

Variable-Rate Distributed Source Coding in the Presence of Byzantine Sensors

Oliver Kosut and Lang Tong
School of Electrical and Computer Engineering
Cornell University, Ithaca, NY 14853
Email: {oek2, lt35}@cornell.edu

Abstract—The distributed source coding problem is considered when the sensors, or encoders, are under Byzantine attack; that is, an unknown number of sensors have been reprogrammed by a malicious intruder to undermine the reconstruction at the fusion center. Three different forms of the problem are considered. The first is a variable-rate setup, in which the decoder adaptively chooses the rates at which the sensors transmit. An explicit characterization of the variable-rate minimum achievable sum rate is stated, given by the maximum entropy over the set of distributions indistinguishable from the true source distribution by the decoder. In addition, two forms of the fixed-rate problem are considered, one with deterministic coding and one with randomized coding. The achievable rate regions are given for both these problems, with a larger region achievable using randomized coding, though both are suboptimal compared to variable-rate coding.

Index Terms—Distributed Source Coding. Byzantine Attack. Sensor Fusion. Network Security.

I. INTRODUCTION

Wireless sensor networks are vulnerable to various forms of attack. A malicious intruder could capture a sensor or a group of sensors and reprogram them, unbeknownst to the other sensors or the fusion center. The intruder could reprogram the sensors to work cooperatively to obstruct or defeat the goal of the network, launching a so-called Byzantine attack.

We refer to sensors that have been reprogrammed as *traitors*, and the rest, which will behave according to the specified procedure, as *honest*. Suppose there are m sensors and at most t traitors. Each time step, sensor i is informed of the value of the random variable X_i . These random variables constitute a discrete memoryless multiple source with probability distribution $p(x_1 \cdots x_m)$. Each sensor encodes its observation independently and transmits the codewords to a common decoder (the fusion center), which attempts to reconstruct the source values with small probability of error based on those transmissions. If there are no traitors, Slepian-Wolf coding [1] can be used to achieve a sum rate as low as

$$H(X_1 \cdots X_m). \quad (1)$$

However, standard Slepian-Wolf coding has no mechanism for handling any deviations from the agreed-upon encoding functions by the sensors. Even a random fault by a single sensor could have devastating consequences for the accuracy of the source estimates produced at the decoder, to say nothing of a Byzantine attack on multiple sensors.

Consider a two sensor example. If sensor 1 transmits at rate $H(X_1)$ and sensor 2 transmits at rate $H(X_2|X_1)$, their source sequences would normally be reconstructable using Slepian-Wolf. Since sensor 2 transmits at a rate below $H(X_2)$, the decoder must use the codeword from sensor 1 to decode X_2 . Thus, if sensor 1 is a traitor, it can manipulate the decoder's estimate of X_2 to cause an error. Generalizing this, it will turn out that for most source distributions, the sum rate given in (1) cannot be achieved if there is even a single traitor. We will present coding schemes that can handle Byzantine attacks, and give explicit characterizations of the achievable rates.

A. Related Work

The notion of Byzantine attack has its root in the Byzantine generals problem [2], [3] in which a clique of traitorous generals conspire to prevent loyal generals from forming consensus. It was shown in [2] that consensus is possible if and only if less than a third of the generals are traitors.

Countering Byzantine attacks in communication networks has also been studied in the past by many authors. See the earlier work of Perlman [4] and also more recent review [5], [6]. An information theoretic network coding approach to Byzantine attack is presented in [7]. The problem of optimal Byzantine attack of sensor fusion for distributed detection is considered in [8]. Sensor fusion with Byzantine sensors was studied in [9]. In that paper, the sensors, having already agreed upon a message, communicate it to the fusion center over a discrete memoryless channel. Quite similar results were shown in [10], in which a malicious intruder takes control of a set of links in the network. The authors show that two nodes can communicate at a nonzero rate as long as less than half of the links between them are Byzantine. This is different from the current paper in that the transmitter chooses its messages, instead of relaying information received from an outside source, but some of the same approaches from [10] are used in the current paper, particularly the use of randomization to fool traitors that have already transmitted.

