

Converses for Distributed Estimation via Strong Data Processing Inequalities

Aolin Xu and Maxim Raginsky

Abstract—We consider the problem of distributed estimation, where local processors observe independent samples conditioned on a common random parameter of interest, map the observations to a finite number of bits, and send these bits to a remote estimator over independent noisy channels. We derive converse results for this problem, such as lower bounds on Bayes risk. The main technical tools include a lower bound on the Bayes risk via mutual information and small ball probability, as well as strong data processing inequalities for the relative entropy. Our results can recover and improve some existing results on distributed estimation with noiseless channels, and also capture the effect of noisy channels on the estimation performance.

Index Terms—Distributed estimation, Bayes risk, strong data processing inequalities.

I. INTRODUCTION

The problem of distributed estimation arises when the estimator does not have direct access to the samples generated according to the parameter of interest, but only to the data received from local processors that observe the samples. In this work, we consider a general model of distributed estimation, where each of the m processors observes n independent samples drawn conditionally on a common d -dimensional parameter, generates a b -bit quantized message, and sends it to a remote estimator with T uses of an independent noisy channel. We derive lower bounds on the Bayes risk and on the minimum b or T needed to achieve a certain Bayes risk. Fundamental limits of similar problems have been studied recently by Duchi et al. [1] and Shamir [2] with the assumption of noiseless channels (cf. also earlier work by Gallager [3] and by Han and Amari [4]).

To some extent, the parameter to be estimated in the problem under consideration can be viewed as a message to be sent in a transmission system, and the samples to be processed and quantized can be viewed as the input data to a compression system. However, a few important features make the problem distinct from data compression and transmission. First, the dimension of the parameter may be fixed and not grow with the number of channel uses. Second, due to communication and computation constraints, the number of bits in the quantized message may not grow with the sample size. For example, as pointed out in [4], the samples can be compressed at asymptotically zero rate, which makes it

impossible to reconstruct the samples, yet still suffices to reliably estimate the parameter. The number of bits in the quantized message may not grow with the number of channel uses either. Due to these features, the conventional coding theorems in information theory cannot be applied here, but we can still use information-theoretic techniques to derive fundamental limits for the general problem of distributed estimation.

One of the major tools we use is a lower bound on the Bayes risk in terms of mutual information and small ball probability, which we derive using techniques introduced in our earlier work [5]. Another major tool is the strong data processing inequality (SDPI) for relative entropy [6]–[8], which lets us quantify the contraction of mutual information caused by communication constraints.

The general results we obtain are non-asymptotic in d , n , b , T and m , and can be used to derive asymptotic results. Examples are given for estimating both discrete and continuous parameters, where the converses closely match achievable performance. Moreover, our results can be naturally applied to minimax lower bounds, since the latter are always lower-bounded by the Bayes risk. We start with the single-processor setting, and then generalize the results to the multi-processor setting. We are able to recover and improve some existing results on distributed estimation with noiseless channels [1] as special cases, while our general results can capture the effect of noisy channels on the estimation performance.

II. MAIN TOOLS

A lower bound on Bayes risk. In the standard Bayesian estimation framework, $\mathcal{P} = \{P_{X|W=w} : w \in W\}$ is a family of distributions on an observation space X , and the parameter space W is endowed with a prior P_W . We estimate W from $X \sim P_{X|W}$ as $\widehat{W} = \psi(X)$, via an estimator ψ . Given a distortion function $\ell : W \times W \rightarrow \mathbb{R}^+$, define the Bayes risk

$$R_B = \inf_{\psi} \mathbb{E}[\ell(W, \widehat{W})].$$

For a given ψ , the excess distortion probability $\mathbb{P}(\ell(W, \widehat{W}) > \rho)$ can be lower bounded in terms of the mutual information $I(W; \widehat{W})$ and the so-called *small ball probability* of W with respect to distortion function ℓ [5], defined as

$$\mathcal{L}(W, \rho) = \sup_{w \in W} \mathbb{P}(\ell(W, w) \leq \rho).$$

It measures the “spread” of the prior distribution P_W . The lower bound on $\mathbb{P}(\ell(W, \widehat{W}) > \rho)$ can be conveniently converted to a lower bound on $\mathbb{E}[\ell(W, \widehat{W})]$ through Markov’s

The authors are with the Department of Electrical and Computer Engineering and the Coordinated Science Laboratory, University of Illinois, Urbana, IL 61801, USA. E-mails: {aolinxu2,maxim}@illinois.edu.

