The Quadratic Gaussian CEO Problem with Byzantine Agents

Oliver Kosut and Lang Tong

***Abstract*— The quadratic Gaussian CEO problem is studied when the agents are under Byzantine attack. That is, an unknown subset of agents is controlled by an adversary that attempts to damage the quality of the estimate at the Central Estimation Officer, or CEO. Inner and outer bounds are presented for the achievable rate region as a function of the fraction of adversarial agents. The inner bound is derived from a generalization of the Berger-Tung quantize-and-bin strategy, which has been shown to be tight in the non-Byzantine case. The outer bound has similarities to the Singleton bound in that the traitorous agents must be prevented from allowing two sources to result in the same transmitted codewords if their values are too far apart for the distortion constraint to be satisfied with a single estimate. The inner and outer bounds on the rate regions are used to find bounds on the asymptotic proportionality constant in the limit of a large number of agents and high sum-rate. These bounds on the proportionality constant differ at most by a factor of 4.**

I. Introduction

Distributed systems are more likely to be susceptible to physical assault. A malicious intruder could seize a group of nodes and reprogram them to cooperate to obstruct the goal of the network, launching a so-called Byzantine attack. Alternatively, nodes may break down and begin transmitting spurious information. In either case, it is necessary to design algorithms and analyze performance in distributed problems when some of the nodes do not behave as they should.

Consider the CEO problem, a special case of multiterminal source coding, in which the fusion center or Central Estimation Officer (CEO), is interested in a sequence $\{X(t)\}_{t=1}^{\infty}$ but cannot observe it directly. Instead, each of L agents observe one of $\{Y_k(t)\}_{t=1}^{\infty}$ for $k = 1, \ldots, L$, where the Y_k are conditionally independent given X. Without cooperating, the agents communicate encoded versions of their measurements to the CEO, which uses these transmissions to produce an estimate of X. We investigate a modification of this problem in which an unknown group of βL agents are traitors. Traitors need not use the stipulated encoders to produce their transmissions to the CEO; indeed they may choose their codewords arbitrarily. We will study how β affects the quality of the CEO's estimate with Gaussian sources and quadratic distortion measure.

The CEO problem was first studied by Berger, Zhang, and Viswanathan [1] for discrete memoryless sources. They showed that for a large number of agents the achievable distortion fell exponentially with increasing sum-rate, and they characterized the associated error exponent. Viswanathan and Berger [2] first studied the quadratic Gaussian version in a similar regime, showing that for many agents the distortion fell as K/R where R is the sum rate and K is a constant independent of R, and they found bounds on the proportionality constant K. Oohama [3] showed that their inner bound was tight, and then [4] simultaneously with Prabhakaran, Tse, and Ramchandran [5] found the rate region for a finite number of agents and heterogeneous observations by the agents. All of these results used only the Berger-Tung inner bound [6], [7] (also known as quantize-and-bin) to prove achievability. These results indicate that the Berger-Tung encoding cannot be improved upon for the quadratic Gaussian CEO problem. It is curious whether the essential Berger-Tung technique remains optimal when some agents are compromised.

The notion of Byzantine attack has its root in the Byzantine Generals Problem [8], [9], in which a clique of traitorous generals conspire to prevent loyal generals from reaching consensus. Byzantine attacks have been applied to many problems in networks, such as network coding [10], [11]. Distributed source coding was investigated in [12], which studied the problem of Slepian-Wolf [13] under Byzantine attack. The discrete memoryless CEO problem under Byzantine attack was investigated in [14], [15], both of which studied the error exponent originally characterized without traitors in [1].

One could consider a range of models for network failures. For example, [10] investigated several different Byzantine models for network coding, and showed that different rates are achievable depending on the insidiousness of the compromised part of the network. For multiterminal source coding, perhaps the simplest model would be one in which failed nodes transmit nothing to the CEO, and therefore their identities are immediately known. Such a model was considered in [16], in which the decoder sought to produce a higher quality estimate when fewer nodes fail; here, we are merely interested in the worst case performance with a limited number of failures. In this case, fully-identified failed nodes can be dealt with simply by decoding based on whatever information is received by the CEO. It is easy to see that the tightness of the Berger-Tung achievable region for the Gaussian CEO problem extends to this form of node failures. Even if failed nodes, instead of making themselves immediately known, transmit random information to the CEO, the problem is easy, because agents are expected to send correlated information, so a node sending a codeword independent from the rest is easy to identify and ignore. On the other end of the spectrum, compromised nodes may be Byzantine, with full access to all the sources, the ability to cooperate, and knowledge of the code. Also, honest

