

Capacity of Cooperative Fusion in the Presence of Byzantine Sensors

Oliver Kosut and Lang Tong

Abstract—The problem of cooperative fusion in the presence of both Byzantine sensors and misinformed sensors is considered. An information theoretic formulation is used to characterize the Shannon capacity of sensor fusion. It is shown that when there are fewer Byzantine sensors than honest sensors, the effect of Byzantine attack can be entirely mitigated, and the fusion capacity is identical to that when all sensors are honest. However, when at least as many sensors are Byzantine as are honest, the Byzantine sensors can completely defeat the sensor fusion so that no information can be transmitted reliably. A capacity achieving transmit-then-verify strategy is proposed for the case that fewer sensors are Byzantine than honest, and its error probability and coding rate is analyzed by using a Markov decision process modeling of the transmission protocol.

Index Terms—Sensor Fusion, Byzantine Attack, Shannon Capacity, Network Security.

I. INTRODUCTION

WIRELESS sensor networks are not physically secure; they are vulnerable to various attacks. For example, sensors may be captured and analyzed such that the attacker gains inside information about the communication scheme and networking protocols. The attacker can then reprogram the compromised sensors and use them to launch the so-called Byzantine attack. This paper presents an information theoretic approach to sensor fusion in the presence of Byzantine sensors.

A. Cooperative Sensor Fusion

We consider the problem of cooperative sensor fusion as illustrated in Fig. 1 where the fusion center extracts information from a sensor field. By cooperative fusion we mean that sensors first reach a consensus among themselves about the fusion message. They then deliver the agreed message to the fusion center collaboratively. We will not be concerned with how sensors reach consensus in this paper; see *e.g.*, [1]. We focus instead on achieving the maximum rate of sensor fusion.

The sensor fusion problem is trivial if the consensus is perfect, *i.e.*, all the sensors agree on the same fusion message.

This work is supported in part by the National Science Foundation under award CCF-0635070, the U. S. Army Research Laboratory under the Collaborative Technology Alliance Program DAAD19-01-2-0011, and TRUST (The Team for Research in Ubiquitous Secure Technology) sponsored by the National Science Foundation under award CCF-0424422, and the following organizations: Cisco, ESCHER, HP, IBM, Intel, Microsoft, ORNL, Qualcomm, Pirelli, Sun and Symantec, and the U. S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0011. The U. S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.




Fig. 1. Cooperative sensor fusion in the presence of Byzantine sensors.

If the fusion center can only communicate with one sensor at a time, and there is no limit on how many times a sensor can transmit (*i.e.*, no energy constraints), there is no difference between having a single sensor delivering the message and having any number of sensors transmitting the message collaboratively. The capacity of such an ideal fusion is given by the classical Shannon theory

$$C = \max_{p(x)} I(X; Y) \quad (1)$$

where X is the transmitted symbol by a sensor, Y the received symbol, and $p(x)$ the distribution used to generate the codebook. Also for this case, even if there is a feedback channel from the fusion center to sensors, the capacity does not increase [2].

Cooperative fusion becomes important if consensus cannot be reached, *i.e.*, there is a probability $\beta > 0$ that a particular sensor is misinformed about what message to transmit. Thus there is a positive probability that a particular sensor communicating with the fusion center is delivering the wrong message. It is no longer obvious what the capacity of sensor fusion is. In [3], a number of sensor fusion models are considered, and the fusion capacity is obtained for several cases. Most relevant to this paper is the fusion model in which there is a feedback channel from the fusion center to individual sensors, and the fusion center polls specific sensors for transmissions. Optimized among all polling strategies, it is shown that, for any $\beta < 1$, the fusion capacity is also given by C in (1). The strategy given in [3] can be characterized as “identify-then-transmit” by first using an asymptotically negligible number of transmissions to identify a sensor that

is correctly informed then letting that sensor transmit the entire codeword.

B. Byzantine Attack and Related Work

The problem considered in this paper is when a fraction β of sensors are Byzantine sensors and a fraction γ are misinformed. The goal of the Byzantine sensors is to disrupt the sensor fusion collaboratively.

As in [3], we assume misinformed sensors behave honestly, but with a randomly chosen message. Byzantine sensors, however, are much more malicious. We assume they have full knowledge of the system and impose no restriction on what they can transmit. In particular, Byzantine sensors know the transmission strategy including the codebook and the polling strategy of the fusion center. They also know, of course, the correct fusion message. Unlike misinformed sensors which are required to pick a message, albeit an incorrect one, and then stick with it, Byzantine sensors can be malicious sometimes and behave in other times as honest sensors in order to evade detection by the fusion center. Furthermore, they can coordinate among themselves (unknown to both the honest sensors and the fusion center) to launch the so-called Byzantine attack. As a result, the capacity achieving coding and transmission strategies developed in [3] are no longer applicable.

The notion of Byzantine attack has its root in the Byzantine generals problem [4], [5] in which a clique of traitorous generals conspire to prevent loyal generals from forming consensus. It was shown in [4] that consensus in the presence of Byzantine attack is possible if and only if less than a third of the generals are traitorous. Relaxing the strict definition of consensus of [4], Pfitzmann and Waidner uses an information theoretic approach to show that the Byzantine generals problem can be solved for an arbitrarily large fraction of traitorous nodes [6]. These and other Byzantine consensus results [1] are relevant to the current paper only in that they deal with the consensus process prior to sensor fusion.

Countering Byzantine attacks in communication networks has also been studied in the past by many authors. See the earlier work of Perlman [7] and also more recent review [8], [9]. An information theoretic network coding approach to Byzantine attack is presented in [10]. Karlof and Wagner [11] consider routing security in wireless sensor networks. They introduce different kinds of attacks and analyze security risks of all major existing sensor network routing protocols. Countermeasures and design considerations for secure routing in sensor networks are also discussed. It is shown that cryptography alone is not enough; careful protocol design is necessary.

There has been limited attempt in dealing with Byzantine attacks for sensor fusion. The problem of optimal Byzantine attack of sensor fusion for distributed detection is considered in [12] where the authors show that exponentially decaying detection error probabilities can still be maintained if and only if the fraction of Byzantine sensors is less than half. A witness-based approach to sensor fusion is proposed by Du

et. al. [13] where the fusion center and a set of witnesses jointly authenticate the fusion data by the use of the Message Authentication Code. The authors of [13] are concerned with the trustworthiness of the fusion center. In contrast, we address the problem of sensor fusion with malicious sensors attacking the fusion center from within.

C. Main Result and Organization

The main result of this paper is to show that, if polling of the fusion center is allowed, and the polling is perfect, the capacity of sensor fusion in the presence of Byzantine attack is again C in (1) when the fraction of Byzantine sensors is less than that of the honest sensors, and 0 otherwise. The condition that there are fewer Byzantine sensors than honest sensors can be written $\beta < 1 - \beta - \gamma$, or $2\beta < 1 - \gamma$.