Research supported in part by the NSF under award no. CCF-1017564, by CAREER award no. CCF-1254041, and by the Center for Science of Information (CSoI), an NSF Science and Technology Center, under grant agreement CCF-0939370.

inequality. Using the techniques from our earlier work [5], we can obtain the following lower bound on the Bayes risk:

Theorem 1. *In the above Bayesian estimation framework,*

$$R_B \geq \sup_{\rho > 0} \rho \left(1 - \frac{I(W; X) + \log 2}{\log(1/\mathcal{L}(W, \rho))} \right).$$

Similar methods to derive Bayes risk lower bounds have been recently proposed by Chen et al. [9], where they obtained lower bounds in terms of general f -informativities [10] and a quantity essentially the same as the small ball probability. Theorem 1 reveals two sources of the intrinsic difficulty of estimating W : the amount of information about W contained in the observation X , captured by $I(W; X)$, and the spread of the prior distribution P_W , captured by $\mathcal{L}(W, \cdot)$. When an estimator does not have direct access to X but only through one or more local processors, the mutual information between W and the estimator's indirect observations will be a contraction of $I(W; X)$. The contraction is caused by the communication constraints between the local processors and the estimator, such as storage limitations of intermediate results, limited transmission blocklength, channel noise, etc.

Contraction of mutual information via SDPI. We quantify the contraction of mutual information using *strong data processing inequalities* for the relative entropy (see [8] and references therein). Given a stochastic kernel (channel) K with input alphabet \mathcal{X} and output alphabet \mathcal{Y} , and a reference input distribution μ on \mathcal{X} , we say that K satisfies an SDPI at μ with constant $c \in [0, 1]$ if $D(\nu K \parallel \mu K) \leq cD(\nu \parallel \mu)$ for any other input distribution ν on \mathcal{X} . Here, μK denotes the marginal distribution of the channel output when the input has distribution μ . The tightest such constants,

$$\eta(\mu, K) \triangleq \sup_{\nu: \nu \neq \mu} \frac{D(\nu K \parallel \mu K)}{D(\nu \parallel \mu)}, \quad \eta(K) \triangleq \sup_{\mu} \eta(\mu, K),$$

are also the maximum contraction ratios of mutual information in a Markov chain [7]: for a Markov chain $W - X - Y$,

$$\sup_{P_{W|X}} \frac{I(W; Y)}{I(W; X)} = \eta(P_X, P_{Y|X}) \quad (1)$$

if the joint distribution $P_{X,Y}$ is fixed, and

$$\sup_{P_{W,X}} \frac{I(W; Y)}{I(W; X)} = \eta(P_{Y|X}) \quad (2)$$

if only the channel $P_{Y|X}$ is fixed. It is generally hard to precisely compute the SDPI constant for an arbitrary pair of μ and K , except for some special cases. One such case is that for binary symmetric channel, $\eta(\text{Bern}(\frac{1}{2}), \text{BSC}(\varepsilon)) = \eta(\text{BSC}(\varepsilon)) = (1 - 2\varepsilon)^2$ [11]. Various upper bounds on SDPI constants have been proposed (see [8] and references therein). We will need one such bound [6]:

Lemma 1. *Define the Dobrushin contraction coefficient of a channel $P_{X|W}$ by $\vartheta(P_{X|W}) = \max_{w,w'} \|P_{X|W=w} - P_{X|W=w'}\|_{\text{TV}}$. Then $\eta(P_{X|W}) \leq \vartheta(P_{X|W})$.*

For product input distributions and product channels, the SDPI constant tensorizes [7] (see [8] for a more general result for other f -divergences):

Lemma 2. *For distributions μ_1, \dots, μ_n on \mathcal{X} and channels K_1, \dots, K_n with input alphabet \mathcal{X} ,*

$$\eta(\mu_1 \otimes \dots \otimes \mu_n, K_1 \otimes \dots \otimes K_n) = \max_{1 \leq i \leq n} \eta(\mu_i, K_i).$$

Finally, motivated by Evans and Schulman [12] and by Polyanskiy and Wu [13], the following lemma characterizes the SDPI constant for multiple uses of a channel:

Lemma 3. *For a stochastic kernel $P_{V|U}$, consider the stationary and memoryless channel $P_{V^T|U^T} = P_{V|U}^T$. The SDPI constant of $P_{V^T|U^T}$ satisfies*

$$\eta(P_{V^T|U^T}) \leq 1 - (1 - \eta(P_{V|U}))^T.$$

III. RESULTS FOR A SINGLE PROCESSOR

Consider the following distributed estimation problem with a single processor, shown schematically in Fig. 1:




Fig. 1. General model (single processor).