B. Fixed-Rate Versus Variable-Rate Coding

In standard multiterminal source coding, each sensor is associated with a rate and an encoding function that transmits information at that rate. We will show that this fixed-rate setup is suboptimal for this problem, in the sense that we can achieve lower sum rates using a variable-rate scheme. By variable-rate

we mean that the number of bits transmitted per source value by a particular sensor will not be fixed. Instead, each sensor has a number of different encoding functions, each with its own rate. The coding session is then made up of a number of transactions. In each transaction, the decoder decides which sensor will transmit information, and which encoding function it should use. Thus we require that the decoder have a reverse channel to transmit information back to the sensors, but it need only send the chosen encoding function index, which will be one of a fixed and small number. In other words, the reverse channel could have arbitrarily small capacity.

C. Honest Sensor Error Requirement

Classical Slepian-Wolf coding requires that the decoder produce perfect estimates of every source value. However, this is no longer possible under Byzantine attack. A traitor could choose to send gibberish to the decoder, in which case the decoder could never correctly decode the associated source values. However, a traitor could also act exactly like an honest sensor, in which case the decoder would never be able to identify it as a traitor. Thus, the decoder will not necessarily be able to produce an accurate estimate for every sensor, but neither will it be able to tell which of its estimates are inaccurate. As a compromise, the decoder will produce an estimate for every source value, but we only require that the estimates corresponding to the honest sensors are correct, even though the decoder may not know which those are. This requirement is reminiscent of that of [2], in which the lieutenants need only perform the order given by the commander if the commander is not a traitor, even though the lieutenants might not know whether he is.

D. Main Results

The main results of this paper give explicit characterizations of the achievable rates for three different setups. The first, discussed in the most depth, is the variable-rate case, for which we give the minimum achievable sum rate. By definition, variable-rate coding involves varying the rates at which different sensors transmit. The choice of these rates will be based on “run time” events such as the source values and the actions of the traitors. Thus, there is no notion of an m -dimensional achievable rate region, since all we can say is that, no matter what happens, the total number of transmitted bits will not exceed a certain value. The second two setups are fixed-rate, divided into deterministic coding and randomized coding, for which we do give m -dimensional achievable rate regions. We show that randomized coding yields a larger achievable rate region than deterministic coding, but we believe that in most cases randomized fixed-rate coding requires an unrealistic assumption. In addition, even randomized fixed-rate coding cannot achieve the same sum rates as variable-rate coding.

For variable-rate coding, the minimum achievable sum rate is given by

$$\sup_{q \in Q} H_q(X_1 \cdots X_m) \quad (2)$$

where H_q is the entropy with respect to the distribution q and Q is a set of distributions which depends on t , the number of allowed traitors. The explicit definition of Q is given later, but intuitively Q is the set of distributions such that if we simulated any distribution $q \in Q$ and handed the resulting source sequences to the decoder as if they had come from the sensors, then it would not be able to correctly identify a single traitor. For example, the source distribution p is always in Q , because if the decoder receives source sequences that appear to come from the true distribution, it will not be able to know which sensors are the traitors. In fact, if $t = 0$, Q is made up of only the source distribution p , so (2) becomes (1). In other words, this result matches the classical Slepian-Wolf result.

On the other hand, if $t = m - 1$, then the decoder knows only that the one honest sensor will report source values distributed according to its single variable marginal distribution, so a traitor will not be detected if it also reports source values distributed according to its marginal distribution. Hence $q \in Q$ if $q(x_i) = p(x_i)$ for all i . It is easy to see that (2) becomes

$$H(X_1) + \cdots + H(X_m). \quad (3)$$

In effect, the decoder must use an independent source code for each sensor.