The converse of the result holds trivially for $2\beta < 1 - \gamma$ because the capacity of the sensor fusion in the absence of Byzantine and misinformed sensors is C . For $2\beta \geq 1 - \gamma$, we show that it is possible for the Byzantine sensors to completely defeat the fusion center and honest sensors by acting in such a way that there are two groups of sensors of exactly the same size, one acting honestly with the true message, the other acting honestly but with a false message. It is thus impossible for the fusion center to distinguish the set transmitting the true message from the set transmitting the false one, so it cannot decode the true message with probability more than $\frac{1}{2}$.

To show the achievability for $2\beta < 1 - \gamma$, we propose a transmission and coding strategy different from that for just misinformed sensors [3], for which the capacity achieving strategy can be called “identify-then-transmit”, where the fusion center first identifies an honest sensor, then receives the entire message from that sensor. Here we must deal with the situation in which a Byzantine sensor may pretend to be an honest sensor. The key idea is one of “transmit-then-verify”. Specifically, we first commit a sensor (Byzantine, misinformed, or honest) to transmit part of a codeword and then verify if the sensor is trustworthy. After a sensor has transmitted, the fusion center verifies the transmission using a random binning procedure. Under this procedure, a Byzantine sensor either has to act honestly or reveal with high probability its identity. We then have to show that the overhead in the verification diminishes as the length of the codeword increases.

This paper is organized as follows. In Section II, we present models for sensors, communication channels, and network setup. The main result is given and sketch of proofs are presented in Section III. We conclude in Section IV.

II. MODEL AND DEFINITIONS

A. Fusion Network and Communication Channels

A sensor is *Byzantine* if it can behave arbitrarily. A sensor is *honest* if it behaves only according to the specified protocol. A sensor is *misinformed* if it behaves exactly like an honest sensor, but with a random message selected uniformly from the set of all possible messages, independent of all other misinformed sensors, and independent of the true message.

Let β be the probability that a randomly selected sensor is Byzantine and γ be the probability that a randomly selected sensor is misinformed. With probability $1 - \beta - \gamma$, a randomly chosen sensor is honest. We assume that the sensor network is large in the sense that there are an infinite number of elements. This assumption ensures that the probability that there are no honest nodes is zero.

Sensors can communicate with the fusion center directly, and the transmissions are time slotted. We assume that the uplink channel from each sensor to the fusion center is a Discrete Memoryless Channel (DMC) $\{\mathcal{X}, \mathcal{Y}, q(y|x)\}$ where \mathcal{X} is the input alphabet, \mathcal{Y} the output alphabet, and $q(y|x)$ the transition probability of the channel. The assumption of identical channel is restrictive and synchronization difficult when the network is large and the fusion center stationary. The assumed model is reasonable, however, if the fusion center is a mobile access point that can travel around the network, and a sensor only transmits to the fusion center when it is activated by and synchronized to the fusion center.

We assume that there is a polling channel from the fusion center to each sensor. Since the fusion center is not power limited, we assume the polling channel is error free with infinite capacity.

B. Transmission Protocol

Before sensor fusion starts, we assume that all honest sensors have, without error, agreed upon a fusion message $W \in \{1, \dots, M\}$ that is uniformly distributed. The code is in general variable length and dynamically generated, so there is no single fixed codebook. However, we assume that the sensors may have any number of fixed codebooks to use as pieces of the code.

The fusion center polls one node to transmit one symbol in each time slot. At time t , the fusion center polls node K_t to transmit a symbol X_t . The symbol received by the fusion center is then Y_t . The fusion center may choose K_t based on previously received symbols Y^{t-1} and polling history K^{t-1} . Since the polling channel has infinite capacity, K_t may choose X_t based on all symbols previously received by the fusion center Y^{t-1} , the polling history K^{t-1} , and anything else the fusion center chooses to send to it. It may also base X_t on all previous transmissions that it has made itself, but not those made by other sensors. Honest sensors, of course, also have access to the message W . We assume that Byzantine sensors also know the fusion message, but misinformed sensors only have access to their own independently generated message.

If a sensor is Byzantine, it may also base its choice of X_t on all transmitted symbols, including those sent by honest sensors, and any additional information the fusion center sends to any sensor. We also assume that the Byzantine sensors know the algorithm the fusion center and honest sensors are using, and that they may communicate securely among themselves with zero error.

After the fusion center receives Y_t , it decides whether to continue polling based on Y^t and K^t . If it decides to continue, then it moves on to the next time slot $t + 1$ and

starts the polling step again. Otherwise, it decodes based on collected observations.

C. Achievable Rates and Capacity

Let N be the random variable representing the total number of symbols sent in a coding session. Once the fusion center decides it is done polling, it decodes the global message based on Y^N and K^N . The decoded message is denoted by $\hat{W} \in \{1, \dots, M\}$. A decoding error occurs if $\hat{W} \neq W$.

The *rate* of a code is defined as

$$R \triangleq \frac{\log(M)}{\mathbb{E}(N)},$$

where M is the number of messages and $\mathbb{E}(N)$ is the expected number of symbols transmitted during a coding session. The probability of error is defined as $P_e \triangleq \Pr(\hat{W} \neq W)$, where W is the message, uniformly selected from $\{1, \dots, M\}$, and \hat{W} is the decoded message. P_e will in general depend on the actions of the Byzantine sensors. A rate R is called *achievable* if for any given error $\epsilon > 0$ and any choice of actions by the Byzantine sensors, there exists a code with rate larger than $R - \epsilon$ and probability of error less than ϵ . The *capacity* of this system is defined as the maximum of all achievable rates.

III. FUSION CAPACITY

The main result of this paper is given by the following theorem that characterizes the capacity for the fusion network described in Sec II.

Theorem: The capacity of this system is

$$C^{\text{byz}} = \begin{cases} C, & \text{if } 2\beta < 1 - \gamma \\ 0, & \text{if } 2\beta \geq 1 - \gamma \end{cases}$$

where C is defined in (1). In particular, if $\gamma = 0$, $C^{\text{byz}} = C$ if and only if $\beta < 1/2$.

The proof of this theorem follows. In Subsection III-A, we prove the converse. In Subsection III-B, we describe the coding strategy used to prove achievability. In Subsection III-C, we define some error events and discuss the error probability. Finally, in Subsection III-D, we discuss the rate of this coding scheme.

