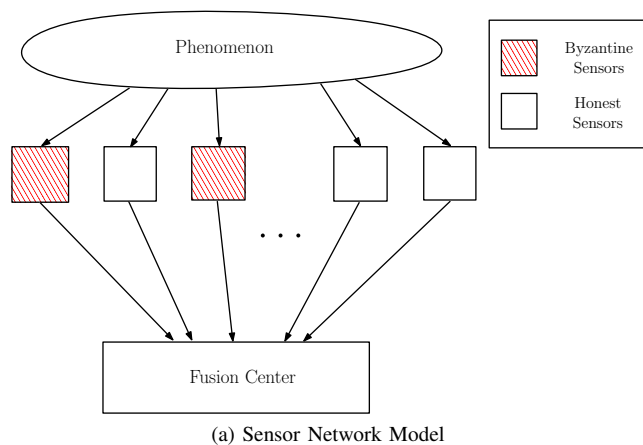


Fig. 1: Distributed Inference Network in the Presence of Byzantine Attacks

phenomenon, regarding which sensors acquire *conditionally-independent* observations. We denote the observation of the i^{th} sensor as r_i . This observation r_i is mapped to one of the M symbols, $u_i \in \{1, \dots, M\}$. In a compromised inference network, since the Byzantine sensors do not transmit their true quantized data, we denote the transmitted symbol as v_i at the i^{th} sensor. If the node i is honest, then $v_i = u_i$. Otherwise, we assume that the Byzantine sensor modifies $u_i = l$ to $v_i = m$ with a probability p_{lm} , as shown in Figure 1b. For the sake of compactness, we denote the transition probabilities depicted in the graph in Figure 1b using a row-stochastic matrix \mathbb{P} , as follows:

$$\mathbb{P} = \begin{bmatrix} p_{11} & p_{12} & \cdots & p_{1M} \\ p_{21} & p_{22} & \cdots & p_{2M} \\ \vdots & \vdots & \ddots & \vdots \\ p_{M1} & p_{M2} & \cdots & p_{MM} \end{bmatrix}. \quad (1)$$

Since the attacker has no knowledge of quantization thresholds employed at each sensor, we assume that \mathbb{P} is independent of the sensor observations. The messages $\mathbf{v} = \{v_1, v_2, \dots, v_N\}$ are transmitted to the fusion center (FC) where a global inference is made about the phenomenon of interest based on \mathbf{v} .

In order to consider the general inference problem, we assume that $\theta \in \Theta$ is the parameter that denotes the phenomenon of interest in the received signal r_i at the i^{th} sensor. If we are considering a detection/classification problem, θ is discrete (finite or countably infinite). In the case of parameter estimation, Θ is a continuous set. Without any loss of generality, we assume $\Theta = \{0, 1, \dots, K-1\}$ if the problem of interest is classification. Hence, detection is a special case of classification with $K = 2$. In the case of estimation, we assume that $\Theta = \mathbb{R}$.

Note that the performance of the FC is determined by the probability distribution (mass function) $P(\mathbf{v}|\theta)$. Therefore, in Section III, we analyze the behavior of $P(\mathbf{v}|\theta)$ in the presence of different attacks and identify the one with the greatest impact on the network.

III. OPTIMAL BYZANTINE ATTACKS

Given the conditional distribution of r_i , $p(r_i|\theta)$, and the sensor quantization thresholds, λ_j for $0 \leq j \leq M$, the conditional distribution of u_i can be found as

$$P(u_i = m|\theta) = \int_{\lambda_{m-1}}^{\lambda_m} p(r_i|\theta) dr_i \quad (2)$$

for all $m = 1, 2, \dots, M$.

If the true quantized symbol at the i^{th} node is $u_i = m$, a compromised node will modify it into $v_i = l$ as depicted in Figure 1b, and transmit it to the FC. Since the FC is not aware of the type of the node (honest or Byzantine), it is natural to assume that node i is compromised with probability α , where α is the fraction of nodes in the network that are compromised. Therefore, we find the conditional distribution of v_i at the FC as follows.

$$\begin{aligned} P(v_i = m|\theta) &= \alpha P(v_i = m|i = \text{Byzantine}, \theta) \\ &\quad + (1 - \alpha) P(v_i = m|i = \text{Honest}, \theta) \\ &= \alpha \sum_{l=1}^M P(u_i = l|\theta) \cdot P(v_i = m|u_i = l, \theta) \\ &\quad + (1 - \alpha) P(u_i = m|\theta) \\ &= \alpha \sum_{l=1}^M p_{lm} P(u_i = l|\theta) + (1 - \alpha) P(u_i = m|\theta) \\ &= \alpha \sum_{l \neq m} p_{lm} P(u_i = l|\theta) + [(1 - \alpha) + \alpha p_{mm}] P(u_i = m|\theta) \\ &= [(1 - \alpha) + \alpha p_{mm}] \\ &\quad + \sum_{l \neq m} \{\alpha p_{lm} - [(1 - \alpha) + \alpha p_{mm}]\} P(u_i = l|\theta). \end{aligned} \quad (3)$$

The goal of a Byzantine attack is to blind the FC with the least amount of effort (minimum α). To totally blind the

FC is equivalent to making $P(v_i = m|\theta) = 1/M$ for all $0 \leq m \leq M - 1$. In Equation (3), the RHS consists of two terms. The first one is based on prior knowledge and the second term conveys information based on the observations. In order to blind the FC, the attacker should make the second term equal to zero. Since the attacker does not have any knowledge regarding $P(u_i = l|\theta)$, it can make the second term of Equation (3) equal to zero by setting

$$\alpha p_{lm} = (1 - \alpha) + \alpha p_{mm}, \quad \forall l \neq m. \quad (4)$$

Then the conditional probability $P(v_i = m|\theta) = (1 - \alpha) + \alpha p_{mm}$ becomes independent of the observations r_i (or its quantized version u_i), resulting in equiprobable symbols at the FC. In other words, the received vector $\mathbf{v} = \{v_1, v_2, \dots, v_N\}$ does not carry any information about θ and, therefore, results in the most degraded performance at the FC. So, the FC now has to solely depend on its prior information about θ in making an inference.

Having identified the condition in Equation (4) under which the Byzantine attack makes the greatest impact on the performance of the network, we identify the strategy that the attacker should employ in order to achieve this condition as follows. Since we need

$$P(v_i = m|\theta) = (1 - \alpha) + \alpha p_{mm} = 1/M,$$

$\alpha = \frac{M-1}{(1-p_{mm})M}$. To minimize α , one needs to make $p_{mm} = 0$. In this paper, we denote the α corresponding to this optimal strategy that minimizes the Byzantine attacker's resources required to blind the FC as α_{blind} . Hence,

$$\alpha_{blind} = \frac{M-1}{M}.$$

Rearranging Equation (4), we have

$$\frac{1}{\alpha} = 1 + (p_{lm} - p_{mm}) = 1 + p_{lm} \quad \forall l \neq m. \quad (5)$$

By setting α to α_{blind} , we have $p_{lm} = 1/(M-1)$, $\forall l \neq m$. That is, the transition probability \mathbb{P} is a highly-symmetric matrix. We summarize the result as a theorem as follows.

Theorem 1. *If the Byzantine attacker has no knowledge of the quantization thresholds employed at each sensor, then the optimal Byzantine attack is given as*

$$p_{lm} = \begin{cases} \frac{1}{M-1} & ; \text{if } l \neq m \\ 0 & ; \text{otherwise} \end{cases} \quad (6)$$

$$\alpha_{blind} = \frac{M-1}{M}.$$

We term Equation (6) as the optimal Byzantine attack, since the FC does not get any information from the data \mathbf{v} it receives from the sensors to perform an inference task. Therefore, the FC has to rely on prior information about the parameter θ , if available.

Theorem 1 can be extended to the case where the channels between sensors (attackers) are not perfect. The result is given in Appendix A.

In Figure 2, we show how α_{blind} scales with increasing quantization alphabet size, M . Since the quantized symbols

Quantization bits	α_{blind}
1	0.5
2	0.75
3	0.875
4	0.9375
5	0.9688
6	0.9844
7	0.9922
8	0.9961

TABLE I: Improvement in α_{blind} with increasing number of quantization bits, $\log_2 M$

are encoded into bits, we also show an exponential increase in α_{blind} as the number of bits needed to encode the M symbols, i.e., $\log_2 M$, increases. This is also shown in Table I. Note that, if the sensors use one additional quantization-bit (2-bit quantization) in their quantization scheme instead of 1-bit quantization (binary quantization), then the α_{blind} increases from 0.5 to 0.75. This trend is observed with increasing number of quantization bits, and when the sensors employ an 8-bit quantizer, then the attacker needs to compromise at least 99.6% of the sensors in the network to blind the FC. Obviously, the improvement in security performance is not free as the sensors incur a communication cost in terms of energy and bandwidth as the number of quantization bits increases. Therefore, in a practical world, the network designer faces a trade-off between the communication cost and the security guarantees.

Also, note that, when $M = 2$ (1-bit quantization), our results coincide with those of Rawat *et al.* in [20], where the focus was on the problem of binary hypotheses testing in a distributed sensor network. On the other hand, our results are more general as they address any inference problem - detection, estimation or classification in a distributed sensor network. Another extreme case to note is when $M \rightarrow \infty$, in which case, $\alpha_{blind} \rightarrow 1$. This means that the Byzantine attacker cannot blind the FC unless all the sensors are compromised.