The fixed-rate achievable regions are based on the Slepian-Wolf achievable region. For randomized coding, the achievable region is such that for every subset of $m - t$ sensors, the rates associated with those sensors fall into the Slepian-Wolf rate region on the corresponding $m - t$ random variables. Note that for $t = 0$, this is identical to the Slepian-Wolf region. For $t = m - 1$, this region is such that for all i , $R_i \geq H(X_i)$, which corresponds to the sum rate in (3). The deterministic region is similar, except that every subset of $m - 2t$ rates is required to fall into the corresponding Slepian-Wolf region.

E. Randomization

Randomization plays a key role in defeating Byzantine attacks. As we have discussed, allowing randomized encoding in the fixed-rate situation expands the achievable region. In addition, the variable-rate coding scheme that we propose relies heavily on randomization to achieve small probability of error. In both fixed and variable-rate coding, randomization is used as follows. Every time a sensor transmits, it randomly chooses from a group of essentially identical encoding functions. The index of the chosen function is transmitted to the decoder along with its output. Without this randomization, a traitor that transmits before an honest sensor i would know exactly the messages that sensor i will send. In particular, it would be able to find fake sequences for sensor i that would produce those same messages. If the traitor tailors the messages it sends to the decoder to match one of those fake sequences, when sensor i then transmits, it would appear to corroborate this fake sequence, causing an error. By randomizing the choice of encoding function, the set of sequences producing the same message is not fixed, so a traitor can no longer know with certainty that a particular fake source sequence will result in the same messages by sensor i as the true

one. This is not unlike Wyner's wiretap channel [11], in which information is kept from the wiretapper by introducing additional randomness.

In both variable-rate and randomized fixed-rate coding, we assume that the traitors know nothing about randomness produced at an honest sensor. Of course, after the randomness has been transmitted, the traitors should have access to that information, which is what we assume in the variable-rate case. However, for the fixed-rate setup, there is no notion of a transmission order, so it would be meaningless to say that the traitors only know about the randomness "after" it has been transmitted. The only choice is to assume that the traitors never find out anything about the randomness. This might be a realistic assumption if the traitors are not able to monitor transmissions to the decoder, but we believe that in most cases it is not. Hence deterministic fixed-rate coding is more realistic.

The rest of the paper is organized as follows. In Section II, we formally give the variable-rate model and present the main result of the paper, which we prove in Section III. In Section IV, we give the rate regions for the fixed-rate setups and illustrate that fixed-rate coding is suboptimal. Finally, in Section V, we offer some future avenues for research.

II. VARIABLE-RATE MODEL AND RESULT

A. Notation

Let X_i be the random variable revealed to sensor i , \mathcal{X}_i the alphabet of that variable, and x_i the corresponding realization. A sequence of random variables revealed to sensor i over n timeslots is denoted X_i^n , and a realization of it $x_i^n \in \mathcal{X}_i^n$. Let $\mathcal{M} \triangleq \{1, \dots, m\}$. For a set $s \subset \mathcal{M}$, let X_s be the set of random variables $\{X_i\}_{i \in s}$, and define x_s and \mathcal{X}_s similarly. By s^c we mean $\mathcal{M} \setminus s$. Let $T_\epsilon^n(X_s)[q]$ be the strongly typical set with respect to the distribution q , or the source distribution p if unspecified. Similarly, $H_q(X_s)$ is the entropy with respect to the distribution q , or p if unspecified. All variations on ϵ , such as ϵ' , ϵ'' , $\dot{\epsilon}$, are assumed to go to 0 as ϵ goes to 0 and may appear without definition. It is meant that either the definition is discernible from context or the existence will be shown.