O. Kosut and L. Tong are with Cornell University, Ithaca, NY 14853 {oek2,lt35}@cornell.edu

This work is supported in part by the National Science Foundation under Award CCF-0635070 and the the Army Research Office under Grant ARO-W911NF-06-1-0346.

nodes are forced to use deterministic encoders. (Alternatively, honest encoders may be random, but this randomness is known to the traitors.) We will consider this model in the present paper, for two reasons. First, assuming very powerful traitors ensures robustness of performance even when they are not. Secondly, after studying some weaker traitor models in [12], [14], [15], such as one in which honest nodes can generate independent randomness that is unknown to the traitors, we have found that the most extreme model, if perhaps overly pessimistic, is the most tractable problem that captures the nature of defeating Byzantine attacks.

In this paper, we present inner and outer bounds on the rate region for the quadratic Gaussian CEO problem under Byzantine attack. Our inner bound is an extension of the Berger-Tung inner bound, and can similarly be applied to a great variety of problems. As it is for the quadratic Gaussion CEO problem without traitors, we conjecture that this inner bound is tight. Our outer bound is a direct generalization of the converses in [3] and [5]. It also has elements of the Singleton bound from coding theory, in that we wish to prevent errors (or, in our case, codeword manipulation by the traitors) from bringing two points together that must remain distinguishable. A similar generalization of the Singleton bound was found in the context of network coding in [11]. We use our bounds on the rate region to bound the proportionality constant originally studied in [2], giving the constant as a function of β to within a factor of 4. We also observe that our bounds on the proportionality constant have dramatically different behaviors for small β, indicating that a small number of traitors may have an unexpectedly harsh effect on performance.

The paper is structured as follows. Section II formally presents the model and states our results. The inner bound is proved in Section III, the outer bound in Section IV, and the bounds on the asymptotic constant in Section V. We conclude in Section VI.

Notation. The superscript n denotes n-length sequences across time (e.g. $X^n = (X(1), \ldots, X(n))$), the superscript L denotes L-length sequences across space (e.g. Y^L), and the superscript nL denotes sequences across time and space (e.g. Y^{nL}). By X_A for $A \subset \{1, \ldots, L\}$ we mean $(X_k)_{k\in A}$. The n-length typical set parameterized by ϵ is written $T_\epsilon^{(n)}(X)$, with joint and condition typical sets written similarly.

II. Model and Results

Let $X(t)$ for $t = 1, \ldots, n$ be an i.i.d. Gaussian sequence with zero mean and variance σ_X^2. For $k = 1, \ldots, L$, let $Y_k(t) = X(t) + N_k(t)$ where $N_k(t)$ is an i.i.d. Gaussian sequence with zero mean and variance $\sigma_{N_k}^2$, and where the N_k are independent of X and of each other. Agent k receives the sequence Y_k^n and encodes it using the encoding function $f_k : \mathbb{R}^n \to \{1, \ldots, 2^{nR_k}\}$. The agents are divided into two groups, labeled *honest agents* and *traitors*. There are βL traitors, where β is assumed to be known to the CEO, but the identity of this group is unknown at the time of the code construction, so the code must be prepared for any possible set of traitors. If agent k is honest, then the codeword C_k that it transmits to the CEO is $C_k = f_k(Y_k^n)$. However, if k is a traitor, then it may choose C_k any way it likes, based on full knowledge of the sources X, Y^L, the code, and cooperation with other traitors. The CEO's decoding function is

$$g : \prod_{k=1}^{L} \{1, \ldots, 2^{nR_k}\} \to \mathbb{R}^n \tag{1}$$

from which it produces an estimate $\hat{X}^n = g(C_1, \ldots, C_L)$. For a given pair (x^n, y^{nL}), we define the maximum possible distortion over all possible actions of the traitors to be

$$D(x^n, y^{nL}) = \max_{\substack{H\subset\{1,\ldots,L\}:\\ |H|=(1-\beta)L}} \max_{C_{H^c}} \frac{1}{n}\sum_{t=1}^{n}(x(t) - \hat{x}(t))^2. \tag{2}$$

In this expression H runs over all possible sets of honest agents, where H^c is the set of traitors. We also maximize over C_{H^c}, the codewords sent by the traitors, ensuring that any potentially traitor actions are considered. Observe that even the choice of which agents to capture may be a function of the source values. Note also that in (2) $\hat{x}^n$ is a function of C^L given by g, and C_H is in turn a function of y_H^n given by the f_i.