- $W = (W_1, \dots, W_d)$ is a random parameter (discrete or continuous) with mutually independent coordinates.
- The $d \times n$ array of observations X^n is generated conditionally on W as follows: For each $1 \leq j \leq d$, given $W_j = w_j$, the n samples in the j th row of X^n , denoted by $X_j^n = (X_{j,1}, \dots, X_{j,n})$, are independently generated according to a given stochastic kernel $P_{X_j|W_j=w_j}$.
- The local processor observes X^n and generates a b -bit message $Y = \varphi_1(X^n)$.
- The encoder maps Y to a codeword $U^T = \varphi_2(Y)$ with blocklength T , and transmits U^T over the noisy channel. The channel is memoryless, with stochastic kernel $P_{V|U}$.
- The remote estimator ψ estimates W from the received codeword V^T , so that $\widehat{W} = \psi(V^T)$.

The Bayes risk in this problem setup is defined as

$$R_B = \inf_{\varphi_1, \varphi_2, \psi} \mathbb{E}[\ell(W, \psi(V^T))].$$

In order to apply Theorem 1, we need an upper bound on the mutual information $I(W; V^T)$ which is independent of φ_1 , φ_2 , and ψ . All logarithms are binary, unless stated otherwise.

Theorem 2. *For any choice of φ_1 , φ_2 , and ψ ,*

$$I(W; V^T) \leq \min \left\{ (H(X^n) \wedge b) \eta_T \max_{1 \leq j \leq d} \eta(P_{X_j^n}, P_{W_j|X_j^n}), \right. \\ \left. I(W; X^n) \eta_T, CT \right\}$$

where $r \wedge s \triangleq \min\{r, s\}$, $\eta_T \triangleq 1 - (1 - \eta(P_{V|U}))^T$, and C is the Shannon capacity of the channel $P_{V|U}$.

Proof: Consider the Markov chain $W - X^n - Y - U^T - V^T$. From Lemma 3, Eq. (2), and the ordinary DPI, we have

$$I(W; V^T) \leq I(W; U^T)\eta_T \leq I(W; Y)\eta_T. \quad (3)$$

On the one hand,

$$I(W; Y) \leq I(X^n; Y)\eta(P_{X^n}, P_{W|X^n}) \quad (4)$$

as a consequence of (1). Lemma 2 and the fact that $(W_1, X_1^n), \dots, (W_d, X_d^n)$ are independent imply that

$$\eta(P_{X^n}, P_{W|X^n}) = \max_{1 \leq j \leq d} \eta(P_{X_j^n}, P_{W_j|X_j^n})$$

Finally, since Y takes values in $[2^b]$,

$$I(X^n; Y) \leq \min\{H(X^n), H(Y)\} \leq \min\{H(X^n), b\}.$$

Using these bounds in (3) and (4), we get

$$I(W; V^T) \leq (H(X^n) \wedge b) \max_{1 \leq j \leq d} \eta(P_{X_j^n}, P_{W_j|X_j^n})\eta_T.$$

Alternatively, using $I(W; Y) \leq I(W; X^n)$ in (3), we get $I(W; V^T) \leq I(W; X^n)\eta_T$. Lastly, because the noisy channel is memoryless, we have $I(W; V^T) \leq I(U^T; V^T) \leq CT$. We complete the proof by taking the minimum of the three resulting estimates to get the tightest bound on $I(W; V^T)$. ■

Next we study a few examples of this problem setup to illustrate the effectiveness of using Theorem 1 and Theorem 2 to derive converse results for the Bayes risk.

Example 1: Transmitting a bit over a BSC. Suppose W is $\text{Bern}(\frac{1}{2})$, $W = X^n = Y$, $P_{V|U}$ is $\text{BSC}(\varepsilon)$, so that $\eta(P_{V|U}) = (1 - 2\varepsilon)^2$, and $\ell(w, \hat{w}) = \mathbf{1}\{w \neq \hat{w}\}$.