B. Communication Protocol

The transmission protocol is composed of L transactions. In each transaction, the decoder selects a sensor to receive information from and selects which of K encoding functions it should use. The sensor then responds by executing that encoding function and transmitting its output back to the decoder. For each sensor $i \in \mathcal{M}$ and encoding function $j \in \{1, \dots, K\}$, there is an associated rate $R_{i,j}$. On the l th transaction, let i_l and j_l be the sensor and encoding function chosen by the decoder, and let h_l be the number of times i_l has transmitted prior to the l th transaction. Note that i_l, j_l, h_l are random variables, since they are chosen by the decoder based on messages it has received, which depend on the source values. The j th encoding function for the i th sensor is given by

$$f_{i,j} : \mathcal{X}_i^n \times \mathcal{Z} \times \{1, \dots, K\}^{h_l} \rightarrow \{1, \dots, 2^{nR_{i,j}}\}$$

where \mathcal{Z} represents randomness generated at the sensor. Let $I_l \in \{1, \dots, 2^{nR_{i_l,j}}\}$ be the message received by the encoder in the l th transaction. If i_l is an honest sensor, then $I_l = f_{i_l,j_l}(X_{i_l}^n, \rho_{i_l}, J_l)$, where $\rho_{i_l} \in \mathcal{Z}$ is the randomness from sensor i_l and $J_l \in \{1, \dots, K\}^{h_l}$ is the history of encoding functions used by sensor i_l so far. If i_l is a traitor, however, it may choose I_l based on all sources X_1^n, \dots, X_m^n , all previous transmissions I_1, \dots, I_{l-1} and polling history i_1, \dots, i_{l-1} and j_1, \dots, j_{l-1} . In particular, it does not have access to the randomness ρ_i for any honest sensor i .

After the decoder receives I_l , if $l < L$ it uses I_1, \dots, I_l to choose the next sensor i_{l+1} and its encoding function index j_{l+1} . After the L th transaction, it decodes according to the decoding function

$$g : \prod_{l=1}^L \{1, \dots, 2^{nR_{i_l,j_l}}\} \rightarrow \mathcal{X}_1^n \times \dots \times \mathcal{X}_m^n.$$

C. Variable-Rate Problem Statement and Main Result

Let $\mathcal{H} \subset \mathcal{M}$ be the set of honest sensors. Define the probability of error $P_e \triangleq \Pr(X_{\mathcal{H}}^n \neq \hat{X}_{\mathcal{H}}^n)$ where $(\hat{X}_1^n, \dots, \hat{X}_m^n) = g(I_1, \dots, I_L)$. This will in general depend on the actions of the traitors. Note again that the only source estimates that matter are those corresponding to the honest sensors.

We define a sum rate R to be ϵ -achievable if for every $\delta > 0$ and sufficiently large n there exists a code such that, for any choice of actions by the traitors, $P_e \leq \epsilon$ and

$$\sum_{l=1}^L R_{i_l,j_l} \leq R + \delta. \quad (4)$$

Note that R_{i_l,j_l} depend on the sensor transmissions, so they are random variables. By (4) we mean that for any messages sent by the sensors, we never exceed a sum rate of $R + \delta$. A sum rate R is *achievable* if it is ϵ -achievable for every $\epsilon > 0$. Let R^* be the minimum achievable sum rate. Certainly then all $R > R^*$ are also achievable.

Some definitions will allow us to state our main result. Let

$$\mathcal{V} \triangleq \{s \subset \mathcal{M} : |s| = m - t\}.$$

This is the collection of all possible sets of honest sensors. For any $V \subset \mathcal{V}$, define

$$Q(V) \triangleq \{q(x_1 \dots x_m) : \forall s \in V, q(x_s) = p(x_s)\}. \quad (5)$$

Let $U(V) \triangleq \bigcup_{s \in V} s$. Finally, define

$$Q \triangleq \bigcup_{V \subset \mathcal{V} : U(V) = \mathcal{M}} Q(V).$$

That is, Q is the set of distributions q such that for each i , there is a marginal distribution of q of $m-t$ variables including X_i that matches the corresponding marginal distribution of p . Thus, those $m-t$ sensors behave as if they were the set of honest sensors, since their sources are distributed correctly. Since every i falls into such a set, every sensor looks like it could be honest.