Let the expected distortion be

$$D = \mathbb{E}(D(X^n, Y^{nL})). \tag{3}$$

We say that a tuple $(R_1, \ldots, R_L, D)$ is *achievable* if for sufficiently large n there exist encoders $(f_1, \ldots, f_L)$ operating at these rates and a decoder g such that the distortion is arbitrarily close to D.

We have assumed above that the decoder is deterministic. In general, that need not be the case, and for certain Byzantine problems it may be that randomization at the decoder can improve performance. However, the convexity of the quadratic distortion function implies it cannot do so for this problem. In particular, given any random decoder, consider the deterministic decoder that simply takes the expectation of the random estimate given the received codewords. This decoder cannot do worse than the random one, even though the traitors may change their behavior based on which decoder is to be used.

We now state our inner bound.

Theorem 1: The tuple $(R_1, \ldots, R_L, D)$ is achievable if there exist r_k for $k = 1, \ldots, L$ and for each matrix $\Sigma \in \mathbb{R}^{L\times L}$ constants $c_k(\Sigma)$ such that

A) for all $S \subset \{1, \ldots, L\}$ with $|S| = (1 - 2\beta)L$ and all $A \subset S$,

$$\sum_{k\in A} R_k \geq \sum_{k\in A} r_k + \frac{1}{2}\log\left(\frac{1}{\sigma_X^2} + \sum_{k\in S}\frac{1-\exp(-2r_k)}{\sigma_{N_k}^2}\right) - \frac{1}{2}\log\left(\frac{1}{\sigma_X^2} + \sum_{k\in S\setminus A}\frac{1-\exp(-2r_k)}{\sigma_{N_k}^2}\right) \tag{4}$$

B) for every $H \subset L$ with $|H| = (1-\beta)L$ and every vector $\lambda \in \mathbb{R}^L$ for which

$$\Sigma_{j,k} = \sigma_X^2 + \frac{\sigma_{N_k}^2}{1-\exp(-2r_k)}\delta_{j,k} \text{ for all } j, k \in H \tag{5}$$

and $\lambda_k = \sigma_X^2$ for $k \in H$,

$$D \geq \mathbb{E}_{\Sigma,\lambda}\left(X - \sum_{k=1}^{L} c_k(\Sigma) U_k\right)^2 \tag{6}$$

where by $\mathbb{E}_{\Sigma,\lambda}$ we mean an expectation taken over a distribution on the variables $(X, U_1, \ldots, U_L)$ with covariance matrix

$$\begin{pmatrix} \sigma_X^2 & \lambda^T \\ \lambda & \Sigma \end{pmatrix}. \tag{7}$$

We offer the following intuition for this result. Agent k will send to the CEO a corrupted version of its measurement represented by U_k. These will be designed so that if all agents were honest, the covariances between them would be

$$\mathbb{E}(U_j U_k) = \sigma_X^2 + \frac{\sigma_{N_k}^2}{1 - \exp(-2r_k)}\delta_{j,k}. \tag{8}$$

However, due to the presence of the traitors, the joint distribution of X, U^L that actually occurs, which is represented by the covariance matrix in (7), may not match the distribution that would result with no traitors. This alternative distribution is parameterized by Σ and λ, where Σ is the covariance matrix of U^L, and λ is the covariance vector between X and U^L. Since the CEO can observe only U^L, it can only recover Σ, from which it must choose an estimator. From Σ, the CEO can identify possible sets of honest agents as the ones satisfying (5), because the honest agents are guaranteed to transmit information using the proper distribution. However, there may be several possible sets that are indistinguishable to the CEO, and for each set many possibilities for λ. The CEO must construct its estimate by choosing constants c_k that satisfy the distortion constraint for each of these possibilities, as (6) stipulates.