A. Converse

Suppose that $\beta = \gamma = 0$ and that all the sensors may communicate with each other with zero error. Certainly these assumptions cannot decrease the capacity for any β and γ . Since the sensors can communicate with each other, we can think of the entire sensor network as a single encoder for the DMC with perfect feedback, because the sensors are allowed to know all previously received symbols by the fusion center. Thus under these assumptions this system reduces to a point-to-point DMC with perfect feedback. In that system, the feedback does not increase capacity [2], so the capacity is C . Thus, the capacity of the sensor network with Byzantine sensors cannot have capacity greater than this, so $C^{\text{byz}} \leq C$ for all β and γ .

Next we show that if $2\beta \geq 1 - \gamma$, then $C^{\text{byz}} = 0$. To do this, we will show that for any algorithm to be used by the fusion center and honest sensors, the Byzantine sensors will be able to make it impossible for the probability of error to be made arbitrarily small. The scheme performed by the Byzantine sensors to accomplish this is as follows. They divide themselves into two groups, one with $\frac{1-\gamma}{2}$ of the sensors, and one with $\beta - \frac{1-\gamma}{2}$ of the sensors. The sensors in the latter group act exactly like honest sensors. Since there is no way for the honest sensors to know anything that the Byzantine sensors do not, it will be impossible to distinguish an honest sensor from a Byzantine sensor acting honestly. The sensors in the former group also act exactly like honest sensors, but with a message different from the true one. There are now three groups of sensors: the misinformed sensors, sensors that act honestly with the true message, and sensors that act honestly with an incorrect message. The second group is made up of the honest sensors and the Byzantine sensors that act honestly. The fusion center will be able to distinguish the misinformed group from the other groups, but since the second and third groups both contain exactly a $\frac{1-\gamma}{2}$ fraction of the sensors, no matter what the fusion center does, it will not be able to determine which group is reporting the true message and which group is reporting the false one, so it will not be able to decode the true message with probability greater than $\frac{1}{2}$. Therefore the converse of the theorem holds.

B. Coding Strategy

To prove the direct part of the theorem, we first describe the coding strategy that will achieve this rate. The coding scheme can be described as a “transmit-then-verify” procedure. In other words, first we ask a sensor to send part of the message to the fusion center. After that, the fusion center polls other sensors to verify whether the received information is correct. Thus, if a Byzantine sensor is selected to transmit the message, it can send erroneous information, but then with high probability it will be discovered to be erroneous in the “verify” step. The Byzantine sensor can send the true information, but then it will be verified, so the fusion center now has that information, and knows it to be correct. As long as the fusion center always verifies any information it receives, the Byzantine sensors can never get any false information through. The best they can do is to prolong the coding process, but we will show that this additional overhead can be made to be negligible.

The coding strategy is as follows. We first break the message up into v chunks, such that each chunk contains an equal part of the information in the message, and the message will be perfectly reconstructible given all the chunks. These chunks could be, for example, the v digits representing the message W when it is written as a number in a particular base. The fusion center will try to obtain the v chunks one at a time, and verify that each chunk obtained is from an honest transmission.

Next we describe the two codebooks to be used in the uplink transmission over the DMC $q(y|x)$. Take any $\epsilon > 0$

and $R < C$. Let the number of possible messages $M = 2^{nR}$, so that the message set is $\{1, \dots, 2^{nR}\}$ and the set of all possible chunks is $\{1, \dots, 2^{nR/v}\}$. The first codebook \mathcal{G}_1 is a $(2^{nR/v}, n/v, \epsilon)$ code to transmit the chunk, where (M, n, ϵ) represents a code over the DMC with M messages, n channel uses, and probability of error less than ϵ . When a sensor is requested to transmit, say, the i th chunk of the message, an honest or misinformed sensor will use \mathcal{G}_1 to transmit the i th chunk of its message. A Byzantine sensor can choose to act honestly and use \mathcal{G}_1 to transmit the correct chunk, or it can transmit any other signal.

The second codebook \mathcal{G}_2 is a (j, l, ϵ) code used by the sensor in the verification process. Specifically, to verify if a transmission represents correct information, the fusion center uses a random binning technique. It distributes all possible chunks into j bins and broadcasts the bin index of each possible chunk to the sensors. The fusion center then asks k sensors to transmit the bin index of the particular chunk that the fusion center is verifying. An honest or misinformed sensor will transmit the bin index of its chunk to the fusion center using this second codebook \mathcal{G}_2 . A Byzantine sensor, if requested for the index, again can transmit arbitrarily, including acting honestly by using \mathcal{G}_2 to transmit the correct index. For fixed j , the code length l is chosen sufficiently long for transmitting the bin index accurately over the DMC. The numbers j and k are functions of decoding error ϵ and are chosen sufficiently large to ensure the fidelity of verification but not large enough to penalize the rate. We comment on the selection of them in Section III-C.

The detailed transmission protocol is as follows.

- 0) The fusion center randomly selects a sensor to transmit the next chunk (starting at the first chunk).
- 1) If the selected sensor is honest, it transmits the entire chunk using the codebook \mathcal{G}_1 . (If the selected sensor is Byzantine, it can act arbitrarily).
- 2) The fusion center randomly places each element in the set of all possible chunks into one of j bins. It then randomly selects k sensors, and sends the binning to each of them. That is, it informs the k sensors which of the $j^{2^{nR/v}}$ possible mappings from $\{1, \dots, 2^{nR/v}\}$ to $\{1, \dots, j\}$ it has selected. Each of those k sensors then sends the bin index of its chunk back to the fusion center using code \mathcal{G}_2 .
- 3) If the plurality of the k received bin indices match the bin index of the chunk that was received in step (1), the fusion center *accepts* that chunk. Otherwise it *declines* it.
- 4) If the chunk was accepted, the fusion center keeps the same sensor selected and moves on to the next chunk (go to step 1). If it was declined, the fusion center randomly selects a new sensor and tries again with the same chunk (step 0).
- 5) Polling stops when all chunks have been received and accepted. To complete the coding process, the fusion center extracts the original message from the v accepted chunks.

Note that each time we run through steps (1) through (4), we use the channel $n/v + kl$ times.

In step (2), we have used a random binning procedure. This is different from the way such a procedure is often used, such as in a common proof of the Slepian-Wolf theorem [2]. There, random binning is used as a technique to show that at least one code satisfying certain properties must exist. Here, we actually construct and use an entirely new random binning *every time* we do step (2). This is necessary because if we used some fixed or deterministic binning, the Byzantine sensors would know the binning to be used beforehand. Thus, if a Byzantine sensor is selected to transmit a chunk in step (1), it could transmit an incorrect chunk falling into the same bin as the real chunk, which would make the verification process useless. The probability that the Byzantine sensor can find such a misleading chunk must be small, so we need *dynamic random binning*.