Theorem 1: The minimum achievable sum rate is

$$R^* = \sup_{q \in Q} H_q(X_1 \cdots X_m). \quad (6)$$

It can be shown that for $t = 1$ and arbitrary m , (6) becomes

$$R^* = H(X_1 \cdots X_m) + \max_{i, i' \in \mathcal{M}} I(X_i; X_{i'} | X_{\{i, i'\}^c}). \quad (7)$$

Relative to the Slepian-Wolf result, we see that we always pay a conditional mutual information penalty for a single traitor. Similar expressions can be found for $t = 2$, $t = m - 2$, and $t = m - 1$ (the last given by (3)). However, analytic expressions do not in general exist for $3 \leq t \leq m - 3$.

III. PROOF OF THE VARIABLE-RATE THEOREM

A. Converse

We first show the converse. Let \tilde{q} be the distribution q that maximizes the entropy in (6). For some s with $|s| = m - t$, we can write $\tilde{q} = p(x_s) \tilde{q}(x_{s^c} | x_s)$. Thus if the s^c sensors are the traitors, they can simulate the conditional distribution $\tilde{q}(x_{s^c} | x_s)$, the outcome of which, when combined with the true values of X_s , will produce a set of $X_1 \cdots X_m$ distributed according to \tilde{q} . Since $\tilde{q} \in Q$, if the traitors act honestly with these fabricated source values, the decoder will not be able to correctly identify a single traitor, so it has no choice but to perfectly decode every value. To do this, it must receive at least $nH_{\tilde{q}}(X_{\mathcal{M}})$ bits, which means $R^* \geq H_{\tilde{q}}(X_{\mathcal{M}})$.

B. Achievability Preliminaries

Now we prove achievability. To do so, we will need the following definitions. For some $V \subset \mathcal{V}$, let

$$S_\epsilon^n(X_{\mathcal{M}})[V] \triangleq \{x_{\mathcal{M}}^n \in \mathcal{X}_{\mathcal{M}}^n : \forall s \in V, x_s^n \in T_\epsilon^n(X_s)\}$$

where T_ϵ^n is the strongly typical set. For $s, s' \subset \mathcal{M}$ and $x_{s'}^n \in \mathcal{X}_{s'}^n$, we define the conditional version

$$S_\epsilon^n(X_s | x_{s'}^n)[V] \triangleq \{x_s^n \in \mathcal{X}_s^n : \exists x_{(s \cup s')^c}^n \in \mathcal{X}_{(s \cup s')^c}^n : (x_s^n x_{s'}^n x_{(s \cup s')^c}^n) \in S_\epsilon^n(X_{\mathcal{M}})[V]\}.$$

The following lemma shows that S_ϵ^n is contained in a union of typical sets.

Lemma 1: Fix $s, s' \subset \mathcal{M}$ and $x_{s'}^n \in \mathcal{X}_{s'}^n$. Then

$$S_\epsilon^n(X_s | x_{s'}^n)[V] \subset \bigcup_{q \in Q(V)} T_{\epsilon'}^n(X_s | x_{s'}^n)[q].$$

C. Coding Scheme Procedure

We propose a multiround coding scheme. Each round is made up of m phases. In the i th phase, transactions are made entirely with sensor i . In addition, all transactions in the first round are based on the first k source values, transactions in the second round on the second k source values, and so on. Each transaction in the i th phase will be associated with a target set chosen by the decoder of the form

$$T_R(\hat{x}_s^k) \triangleq \bigcup_{q: H_q(X_i | X_s) \leq R} T_{\epsilon'}^k(X_i | \hat{x}_s^k)[q] \quad (8)$$

with $s \subset \mathcal{M}$ to be defined, and ϵ' is as defined in Lemma 1. It takes about kR bits to encode any sequence in this set, so we

can think of $T_R(\hat{x}_s^k)$ as the set of all the sequences that can be decoded if a sensor has only sent kR bits so far in the current phase. The strategy will be to slowly increase R , expanding $T_R(\hat{x}_s^k)$ until it contains the relevant source sequence.