This inner bound is a natural extension of the Berger-Tung inner bound for the non-Byzantine setting. This bound is tight in the non-Byzantine setting, and we conjecture that our inner bound is likewise tight, in which case the rate region would be given by conditions (A) and (B) above. However, this region does not match that of our outer bound, stated as follows.

Theorem 2: If the tuple $(R_1, \ldots, R_L, D)$ is achievable, then there exist r_k for $k = 1, \ldots, L$ such that for all $S \subset \{1, \ldots, L\}$ with $|S| = (1-2\beta)L$ and all $A \subset S$,

$$\sum_{k\in A} R_k \geq \sum_{k\in A} r_k + \frac{1}{2}\log\frac{1}{D} - \frac{1}{2}\log\left(\frac{1}{\sigma_X^2} + \sum_{k\in S\setminus A}\frac{1-\exp(-2r_k)}{\sigma_{N_k}^2}\right), \tag{9}$$

$$\frac{1}{D} \leq \frac{1}{\sigma_X^2} + \sum_{k\in S}\frac{1-\exp(-2r_k)}{\sigma_{N_k}^2}. \tag{10}$$

The region specified in our outer bound in Theorem 2 is identical to the rate region for the non-Byzantine problem given in [4], [5] except that the two conditions on $\{1, \ldots, L\}$ have been replaced with conditions on S for all sets of size $(1-2\beta)L$. In fact, Theorems 1 and 2 reduce to the result in [4], [5] when $\beta = 0$.

When $\beta \geq 1/2$, (10) reduces to $D \geq \sigma_X^2$. That is, no matter how high the communication rate from the agents, the quality of the CEO's estimate can never improve over the a priori variance. The reason for this is that once traitors control at least half of the network, it is impossible for the CEO to distinguish the group of traitors from the group of honest agents, so the traitors can simply report a completely different value of X than the true one, and the CEO will never know which one is the truth. For this reason, we focus mainly on the nontrivial regime $\beta < 1/2$.

We now define the asymptotic proportionality constant. Let $D(R, L)$ be the minimum achievable distortion for L agents where the sum-rate is at most R. In the case that all agents have the same quality of observation (i.e. $\sigma_{N_k}^2 = \sigma_N^2$ for all k), let $D(R) = \lim_{L\to\infty} D(R, L)$. Finally define

$$K(\sigma_X^2, \sigma_N^2, \beta) = \lim_{R\to\infty} R\frac{D(R)}{\sigma_X^2}. \tag{11}$$

That is, $D(R)$ goes like $K\sigma_X^2/R$ for large R. The following theorem bounds K.

Theorem 3:

$$\frac{\sigma_N^2}{2\sigma_X^2}\frac{1}{1-2\beta} \leq K(\sigma_X^2, \sigma_N^2, \beta) \leq \frac{\sigma_N^2}{2\sigma_X^2}\frac{\sqrt{1-\beta}+\sqrt{\beta}}{(1-\beta)(\sqrt{1-\beta}-\sqrt{\beta})}. \tag{12}$$

At $\beta = 0$, the two bounds meet at $\sigma_N^2/(2\sigma_X^2)$, matching the result of [3]. They also both diverge at $\beta = 1/2$. The ratio between them is monotonically increasing in β and is never more than 4.

We do not make the same conjecture for our upper bound on K as we did for our inner bound on the rate region. The complexity of the statement of Theorem 1 makes it difficult to calculate the best value of K that would result from it, and it may be possible to improve on the upper bound in (12). However, though we omit the proof of this in the interest of space, it is true that any upper bound on K resulting from Theorem 1 would have no better behavior for small β than the upper bound in (12). That is, if our conjecture on the tightness of Theorem 1 holds, then for $\beta \ll 1$,

$$K(\sigma_X^2, \sigma_{N_k}^2, \beta) \simeq \frac{\sigma_N^2}{2\sigma_X^2}(1 + 2\sqrt{\beta}). \tag{13}$$

Compare this to our lower bound from Theorem 3, which states that for small β,

$$K(\sigma_X^2, \sigma_{N_k}^2, \beta) \gtrsim \frac{\sigma_N^2}{2\sigma_X^2}(1 + 2\beta). \tag{14}$$

Observe that (13) increases rapidly with β near $\beta = 0$ as compared to (14).