Corollary 1. *The minimum blocklength T^* to achieve $R_B \leq p$ satisfies*

$$T^* \geq \frac{\log \frac{1}{h(\frac{p})}}{\log \frac{1}{4\varepsilon\bar{\varepsilon}}} \geq \frac{\log \frac{1}{p} - \log \log \frac{e}{p}}{\log \frac{1}{4\varepsilon\bar{\varepsilon}}} \sim \frac{\log \frac{1}{p}}{\log \frac{1}{4\varepsilon\bar{\varepsilon}}} \text{ as } p \rightarrow 0,$$

where $h(\cdot)$ is the binary entropy function, and $\bar{\varepsilon} \triangleq 1 - \varepsilon$.

Proof: In this case, we can bypass Theorem 1 by using the bound $1 - h(\mathbb{P}(\widehat{W} \neq W)) \leq I(W; V^T)$. Theorem 2 gives $I(W; V^T) \leq I(W; X^n)\eta_T \leq 1 - (4\varepsilon\bar{\varepsilon})^T$. We obtain the lower bound using the fact that $h(p) \leq p \log \frac{e}{p}$. ■

The blocklength of a repetition code with error probability $\leq p$ gives an upper bound on T^* . By the Chernoff bound [14], a blocklength- T repetition code can achieve $\mathbb{P}(\widehat{W} \neq W) \leq 2^{-\frac{T}{2} \log \frac{1}{4\varepsilon\bar{\varepsilon}}}$. Thus

$$T^* \leq 2 \log \frac{1}{p} / \log \frac{1}{4\varepsilon\bar{\varepsilon}}.$$

We see that the upper and lower bounds on T^* only differ by a factor of 2 as $p \rightarrow 0$, and have the same dependence on ε .

Example 2: Estimating a discrete parameter. Consider the case where W is uniformly distributed on $\{\pm 1\}^d$ and $X^n \in \{\pm 1\}^{d \times n}$. Given some fixed $\delta \in [0, 1]$, $P_{X_j|W_j}(x_{j,k}|w_j) = (1 + x_{j,k}w_j\delta)/2$ for $j \in \{1, \dots, d\}$ and $k \in \{1, \dots, n\}$. In other words, $P_{X_j|W_j}$ is $\text{BSC}(\frac{1-\delta}{2})$. It follows that $X_{j,k}$ is

uniform on $\{\pm 1\}$, and $P_{W_j|X_{j,k}}$ is $\text{BSC}(\frac{1-\delta}{2})$ as well. Channel $P_{V|U}$ is assumed to be arbitrary.

Corollary 2. *In Example 2, for $n = 1$,*

$$I(W; V^T) \leq \min \left\{ (d \wedge b)\delta^2\eta_T, CT \right\}. \quad (5)$$

For $n > 1$, with $\beta \triangleq \frac{1-\delta}{1+\delta}$ and $\xi_n \triangleq \frac{1-\beta^n}{1+\beta^n}$,

$$I(W; V^T) \leq \min \left\{ (d(1 + nh(\frac{1-\delta}{2})) \wedge b) \xi_n \eta_T, d\eta_T, CT \right\}.$$

Proof: For $n = 1$, we have the exact SDPI constant $\eta(P_{X_j}, P_{W_j|X_j}) = \delta^2$, due to the fact that X_j is uniform on $\{\pm 1\}$ and $P_{W_j|X_j}$ is $\text{BSC}(\frac{1-\delta}{2})$.

For $n > 1$, by Lemma 1,

$$\eta(P_{X_j^n}, P_{W_j|X_j^n}) \leq \eta(P_{W_j|X_j^n}) \leq \vartheta(P_{W_j|X_j^n}), \quad (6)$$

where a simple computation gives $\vartheta(P_{W_j|X_j^n}) = (1 - \beta^n)/(1 + \beta^n)$. We also have $I(W; X^n) \leq d$, and

$$\begin{aligned} H(X^n) &= dH(X_1^n) \leq dH(W_1, X_1^n) \\ &= d(H(W_1) + H(X_1^n|W_1)) = d(1 + nh(\frac{1-\delta}{2})). \end{aligned}$$

The results then follow from Theorem 2. ■

Duchi et al. [1] considered the same problem with $n = 1$ and noiseless $P_{V|U}$. Their result (Lemma 3 in [1]), proved in a much more complicated way, shows that

$$I(W; Y) \leq \min\{d, b\}32\delta^2/(1 - \delta)^4 \quad (7)$$

where the contraction coefficient is less than 1 only when $\delta < 0.133$. In contrast, the contraction coefficient in (5) can never go greater than 1, and it considerably improves the contraction coefficient in (7) over all $\delta \in [0, 1]$, especially for large δ , under the same noiseless channel assumption. Combined with Theorem 1, Corollary 2 can be applied to derive lower bounds on the minimax risk in estimating the mean of an arbitrary probability distribution on the cube $[-1, 1]^d$. We discuss this application in Sec. IV, in the multi-processor setting.