In the following sections, we consider distributed detection and estimation problems in sensor networks and analyze the impact of the optimal Byzantine attack on these systems. For the sake of tractability, we consider a noiseless channel ($\mathbb{Q} = \mathbb{I}$) at the FC in the framework of resource-constrained Byzantine attack. Therefore, according to Theorem 1, we restrict our attention to the set of highly-symmetric \mathbb{P} for the sake of tractability. In other words, we assume that

$$p_{lm} = \begin{cases} p & \text{if } l \neq m \\ 1 - (M-1)p & \text{otherwise.} \end{cases} \quad (7)$$

IV. DISTRIBUTED DETECTION IN THE PRESENCE OF RESOURCE-CONSTRAINED BYZANTINE ATTACKS

In this section, we consider a resource-constrained Byzantine attack on binary hypotheses testing in a distributed sensor network where the phenomenon of interest is denoted as θ and is modeled as follows:

$$\theta = \begin{cases} 0; & \text{if } H_0 \\ 1; & \text{if } H_1 \end{cases} \quad (8)$$

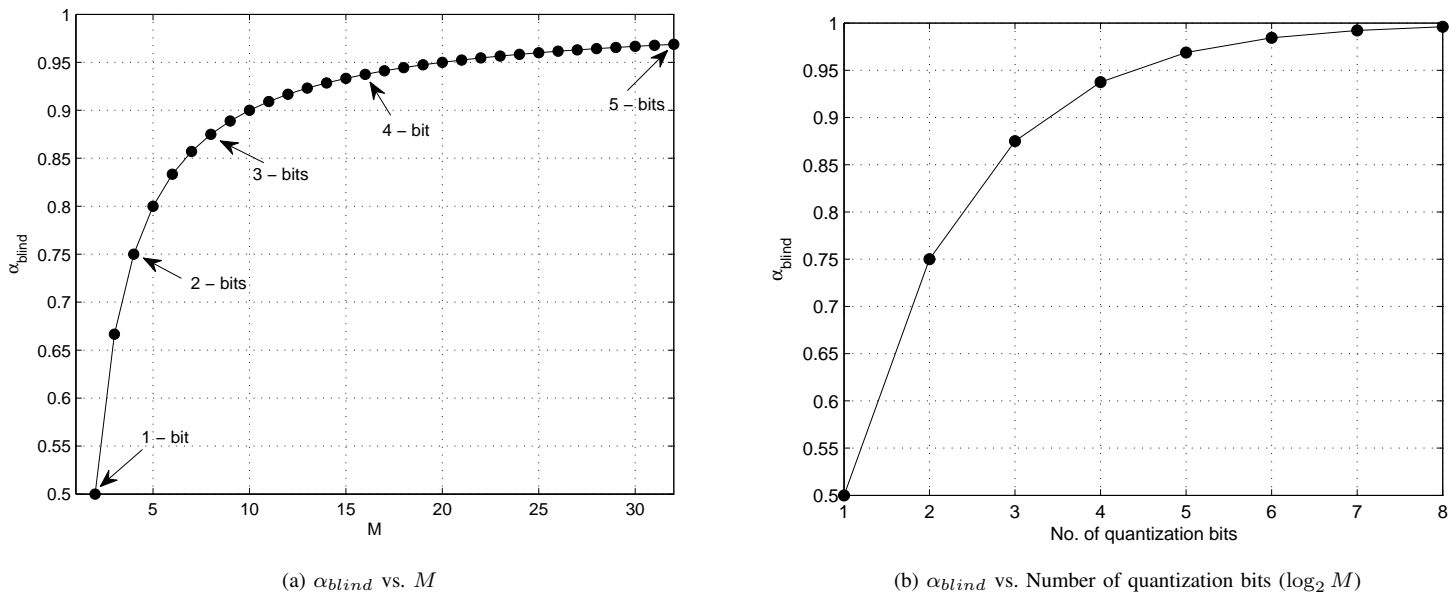


Fig. 2: Improvement in α_{blind} with increasing number of quantization levels

In order to characterize the performance of the FC, we consider Kullback-Leibler Divergence (KLD) as the performance metric. Note that KLD can be interpreted as the error exponent in the Neyman-Pearson detection framework [25], which means that the probability of missed detection goes to zero exponentially with the number of sensors at a rate equal to KLD computed at the FC. We denote KLD at the FC by \mathbb{D}_{FC} and define it as follows:

$$\begin{aligned} \mathbb{D}_{FC} &= \mathbb{E}_{H_0} \left[\log \left(\frac{P(\mathbf{v}|H_0)}{P(\mathbf{v}|H_1)} \right) \right] \\ &= \sum_{\mathbf{m} \in \{1, \dots, M\}^N} P(\mathbf{v} = \mathbf{m}|H_0) \cdot \log \left(\frac{P(\mathbf{v} = \mathbf{m}|H_0)}{P(\mathbf{v} = \mathbf{m}|H_1)} \right) \end{aligned} \quad (9)$$

Since we have assumed that the sensor observations are conditionally independent,² KLD can be expressed as

$$\mathbb{D}_{FC} = N D_{FC}, \quad (10)$$

where

$$D_{FC} = \sum_{m=1}^M P(v = m|H_0) \cdot \log \left(\frac{P(v = m|H_0)}{P(v = m|H_1)} \right).$$

Note that the optimal Byzantine attack, as given in Equation (6), results in equiprobable symbols at the FC irrespective of the hypotheses. Therefore, $\mathbb{D}_{FC} = 0$ under optimal Byzantine attack, resulting in the blinding of the FC.

On the other hand, if the attacker does not have enough resources to compromise α_{blind} fraction of sensors in the network (i.e. $\alpha < \alpha_{blind}$), an optimal strategy for the Byzantine node is to use an appropriate \mathbb{P} matrix that deteriorates the performance of the sensor network to the maximal extent.

²For notational convenience, sensor index i is ignored in the rest of the paper.

As mentioned earlier in Section III, we restrict our search to finding the optimal \mathbb{P} within a space of highly symmetric row-stochastic matrices, as given in Equation (7). Thus, we formulate the problem as follows.

Problem 1. Given the value of $\alpha < \alpha_{blind}$, find the optimal \mathbb{P} within a space of highly symmetric row-stochastic matrices, as given in Equation (7), such that

$$\begin{aligned} &\text{minimize}_p \quad D_{FC} \\ &\text{subject to} \quad 0 \leq p \leq \frac{1}{M-1} \end{aligned}$$

Theorem 2 presents the optimal flipping probability that provides the solution to Problem 1. Note that this result is independent of the design of the sensor network and, therefore, can be employed when the Byzantine has no knowledge about the network.

Theorem 2. Given a fixed $\alpha < \frac{M-1}{M}$, the probability p that optimizes \mathbb{P} within a space of highly symmetric row-stochastic matrices, as given in Equation (7), such that D_{FC} is minimized, is given by

$$p^* = \frac{1}{M-1}. \quad (11)$$

Proof. See Appendix B. □

Note that this solution is of particular interest to the Byzantine attacker since the solution does not require any knowledge about the sensor network design. Also, the attacker's strategy is very simple to implement.

Numerical Results

For illustration purposes, let us consider the following example, where the inference network is deployed to aid the

opportunistic spectrum access for a cognitive radio network (CRN). In other words, the CRs are sensing a licensed spectrum band to find the vacant band for the operation of the CRN.

Let the observation model at the i^{th} sensor be defined as follows.

$$r_i = s(\theta) + n_i, \quad (12)$$

where $\theta \in \{0, 1\}$, $s(\theta) = \mu \cdot (-1)^{1+\theta}$ is a BPSK-modulated symbol transmitted by the licensed (or the primary) user transmitter, and the noise n_i is the AWGN at the i^{th} sensor with probability distribution $\mathcal{N}(0, \sigma^2)$.

Therefore, the conditional distribution of r_i under H_0 and H_1 can be given as $\mathcal{N}(-\mu, \sigma^2)$ and $\mathcal{N}(\mu, \sigma^2)$ respectively. The range of r_i spans the entire real line (\mathbb{R}). However, we assume that the quantizer restricts the support by limiting the range of output values to a smaller range, say $[-A, A]$. This parameter A is called the *overloading* parameter [26] because the choice of A dictates the amount of overloading distortion caused by the quantizer. Within this restricted range of observations, we assume a uniform quantizer with a step size (called the *granularity* parameter) given by $\Delta = \frac{2}{M-2}$, which dictates the granularity distortion of the quantizer. In other words, the observation r_i is quantized using the following quantizer:

$$u_i = \begin{cases} 0; & \text{if } -\infty < r_i \leq \lambda_1 \\ 1; & \text{if } \lambda_1 < r_i \leq \lambda_2 \\ \vdots & \\ M-1; & \text{if } \lambda_{M-1} < r_i \leq \infty \end{cases}, \quad (13)$$

where

$$\lambda_i = A \cdot \left[\frac{2(i-1)}{M-2} - 1 \right].$$

Note that, $\lambda_1 = -A$ and $\lambda_{M-1} = A$ represent the restricted range of the quantizer, as discussed earlier. The i^{th} sensor transmits a symbol v_i to the FC, where $v_i = u_i$ if it is honest. In the case of the i^{th} sensor being a Byzantine node, the decision u_i is modified into v_i using the flipping probability matrix \mathbb{P} as given in Equation (6).

Although the performance of a given sensor network is quantified by the probability of error at the FC, we use a surrogate metric, as described earlier, called the KLD at the FC (Refer to Equation (9)) for the sake of tractability. In an asymptotic sense, Stein's Lemma [25] states that the KLD is the rate at which the probability of missed detection converges to zero under a constrained probability of false alarm. Therefore, in our numerical results, we present how KLD at the FC varies with the fraction of Byzantine nodes α , in the network.