III. Inner Bound Proof

Fix R_k, D, r_k, and $c_k(\Sigma)$ satisfying conditions (A) and (B) in the statement of Theorem 1. We will present a coding scheme to show that $(R_1, \ldots, R_L, D)$ is achievable. First define for each agent k an auxiliary random variable $U_k = Y_k + W_k$, where $W_k \sim \mathcal{N}(0, \sigma_{W_k}^2)$ and the W_k are independent from each other and X, Y^L. The variance $\sigma_{W_k}^2$ is chosen so that

$$r_k = I(Y_k; U_k|X) = \frac{1}{2}\log\frac{\sigma_{N_k}^2 + \sigma_{W_k}^2}{\sigma_{W_k}^2}. \tag{15}$$

Descriptions of the codebook, and the encoding and decoding rules follow.

1) Random Code Structure: Each agent k forms its codebook in the following way. It generates $2^{n(I(Y_k;U_k)+\epsilon)}$ n-length codewords from the marginal distribution of U_k. Let $\mathcal{C}_k^{(n)}$ be the codeword set. These codewords are then placed into 2^{nR_k} bins uniformly at random.

2) Encoding Rule: Upon receiving Y_k^n, agent k selects at random an element of

$$\mathcal{C}_k^{(n)} \cap T_\epsilon^{(n)}(U_k|Y_k^n) \tag{16}$$

which it denotes U_k^n. Agent k then sends to the CEO the index of the bin containing U_k^n.

3) Decoding Rule: For each $H \subset \{1,\dots,L\}$ with $|H| = (1-\beta)L$, the CEO looks for a group of codewords in $T_\epsilon^{(n)}(U_H)$ that matches the received bins from all agents in H. If there is exactly one such a sequence, call it $\hat{U}_k^n[H]$ for all $k \in H$. If there is no such sequence or more than one, define this to be null.

For all k, if there is exactly one non-null value of $\hat{U}_k^n[H]$ among all $H \ni k$, then call this sequence $\hat{U}_k^n$. If the values of $\hat{U}_k^n[H]$ are all null or they are inconsistent, then set $\hat{U}_k^n = 0^n$. Let Σ be the sample covariance of $\hat{U}^{nL}$, with the exception that if $\hat{U}_H^n$ is jointly typical, we reset

$$\Sigma_{j,k} = \sigma_X^2 + \frac{\sigma_{N_k}^2}{1-\exp(-2r_k)}\delta_{j,k} \tag{17}$$

for all $j,k \in H$. This value is identical to $\mathbb{E}(U_jU_k)$, so if the pair (U_j,U_k) is already jointly typical, this revision does not change Σ much. Finally, the CEO chooses for its estimate

$$\hat{X}^n = \sum_{k=1}^{L} c_k(\Sigma)\hat{U}_k^n. \tag{18}$$

A. Error Analysis

Let H be the true set of honest agents. It can be shown that if (4) is satisfied, then with high probability $(X^n,\hat{U}_H^n)$ are jointly typical. Let λ be the sample covariance between $\hat{U}^{nL}$ and X^n, but again with the exception that we reset $\lambda_k = \sigma_X^2$ for all $k \in H$. As with our construction of Σ, because $\mathbb{E}(XU_k) = \sigma_X^2$, the joint typicality of $(X^n,\hat{U}_H^n)$ ensures that this revision is slight. Hence

$$\frac{1}{n}\sum_{t=1}^{n}(X(t)-\hat{X}(t))^2 \le \mathbb{E}_{\Sigma,\lambda}\left(X - \sum_{k=1}^{L} c_k(\Sigma)U_k\right)^2 + \dot{\epsilon} \tag{19}$$

where $\dot{\epsilon} \to 0$ as $\epsilon \to 0$. Therefore, by (6)

$$\frac{1}{n}\sum_{t=1}^{n}(X(t)-\hat{X}(t))^2 \le D + \dot{\epsilon}. \tag{20}$$

Taking the limit as $\epsilon \to 0$ proves achievability.