If $\beta + \gamma < 1/2$ (in particular, if $\gamma = 0$), then the following simplification can be made to step (2). Instead of constructing a random binning, each of the k sensors sends to the fusion center whether or not the chunk that the fusion center has just received matches the true chunk. Thus \mathcal{G}_2 can be reduced to a $(2, l, \epsilon)$ code, since only one of two messages—that the chunk is correct or not—must be communicated. This only works when more than half of the sensors are honest, since all the misinformed and Byzantine sensors will likely report that the true chunk is incorrect, and we expect more than half of the k sensors to be honest only if more than half of all the sensors are honest. The rest of the proof will assume use of the above scheme and not the simplification, but the argument is almost entirely identical.

C. Error Events and Error Probability Analysis

We show next that, with appropriately chosen n, v, j, l, k in the two codebooks, the probability that a message is decoded incorrectly goes to zero, and the decoding process will end with an average number of transmissions approximately $n + O(\epsilon n)$. Thus with a message set of size 2^{nR} , and $R \geq C - \epsilon$, we have the proof of the main theorem.

To analyze the probability of error, we need to define some events. Events $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ are the most basic ways in which errors can occur. $\mathcal{B}_1, \mathcal{B}_2, \mathcal{C}$ have to do with the conclusion the fusion center reaches, and thus determine how the coding will progress.

- \mathcal{A}_1 : A coding error occurs in step (1), *i.e.*, the transmitted chunk is different from the decoded one.
- \mathcal{A}_2 : Of the k bin indices that are decoded in step (2), the plurality do not equal the bin index for the true chunk.
- \mathcal{A}_3 : For a given pair of distinct chunks, they are both put into the same bin in step (2).
- \mathcal{B}_1 : The chunk is declined in step (3).
- \mathcal{B}_2 : A chunk is accepted in step (3) and that chunk is not the true one.
- \mathcal{C} : The true chunk is transmitted in step (1).

The following lemma bounds the probabilities of events relevant to the error analysis.

Lemma 1: Define

$$\begin{aligned} p_1 &\triangleq \Pr(\mathcal{B}_1|\mathcal{C}), & p_2 &\triangleq \Pr(\mathcal{B}_2|\mathcal{C}^c), \\ p_3 &\triangleq \Pr(\mathcal{B}_2|\mathcal{C}). \end{aligned}$$

For sufficiently large j and k , and no matter what the Byzantine sensors do, $\Pr(\mathcal{A}_i) \leq \epsilon$ for $i = 1, 2, 3$, and

$$\begin{aligned} p_1 &\leq \Pr(\mathcal{A}_1) + \Pr(\mathcal{A}_2), & p_2 &\leq \Pr(\mathcal{A}_2) + \Pr(\mathcal{A}_3), \\ p_3 &\leq \Pr(\mathcal{A}_1)(\Pr(\mathcal{A}_2) + \Pr(\mathcal{A}_3)). \end{aligned}$$

Proof: Since \mathcal{G}_1 was constructed to have error probability less than ϵ , $\Pr(\mathcal{A}_1) \leq \epsilon$.

Now we show $\Pr(\mathcal{A}_2) \leq \epsilon$ for sufficiently large k . Consider one of the k sensors polled in step (2). Let α be the probability that the sensor is honest and a \mathcal{G}_2 error does not occur when it sends its bin index in step (2). Thus $\alpha = (1 - \beta - \gamma)(1 - \epsilon)$. Let β' be the probability that the sensor is Byzantine or a \mathcal{G}_2 error occurs, so $\beta' = \beta + \epsilon(1 - \beta)$. Let γ' be the probability that the sensor is misinformed and a \mathcal{G}_2 error does not occur, so $\gamma' = \gamma(1 - \epsilon)$. Let the random variable A be the number of the k sensors that are honest without a \mathcal{G}_2 error, B be the number that are Byzantine or have a \mathcal{G}_2 error, and C be the number that are misinformed without a \mathcal{G}_2 error. By the law of large numbers, there is some k large enough such that

$$\Pr\left(\left|\frac{A}{k} - \alpha\right| < \epsilon, \left|\frac{B}{k} - \beta'\right| < \epsilon, \left|\frac{C}{k} - \gamma'\right| < \epsilon\right) > 1 - \frac{\epsilon}{2}. \quad (2)$$

If $\beta + \gamma < 1/2$, then for sufficiently small ϵ , $\alpha - \epsilon > 1/2$, so by (2) $\Pr(A > k/2) > 1 - \epsilon/2$, which means the probability of the majority—and hence the plurality—of the decoded bin indices being incorrect is less than ϵ .

Now we consider the $\beta + \gamma \geq 1/2$ case. In order for an incorrect bin index to achieve the plurality, some combination of the Byzantine and misinformed sensors all transmitting the same bin index must form a group larger than the group of honest sensors that are polled. If we assume the Byzantine sensors know the misinformed sensors' messages, the worst they can do is all choose to send the same bin index as that of the largest group of misinformed sensors with messages in the same bin. Thus, \mathcal{A}_2 will occur if there is a set of $A - B$ of the C misinformed sensors that all have the same message. Let $q(i, c)$ be the probability that, of a set of c misinformed sensors, at least one subset of i sensors have messages in the same bin. There are $\binom{c}{i}$ subsets of size i from a set of size c , and the chunks of any such subset of misinformed sensors have probability $1/j^{i-1}$ of all falling into the same bin. Thus

$$q(i, c) \leq \binom{c}{i} \frac{1}{j^{i-1}}. \quad (3)$$

From (2), $A \leq (\alpha - \epsilon)k$, $B \geq (\beta' + \epsilon)k$, or $C \geq (\gamma' + \epsilon)k$ with probability at most $\epsilon/2$. If not, recall $\Pr(\mathcal{A}_2) = q(A - B, C)$, and $A - B > (\alpha - \epsilon)k - (\beta' + \epsilon)k = (\alpha - \beta' + 2\epsilon)k$ while $C < (\gamma' + \epsilon)k$. Note that $q(i, c)$ is monotonically decreasing

in i and monotonically increasing in c . Thus

$$\begin{aligned} \Pr(\mathcal{A}_2) &\leq \frac{\epsilon}{2} + q(\lfloor(\alpha - \beta' - 2\epsilon)k\rfloor, \lceil(\gamma' + \epsilon)k\rceil) \\ &\leq \frac{\epsilon}{2} + \binom{\lceil(\gamma' + \epsilon)k\rceil}{\lfloor(\alpha - \beta' - 2\epsilon)k\rfloor} j^{-\lfloor(\alpha - \beta' - 2\epsilon)k\rfloor + 1} \quad (4) \\ &\leq \frac{\epsilon}{2} + j \binom{\lceil(\gamma' + \epsilon)k\rceil}{\lfloor(\alpha - \beta' - 2\epsilon)k\rfloor} \\ &\quad \binom{\lfloor(\alpha - \beta' - 2\epsilon)k\rfloor}{-\lfloor(\alpha - \beta' - 2\epsilon)k\rfloor + 1} \quad (5) \end{aligned}$$