Using Corollary 2, we can obtain lower bounds on the bit error probability for estimating W and on the number of bits to quantize the message Y .

Corollary 3. *In Example 2, let $\ell(w, \hat{w}) = \frac{1}{d} \sum_{j=1}^d \mathbf{1}\{w_j \neq \hat{w}_j\}$. Then, for $n = 1$,*

$$R_B \geq h^{-1} \left(1 - \frac{1}{d} \min \{b\delta^2\eta_T, CT\} \right),$$

provided b, d , and T are such that the argument of $h^{-1}(\cdot)$ lies in $[0, 1]$.

Proof: Let $d_2(\cdot|\cdot)$ be the binary divergence function; then, choosing $\varphi_1, \varphi_2, \psi$ that attain R_B , we have

$$\begin{aligned} 1 - h(R_B) &= d_2(R_B|\frac{1}{2}) \leq \frac{1}{d} \sum_{j=1}^d d_2(\mathbb{P}(W_j \neq \widehat{W}_j)|\frac{1}{2}) \\ &\leq \frac{1}{d} \sum_{j=1}^d I(W_j; \widehat{W}_j) \leq \frac{1}{d} I(W; \widehat{W}) \leq \frac{1}{d} (b\delta^2\eta_T \wedge CT), \end{aligned}$$

where the first line uses the convexity of divergence, and the second line uses the data processing inequality for divergence, the fact that W_j 's are i.i.d., and Corollary 2. Applying h^{-1} to both sides, we get the result. ■

Corollary 4. *In Example 2, for $n = 1$, to achieve $R_B \leq p$, it is necessary that*

$$\frac{b}{d} \geq \frac{1 - h(p)}{\delta^2 \eta_T} = \frac{1 - h(p)}{\delta^2 (1 - (1 - \eta(P_{V|U}))^T)}.$$

In Fig. 2, this lower bound is compared with the asymptotic compression ratio $\tilde{R}(p) = 1 - h(\frac{2p + \delta - 1}{2\delta})$, $0 \leq \frac{1 - \delta}{2} \leq p \leq \frac{1}{2}$, of noisy lossy coding of an i.i.d. Bern($\frac{1}{2}$) source over BSC($\frac{1 - \delta}{2}$), and also with the rate-distortion function $R(p) = 1 - \tilde{h}(p)$ of Bern($\frac{1}{2}$).

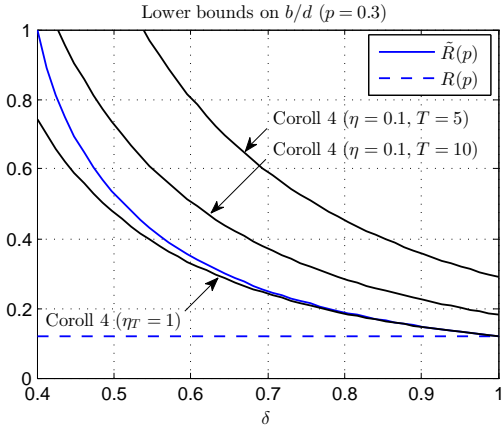


Fig. 2. Comparison of lower bounds on b/d ($p = 0.3$).

Example 3: Estimating a continuous parameter. Consider the problem of estimating the bias of a Bernoulli random variable through a BSC. In this case, W is assumed to be uniformly distributed on $[0, 1]$, $P_{X|W=w}$ is Bern(w), and $P_{V|U}$ is BSC(ε). We are interested in lower-bounding the Bayes risk with respect to the absolute loss $\ell(w, \hat{w}) = |w - \hat{w}|$.

Corollary 5. *In Example 3, let $I^* \triangleq \sup_{\varphi_1, \varphi_2, \psi} I(W; V^T)$. Then the Bayes risk can be lower-bounded by*

$$R_B \geq \frac{1}{16} 2^{-2I^*} \quad (8)$$

for all values of I^* , and by

$$R_B \gtrsim \frac{1}{8I^*} 2^{-I^*} \quad (9)$$

for $I^* \rightarrow \infty$. The notation $y \gtrsim g(x)$ means that there exists some function f such that $y \geq f(x)$ for all x , and $\lim_{x \rightarrow \infty} f(x)/g(x) = 1$.