IV. Outer Bound Proof

Assume $(R_1,\dots,R_L,D)$ is achievable, and consider a code that achieves it with codewords $(C_1,\dots,C_L)$. We may assume without loss of generality that the code achieves distortion D with probability at least $1-\epsilon$, because we can always repeat the code multiple times and apply the law of large numbers. Fix $S \subset \{1,\dots,L\}$ with $|S| = (1-2\beta)L$, and $A \subset S$. A standard inequality train yields

$$\sum_{k\in A} R_k \ge \frac{1}{n}I(X^n;C_S) - \frac{1}{n}I(X^n;C_{S\setminus A}) + \sum_{k\in A} r_k \tag{21}$$

where $r_k = \frac{1}{n}I(Y_k^n;C_k|X^n)$. Lemma 3.1 in [5] states that for any $B \subset \{1,\dots,L\}$,

$$\frac{1}{\sigma_X^2}\exp\left(\frac{2}{n}I(X^n;C_B)\right) \le \frac{1}{\sigma_X^2} + \sum_{k\in B}\frac{1-\exp(-2r_k)}{\sigma_{N_k}^2} \tag{22}$$

which allows us to bound the second term in (21). We will proceed to show that

$$\frac{1}{n}h(X^n|C_S) \le \frac{1}{2}\log 2\pi e D \tag{23}$$

which, applied to (21) along with (22), gives (9). Taking $A = \emptyset$ yields (10).

We now prove (23). Let H_1, H_2 be sets of size $(1-\beta)L$ such that $S = H_1 \cap H_2$. If H_i is the true set of honest agents, for $i = 1$ or 2, then they use the deterministic encoding functions f_k to get C_{H_i} from $Y_{H_i}^n$. Meanwhile, the traitors, H_i^c, choose $C_{H_i^c}$. The CEO's estimate $\hat{X}^n$ is effectively a deterministic function of $Y_{H_i}^n$ and $C_{H_i^c}$. Thus we can define the set

$$S_D(X,Y_{H_i}) = \Big\{(x^n,y_{H_i}^n) : \forall c_{H_i^c}, \frac{1}{n}d(x^n,\hat{X}^n(y_{H_i}^n,c_{H_i^c})) \le D\Big\}. \tag{24}$$

This is the set of all $(x^n,y_{H_i}^n)$ pairs for which $\hat{X}^n$ achieves the distortion constraint no matter what the traitors do. Because we assume that distortion D is achieved with probability nearly one, the probability of the set $S_D(X,Y_{H_i})$ is also nearly one.

Now define

$$Q_D(X,Y_S) = \big\{(x^n,y_S^n) : \exists y_{H_1\setminus H_2}^n, y_{H_2\setminus H_1}^n : (x^n,y_{H_1}^n) \in S_D(X,Y_{H_1}), (x^n,y_{H_2}^n) \in S_D(X,Y_{H_2})\big\}. \tag{25}$$

That is, $Q_D(X,Y_S)$ is the set of pairs (x^n,y_S^n) such that $\hat{X}^n$ may achieve the distortion constraint (depending on the Y values) if either H_1 or H_2 is the set of honest agents. Because the S_D sets have probability nearly one, so does Q_D.

Given a codeword c_S, let $Q_D(X|c_S)$ be the set of x^n such that $(x^n,y_S^n) \in Q_D(X,Y_S)$ for some y_S^n for which $f_S(y_S^n) = c_S$. It follows from the high probability property of $Q_D(X,Y_S)$ that $Q_D(X|c_S)$ also has high probability conditioned on c_S being sent. Hence

$$\frac{1}{n}h(X^n|C_S) \le \frac{1}{n}\max_{c_S}\log \mathrm{Vol}(Q_D(X|c_{H\cap S})) + \ddot{\epsilon} \tag{26}$$

where $\ddot{\epsilon} \to 0$ as $\epsilon \to 0$.