For the above sensor network, we assume that $\mu = 1$, $\sigma^2 = 1$ and $A = 2$. In Figure 3, we plot the contribution of each sensor in terms of KLD at the FC as a function of α , for 1-bit, 2-bit, 3-bit and 4-bit quantizations, i.e., $M = 2, 4, 8$ and 16 respectively, at the sensors. As per our intuition, we observe an improvement in both the detection performance (KLD) as well as security performance (α_{blind}). Therefore, for a given α , the Byzantine attack can be mitigated by employing finer quantization at the sensors. Of course, the best that the




Fig. 3: Contribution of a sensor to the overall KLD at the fusion center as a function of α , for different number of quantization levels. The pentagrams on the x-axis correspond to the α_{blind} for 1-bit, 2-bit, 3-bit and 4-bit quantizations respectively from left to right.

designer can do is to let the sensors transmit unquantized data to the FC, whether in the form of observation samples or their sufficient statistic (likelihood ratio). In this case, we can see that $\alpha_{blind} = 1$, since $\lim_{M \rightarrow \infty} \frac{M-1}{M} = 1$.

V. DISTRIBUTED ESTIMATION IN THE PRESENCE OF RESOURCE-CONSTRAINED BYZANTINE ATTACKS

In this section, we consider the problem of estimating a scalar parameter of interest, denoted by $\theta \in \mathbb{R}$, in a distributed sensor network. As described in the system model, we assume that the i^{th} sensor quantizes its observation r_i into an M -ary symbol u_i , and transmits v_i to the FC. If the i^{th} node is honest, then $v_i = u_i$. Otherwise, we assume that the sensor is compromised and flips u_i into v_i using a flipping probability matrix \mathbb{P} . Under the assumption that the FC receives the symbols \mathbf{v} over an ideal channel, the estimation performance at the FC depends on the probability mass function $P(\mathbf{v}|\theta)$.

The performance of a distributed estimation network can be expressed in terms of the mean-squared error, defined as $\mathbb{E}[(\hat{\theta} - \theta)^2]$. In the case of unbiased estimators, this mean-squared error is lower bounded by the *Cramer-Rao lower bound* (CRLB) [27], which provides a benchmark for the design of an estimator at the FC. We present this result in Equation (14):

$$\mathbb{E}[(\hat{\theta}(\mathbf{v}) - \theta)^2] \geq \frac{1}{I_{FC}}, \quad (14)$$

where

$$I_{FC} = \mathbb{E} \left[\left(\frac{\partial \log P(\mathbf{v}, \theta)}{\partial \theta} \right)^2 \right]. \quad (15)$$

The term I_{FC} is well known as the Fisher information (FI), and is, therefore, a performance metric that captures the performance of the optimal estimator at the FC. Note that, as shown in Equation (16), I_{FC} can be further decomposed into two parts, one corresponding to the prior knowledge about θ at the FC, and the other (denoted as J_{FC}) representing the information about θ , in the sensor transmissions \mathbf{v} :

$$I_{FC} = J_{FC} + \mathbb{E} \left[\left(\frac{\partial \log p(\theta)}{\partial \theta} \right)^2 \right], \quad (16)$$

where

$$J_{FC} = \mathbb{E} \left[\left(\frac{\partial \log P(\mathbf{v}|\theta)}{\partial \theta} \right)^2 \right]. \quad (17)$$

In most cases, a closed form expression for the mean-squared error is intractable and, therefore, conditional Fisher information (FI) is used as a surrogate metric to quantify the performance of a distributed estimation network. In this paper, we also use conditional FI of the received data \mathbf{v} as the performance metric. Since the sensor observations are conditionally independent resulting in independent \mathbf{v} , we denote the conditional FI as \mathbb{J}_{FC} and is defined as follows:

$$\mathbb{J}_{FC} = N J_{FC}, \quad (18)$$

where

$$J_{FC} = \mathbb{E} \left[\frac{\partial}{\partial \theta} \log P(\mathbf{v}|\theta) \right]^2 = -\mathbb{E} \left[\frac{\partial^2}{\partial \theta^2} \log P(\mathbf{v}|\theta) \right]. \quad (19)$$

Following the same approach as in Section IV, we consider the problem of finding an optimal resource-constrained Byzantine attack when $\alpha < \alpha_{blind}$, by finding the symmetric transition matrix \mathbb{P} that minimizes the conditional FI at the FC. This can be formulated as follows.

Problem 2. *Given the value of α , determine the optimal \mathbb{P} within a space of highly symmetric row-stochastic matrices, as given in Equation (7), such that*

$$\begin{aligned} & \underset{p}{\text{minimize}} && J_{FC} \\ & \text{subject to} && 0 \leq p \leq \frac{1}{M-1}. \end{aligned}$$

Theorem 3 presents the optimal flipping probability that provides a solution to Problem 2. Note that this result is independent of the design of the sensor network and, therefore, can be employed when the Byzantine has no knowledge about the network.

Theorem 3. *Given a fixed $\alpha < \frac{M-1}{M}$, the flipping probability p that optimizes \mathbb{P} over a space of highly symmetric row-stochastic matrices, as given in Equation (7), by minimizing J_{FC} is given by*

$$p^* = \frac{1}{M-1}.$$

Proof. See Appendix C. □

Numerical Results

As an illustrative example, we consider the problem of estimating $\theta = 1$ at the FC based on all the sensors' transmitted messages. Let the observation model at the i^{th} sensor be defined as follows:

$$r_i = \theta + n_i, \quad (20)$$

where the noise n_i is the AWGN at the i^{th} sensor with probability distribution $\mathcal{N}(0, \sigma^2)$. The sensors employ the same quantizer as the one presented in Equation (13). The quantized symbol, denoted as u_i at the i^{th} sensor, is then modified into v_i using the flipping probability matrix \mathbb{P} , as given in Equation (6).

Figure 4 plots the conditional FI corresponding to one sensor, for different values of α and M , when the uniform quantizer is centered around the true value of θ . Note that as SNR increases ($\sigma \rightarrow 0$), we observe that it is better for the network to perform as much finer quantization as possible to mitigate the Byzantine attackers. On the other hand, if SNR is low, coarse quantization performs better for lower values of α . This phenomenon of coarse quantization performing better under low SNR scenarios, can be attributed to the fact that more noise gets filtered as the quantization gets coarser (decreasing M) than the signal itself. On the other hand, in the case of high SNR, since the signal level is high, coarse quantization cancels out the signal component significantly, thereby resulting in a degradation in performance.

VI. MITIGATION OF BYZANTINE ATTACKS IN A BANDWIDTH-CONSTRAINED INFERENCE NETWORK

Given that the distributed inference network is under Byzantine attack, we showed that the performance of the network can be improved by increasing the quantization alphabet size of the sensors. Obviously, in a bandwidth-constrained distributed inference network, the sensors can only transmit with the maximum possible M , which is finite. In this section, we assume that the network cannot further increase the quantization alphabet size due to this bandwidth constraint. Therefore, we present a reputation-based Byzantine identification/mitigation scheme, which is an extension of the one proposed by Rawat *et al.* in [20], in order to improve the inference performance of the network.

A. Reputation-Tagging at the Sensors

As proposed by Rawat *et al.* in [20], the FC identifies the Byzantine nodes by iteratively updating a reputation-tag for each node as time progresses. We extend the scheme to include fine quantization scenarios, i.e., $M > 2$, and analyze its performance through simulation results.

As mentioned earlier in the paper, the FC receives a vector \mathbf{v} of received symbols from the sensors and fuses them to yield a global decision, denoted as $\hat{\theta}$. We assume that the observation model is known to the network designer, and is given as follows:

$$r_i = f_i(\theta) + n_i, \quad (21)$$

where $f_i(\cdot)$ denotes the known observation model. We denote the quantization rule employed at the sensor as γ . Therefore,




Fig. 4: Contribution of a sensor to the overall conditional FI at the FC as a function of α , for different number of quantization levels when $\theta = 0$ and $A = 2$. The pentagrams on the x-axis correspond to the α_{blind} for 1-bit, 2-bit, 3-bit and 4-bit quantizations respectively from left to right.

the quantized message at the sensor is given by $u_i = \gamma(r_i)$. As discussed earlier, the i^{th} sensor flips u_i into v_i using a flipping probability matrix \mathbb{P} . Since the FC makes a global inference $\hat{\theta}$, it can calculate the squared-deviation d_i of each sensor from the expected message that it is to nominally transmit as follows:

$$d_i = \left(\gamma^{-1}(v_i) - f_i(\hat{\theta}) \right)^2, \quad (22)$$

where $\gamma^{-1}(v_i)$ is the inverse of the sensor quantizer $\gamma(v_i)$ and it is assumed to be the centroid of the corresponding decision region of the quantizer v_i .

Note that v_i is the received symbol which characterizes the behavior (honest or Byzantine) of the i^{th} sensor, while $f_i(\hat{\theta})$ is the signal that the FC expects the sensor to observe. If the i^{th} sensor is honest, we expect the mean of d_i to be small. On the other hand, if the i^{th} sensor is a compromised node, then the mean of d_i is expected to be large. Therefore, we accumulate the squared-deviations $\mathbf{d}_i = \{d_i(1), \dots, d_i(T)\}$ over T time intervals and compute a reputation tag $\Lambda_i(\mathbf{d}_i)$, as a time-average for the i^{th} node as follows:

$$\Lambda_i = \frac{1}{T} \sum_{t=1}^T d_i(t). \quad (23)$$

The i^{th} sensor is declared honest/Byzantine using the following threshold-based tagging rule

$$\Lambda_i \underset{\text{Honest}}{\overset{\text{Byzantine}}{\geq}} \eta. \quad (24)$$

The performance of the above tagging rule depends strongly on the choice of η . Note that the threshold η should be chosen based on two factors. Firstly, η should be chosen in such a way that the probability with which a malicious node is tagged

Byzantine is high. Higher the value of η , lower is the chance of tagging a node to be Byzantine and vice-versa. This results in a tradeoff between the probability of detecting a Byzantine vs. the probability of falsely tagging an honest node as a Byzantine. Secondly, the value of M also plays a role in the choice of η , and therefore, the performance of the tagging rule. We illustrate this phenomenon in our simulation results.