$$\begin{aligned} &\leq \frac{\epsilon}{2} + j \left(\frac{(\gamma' - \alpha + \beta' + 3\epsilon)k + 3}{j} \right)^{(\alpha - \beta' - 2\epsilon)k} \\ &= \epsilon + j \left(\frac{(1 - 2\alpha + 3\epsilon)k + 3}{j} \right)^{(\alpha - \beta' - 2\epsilon)k} \quad (6) \end{aligned}$$

where (4) is from (3), (5) because $\binom{a}{b} \leq (a - b + 1)^b$, and (6) since $\alpha + \beta' + \gamma' = 1$. Note that $1 - 2\alpha > 0$ because $\beta + \gamma > 1/2$. For sufficiently small ϵ , $1 - 2\alpha + 3\epsilon < 1$, so for sufficiently large k , $(1 - 2\alpha + 3\epsilon)k + 3 \leq k$. Thus

$$\Pr(\mathcal{A}_2) \leq \frac{\epsilon}{2} + j \left(\frac{k}{j} \right)^{(\alpha - \beta' - 2\epsilon)k}.$$

Since $2\beta < 1 - \gamma$, for sufficiently small ϵ , $\alpha - \beta' - 2\epsilon > 0$. Thus if we let $2k \leq j \leq 3k$, then for large enough k , the second term above is less than $\epsilon/2$, so $\Pr(\mathcal{A}_2) \leq \epsilon$.

Next we show $\Pr(\mathcal{A}_3) \leq \epsilon$ for sufficiently large j . Since there are j bins, the probability that two different chunks are put into the same bin in step (2) is $1/j$. Thus if $j \geq 1/\epsilon$, $\Pr(\mathcal{A}_3) \leq \epsilon$. Note that achieving this condition on j as well as that above requires $k \geq \frac{1}{3\epsilon}$.

Note that p_1 is the probability that the received chunk is declined in step (3) given the true chunk was transmitted in step (1). One way for this to happen is for there to be a coding error in step (1), *i.e.*, \mathcal{A}_1 occurs, so the received chunk will not be the true chunk, so the polled sensors may not confirm it. Note that a coding error does not necessitate the chunk being declined, but it does cover a large set of the ways it could happen. If \mathcal{A}_1 does not occur, then the received chunk is the true one, so the chunk could only be declined if the majority of the bin indices received in step (2) do not match the true chunk, *i.e.*, \mathcal{A}_2 occurs. Thus

$$p_1 \leq \Pr(\mathcal{A}_1 \cup \mathcal{A}_2) \leq \Pr(\mathcal{A}_1) + \Pr(\mathcal{A}_2).$$

Next, p_2 is the probability that an incorrect chunk is accepted given that an incorrect chunk is transmitted in step (1). If more than half of the decoded bin indices are incorrect (\mathcal{A}_2), then those incorrect bin indices might confirm the incorrect chunk. If not, then the only way for the incorrect chunk to be accepted is for it to fall into the same bin as the true chunk (\mathcal{A}_3). Thus

$$p_2 \leq \Pr(\mathcal{A}_2 \cup \mathcal{A}_3) \leq \Pr(\mathcal{A}_2) + \Pr(\mathcal{A}_3).$$

Finally, p_3 is the probability that an incorrect chunk is accepted given that the correct chunk is transmitted in step (1). In order for this to happen, the decoded chunk must

not be the true one, so a coding error must occur (\mathcal{A}_1). In addition, for that decoded incorrect chunk to be accepted, more than half of the decoded bin indices must be incorrect (\mathcal{A}_2) or the incorrect chunk must fall into the same category as the real one (\mathcal{A}_3). Thus

$$p_3 \leq \Pr(\mathcal{A}_1) \Pr(\mathcal{A}_2 \cup \mathcal{A}_3) \leq \Pr(\mathcal{A}_1) (\Pr(\mathcal{A}_2) + \Pr(\mathcal{A}_3)).$$

■

For the following part of the analysis, we will assume that all misinformed sensors behave as Byzantine sensors. Certainly this cannot reduce the error probability, and as long as we continue using the bounds on the p_i s from Lemma 1, the analysis becomes simpler. Thus we will assume that a sensor is Byzantine with probability $\tilde{\beta} \triangleq \beta + \gamma$ and honest with probability $1 - \tilde{\beta}$.

As the coding scheme commences, it moves through a number of different states, depending on the number chunks the fusion center has received thus far, and whether the selected sensor is Byzantine. Depending on the exact sequence of events, the fusion center might remain at a certain state for some time, requesting the same chunk several times until it finds an honest sensor. The progress is probabilistic because every time the fusion center selects a sensor it might be Byzantine or honest, and every time it receives a transmission, a transmission error might or might not occur. In fact, the progress of the coding scheme can be modeled as a Markov process. In particular, it will be a Markov decision process, because a Byzantine sensor, if it is selected to transmit a chunk, has some choice about what to transmit. That choice will influence the probabilities of future events. The Markov decision process that we will use to analyze the error probability of this scheme is diagrammed in Fig. 2.

The process will have $2v + 3$ states. State i , for $i = 0, \dots, v$ represents the fusion center having successfully received i true chunks and the currently selected sensor is honest. State i' is the same except the currently selected sensor is Byzantine. Finally, state e represents the fusion center having accepted at least one false chunk. The decision for the Markov decision process will be whether a Byzantine sensor, if it is asked to send a chunk in step (1), chooses to send the true chunk or not. Thus a decision will only be made when a Byzantine sensor has been selected, *i.e.*, we are in one of the i' states.

States v , v' , and e will be terminal states, so an error will occur if we reach state e before state v or v' . Define

$$\begin{aligned} e_i &\triangleq \Pr(\text{error occurs starting from state } i), \\ e'_i &\triangleq \Pr(\text{error occurs starting from state } i'). \end{aligned}$$

In executing the Markov decision process, the Byzantine sensors make decisions to maximize the probability of error. At the very beginning of the coding scheme, we select a sensor which will be with probability $1 - \tilde{\beta}$ honest and probability $\tilde{\beta}$ Byzantine. Thus, the total probability of error is

$$P_e = (1 - \tilde{\beta})e_0 + \tilde{\beta}e'_0.$$




Fig. 2. The Markov decision process used to find the error probability. Dashed lines from a state represent the Byzantine sensor choosing to send erroneous information, and dotted lines represent the Byzantine sensor choosing to send true information.