Consider two elements x^n, x'^n of $Q_D(X|c_S)$. From the definition of Q_D, there must be $y_{H_1}^n$ and $y_{H_2}'^n$ such that

$$(x^n,y_{H_1}^n) \in S_D(X,Y_{H_1}), \quad (x'^n,y_{H_2}'^n) \in S_D(X,Y_{H_2}). \tag{27}$$

Consider the case that c_S, $c_{H_1\setminus H_2} = f_{H_1\setminus H_2}(y_{H_1\setminus H_2}^n)$, and $c_{H_2\setminus H_1} = f_{H_2\setminus H_1}(y_{H_2\setminus H_1}'^n)$ are sent. First observe that this

set of messages could have been produced if $X^n = x^n$, $Y^n_{H_1} = y^n_{H_1}$, and H_1 were the set of honest agents. Then the agents in $H_2 \setminus H_1$, which are all traitors, could send $c_{H_2 \setminus H_1}$. Since $(x^n, y^n_{H_1}) \in S_D(X, Y_{H_1})$, by definition the estimate $\hat{x}^n$ produced at the CEO must satisfy $\frac{1}{n} d(x^n, \hat{x}^n) \leq D$. However, the same set of messages could have been produced if $X^n = x'^n$, $Y^n_{H_2} = y^n_{H_2}$, and H_2 were the set of honest agents, where $H_1 \setminus H_2$ decide to send $c_{H_1 \setminus H_2}$. Since the CEO produces just one estimate given a set of input messages, the very same estimate $\hat{x}^n$, must satisfy $\frac{1}{n} d(x'^n, \hat{x}^n) \leq D$. Hence, by the triangle inequality, for any $x^n, x'^n \in Q_D(X|c_S)$,

$$\|x^n - x'^n\|_2 \leq 2\sqrt{nD}. \tag{28}$$

That is, $Q_D(X|c_S)$ has diameter at most $2\sqrt{nD}$. The following lemma bounds the volume of subsets of $\mathbb{R}^n$ as a function of their diameter. The proof, which we omit, makes use of the Brunn-Minkowski inequality.

Lemma 1: The volume of any subset of $\mathbb{R}^n$ is no more than that of the n-ball with the same diameter.

Hence, the volume of $Q_D(X|c_S)$ is no more than that of an n-ball with radius $\sqrt{nD}$, which is less than $(2\pi eD)^{n/2}$. Applying this to (26) gives (23), completing the proof.

V. Proof of Asymptotic Bounds

The proof of the lower bound in (12) using Theorem 2 is straightforward and we omit it. We proceed to prove the upper bound in (12) using Theorem 1.

For a given sum-rate R, we must specify r_k and c_k to satisfy conditions (A) and (B) in the statement of Theorem 1. Let $R_k = R/L$ for all k. Let r be the largest possible value satisfying (4) where $r_k = r/L$. It is not hard to show that for large L and R, r is nearly equal to R.

For all $A \subset \{1, \ldots, L\}$, let $\hat{X}_A = \mathbb{E}(X|U_A)$. When X and U_A are related according to the nominal distribution, for fixed $|A|/L$ and large L and R,

$$\mathbb{E}(X - \hat{X}_A)^2 \simeq \frac{\sigma_N^2}{2R} \frac{L}{|A|}. \tag{29}$$

Also observe that if $B \subset A$,

$$\mathbb{E}(\hat{X}_A - \hat{X}_B)^2 = \mathbb{E}(X - \hat{X}_B)^2 - \mathbb{E}(X - \hat{X}_A)^2. \tag{30}$$

We choose the c_k in the following way. Given Σ, we look for a set $\hat{H} \subset \{1, \ldots, L\}$ of size $(1-\beta)L$ that has the anticipated distribution if $\hat{H}$ were the set of honest agents. That is, (5) holds for $\hat{H}$. If there is more than one such $\hat{H}$, choose between them arbitrarily. Define c_k such that

$$\sum_{k=1}^{L} c_k U_k = \hat{X}_{\hat{H}}. \tag{31}$$

Now we show that this choice satisfies condition (2) for a value of D giving the upper bound in (12). To do this, we must consider all possible values of λ. In the worst case, the true set of honest agents H shares just $(1-2\beta)L$ agents with $\hat{H}$. Because $U_{\hat{H}}$ is distributed according to the nominal distribution,

$$\mathbb{E}_{\Sigma,\lambda}(\hat{X}_{\hat{H}} - \hat{X}_{\hat{H}\cap H})^2 = \mathbb{E}(X - \hat{X}_{\hat{H}\cap H})^2 - \mathbb{E}(X - \hat{X}_{\hat{H}})^2. \tag{32}$$