B. Optimal Choice of the Tagging Threshold as $T \rightarrow \infty$

In this paper, we denote the true type of the i^{th} node as \mathcal{S}_i , where $\mathcal{S}_i = H$ corresponds to honest behavior, while $\mathcal{S}_i = B$ corresponds to Byzantine behavior, for all $i = 1, \dots, N$. Earlier, in this section, we presented Equation (24) which allows us to make inferences about the true type. But, the performance of the Byzantine tagging scheme corresponding the i^{th} sensor is quantified by the conditional probabilities $P(\Lambda_i \geq \eta | \mathcal{S}_i = \mathcal{S})$, for both $\mathcal{S} = H, B$. In order to find the optimal choice of η in Equation (24), we continue with the Neyman-Pearson framework even in the context of Byzantine identification, where the goal is to maximize $P(\Lambda_i \geq \eta | \mathcal{S}_i = B)$, subject to the condition that $P(\Lambda_i \geq \eta | \mathcal{S}_i = H) \leq \xi$.

To find these two conditional probabilities $P(\Lambda_i \geq \eta | \mathcal{S}_i = H)$ and $P(\Lambda_i \geq \eta | \mathcal{S}_i = B)$, we need a closed form expression of the conditional distributions, $P(\Lambda_i | \mathcal{S}_i = H)$ and $P(\Lambda_i | \mathcal{S}_i = B)$ respectively. In practice, where T is finite, it is intractable to determine the conditional distribution of Λ_i , which is necessary to come up with the optimal choice of η . Therefore, in this paper, we assume that $T \rightarrow \infty$ and present an asymptotic choice of the tagging threshold η used in Equation (24).

As $T \rightarrow \infty$, since $d_i(t)$ is independent across $t = 1, \dots, T$, due to central-limit theorem, $(\Lambda_i | \mathcal{S}_i = \mathcal{S}) \sim \mathcal{N}(\mu_{i,\mathcal{S}}, \sigma_{i,\mathcal{S}})$,

where

$$\begin{aligned}\mu_{i,\mathcal{T}} &= \mathbb{E}(\Lambda_i | \mathcal{T}_i = \mathcal{T}) \\ &= \mathbb{E} \left[\left(\gamma^{-1}(v_i(t)) - \hat{\theta}(t) \right)^2 | \mathcal{T}_i = \mathcal{T} \right]\end{aligned}\quad (25)$$

and

$$\begin{aligned}\sigma_{i,\mathcal{T}}^2 &= \text{Var}(\Lambda_i | \mathcal{T}_i = \mathcal{T}) \\ &= \frac{1}{T} \text{Var} \left[\left(\gamma^{-1}(v_i(t)) - \hat{\theta}(t) \right)^2 | \mathcal{T}_i = \mathcal{T} \right].\end{aligned}\quad (26)$$

In this paper, we do not present the final form of $\mu_{i,\mathcal{T}}$ and $\sigma_{i,\mathcal{T}}$ in order to preserve generality. Assuming that $v_i(t)$ is independent across sensors as well as time, the moments of d_i can be computed for any given FC's inference $\hat{\theta}(t)$ at time t about a given phenomenon. Although the final form of $\mu_{i,\mathcal{T}}$ and $\sigma_{i,\mathcal{T}}$ is not presented, since $d_i(t)$ is a function of \mathbf{v} , we present the conditional probability of $(v_j | \mathcal{T}_i = \mathcal{T})$ in Equation (27), which is necessary for the computation of $\mu_{i,\mathcal{T}}$ and $\sigma_{i,\mathcal{T}}$.

$$P(v_j | \mathcal{T}_i = \mathcal{T}) = \int P(v_j | \theta, \mathcal{T}_i = \mathcal{T}) p(\theta) d\theta, \quad (27)$$

where $P(v_j | \theta, \mathcal{T}_i = \mathcal{T})$ can be calculated as follows:

$$P(v_j = m | \theta, \mathcal{T}_i = H) = \begin{cases} P(u_j = m | \theta), & \text{if } j = i \\ (1 - \pi_{BH})P(u_j = m | \theta) + \\ \pi_{BH} \sum_{k=1}^M p_{km} P(u_j = k | \theta), & \text{if } j \neq i \end{cases} \quad (28)$$

and

$$P(v_j = m | \theta, \mathcal{T}_i = B) = \begin{cases} \sum_{k=1}^M p_{km} P(u_j = k | \theta), & \text{if } j = i \\ (1 - \pi_{BB})P(u_j = m | \theta) + \\ \pi_{BB} \sum_{k=1}^M p_{km} P(u_j = k | \theta), & \text{if } j \neq i \end{cases}, \quad (29)$$

where $\pi_{BH} = P(\mathcal{T}_j = B | \mathcal{T}_i = H)$ and $\pi_{BB} = P(\mathcal{T}_j = B | \mathcal{T}_i = B)$ are conditional probabilities of the j^{th} node's type, given the type of the i^{th} node. Since there are α fraction of nodes in the network, given that the FC knows the type of i^{th} node as H , the conditional probability of the j^{th} node belonging to a type \mathcal{T} is given by $\pi_{BH} = \frac{N\alpha}{N-1}$ and $\pi_{BB} = \frac{N\alpha-1}{N-1}$.

Given the conditional distributions $P(\Lambda_i | \mathcal{T}_i = H)$ and $P(\Lambda_i | \mathcal{T}_i = B)$, we find the performance of the Byzantine

identification scheme as follows:

$$\begin{aligned}P(\Lambda_i \geq \eta | \mathcal{T}_i = H) &= Q \left(\frac{\eta - \mu_{i,H}}{\sigma_{i,H}} \right) \\ P(\Lambda_i \geq \eta | \mathcal{T}_i = B) &= Q \left(\frac{\eta - \mu_{i,B}}{\sigma_{i,B}} \right).\end{aligned}\quad (30)$$

Under the NP framework, the optimal η can be chosen by letting $P(\Lambda_i \geq \eta | i = H) = \beta$, when Λ_i is normally distributed conditioned on the true type of a given node. In other words,

$$Q \left(\frac{\eta - \mu_{i,H}}{\sigma_{i,H}} \right) = \xi \quad (31)$$

or equivalently,

$$\eta_{\text{optimal}} = \mu_{i,H} + \sigma_{i,H} Q^{-1}(\xi). \quad (32)$$

Note that, since $P(v_i | \mathcal{T}_i = H)$ is a function of α , it follows that both $\mu_{i,H}$ and $\sigma_{i,H}$ are functions of α . Although we do not provide a closed-form expression for η as a function of α , we provide the following example to portray how η varies with different values of α .

1) *Example: Variation of η as a function of α :* Consider a distributed estimation network with $N = 5$ identical nodes. Let the prior distribution of the true phenomenon θ be the uniform distribution $\mathcal{U}(0, 1)$. We assume that the sensing channel is an AWGN channel where the sensor observations is given by $r_i = \theta + n_i$. Therefore, the conditional distribution of the sensor observations is $\mathcal{N}(\theta, \sigma^2)$, when conditioned on θ . We assume that the sensors employ the quantizer rule shown in Equation (13) on their observations r_i . At the FC, we let $\gamma^{-1}(\cdot)$ be defined as the centroid function that returns $c_i = \frac{\lambda_{i-1} + \lambda_i}{2}$.

Let $\hat{\theta} = \frac{1}{M} \sum_{i=1}^N \gamma^{-1}(v_i(t))$ be the fusion rule employed at the FC to estimate θ .

Since the network comprises of identical nodes, without any loss of generality, we henceforth focus our attention on the reputation-based identification rule at sensor-1. Substituting the above mentioned fusion rule in the squared-deviation d_1 corresponding to sensor-1 in Equation (22), we have

$$\begin{aligned}d_1 &= \left(\gamma^{-1}(v_1) - \frac{1}{M} \sum_{i=1}^5 \gamma^{-1}(v_i(t)) \right)^2 \\ &= \left(\frac{M-1}{M} \gamma^{-1}(v_1) - \frac{1}{M} \sum_{i=2}^5 \gamma^{-1}(v_i(t)) \right)^2.\end{aligned}\quad (33)$$

Let us denote $\phi_{ij} = \mathbb{E} \left\{ (\gamma^{-1}(v_i))^j | \mathcal{T}_1 = H \right\} = \sum_{v_i=1}^M \left[(\gamma^{-1}(v_i))^j P(v_i | \mathcal{T}_1 = H) \right]$, for all $i = 1, \dots, 5$ and $j = 1, 2, \dots, \infty$. Here, $P(v_i | \mathcal{T}_1 = H)$ can be computed using

Equation (28) as follows:

$$\begin{aligned}
 P(v_i = m | \mathcal{F}_1 = H) &= \int_{-\infty}^{\infty} P(v_i = m | \theta, \mathcal{F}_1 = H) p(\theta) d\theta \\
 &= \int_0^1 P(v_i = m | \theta, \mathcal{F}_1 = H) d\theta \\
 &= \begin{cases} a_{1,m} & \text{if } i = 1 \\ \left(\frac{N\alpha}{(N-1)(M-1)} + \left(1 - \frac{MN\alpha}{(N-1)(M-1)} \right) a_{i,m} \right) & \text{otherwise.} \end{cases}
 \end{aligned} \tag{34}$$

where $a_{i,m} = \int_0^1 P(u_i = m | \theta) d\theta$, for all $i = 1, \dots, N$.