From state i , with probability p_1 the chunk will be declined. The fusion center then selects a new sensor, which will be Byzantine with probability $\tilde{\beta}$ and honest with probability $1 - \tilde{\beta}$. Thus we transition to state i' with probability $p_1\tilde{\beta}$ and back to state i with probability $p_1(1 - \tilde{\beta})$. With probability p_3 , an incorrect chunk is accepted, so we transition to state e . Finally, with probability $1 - p_1 - p_3$ the true chunk is accepted, so we transition to state $i + 1$.

From state i' , the transition probabilities depend on the decision. If the Byzantine sensor chooses not to send the true chunk, then with probability p_2 the false chunk will be accepted, so we transition to state e . Otherwise, the fusion center selects a new sensor. Thus with probability $(1 - p_2)\tilde{\beta}$ we return to state i' , and with probability $(1 - p_2)(1 - \tilde{\beta})$ we transition to state i . If the Byzantine sensor decides to send the true chunk, then the transition probabilities are essentially the same as they were from state i ; with probability $p_1\tilde{\beta}$ we return to state i' , with probability $p_1(1 - \tilde{\beta})$ we transition to state i , with probability p_3 we transition to state e , and with probability $1 - p_1 - p_3$ we transition to state $i + 1'$.

From these transition probabilities, we see that

$$e_i = p_3 + (1 - p_1 - p_3)e_{i+1} + p_1\tilde{\beta}e'_i + p_1(1 - \tilde{\beta})e_i,$$

$$e'_i = \max\{p_2 + (1 - p_2)\tilde{\beta}e'_i + (1 - p_2)(1 - \tilde{\beta})e_i,$$

$$p_3 + (1 - p_1 - p_3)e'_{i+1} + p_1\tilde{\beta}e'_i + p_1(1 - \tilde{\beta})e_i\}.$$

The maximum represents the Byzantine sensors always making the decision that maximizes the error probabilities. In addition, if we arrive at either state v or v' , the fusion center has received the entire message without error, so $e_v = e_{v'} = 0$.

D. Code Rate

We also need to consider the rate of this code. To show that the rate can be made arbitrarily close to C , we need to show that the expected number of channel uses $\mathbb{E}(N)$ converges to n as ϵ goes to zero. Each time a chunk is transmitted (*i.e.*, each time we run through steps (1) to (4)), the channel is used $n/v + kl$ times. All we need to know is the expected number of chunks that are transmitted in the entire coding scheme. To find this, we will use a similar Markov decision process as the one described above. The only differences lie in the fact that we are not interested in whether an error occurs, only in how long it takes to finish. Thus we remove state e and redefine states i and i' to represent the fusion center having accepted i states, but with all of them not necessarily correct. Thus every time we would transition to state e , we actually transition somewhere else. For instance, if we are in state i' and the Byzantine sensors choose to send erroneous information, then with probability p_2 , the chunk is accepted, so we transition to state $i + 1'$ instead of e . Let q_i and q'_i be the expected number of steps made in the Markov decision process before reaching one of the terminal states (v or v') given that we start at state i or i' respectively and the Byzantine sensors make decisions that maximize the expected number of steps. Then

$$q_i = 1 + (1 - p_1)q_{i+1} + p_1\tilde{\beta}q'_i + p_1(1 - \tilde{\beta})q_i,$$

$$q'_i = \max\{1 + p_2q'_{i+1} + (1 - p_2)\tilde{\beta}q'_i + (1 - p_2)(1 - \tilde{\beta})q_i,$$

$$1 + (1 - p_1)q'_{i+1} + p_1\tilde{\beta}q'_i + p_1(1 - \tilde{\beta})q_i\}.$$
(7)

Again, $q_v = q'_v = 0$.

Lemma 2 (Average Code Length): There exist n , v , j , and k as functions of ϵ such that the error probability $P_e \rightarrow 0$ and the expected number of channel uses $\mathbb{E}(N) \rightarrow n$ as $\epsilon \rightarrow 0$.

Proof: Take j and k large enough for Lemma 1 to hold, and n and v such that

$$\frac{2}{\epsilon} \geq v \geq \frac{1}{\epsilon}, \quad n \geq \frac{klv}{\epsilon}.$$
(8)

We define f_i, f'_i for $i = 0, \dots, v$ as follows. Let $f_v \triangleq f'_v \triangleq 0$ and for $i < v$,

$$f_i \triangleq p_3 + (1 - p_1 - p_3)f_{i+1} + p_1\tilde{\beta}f'_i + p_1(1 - \tilde{\beta})f_i, \quad (9)$$

$$f'_{i,a} \triangleq p_2 + \tilde{\beta}f'_{i+1} + (1 - \tilde{\beta})f_i, \quad (10)$$

$$f'_{i,b} \triangleq p_3 + (1 - p_1 - p_3)f'_{i+1} + p_1\tilde{\beta}f'_i, \quad (11)$$

$$f'_i \triangleq \max\{f'_{i,a}, f'_{i,b}\}. \quad (12)$$

The only difference between f_i, f'_i and e_i, e'_i is that the $(1 - p_2)$ factors have been dropped from the second two terms in (10). Thus $e_i \leq f_i, e'_i \leq f'_i$, for all i . Fix some $i \in \{0, \dots, v-1\}$. If $f'_i = f'_{i,a}$, then by (10)

$$f'_i = \frac{p_2}{1 - \tilde{\beta}} + f_i. \quad (13)$$

Combining this with (9) gives

$$f_i = f_{i+1} + \frac{p_3}{1 - p_1} + \frac{p_1 p_2 \tilde{\beta}}{(1 - p_1)(1 - \tilde{\beta})}, \quad (14)$$

which with (13) produces

$$\begin{aligned} f'_i &= f_{i+1} + \frac{p_3}{1 - p_1} + \frac{p_1 p_2 \tilde{\beta}}{(1 - p_1)(1 - \tilde{\beta})} + \frac{p_2}{1 - \tilde{\beta}} \\ &= f_{i+1} + \frac{p_3}{1 - p_1} + \frac{p_2(1 - p_1(1 - \tilde{\beta}))}{(1 - p_1)(1 - \tilde{\beta})}. \end{aligned} \quad (15)$$

If $f'_i = f'_{i,b}$, then combining (9) with (11) gives

$$f'_i = \frac{p_3}{1 - p_1} + p_1(1 - \tilde{\beta})f_{i+1} + (1 - p_1(1 - \tilde{\beta}))f'_{i+1}. \quad (16)$$

Note that (15) and (16) are what f'_i would be if f'_i equaled $f'_{i,a}$ or $f'_{i,b}$ respectively. However, these expressions are not necessarily equal to $f'_{i,a}$ and $f'_{i,b}$, because we have used (9) to derive both of them, which contains the real value of f'_i . Still, because of the definition of f'_i in (12), the larger of (15) and (16) will be the true value of f'_i .