Furthermore, since $\hat{H} \cap H$ contains only honest agents,

$$\mathbb{E}_{\Sigma,\lambda}(\hat{X}_{\hat{H}\cap H} - X)^2 = \mathbb{E}(\hat{X}_{\hat{H}\cap H} - X)^2. \tag{33}$$

The Cauchy-Shwartz inequality and (29) can now be used to show

$$\mathbb{E}_{\Sigma,\lambda}(\hat{X}_{\hat{H}} - X)^2 \lesssim \frac{\sigma_N^2}{2R} \frac{\sqrt{1-\beta} + \sqrt{\beta}}{(1-\beta)(\sqrt{1-\beta} - \sqrt{\beta})}. \tag{34}$$

This proves the upper bound in (12).

VI. Conclusion

We presented inner and outer bounds for the rate region of the quadratic Gaussian CEO problem as a function of the size of a subset of agents reprogrammed by an adversary. These bounds were used to bound the asymptotic proportionality constant for many agents and high sum-rates.

We conjectured that our inner bound on the rate region is tight, which would indicate that our upper bound on the proportionality constant is approximately tight for a small fraction of traitors. In particular, if our conjecture holds then the proportionality constant goes like $1 + 2\sqrt{\beta}$ near $\beta = 0$, as opposed to our lower bound of $1 + 2\beta$. Hence, if our conjecture is true, then a small number of traitors would have a surprisingly damaging effect on the quality of the CEO's estimate.

References

[1] T. Berger, Z. Zhang, and H. Viswanathan, "The CEO problem," *IEEE Trans. Inform. Theory*, vol.-42, pp. 887–902, 1996.

[2] H. Viswanathan and T. Berger, "The quadratic Gaussian CEO problem," *IEEE Trans. Inform. Theory*, vol. 43, pp. 1549–1559, 1997.

[3] Y. Oohama, "The rate-distortion function for the quadratic Gaussian CEO problem," *IEEE Trans. Inform. Theory*, vol 44, pp. 55–67, 1998.

[4] Y. Oohama, "Rate-distortion theory for Gaussian multiterminal source coding systems with several side informations at the decoder," *IEEE Trans. Inform. Theory*, vol. 51, pp. 2577-2593, 2005.

[5] V. Prabhakaran, D. Tse, and K. Ramchandran, "Rate region of the quadratic Gaussian CEO problem'," in *Proc. Int. Symp. Inf. Theory*, 2004.

[6] S. Y. Tung, *Multiterminal source coding*. PhD thesis, Cornell University, Ithaca, NY, 1978.

[7] T. Berger, "Multiterminal source coding" in *The Information Theory Approach to Communications* (CISM Courses and Lectures, no. 229), G. Longo, Ed. Vienna and New York: Springer Verlag, pp. 171–231, 1978.

[8] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," *ACM Transactions on Programming Languages and Systems*, vol. 4, pp. 382–401, 1982.

[9] D. Dolev, "The Byzantine generals strike again," *Journal of Algorithms*, vol. 3, no. 1, pp. 14–30, 1982.

[10] S. Jaggi, et al, "Resilient Network Coding in the Presence of Byzantine Adversaries," in *Proc. INFOCOM*, pp. 616–624, 2007.

[11] N. Cai and R. W. Yeung, "Network error correction, part I: Basic concepts and upper bounds," *Comm. in Inf. and Syst.*, vol. 6, no. 1, pp 19–36, 2006.

[12] O. Kosut and L. Tong, "Distributed source coding in the presence of Byzantine sensors," *IEEE Trans. Inform. Theory*, vol. 54, pp. 2550–2565, 2008.

[13] D. Slepian and J. Wolf, "Noiseless coding of correlated information sources," *IEEE Trans. Inform. Theory*, vol. 19, pp. 471–480, 1973.

[14] O. Kosut and L. Tong, "The CEO problem," in *Proc. Int. Symp. Inf. Theory*, 2008.

[15] O. Kosut and L. Tong, "A characterization of the error exponent for the Byzantine CEO problem," in *Proc. 46th Allerton Conf. on Comm., Control and Comp.*, 2008.

[16] J. Chen and T. Berger, "Robust Distributed Source Coding," *IEEE Trans. Inform. Theory*, vol. 54, pp. 3385–3398, 2008.