Note that, since all the nodes in the network are identical, $P(u_i | \theta)$ is independent of the node-index i , and therefore, $\phi_{ij} = \phi_{2j}$, for all $i \neq 1$.

Thus, the conditional mean and variance, μ_{1H} and σ_{1H}^2 , are given as follows for the special case of $N = 5$:

$$\begin{aligned}
 \mu_{1H} &= \mathbb{E} \left[\left(\frac{M-1}{M} \gamma^{-1}(v_1) - \frac{1}{M} \sum_{i=2}^5 \gamma^{-1}(v_i(t)) \right)^2 \mid \mathcal{F}_i = H \right] \\
 &= \frac{1}{M^2} \mathbb{E} \left[\left((M-1) \gamma^{-1}(v_1) - \sum_{i=2}^5 \gamma^{-1}(v_i(t)) \right)^2 \mid \mathcal{F}_i = H \right] \\
 &= \frac{1}{M^2} [(M-1)^2 \phi_{12} + 4\phi_{22} + 12\phi_{21}^2 - 8(M-1)\phi_{11}\phi_{21}]
 \end{aligned} \tag{35}$$

and

$$\begin{aligned}
 \sigma_{1H}^2 &= \frac{1}{T} \text{Var} \left[\left(\gamma^{-1}(v_i(t)) - \hat{\theta}(t) \right)^2 \mid \mathcal{F}_i = H \right] \\
 &= \frac{1}{T} \{ \Delta - \mu_{1H}^2 \},
 \end{aligned} \tag{36}$$

where

$$\begin{aligned}
 \Delta &= \mathbb{E} \left[\left(\frac{M-1}{M} \gamma^{-1}(v_1) - \frac{1}{M} \sum_{i=2}^5 \gamma^{-1}(v_i(t)) \right)^4 \mid \mathcal{F}_i = H \right] \\
 &= \frac{1}{M^4} [(M-1)^4 \phi_{14} - 16(M-1)^3 \phi_{13}\phi_{21} \\
 &\quad + 6(M-1)^2 \phi_{12} \{ 4\phi_{22} + 12\phi_{21}^2 \} \\
 &\quad - 4(M-1)\phi_{11} (4\phi_{23} + 36\phi_{22}\phi_{21} + 24\phi_{21}^3) + 4\phi_{24} \\
 &\quad + 12\phi_{23}\phi_{21} + 36(\phi_{23}\phi_{21} + \phi_{22}^2 + 2\phi_{22}\phi_{21}^2) \\
 &\quad + 24(\phi_{21}^4 + 3\phi_{22}\phi_{21}^2)].
 \end{aligned} \tag{37}$$

Thus, for $\xi = 0.01$, we compute the tagging threshold η numerically as shown in Equation (32), and plot the variation

of η as a function of α in Figure 5. Note that, in our numerical results, we observe that the optimal choice of η is a convex function of α , where the curvature of the convexity decreases with increasing M . This can be clearly seen from Figure 5b, where we only plot the case of $M = 7$. We observe a similar behavior for all the other values of M , and therefore, present the case of $M = 7$ to illustrate the convex behavior of η . In other words, for very large values of M , the choice of η becomes independent of α , for any fixed $\alpha \leq \alpha_{blind}$.

C. Simulation Results

In order to illustrate the performance of the proposed reputation-based scheme, we consider a sensor network with a total of 100 sensors in the network, out of which 20 are Byzantine sensors. Let the sensor quantizers be given by Equation (13) and the fusion rule at the FC be the MAP rule, given as follows:

$$\sum_{i=1}^N \log \left(\frac{P(v_i | H_1)}{P(v_i | H_0)} \right) \underset{\hat{\theta}=0}{\overset{\hat{\theta}=1}{\geq}} \log \frac{p_0}{p_1}. \tag{38}$$

Figure 6 plots the rate of identification of the number of Byzantine nodes in the network for the proposed reputation-based scheme for different sizes of the quantization alphabet set. Note that the convergence rate deteriorates as M increases. This is due to the fact that the Byzantine nodes have increasing number of symbol options to flip to, because of which a greater number of time-samples are needed to identify the malicious behavior. In addition, we also simulate the evolution of mislabelling an honest node as a Byzantine node in time, and plot the probability of the occurrence of this event in Figure 7. Just as the convergence deteriorates with increasing M , we observe a similar behavior in the evolution of the probability of mislabelling honest nodes. Another important observation in Figure 7 is that the probability of mislabelling a node always converges to zero in time. Similarly, we simulate the evolution of mislabelling a Byzantine node as an honest one in time in Figure 8. We observe similar convergence of the probability of mislabelling a Byzantine node as an honest node to zero, with a rate that decreases with increasing number of quantization levels, M . Therefore, Figures 6, 7 and 8 demonstrate that, after a sufficient amount of time, the reputation-based scheme always identifies the true behavior of a node within the network, with negligible number of mislabels.

VII. CONCLUDING REMARKS

In summary, we modelled the problem of distributed inference with M-ary quantized data in the presence of Byzantine attacks, under the assumption that the attacker does not have knowledge about either the true hypotheses or the quantization thresholds at the sensors. We found the optimal Byzantine attack that *blinds* the FC in the case of any inference task for both noiseless and noisy FC channels. We also considered the problem of resource-constrained Byzantine attack ($\alpha < \alpha_{blind}$) for distributed detection and estimation in the presence of resource-constrained Byzantine attacker for




Fig. 5: Variation of the optimal tagging threshold η (in the asymptotic sense, where $T \rightarrow \infty$) as a function of α




Fig. 6: Rate of identification of the number of Byzantine nodes in time for different number of quantization levels




Fig. 7: Evolution of the probability of mislabelling an honest node as a Byzantine in time for different number of quantization levels

the special case of highly symmetric attack strategies in the presence of noiseless channels at the FC. From the inference network's perspective, we presented a mitigation scheme that identifies the Byzantine nodes through reputation-tagging. We also showed how the optimal tagging threshold can be found when the time-window $T \rightarrow \infty$. Finally, we also investigated the performance of our reputation-based scheme in our simulation results and show that our scheme always converges to finding all the compromised nodes, given sufficient amount of time. In our future work, we will investigate the optimal resource-constrained Byzantine attack in the space of all row-stochastic flipping probability matrices, and if possible, find schemes that mitigate the Byzantine attack more effectively.

APPENDIX A OPTIMAL BYZANTINE ATTACK IN THE PRESENCE OF A DISCRETE NOISY CHANNEL AT THE FC

Given that the messages $\mathbf{v} = \{v_1, v_2, \dots, v_N\}$ are transmitted to the fusion center (FC), we assume a discrete noise channel $\mathbb{Q} = [q_{mn}]$ between the sensors and the FC, where q_{mn} is the probability with which $v_i = m$ is transformed to symbol $z_i = n$ at the i^{th} sensor. Based on the received \mathbf{z} at the FC, a global inference is made about the phenomenon of interest. In this paper, we assume that the row-stochastic channel matrix \mathbb{Q} is invertible for the sake of tractability.

Given the transition probability matrix \mathbb{Q} for the channel between the sensors and the FC, we assume that the FC receives $z_i = n$ when the i^{th} sensor transmits $v_i = m$, with a probability q_{mn} . The conditional distribution of $z_i = n$




Fig. 8: Evolution of the probability of mislabelling a Byzantine node as an honest node in time for different number of quantization levels

under a given phenomenon θ , is given as

$$P(z_i = n|\theta) = \sum_{m=1}^M q_{mn} P(v_i = m|\theta). \quad (39)$$

Note that if \mathbb{Q} is a doubly stochastic matrix, since $\sum_{m=1}^M q_{mn} = 1$, it is sufficient for the Byzantine attacker to ensure $P(v_i = m|\theta) = \frac{1}{M}$. Thus, by Theorem 1, we have the following theorem when \mathbb{Q} is a doubly stochastic matrix.

Theorem 4. *If the channel matrix \mathbb{Q} is doubly-stochastic, and if the Byzantine attacker has no knowledge about the sensors' quantization thresholds, then the optimal Byzantine attack is given as*

$$p_{lm} = \begin{cases} \frac{1}{M-1} & ; \text{if } l \neq m \\ 0 & ; \text{otherwise} \end{cases} \quad (40)$$

$$\alpha_{blind} = \frac{M-1}{M}.$$

Therefore, we focus our attention to any general row-stochastic channel matrix \mathbb{Q} , where $\sum_{m=1}^M q_{mn}$ need not necessarily sum to unity for all $n = 1, \dots, M$. In other words, the Byzantine attacker has to find an alternative strategy to blind the FC, where $P(z_i = n|\theta) = \frac{1}{M}$. Substituting Equation (3) in Equation (39) and rearranging the terms, we have the

following.

$$\begin{aligned} P(z_i = n|\theta) &= \sum_{m=1}^M q_{mn} P(v_i = m|\theta) \\ &= \sum_{m=1}^M q_{mn} [(1-\alpha) + \alpha p_{mm}] \\ &+ \sum_{m=1}^M q_{mn} \left\{ \sum_{l \neq m} \{\alpha p_{lm} - [(1-\alpha) + \alpha p_{mm}]\} P(u_i = l|\theta) \right\} \\ &= \sum_{m=1}^M q_{mn} [(1-\alpha) + \alpha p_{mm}] \\ &+ \sum_{l=1}^M \left[\sum_{m \neq l} q_{mn} \{\alpha p_{lm} - [(1-\alpha) + \alpha p_{mm}]\} \right] P(u_i = l|\theta). \end{aligned} \quad (41)$$

The goal of a Byzantine attack is to blind the FC with the least amount of effort (minimum α). To totally blind the FC is equivalent to making $P(z_i = n|\theta) = 1/M$ for all $0 \leq n \leq M-1$. In Equation (41), the RHS consists of two terms. The first one is based on prior knowledge and the second term conveys information based on the observations. In order to blind the FC, the attacker should make the second term equal to zero. Since the attacker does not have any knowledge regarding $P(u_i = l|\theta)$, it can make the second term of Equation (41) equal to zero by setting

$$\sum_{m \neq l} q_{mn} \{\alpha p_{lm} - [(1-\alpha) + \alpha p_{mm}]\} = 0 \text{ for all } 1 \leq n, l \leq M. \quad (42)$$

Then the conditional probability $P(z_i = n|\theta) = \sum_{m=1}^M q_{mn} [(1-\alpha) + \alpha p_{mm}]$ becomes independent of the observations r_i (or its quantized version u_i), resulting in equiprobable symbols at the FC. In other words, the received vector $\mathbf{z} = \{z_1, z_2, \dots, z_N\}$ does not carry any information about $\mathbf{u} = \{u_1, u_2, \dots, u_N\}$, thus making FC solely dependent on its prior information about θ in making an inference.