We will now show by induction that $f'_i = f'_{i,a}$ for $i = 0, \dots, v-1$. For $i = v-1$, since $f_v = f'_v = 0$, it is clear that the expression in (15) is larger than that in (16), so $f'_{v-1} = f'_{v-1,a}$. Now we assume that $f'_{i+1} = f'_{i+1,a}$ and show that $f'_i = f'_{i,a}$. By (13),

$$f'_{i+1} = \frac{p_2}{1 - \tilde{\beta}} + f_{i+1}.$$

Thus, if $f'_i = f'_{i,b}$, (16) becomes

$$\begin{aligned} f'_i &= \frac{p_3}{1 - p_1} + p_1(1 - \tilde{\beta})f_{i+1} \\ &\quad + (1 - p_1(1 - \tilde{\beta})) \left(\frac{p_2}{1 - \tilde{\beta}} + f_{i+1} \right) \\ &= \frac{p_3}{1 - p_1} + f_{i+1} + \frac{p_2(1 - p_1(1 - \tilde{\beta}))}{1 - \tilde{\beta}}. \end{aligned}$$

Since the expression in (15) is larger than this, $f'_i = f'_{i,a}$.

Therefore (14) holds for $i = 0, \dots, v-1$, so

$$f_i = \left(\frac{p_3}{1 - p_1} + \frac{p_1 p_2 \tilde{\beta}}{(1 - p_1)(1 - \tilde{\beta})} \right) (v - i). \quad (17)$$

Thus

$$\begin{aligned} P_e &= (1 - \tilde{\beta})e_0 + \tilde{\beta}e'_0 \\ &\leq (1 - \tilde{\beta})f_0 + \tilde{\beta}f'_0 \\ &= f_0 + \frac{p_2 \tilde{\beta}}{1 - \tilde{\beta}} \end{aligned} \quad (18)$$

$$= \frac{p_1 p_2 \tilde{\beta} + p_3(1 - \tilde{\beta})}{(1 - p_1)(1 - \tilde{\beta})} v + \frac{p_2 \tilde{\beta}}{1 - \tilde{\beta}} \quad (19)$$

$$\leq \frac{4\epsilon^2 \tilde{\beta} + 2\epsilon^2(1 - \tilde{\beta})}{(1 - 2\epsilon)(1 - \tilde{\beta})} \left(\frac{2}{\epsilon} \right) + \frac{2\epsilon \tilde{\beta}}{1 - \tilde{\beta}} \quad (20)$$

$$= \left(\frac{8\tilde{\beta} + 4(1 - \tilde{\beta})}{(1 - 2\epsilon)(1 - \tilde{\beta})} + \frac{2\tilde{\beta}}{1 - \tilde{\beta}} \right) \epsilon$$

where (18) is from (13), (19) is from (17), and (20) is from Lemma 1 and (8). Thus $P_e \rightarrow 0$ as $\epsilon \rightarrow 0$.

Now we analyze q_i, q'_i to find $\mathbb{E}(N)$. Combining the expression for q_i in (7) with either expression for q'_i in the maximum in (7) yields expressions of the form

$$q_i = 1 + \gamma + \delta q_{i+1} + (1 - \delta)q'_{i+1}, \quad (21)$$

$$q'_i = 1 + \gamma' + \delta' q_{i+1} + (1 - \delta')q'_{i+1}, \quad (22)$$

where $\gamma, \gamma' \geq 0$ and $\delta, \delta' \in [0, 1]$. The quantity γ represents the expected number of state transitions between states i and i' before moving on to state $i+1$ or $i+1'$, given that we start at state i' , and δ represents the probability that when we do transition away from states i and i' , we go to state $i+1$ and not $i+1'$. The quantities γ' and δ' are the same except starting at state i' . Obviously, the values of these will depend on which element of the maximum is larger, but for our current purposes it only matters that the expressions will have this form.

We will now show by induction that $q_i - q_{i+1} \geq 1$ and $q'_i - q'_{i+1} \geq 1$ for $i = 0, \dots, v-1$. First consider $i = v-1$. $q_v = q'_v = 0$, so by (21) and (22), $q_{v-1} = 1 + \gamma$ and $q'_{v-1} = 1 + \gamma'$. Thus $q_{v-1} - q_v \geq 1$ and $q'_{v-1} - q'_v \geq 1$. Now we assume that $q_{i+1} - q_{i+2} \geq 1$ and $q'_{i+1} - q'_{i+2} \geq 1$ and show that $q_i - q_{i+1} \geq 1$ and $q'_i - q'_{i+1} \geq 1$. By assumption and (21),

$$\begin{aligned} q_i - q_{i+1} &= \delta(q_{i+1} - q_{i+2}) + (1 - \delta)(q'_{i+1} - q'_{i+2}) \\ &\geq \delta + (1 - \delta) \\ &= 1. \end{aligned}$$

Similarly by (22),

$$\begin{aligned} q'_i - q'_{i+1} &= \delta'(q_{i+1} - q_{i+2}) + (1 - \delta')(q'_{i+1} - q'_{i+2}) \\ &\geq \delta' + (1 - \delta') \\ &= 1. \end{aligned}$$

Thus $q_i - q_{i+1} \geq 1$ and $q'_i - q'_{i+1} \geq 1$ for $i = 0, \dots, v-1$. In particular, $q'_{i+1} \leq q'_i - 1$.

Suppose the first element of the maximum is larger in (7). Then

$$\begin{aligned} q'_i &= 1 + p_2 q'_{i+1} + (1 - p_2) \tilde{\beta} q'_i + (1 - p_2)(1 - \tilde{\beta}) q_i \\ &\leq 1 + p_2 (q'_i - 1) + (1 - p_2) \tilde{\beta} q'_i + (1 - p_2)(1 - \tilde{\beta}) q_i. \end{aligned}$$

This can be rewritten

$$q'_i \leq \frac{1}{1-\tilde{\beta}} + q_i. \quad (23)$$

Now suppose the second element of the maximum is larger in (7). Then

$$\begin{aligned} q'_i &= 1 + (1-p_1)q'_{i+1} + p_1\tilde{\beta}q'_i + p_1(1-\tilde{\beta})q_i \\ &\leq 1 + (1-p_1)(q'_i - 1) + p_1\tilde{\beta}q'_i + p_1(1-\tilde{\beta})q_i. \end{aligned}$$

This can also be rewritten to (23), so (23) must hold no matter which value is larger in the maximum in (7). Thus

$$q_i \leq 1 + (1-p_1)q_{i+1} + p_1\tilde{\beta}\left(\frac{1}{1-\tilde{\beta}} + q_i\right) + p_1(1-\tilde{\beta})q_i.$$