The decoder will attempt to determine whether the source sequence is contained in $T_R(\hat{x}_s^k)$, and if so to decode it. Sensor i will randomly choose from a number of encoding functions f_1, \dots, f_C . Each of these encoding functions will be created by means of a random binning procedure and the codebooks revealed to both the sensor and decoder. Sensor i will transmit up to $k(R + \epsilon)$ bits containing the index of the randomly chosen encoding function and its output. If there is exactly one source sequence in the target set that matches every value received so far from sensor i in this round, call it \hat{x}_i^k . If there is more than one such sequence, we declare an error. If there is no such sequence, we conclude that the source sequence is not contained in the target set, increase R by ϵ , and do another transaction. Note that when $R \geq \log |\mathcal{X}_i|$, every sequence will be in $T_R(\hat{x}_s^k)$, so we will definitely decode the sequence or declare an error.

The collection $V \subset \mathcal{V}$ will always contain only those sets that could be the set of honest sensors. We begin by setting $V = \mathcal{V}$, and pare it down after each round based on new information. Define $s_i \triangleq \{1, \dots, i\} \cap U(V)$. Phase i of any round is made up of the following steps.

- 1) If $i \notin U(V)$, ignore i and go to the next phase.
- 2) Otherwise, let $R = \epsilon$.
- 3) Receive up to $k(R + \epsilon)$ bits from sensor i , with target set $T_R(\hat{x}_{s_{i-1}}^k)$. If possible, decode the sequence to \hat{x}_i^k and go to the next phase. If not, increase R by ϵ and repeat.
- 4) After phase m , let $V' \in \mathcal{V}$ be the largest subset of V such that $\hat{x}_{U(V')} \in S_\epsilon^n(X_{U(V')})[V']$. Use V' as V in the next round. If there is no such V' , declare an error.

D. Code Rate

It can be shown that the probability of error can be made arbitrarily small if C , the number of encoding functions from which each sensor chooses randomly during each transaction, is sufficiently large. We can then make k large enough that transmitting the index of the chosen encoding function takes negligible rate compared to transmitting its output. Thus in each phase we need only transmit $R + \epsilon$ bits per symbol. Let $q_{\hat{x}}$ be the type of $\hat{x}_{U(V)}^k$. The total number of bits sent per symbol for the entire round is therefore at most

$$\sum_{i=1}^m \inf_{q: \hat{x}_i^k \in T_{\epsilon'}^n(X_i | \hat{x}_{s_{i-1}}^k)[q]} H_q(X_i | X_{s_{i-1}}) + \epsilon + \epsilon$$

$$\leq \inf_{q: \hat{x}_{U(V)}^k \in T_{\epsilon'}^n(X_{U(V)})[q]} \sum_{i=1}^m H_q(X_i | X_{s_i}) + m(\epsilon + \epsilon) \quad (9)$$

$$\leq H_{q_{\hat{x}}}(X_{U(V)}) + m(\epsilon + \epsilon) \quad (10)$$

$$\leq \sup_{q \in Q(V')} H_q(X_{U(V)}) + \epsilon \quad (11)$$

$$\leq \sup_{q \in Q} H_q(X_{\mathcal{M}}) + \log |\mathcal{X}_{U(V) \setminus U(V')}| + \epsilon \quad (12)$$