In order to identify the strategy that the attacker should employ to achieve the condition in Equation (42), for all $n = 1, \dots, M$, we need

$$P(z_i = n|\theta) = \frac{1}{M}, \quad (43)$$

$$\text{or, } \sum_{m=1}^M q_{mn} \{(1-\alpha) + \alpha p_{mm}\} = \frac{1}{M}.$$

In matrix form, we can rewrite Equation (43) as

$$(1-\alpha)\mathbf{1}^T \mathbb{Q} + \alpha \mathbf{p}^T \mathbb{Q} = \frac{1}{M} \mathbf{1}^T,$$

where $\mathbf{1}$ is an all-one column-vector and $\mathbf{p} = [p_{11}, \dots, p_{MM}]^T$ is the column-vector of all diagonal elements of \mathbb{P} . In other words,

$$\alpha(\mathbf{1} - \mathbf{p}) = \mathbf{1} - \frac{1}{M} (\mathbb{Q}^T)^{-1} \mathbf{1} \quad (44)$$

Note that every element in the LHS of Equation (44) always lies between 0 and 1. Therefore, the existence of the Byzantine's optimal strategy relies on the following condition. In other words,

$$\mathbf{0} \leq (\mathbb{Q}^T)^{-1} \mathbf{1} \leq M \mathbf{1}. \quad (45)$$

If (45) does not hold, there does not exist an optimal strategy. Given that the condition in Equation (45) holds, the minimum α can be found as follows.

$$\begin{aligned} \alpha_{blind} &= \min \left\{ \mathbf{1} - \frac{1}{M} (\mathbb{Q}^T)^{-1} \mathbf{1} \right\} \\ &= 1 - \frac{1}{M} \max \left\{ (\mathbb{Q}^T)^{-1} \mathbf{1} \right\}. \end{aligned} \quad (46)$$

Therefore, \mathbf{p} can be calculated as

$$\begin{aligned} \mathbf{p} &= \mathbf{1} - \frac{1}{\alpha_{blind}} \left(\mathbf{1} - \frac{1}{M} (\mathbb{Q}^T)^{-1} \mathbf{1} \right) \\ &= \frac{1}{\alpha_{blind} M} (\mathbb{Q}^T)^{-1} \mathbf{1} - \frac{1 - \alpha_{blind}}{\alpha_{blind}} \mathbf{1}. \end{aligned} \quad (47)$$

Next, in order to find the rest of the \mathbb{P} matrix, let us consider Equation (42). Adding $q_{ln} \{ \alpha p_{lu} - [1 - \alpha + \alpha p_{lu}] \}$ on both sides to Equation (42), we have

$$\sum_{m=1}^M q_{mn} \{ \alpha p_{lm} - [(1 - \alpha) + \alpha p_{mm}] \} = -q_{ln} (1 - \alpha) \quad \text{for all } 1 \leq n, l \leq M.$$

$$\text{or, } \alpha \sum_{m=1}^M q_{mn} p_{lm} = \frac{1}{M} - q_{ln} (1 - \alpha) \quad \text{for all } 1 \leq n, l \leq M. \quad (48)$$

In matrix form, we have

$$\alpha \mathbb{P} \mathbb{Q} = \frac{1}{M} \mathbf{1} - (1 - \alpha) \mathbb{Q}, \quad (49)$$

where $\mathbf{1}$ is an all-one matrix. Equivalently, we have

$$\mathbb{P} = \frac{1}{\alpha M} \mathbf{1} \mathbb{Q}^{-1} - \frac{1 - \alpha}{\alpha} \mathbb{I}, \quad (50)$$

where \mathbb{I} is the identity matrix. Note that the vector \mathbf{p} (comprising the diagonal elements of \mathbb{P}) obtained from Equation (50) is verified to be same as that from Equation (47).

In summary, we have the following theorem that provides the optimal Byzantine strategy in the presence of noisy FC channels:

Theorem 5. *Let the Byzantine attacker have no knowledge about the sensors' quantization thresholds, and, the FC's channel matrix be \mathbb{Q} . If \mathbb{Q} is non-singular, and, if $\mathbf{0} \leq (\mathbb{Q}^T)^{-1} \mathbf{1} \leq M \mathbf{1}$, then the optimal Byzantine attack is given as*

$$\begin{aligned} \alpha_{blind} &= 1 - \frac{1}{M} \max \left\{ (\mathbb{Q}^T)^{-1} \mathbf{1} \right\} \\ \mathbb{P} &= \frac{1}{\alpha_{blind} M} \mathbf{1} \mathbb{Q}^{-1} - \frac{1 - \alpha_{blind}}{\alpha_{blind}} \mathbb{I}. \end{aligned} \quad (51)$$

Note that, if the channel matrix \mathbb{Q} is doubly-stochastic, we have $\mathbb{Q} \mathbf{1} = \mathbf{1}$ and $\mathbb{Q}^T \mathbf{1} = \mathbf{1}$. Substituting these conditions in Equation (51), Theorem 5 reduces to Theorem 4.

Having identified the optimal Byzantine attack, one can observe that the attacker needs to compromise a huge number of sensors ($\alpha_{blind} = 1 - \frac{1}{M} \max \left\{ (\mathbb{Q}^T)^{-1} \mathbf{1} \right\}$) in the network to blind the FC. Therefore, it is obvious that, in the case of a resource-constrained attacker, the attacker compromises a fixed fraction of nodes $\alpha \leq \alpha_{blind}$ in such a way that the performance degradation at the FC is maximized. In our future work, we will investigate the problem of finding the optimal strategy in the context of resource-constrained Byzantine attacks in the presence of noisy FC channels.

APPENDIX B PROOF FOR THEOREM 2

For the sake of notational simplicity, let us denote $x_m = P(u = m|H_0)$ and $y_m = P(u = m|H_1)$. Similarly, $\tilde{x}_m = P(v = m|H_0)$ and $\tilde{y}_m = P(v = m|H_1)$.

Rewriting Equation (3) in our new notation, we have

$$\tilde{x}_m = \alpha \sum_{l \neq m} p_{xl} + (1 - \alpha(M - 1)p)x_m = \alpha p + (1 - M\alpha p)x_m \quad (52)$$

and

$$\tilde{y}_m = \alpha \sum_{l \neq m} p_{yl} + (1 - \alpha(M - 1)p)y_m = \alpha p + (1 - M\alpha p)y_m. \quad (53)$$

Therefore, the KLD at the FC can be rewritten as

$$D_{FC} = \sum_{m=1}^M \tilde{x}_m \log \left(\frac{\tilde{x}_m}{\tilde{y}_m} \right). \quad (54)$$

On partially differentiating D_{FC} with respect to p , we have

$$\begin{aligned} \frac{\partial D_{FC}}{\partial p} &= \frac{\partial}{\partial p} \sum_{m=1}^M \tilde{x}_m \log \left(\frac{\tilde{x}_m}{\tilde{y}_m} \right) \\ &= \alpha \sum_{m=1}^M \left[(1 - Mx_m) \left(1 + \log \frac{\tilde{x}_m}{\tilde{y}_m} \right) - (1 - My_m) \frac{\tilde{x}_m}{\tilde{y}_m} \right] \\ &= \alpha \sum_{m=1}^M (1 - Mx_m) + \alpha \sum_{m=1}^M (1 - Mx_m) \log \frac{\tilde{x}_m}{\tilde{y}_m} \\ &\quad - \alpha \sum_{m=1}^M (1 - My_m) \frac{\tilde{x}_m}{\tilde{y}_m}. \end{aligned} \quad (55)$$

Consider the first term in the RHS of Equation (55). Note that, since $\mathbf{x} = \{x_1, \dots, x_M\}$ is a probability mass function, we have

$$\sum_{m=1}^M (1 - Mx_m) = M - M \sum_{m=1}^M x_m = M - M = 0.$$

Therefore, Equation (55) reduces to

$$\frac{\partial D_{FC}}{\partial p} = \alpha \sum_{m=1}^M (1 - Mx_m) \log \frac{\tilde{x}_m}{\tilde{y}_m} - \alpha \sum_{m=1}^M (1 - My_m) \frac{\tilde{x}_m}{\tilde{y}_m}. \quad (56)$$

Rearranging the terms in Equation (56), we have

$$\begin{aligned} \frac{\partial D_{FC}}{\partial p} = & \alpha \sum_{m=1}^M \left[\log \frac{\tilde{x}_m}{\tilde{y}_m} - \frac{\tilde{x}_m}{\tilde{y}_m} \right] - \alpha M \sum_{m=1}^M x_m \log \frac{\tilde{x}_m}{\tilde{y}_m} \\ & + \alpha M \sum_{m=1}^M y_m \frac{\tilde{x}_m}{\tilde{y}_m}. \end{aligned} \quad (57)$$