This can be rewritten

$$q_i \leq 1 + \frac{p_1}{(1-p_1)(1-\tilde{\beta})} + q_{i+1},$$

so

$$q_i \leq \left(1 + \frac{p_1}{(1-p_1)(1-\tilde{\beta})}\right)(v-i). \quad (24)$$

Let V be the random variable denoting the total number of chunks that are requested in the entire coding session. Since we start at state 0 with probability $1-\tilde{\beta}$ and at state $0'$ with probability $\tilde{\beta}$,

$$\begin{aligned} \mathbb{E}(V) &= (1-\tilde{\beta})q_0 + \tilde{\beta}q'_0 \\ &\leq q_0 + \frac{\tilde{\beta}}{1-\tilde{\beta}} \end{aligned} \quad (25)$$

$$\leq \left(1 + \frac{p_1}{(1-p_1)(1-\tilde{\beta})}\right)v + \frac{\tilde{\beta}}{1-\tilde{\beta}} \quad (26)$$

where (25) is from (23) and (26) is from (24). Thus

$$\begin{aligned} \mathbb{E}(N) &= \mathbb{E}(V) \left(\frac{n}{v} + kl\right) \\ &\leq \left[\left(1 + \frac{p_1}{(1-p_1)(1-\tilde{\beta})}\right)v + \frac{\tilde{\beta}}{1-\tilde{\beta}}\right] \left(\frac{n}{v} + kl\right) \\ &= n \left[1 + \frac{p_1}{(1-p_1)(1-\tilde{\beta})} + \frac{\tilde{\beta}}{1-\tilde{\beta}} \frac{1}{v}\right. \\ &\quad \left. + \left(1 + \frac{p_1}{(1-p_1)(1-\tilde{\beta})}\right) \frac{klv}{n} + \frac{\tilde{\beta}}{1-\tilde{\beta}} \frac{kl}{n}\right] \\ &\leq n \left[1 + \frac{2\epsilon}{(1-2\epsilon)(1-\tilde{\beta})} + \frac{\tilde{\beta}}{1-\tilde{\beta}} \epsilon\right. \\ &\quad \left. + \left(1 + \frac{2\epsilon}{(1-2\epsilon)(1-\tilde{\beta})}\right) \epsilon + \frac{\tilde{\beta}}{1-\tilde{\beta}} \epsilon^2\right] \quad (27) \\ &= n \left[1 + \left(\frac{2(1+\epsilon)}{(1-2\epsilon)(1-\tilde{\beta})} + \frac{\tilde{\beta}(1+\epsilon)}{1-\tilde{\beta}} + 1\right) \epsilon\right] \end{aligned}$$

where (27) is from Lemma 1 and (8). Thus $\mathbb{E}(N) \rightarrow n$ as $\epsilon \rightarrow 0$. ■

Therefore the rate of this code,

$$\frac{nR}{\mathbb{E}(N)},$$

converges to R as ϵ goes to 0. Thus C is achievable.

IV. CONCLUSION

We showed in this paper that, by cooperative sensor fusion, the presence of Byzantine and misinformed sensors can be completely mitigated when the Byzantine sensor population is less than that of the honest sensors, but no information can be transmitted when at least as many sensors are Byzantine as are honest. We proposed a “transmit-then-verify” scheme that forces a Byzantine sensor to either act honestly or reveal its Byzantine identity. The key of this idea is the use of random binning in sensor polling. Note that the random binning in our strategy is not a random coding argument; it is an actual randomized transmission protocol. However, this random binning is not needed when more than half of the sensors are honest.

The network does not have to contain infinite number of sensors. For a finite size network, we will assume that a deterministic β fraction of the sensors are Byzantine and γ fraction are misinformed. In that case, all the sensors can be polled when verifying a transmission (*i.e.*, k can be set to the total number of sensors). Thus information will always be correctly verified, because the polled honest sensors will outnumber the polled Byzantine sensors. This requires a constant and hence asymptotically negligible number of channel uses, so polling every sensor instead of a random subset does not effect the rate.

REFERENCES

- [1] M. Barborak, M. Malek, and A. Dabhura, “The consensus problem in fault-tolerant computing,” *ACM Computing Surveys*, vol. 25, no. 2, pp. 171–220, 1993.
- [2] T. Cover and J. Thomas, *Elements of Information Theory*. John Wiley & Sons, Inc., 1991.
- [3] Z. Yang and L. Tong, “Cooperative sensor networks with misinformed nodes,” *IEEE Trans. Inform. Theory*, vol. 51, pp. 4118–4133, Dec. 2005.
- [4] L. Lamport, R. Shostak, and M. Pease, “The byzantine generals problem,” *ACM Transactions on Programming Languages and Systems*, vol. 4, pp. 382–401, July 1982.
- [5] D. Dolev, “The Byzantine generals strike again,” *Journal of Algorithms*, vol. 3, no. 1, pp. 14–30, 1982.
- [6] B. Pfitzmann and M. Waidner, “Information theoretic pseudosignatures and Byzantine agreement for $t \geq n/3$,” Tech. Rep. RZ2882, IBM Research Report, Nov 1996.
- [7] R. Perlman, *Network Layer Protocols with Byzantine Robustness*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, August 1988.
- [8] L. Zhou and Z. J. Haas, “Securing ad hoc networks,” *IEEE Network Magazine*, vol. 13, pp. 24–30, Nov/Dec 1999.
- [9] Y. Hu and A. Perrig, “Security and privacy in sensor networks,” *IEEE Security and Privacy Magazine*, vol. 2, pp. 28–39, 2004.
- [10] T. Ho, B. Leong, R. Koetter, M. Médard, M. Effron, and D. Karger, “Byzantine modification detection in multicast networks using randomized network coding,” in *IEEE Proc. Intl. Sym. Inform. Theory*, p. 143, June 27–July 2 2004.
- [11] C. Karlof and D. Wagner, “Secure routing in wireless sensor networks: attacks and countermeasures,” in *Proceedings of the 2003 IEEE International Workshop on Sensor Network Protocols and Applications*, pp. 113–127, May 2003.
- [12] S. Marano, V. Matta, and L. Tong, “Distributed inference in the presence of Byzantine sensors,” in *Proc. 40th Annual Asilomar Conf. on Signals, Systems, and Computers*, (Pacific Grove, CA), Oct 29–Nov 1 2006.
- [13] W. Du, J. Deng, Y. S. Han, and P. Varshney, “A witness-based approach for data fusion assurance in wireless sensor networks,” in *IEEE Global Telecommunications Conference 2003*, vol. 3, pp. 1435–1439, December 2003.