where (9) holds because the set of distributions q such that $\hat{x}_{s_i}^k \in T_{\epsilon'}^n(X_{s_i})[q]$ contains the set of distributions q such that $\hat{x}_{U(V)}^k \in T_{\epsilon'}^n(X_{U(V)})[q]$, and (10) holds because $\hat{x}_{U(V)}$ is typical with respect to its own type. Because $\hat{x}_{U(V)} \in S_{\epsilon'}^n(X_{U(V)})[V']$, by Lemma 1, for some $q \in Q(V')$, $\hat{x}_{U(V)} \in T_{\epsilon'}^n(X_{U(V)})[q]$. For this q , for all $x_{U(V)} \in \mathcal{X}_{U(V)}$, $|q_{\hat{x}}(x_{U(V)}) - q(x_{U(V)})| \leq \frac{\epsilon'}{|\mathcal{X}_{U(V)}|}$. Since the distributions are arbitrarily close, the entropies with respect to these distributions will be arbitrarily close, so (11) holds.

If $U(V') = U(V)$, then the second term in (12) is 0, so we can bound (12) by $\sup_{q \in Q} H_q(X_{\mathcal{M}}) + \tilde{\epsilon}$. However, if $U(V) \setminus U(V') \neq \emptyset$, we cannot. Even so, since at least one sensor is eliminated whenever $U(V) \setminus U(V') \neq \emptyset$, this can only happen for at most t rounds, after which we will have eliminated every traitor. Thus with enough rounds, we can always bound the sum rate by $\sup_{q \in Q} H_q(X_{\mathcal{M}}) + \tilde{\epsilon}$.

IV. FIXED-RATE RESULTS

Consider an m -tuple of rates (R_1, \dots, R_m) , encoding functions $f_i : \mathcal{X}_i^n \rightarrow \{1, \dots, 2^{nR_i}\}$ for $i \in \mathcal{M}$, and decoding function

$$g : \prod_{i=1}^m \{1, \dots, 2^{nR_i}\} \rightarrow \mathcal{X}_1^n \times \dots \times \mathcal{X}_m^n.$$

Let $I_i \in \{1, \dots, 2^{nR_i}\}$ be the message transmitted by sensor i . If sensor i is honest, $I_i = f_i(X_i^n)$. If it is a traitor, it may choose I_i arbitrarily, based on all the sources $X_{\mathcal{M}}^n$. Define the probability of error $P_e \triangleq \Pr(X_{\mathcal{J}_c}^n \neq \hat{X}_{\mathcal{J}_c}^n)$ where $(\hat{X}_1^n, \dots, \hat{X}_m^n) = g(I_1, \dots, I_m)$.

We say an m -tuple (R_1, \dots, R_m) is *deterministic-fixed-rate achievable* if for any $\epsilon > 0$ and sufficiently large n , there exist coding functions f_i and g such that, for any choice of actions by the traitors, $P_e \leq \epsilon$. Let $\mathcal{R}_{\text{dfr}} \subset \mathbb{R}^m$ be the set of deterministic-fixed-rate achievable m -tuples.

Define an m -tuple to be *randomized-fixed-rate achievable* in the same way as above, except we allow the encoding functions f_i to be randomized. Let $\mathcal{R}_{\text{rfr}} \subset \mathbb{R}^m$ be the set of randomized-fixed-rate achievable rate vectors.

For any $s \subset \mathcal{M}$, let $\text{SW}(X_s)$ be the Slepian-Wolf rate region for the random variables X_s . For any integer $k \leq m$, define

$$\mathcal{R}_k \triangleq \{(R_1 \cdots R_m) : \forall s \subset \mathcal{M}, |s| = k : (R_i)_{i \in s} \in \text{SW}(X_s)\}.$$

The following theorem gives the rate regions explicitly.