Let us denote the first term as T_1 . In other words,

$$T_1 = \alpha \sum_{m=1}^M \left[\log \frac{\tilde{x}_m}{\tilde{y}_m} - \frac{\tilde{x}_m}{\tilde{y}_m} \right].$$

Let us now focus our attention on the other terms in the RHS of Equation (57). Substituting Equations (52) and (53) in the second and third terms of the RHS of Equation (57), we have

$$\begin{aligned} \frac{\partial D_{FC}}{\partial p} = & T_1 - \frac{M\alpha}{1 - M\alpha p} \sum_{m=1}^M (\tilde{x}_m - \alpha p) \log \frac{\tilde{x}_m}{\tilde{y}_m} \\ & + \frac{M\alpha}{1 - M\alpha p} \sum_{m=1}^M (\tilde{y}_m - \alpha p) \frac{\tilde{x}_m}{\tilde{y}_m} \\ = & T_1 - \frac{M\alpha}{1 - M\alpha p} D(\tilde{\mathbf{x}}||\tilde{\mathbf{y}}) \\ & + \frac{M\alpha}{1 - M\alpha p} \left\{ \sum_{m=1}^M \alpha p \log \frac{\tilde{x}_m}{\tilde{y}_m} - \sum_{m=1}^M \alpha p \frac{\tilde{x}_m}{\tilde{y}_m} + \sum_{m=1}^M \tilde{x}_m \right\}, \end{aligned} \quad (58)$$

where $D(\tilde{\mathbf{x}}||\tilde{\mathbf{y}})$ is the KLD between $\tilde{\mathbf{x}}$ and $\tilde{\mathbf{y}}$ and is, therefore, non-negative. Also, note that in Equation (58), since $\tilde{\mathbf{x}} = \{\tilde{x}_1, \dots, \tilde{x}_M\}$ is a probability mass function, $\sum_{m=1}^M \tilde{x}_m = 1$.

Therefore, Equation (58) reduces to

$$\begin{aligned} \frac{\partial D_{FC}}{\partial p} = & T_1 - \frac{M\alpha}{1 - M\alpha p} D(\tilde{\mathbf{x}}||\tilde{\mathbf{y}}) + \frac{M\alpha}{1 - M\alpha p} \\ & + \frac{M\alpha^2 p}{1 - M\alpha p} \sum_{m=1}^M \left[\log \frac{\tilde{x}_m}{\tilde{y}_m} - \frac{\tilde{x}_m}{\tilde{y}_m} \right]. \end{aligned} \quad (59)$$

Note that the last term in the RHS of Equation (59),

$$\frac{M\alpha^2 p}{1 - M\alpha p} \sum_{m=1}^M \left[\log \frac{\tilde{x}_m}{\tilde{y}_m} - \frac{\tilde{x}_m}{\tilde{y}_m} \right] = \frac{M\alpha p}{1 - M\alpha p} T_1.$$

In other words,

$$\begin{aligned} \frac{\partial D_{FC}}{\partial p} = & \left(1 + \frac{M\alpha p}{1 - M\alpha p} \right) T_1 - \frac{M\alpha}{1 - M\alpha p} D(\tilde{\mathbf{x}}||\tilde{\mathbf{y}}) \\ & + \frac{M\alpha}{1 - M\alpha p} \\ = & \frac{1}{1 - M\alpha p} T_1 - \frac{M\alpha}{1 - M\alpha p} D(\tilde{\mathbf{x}}||\tilde{\mathbf{y}}) + \frac{M\alpha}{1 - M\alpha p}. \end{aligned} \quad (60)$$

Rearranging the terms in Equation (60) and expanding T_1 , we have

$$\begin{aligned} \frac{\partial D_{FC}}{\partial p} = & -\frac{M\alpha}{1 - M\alpha p} D(\tilde{\mathbf{x}}||\tilde{\mathbf{y}}) + \frac{M\alpha}{1 - M\alpha p} \\ & + \frac{\alpha}{1 - M\alpha p} \sum_{m=1}^M \left[\log \frac{\tilde{x}_m}{\tilde{y}_m} - \frac{\tilde{x}_m}{\tilde{y}_m} \right] \\ = & -\frac{M\alpha}{1 - M\alpha p} D(\tilde{\mathbf{x}}||\tilde{\mathbf{y}}) \\ & + \frac{\alpha}{1 - M\alpha p} \sum_{m=1}^M \left[\log \frac{\tilde{x}_m}{\tilde{y}_m} - \left(\frac{\tilde{x}_m}{\tilde{y}_m} - 1 \right) \right]. \end{aligned} \quad (61)$$

Since $\log x \leq x - 1$ for all x , we find that the second term in the RHS of Equation (56) is negative. Therefore, we have

$$\frac{\partial D_{FC}}{\partial p} \leq 0. \quad (62)$$

Since D_{FC} is a non-increasing function of p , the optimal p , p^* , takes the maximum value $1/(M - 1)$.

APPENDIX C PROOF FOR THEOREM 3

For the sake of notational simplicity, we let $z_m = P(u = m|\theta)$. Similarly, $\tilde{z}_m = P(v = m|\theta)$. Using this notation in Equation (19), we have

$$\begin{aligned} J_{FC} = & \sum_{m=1}^M P(v = m|\theta) \left(\frac{\partial \log P(v = m|\theta)}{\partial \theta} \right)^2 \\ = & \sum_{m=1}^M \tilde{z}_m \left(\frac{\partial \log \tilde{z}_m}{\partial \theta} \right)^2 \\ = & (1 - M\alpha p)^2 \sum_{m=1}^M \frac{1}{\tilde{z}_m} \left(\frac{\partial z_m}{\partial \theta} \right)^2. \end{aligned} \quad (63)$$

Partially differentiating J_{FC} with respect to p , we have

$$\begin{aligned}
 \frac{\partial J_{FC}}{\partial p} &= 2(1 - M\alpha p)(-M\alpha) \sum_{m=1}^M \frac{1}{\tilde{z}_m} \left(\frac{\partial z_m}{\partial \theta} \right)^2 \\
 &\quad + (1 - M\alpha p)^2 \sum_{m=1}^M \left(-\frac{1}{\tilde{z}_m^2} \right) (\alpha - M\alpha z_m) \left(\frac{\partial z_m}{\partial \theta} \right)^2 \\
 &= -(1 - M\alpha p) \left[2M\alpha \sum_{m=1}^M \tilde{z}_m \left(\frac{1}{\tilde{z}_m} \frac{\partial z_m}{\partial \theta} \right)^2 \right. \\
 &\quad \left. + (1 - M\alpha p) \sum_{m=1}^M \alpha \left(\frac{1}{\tilde{z}_m} \frac{\partial z_m}{\partial \theta} \right)^2 \right. \\
 &\quad \left. - (1 - M\alpha p) \sum_{m=1}^M M\alpha z_m \left(\frac{1}{\tilde{z}_m} \frac{\partial z_m}{\partial \theta} \right)^2 \right] \\
 &= -(1 - M\alpha p) \left[\alpha(1 - M\alpha p) \sum_{m=1}^M \left(\frac{1}{\tilde{z}_m} \frac{\partial z_m}{\partial \theta} \right)^2 \right. \\
 &\quad \left. + M\alpha(1 + M\alpha p) \sum_{m=1}^M z_m \left(\frac{1}{\tilde{z}_m} \frac{\partial z_m}{\partial \theta} \right)^2 \right] \quad (64)
 \end{aligned}$$

In Equation (64), we have a negative term multiplied by a non-negative term, and hence we have

$$\frac{\partial J_{FC}}{\partial p} \leq 0. \quad (65)$$

Since J_{FC} is a non-increasing function of p , $p^* = \frac{1}{M-1}$, being the maximum value, is the optimal solution to Problem 2.

ACKNOWLEDGEMENT

This work was supported in part by AFOSR under Grants FA9550-10-1-0458, FA9550-10-1-0263, FA 9550-10-C-0179 and by CASE at Syracuse University and National Science Council of Taiwan, under grants NSC 99-2221-E-011-158-MY3, NSC 101-2221-E-011-069 -MY3. Han's work was completed during his visit to Syracuse University from 2012 to 2013.

REFERENCES

- [1] P. K. Varshney, *Distributed Detection and Data Fusion*. Springer, New York, 1997.
- [2] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "Wireless sensor networks: a survey," *Computer networks*, vol. 38, no. 4, pp. 393–422, 2002.
- [3] V. V. Veeravalli and P. K. Varshney, "Distributed inference in wireless sensor networks," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 370, no. 1958, pp. 100–117, 2012.
- [4] J. N. Tsitsiklis, "Decentralized detection," in *Advances in Signal Processing*, H. V. Poor and J. B. Thomas, Eds. JAI Press, 1993, vol. 2, pp. 297–344.
- [5] R. Viswanathan and P. K. Varshney, "Distributed detection with multiple sensors i. fundamentals," *Proc. IEEE*, vol. 85, no. 1, pp. 54–63, 1997.