Proof: We have $\mathcal{L}(W, \rho) = \sup_{w \in [0, 1]} \mathbb{P}(|W - w| < \rho) \leq \min\{2\rho, 1\}$. From Theorem 1,

$$R_B \geq \sup_{0 < \rho < \frac{1}{2}} \rho \left(1 - \frac{I^* + \log 2}{\log(1/2\rho)}\right) \geq \frac{1}{2} \sup_{0 < s < 1} s 2^{-\frac{I^* + 1}{1-s}} \quad (10)$$

where the last inequality is obtained by requiring $1 - \frac{I^* + \log 2}{\log(1/2\rho)} > s$ for each $s \in (0, 1)$. In (10), taking $s = \frac{1}{2}$,

we get (8), while optimizing over s and sending $I^* \rightarrow \infty$, we get (9). ■

Corollary 6. *In Example 3, for any choice of $\varphi_1, \varphi_2, \psi$,*

$$I(W; V^T) \leq \min \left\{ b(1 - 2^{-n})(1 - (4\varepsilon\bar{\varepsilon})^T), \left(\frac{1}{2} \log n + \gamma_n\right)(1 - (4\varepsilon\bar{\varepsilon})^T), (1 - h(\varepsilon))T \right\},$$

where $\lim_{n \rightarrow \infty} \gamma_n = c$ with some absolute constant c .

Proof: From Lemma 1 and an explicit computation,

$$\eta(P_{W|X^n}) \leq \vartheta(P_{W|X^n}) = 1 - 2^{-n}.$$

Moreover (see, e.g., [15])

$$I(W; X^n) = \frac{1}{2} \log n + \gamma_n.$$

With these facts, the result follows from Theorem 2. ■

Now we apply Corollaries 5 and 6 to two specific cases:

Case 1: $\varepsilon = 0$, $T \geq b$. We have

$$R_B \gtrsim \frac{1}{8(1 - 2^{-n})b} 2^{-(1 - 2^{-n})b} \sim \frac{1}{4\sqrt{n} \log n}$$

for $b = \frac{1}{2} \log n$, and $n \rightarrow \infty$. To obtain an upper bound on R_B , consider the scheme where the local processor quantizes the sample mean $\bar{X}^n = n^{-1} \sum_{j=1}^n X_j$ into \tilde{W} using a uniform b -bit quantization of $[0, 1]$, and the remote estimator sets $\hat{W} = \tilde{W}$. By the triangle inequality,

$$\mathbb{E}|W - \hat{W}| \leq \mathbb{E}|W - \bar{X}^n| + \mathbb{E}|\bar{X}^n - \tilde{W}| \leq \frac{1}{\sqrt{6n}} + 2^{-b}.$$

Thus, for $b = \frac{1}{2} \log n$, $R_B \leq 1.41/\sqrt{n}$, which only differs from the lower bound by a logarithmic factor as $n \rightarrow \infty$.

Case 2: $\varepsilon > 0$, $b \geq \log(n + 1)$. We have

$$\begin{aligned} R_B &\gtrsim \max \left\{ \frac{c_1}{\eta_T \sqrt{n}^{\eta_T} \log n}, \frac{2^{-(1-h(\varepsilon))T}}{8(1-h(\varepsilon))T} \right\} \\ &\geq \frac{\alpha c_1}{\eta_T \sqrt{n}^{\eta_T} \log n} + \frac{\bar{\alpha} 2^{-(1-h(\varepsilon))T}}{8(1-h(\varepsilon))T}, \quad \forall \alpha \in [0, 1], \end{aligned}$$

where $\eta_T = 1 - (4\varepsilon\bar{\varepsilon})^T$, and c_1 is an absolute constant. Consider the scheme where the local processor computes the sample sum S^n , which is uniformly distributed on $\{0, \dots, n\}$, represents it with $\log(n + 1)$ bits, and transmits these bits over the channel using a blocklength- T code. The remote estimator decodes S^n as \hat{S}^n , and sets $\hat{W} = \hat{S}^n/n$. Then