Theorem 2: The fixed-rate achievable regions are given by

$$\mathcal{R}_{\text{dfr}} = \mathcal{R}_{\max\{1, m-2t\}} \quad \text{and} \quad \mathcal{R}_{\text{rfr}} = \mathcal{R}_{m-t}.$$

We omit the proof of this, but we briefly illustrate that circumstances exist for which fixed-rate coding is suboptimal compared to variable-rate coding. Suppose $m = 3$ and $t = 1$. Recall from (7) that the variable-rate minimum achievable sum rate is given by

$$R^* = H(X_1 X_2 X_3) + \max\{I(X_1; X_2|X_3), I(X_1; X_3|X_2), I(X_2; X_3|X_1)\}. \quad (13)$$

Suppose that $I(X_1; X_2|X_3)$ achieves this maximum. If the rate triple (R_1, R_2, R_3) is randomized fixed-rate achievable, then

$(R_1, R_2, R_3) \in \mathcal{R}_2$, which means $R_i + R_j \geq H(X_i X_j)$ for all $i, j \in \{1, 2, 3\}$. Thus

$$\begin{aligned} R_1 + R_2 + R_3 &\geq \frac{1}{2} [H(X_1 X_2) + H(X_1 X_3) + H(X_2 X_3)] \\ &= H(X_1 X_2 X_3) + \frac{1}{2} [I(X_1; X_2|X_3) + I(X_1 X_2; X_3)]. \end{aligned} \quad (14)$$

If $I(X_1 X_2; X_3) > I(X_1; X_2|X_3)$, (14) is larger than (13). Hence, for some source distributions, a larger sum rate is required for fixed-rate coding than variable-rate coding.

V. FUTURE WORK

Much more work could be done in the area of Byzantine network source coding. In this paper, we assumed that the traitors have access to all the source values, an assumption that was vital in our converse proofs. This is a significant assumption that may not be all that realistic. It would be worthwhile, though perhaps more difficult, to characterize the achievable rate region without this assumption, assuming that the traitors have access only to their own source values, or possibly degraded versions of those of the honest sensors.

Finally, we could consider Byzantine attacks on other sorts of multi-terminal source coding problems, such as the rate distortion problem [12], [13] or the CEO problem [14].

REFERENCES

- [1] D. Slepian and J. Wolf, "Noiseless coding of correlated information sources," *IEEE Trans. Information Theory*, vol. IT-19, pp. 471–480, 1973.
- [2] 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.
- [3] D. Dolev, "The Byzantine generals strike again," *Journal of Algorithms*, vol. 3, no. 1, pp. 14–30, 1982.
- [4] R. Perlman, *Network Layer Protocols with Byzantine Robustness*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, August 1988.
- [5] L. Zhou and Z. J. Haas, "Securing ad hoc networks," *IEEE Network Magazine*, vol. 13, pp. 24–30, Nov/Dec 1999.
- [6] Y. Hu and A. Perrig, "Security and privacy in sensor networks," *IEEE Security and Privacy Magazine*, vol. 2, pp. 28–39, 2004.
- [7] T. Ho, B. Leong, R. Koetter, M. Médard, M. Effros, 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.
- [8] 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.
- [9] O. Kosut and L. Tong, "Capacity of cooperative fusion in the presence of Byzantine sensors," in *Proc. 44th Annual Allerton Conf. on Commun., Control and Comp.*, (Monticello, IL), Sep 27–29 2006.
- [10] T. H. S. Jaggi, M. Langberg and M. Effros, "Correction of adversarial errors in networks," in *Proceedings of International Symposium in Information Theory and its Applications*, (Adelaide, Australia), 2005.
- [11] A. Wyner, "The wiretap channel," *Bell Syst. Tech. J.*, vol. 54, pp. 1355–1387, 1975.
- [12] S. Y. Tung, *Multiterminal Source Coding*. PhD thesis, Cornell University, Ithaca, NY, 1978.
- [13] T. Berger, *The Information Theory Approach to Communications* (G. Longo, ed.), chapter Multi-terminal source coding. Springer-Verlag, 1978.
- [14] T. Berger, Z. Zhang, and H. Viswanathan, "The CEO problem [multiterminal source coding]," *IEEE Trans. Inform. Theory*, vol. 42, pp. 887–902, May. 1996.