- [6] R. S. Blum, S. A. Kassam, and H. V. Poor, "Distributed detection with multiple sensors ii. advanced topics," *Proc. IEEE*, vol. 85, no. 1, pp. 64–79, 1997.
- [7] B. Chen, L. Tong, and P. K. Varshney, "Channel-aware distributed detection in wireless sensor networks," *IEEE Signal Processing Magazine*, vol. 23, no. 4, pp. 16–26, 2006.
- [8] J. Chamberland and V. Veeravalli, "Wireless sensors in distributed detection applications," *IEEE Signal Processing Magazine*, vol. 24, no. 3, pp. 16–25, 2007.
- [9] L. Cheng, C. Wu, Y. Zhang, H. Wu, M. Li, and C. Maple, "A survey of localization in wireless sensor network," *International Journal of Distributed Sensor Networks*, vol. 2012, pp. 1–12, 2012.
- [10] N. Patwari, J. Ash, S. Kyperountas, A. Hero, R. Moses, and N. Correal, "Locating the nodes: cooperative localization in wireless sensor networks," *IEEE Signal Processing Magazine*, vol. 22, no. 4, pp. 54–69, 2005.
- [11] R. R. Brooks, P. Ramanathan, and A. M. Sayeed, "Distributed target classification and tracking in sensor networks," *Proceedings of the IEEE*, vol. 91, no. 8, pp. 1163–1171, 2003.
- [12] J. Tsitsiklis, "Decentralized detection by a large number of sensors," *Math. Control, Signals, Systems*, vol. 1, no. 2, pp. 167–182, 1988.
- [13] A. Perrig, R. Szewczyk, J. D. Tygar, V. Wen, and D. E. Culler, "Spins: security protocols for sensor networks," *Wirel. Netw.*, vol. 8, no. 5, pp. 521–534, Sep. 2002. [Online]. Available: <http://dx.doi.org/10.1023/A:1016598314198>
- [14] A. Perrig, J. Stankovic, and D. Wagner, "Security in wireless sensor networks," *Communications of the ACM*, vol. 47, no. 6, pp. 53–57, 2004.
- [15] C. Karlof, N. Sastry, and D. Wagner, "Tinysec: a link layer security architecture for wireless sensor networks," in *Proceedings of the 2nd international conference on Embedded networked sensor systems*, ser. SenSys '04. New York, NY, USA: ACM, 2004, pp. 162–175. [Online]. Available: <http://doi.acm.org/10.1145/1031495.1031515>
- [16] A. Vempaty, L. Tong, and P. Varshney, "Distributed inference with byzantine data," to appear in *IEEE Signal Process. Mag., Special Issue: Signal Processing for Cyber-security and Privacy*.
- [17] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, Jul. 1982. [Online]. Available: <http://doi.acm.org/10.1145/357172.357176>
- [18] G. N. Nayak and S. Samaddar, "Different flavours of man-in-the-middle attack, consequences and feasible solutions," in *Computer Science and Information Technology (ICCSIT), 2010 3rd IEEE International Conference on*, vol. 5, 2010, pp. 491–495.
- [19] S. Marano, V. Matta, and L. Tong, "Distributed detection in the presence of byzantine attacks," *IEEE Trans. Signal Process.*, vol. 57, no. 1, pp. 16–29, 2009.
- [20] A. S. Rawat, P. Anand, H. Chen, and P. K. Varshney, "Collaborative spectrum sensing in the presence of byzantine attacks in cognitive radio networks," *IEEE Trans. Signal Process.*, vol. 59, no. 2, pp. 774–786, 2011.
- [21] A. Vempaty, K. Agrawal, H. Chen, and P. Varshney, "Adaptive learning of byzantines' behavior in cooperative spectrum sensing," in *Proc. IEEE Wireless Communications and Networking Conf. (WCNC)*, 2011, pp. 1310–1315.
- [22] B. Kailkhura, S. Brahma, and P. K. Varshney, "Optimal byzantine attack on distributed detection in tree based topologies," in *Proc. of International Conference on Computing, Networking and Communications Workshops (ICNC-CPS)*, San Diego, USA, January 2013.
- [23] E. Soltanmohammadi, M. Orooji, and M. Naraghi-Pour, "Decentralized hypothesis testing in wireless sensor networks in the presence of misbehaving nodes," *IEEE Transactions on Information Forensics and Security*, vol. 8, no. 1, pp. 205–215, 2013.
- [24] A. Vempaty, O. Ozdemir, K. Agrawal, H. Chen, and P. Varshney, "Localization in wireless sensor networks: Byzantines and mitigation techniques," *IEEE Transactions on Signal Processing*, vol. 61, no. 6, pp. 1495–1508, 2013.
- [25] J.-F. Chamberland and V. V. Veeravalli, "Decentralized detection in sensor networks," *IEEE Trans. Signal Process.*, vol. 51, no. 2, pp. 407–416, 2003.
- [26] J. G. Proakis, *Digital signal processing: principles algorithms and applications*. Pearson Education India, 2001.
- [27] A. Ribeiro and G. B. Giannakis, "Bandwidth-constrained distributed estimation for wireless sensor networks - part i: Gaussian case," *IEEE Transactions on Signal Processing*, vol. 54, no. 3, pp. 1131–1143, 2006.



V. Sriram Siddharth (Sid) Nadendla V. Sriram Siddharth (Sid) Nadendla was born in Rajahmundry, India, 1986. He received B.E. degree in electronics and communication engineering from Sri Chandrasekharendra Saraswati Viswa Mahavidyalaya (SCSVMV University), Kanchipuram, India, in 2007, his M.S. degree in electrical engineering from Louisiana State University, Baton Rouge, LA, USA, in 2009. He is currently pursuing his Ph.D. in Electrical and Computer Engineering in the Department of Electrical Engineering and Computer

Science, Syracuse University, Syracuse, NY, USA. His current research interests are in security, distributed inference, wireless communications and networks, resource allocation, game theory and mechanism design.

During his stay at SCSVMV University, he was awarded the University Gold-Medal along with his Bachelors degree in 2007. In addition to annual merit-scholarships, he was also a recipient of Dr. S. Subbulakshmi Endowment Cash Prize and Dr. S. Suryanarayanan Endowment Cash Prize for securing the Highest Marks in the university examinations, SCSVMV(DU), India in the period 2003-07. During 2007-09, he worked as a teaching assistant in the Department of Electrical and Computer Engineering, Louisiana State University. From August 2009, he is working as a research assistant at the Center for Advanced Systems Engineering (CASE), Syracuse University. He was a research intern at ANDRO Computational Solutions, LLC, Rome, NY in Summer 2013. He is a graduate student member of IEEE and serves as a reviewer to IEEE Transactions in Wireless Communications and IET Communications.



Pramod K. Varshney Pramod K. Varshney was born in Allahabad, India, on July 1, 1952. He received the B.S. degree in electrical engineering and computer science (with highest honors), and the M.S. and Ph.D. degrees in electrical engineering from the University of Illinois at Urbana-Champaign in 1972, 1974, and 1976 respectively.

During 1972-76, he held teaching and research assistantships at the University of Illinois. Since 1976 he has been with Syracuse University, Syracuse, NY where he is currently a Distinguished Professor of Electrical Engineering and Computer Science and the Director of CASE: Center for Advanced Systems and Engineering. He served as the Associate Chair of the department during 1993-96. He is also an Adjunct Professor of Radiology at Upstate Medical University in Syracuse, NY. His current research interests are in distributed sensor networks and data fusion, detection and estimation theory, wireless communications, physical layer security and image processing. He has published extensively. He is the author of *Distributed Detection and Data Fusion*, published by Springer-Verlag in 1997. He has served as a consultant to several major companies.

While at the University of Illinois, Dr. Varshney was a James Scholar, a Bronze Tablet Senior, and a Fellow. He is a member of Tau Beta Pi and is the recipient of the 1981 ASEE Dow Outstanding Young Faculty Award. He was elected to the grade of Fellow of the IEEE in 1997 for his contributions in the area of distributed detection and data fusion. He was the guest editor of the special issue on data fusion of the *Proceedings of the IEEE*, January 1997. In 2000, he received the Third Millennium Medal from the IEEE and Chancellor's Citation for exceptional academic achievement at Syracuse University. He is the recipient of the IEEE 2012 Judith A. Resnik Award. He is on the editorial board of *Journal on Advances in Information Fusion*. He was the President of International Society of Information Fusion during 2001.



Yunghsiang S. Han Yunghsiang S. Han (S'90-M'93-SM'08-F'11) was born in Taipei, Taiwan, 1962. He received B.Sc. and M.Sc. degrees in electrical engineering from the National Tsing Hua University, Hsinchu, Taiwan, in 1984 and 1986, respectively, and a Ph.D. degree from the School of Computer and Information Science, Syracuse University, Syracuse, NY, in 1993. He was from 1986 to 1988 a lecturer at Ming-Hsin Engineering College, Hsinchu, Taiwan. He was a teaching assistant from 1989 to 1992, and a research associate in the School

of Computer and Information Science, Syracuse University from 1992 to 1993. He was, from 1993 to 1997, an Associate Professor in the Department of Electronic Engineering at Hua Fan College of Humanities and Technology, Taipei Hsien, Taiwan. He was with the Department of Computer Science and Information Engineering at National Chi Nan University, Nantou, Taiwan from 1997 to 2004. He was promoted to Professor in 1998. He was a visiting scholar in the Department of Electrical Engineering at University of Hawaii at Manoa, HI from June to October 2001, the SUPRIA visiting research scholar in the Department of Electrical Engineering and Computer Science and CASE center at Syracuse University, NY from September 2002 to January 2004 and July 2012 to June 2013, and the visiting scholar in the Department of Electrical and Computer Engineering at University of Texas at Austin, TX from August 2008 to June 2009. He was with the Graduate Institute of Communication Engineering at National Taipei University, Taipei, Taiwan from August 2004 to July 2010. From August 2010, he is with the Department of Electrical Engineering at National Taiwan University of Science and Technology as Chair professor. His research interests are in error-control coding, wireless networks, and security.

Dr. Han was a winner of the 1994 Syracuse University Doctoral Prize and a Fellow of IEEE. One of his papers won the prestigious 2013 ACM CCS Test-of-Time Award in cybersecurity.