$$\begin{aligned} \mathbb{E}|W - \hat{W}| &\leq \mathbb{E}|W - S^n/n| + \mathbb{E}|S^n/n - \hat{S}^n/n| \\ &\leq \frac{1}{\sqrt{6n}} + \mathbb{P}(S^n \neq \hat{S}^n) \leq \frac{1}{\sqrt{6n}} + 2^{-E_r(\frac{\log(n+1)}{T})T}, \end{aligned}$$

where $E_r(\cdot)$ is the random coding error exponent of BSC(ε):

$$E_r\left(\frac{1}{T} \log(n + 1)\right) = 1 - \log(1 + \sqrt{4\varepsilon\bar{\varepsilon}}) - \frac{1}{T} \log(n + 1)$$

when $\frac{1}{T} \log(n + 1) \leq 1 - h(\frac{\sqrt{\varepsilon}}{\sqrt{\varepsilon} + \sqrt{\bar{\varepsilon}}})$ [14, p. 146]. Note that

$$1 \leq \frac{1 - h(\varepsilon)}{1 - \log(1 + \sqrt{4\varepsilon\bar{\varepsilon}})} \leq 2, \quad \forall \varepsilon \in (0, \frac{1}{2}),$$

which implies that the error exponent in the lower bound can closely match that in the upper bound at low transmission rate.

IV. A RESULT FOR MULTIPLE PROCESSORS

We now consider a set-up with m local processors. Each processor observes an independent set of samples generated from a common random parameter W , and communicates with the remote estimator over an independent noisy channel. For notational simplicity, we assume that each processor applies the same local transformation to its samples, and that the channels between each processor and the remote estimator have the same transition probabilities. The results can be straightforwardly generalized to the case where the local encoders and the channels are different across the processors.

Theorem 3. *In the multi-processor setup described above,*

$$I(W; V^{m \times T}) \leq m \min \left\{ I(W; X^n) \eta_T, CT, \right. \\ \left. (H(X^n) \wedge b) \max_{1 \leq j \leq d} \eta(P_{X_j^n}, P_{W_j | X_j^n}) \eta_T \right\}.$$

Proof: Due to the independence assumption, the codewords $V^{m \times T} \triangleq (V_{(1)}^T, \dots, V_{(m)}^T)$ received by the remote estimator from the processors $\{1, \dots, m\}$ are conditionally independent given W . This implies that $I(W; V^{m \times T}) \leq \sum_{i=1}^m I(W; V_{(i)}^T)$ (see, e.g., [1, Lemma 4]). Using Theorem 2 to upper-bound each term, we obtain the result of Theorem 3. ■

Using Theorem 3 with Theorem 1 and Corollary 2, we can obtain a lower bound on the minimax risk for estimating the mean of an unknown distribution P on $X = [-1, 1]^d$, where each processor $i \in \{1, \dots, m\}$ only observes a single independent sample $X_{(i)} \sim P$. Let \mathcal{P} denote family of all probability distributions on $[-1, 1]^d$. For $P \in \mathcal{P}$, the parameter of P in this example is formally defined as $\theta(P) = \mathbb{E}_P[X]$. The minimax risk is defined as

$$R_M = \inf_{\varphi_1^m, \varphi_2^m, \psi} \sup_{P \in \mathcal{P}} \mathbb{E}_P[\|\theta(P) - \psi(V^{m \times T})\|^2],$$

where ψ is an estimator of $\theta \in [-1, 1]^d$.

Corollary 7. *For the above minimax estimation problem,*

$$R_M \geq \frac{d}{6} \min \left\{ 1, \frac{d}{24\eta_T m(d \wedge b)} \right\}, \text{ for } d \geq 12.$$

Proof: At a high level, the proof strategy follows that in [1]. But we use Theorem 1 instead of their distance-based Fano inequality, and we improve their mutual information upper bound by Corollary 2, which also captures the influence of noisy channels between the processors and the remote estimator. Let W , δ , and $P_{X_j | W_j}$ be defined as in Example 2 with $n = 1$. Given $W = w$, each processor observes a sample X with its coordinates drawn according to $P_{X_j | W_j = w_j}$. Hence $P_{X | W = w} \in \mathcal{P}$ for all $w \in \{\pm 1\}^d$. Let $\theta_w \triangleq \theta(P_{X | W = w}) = \delta w$, so that $\|\theta_w - \theta_{w'}\|^2 = 4\delta^2 \ell_H(w, w')$, where ℓ_H denotes the Hamming distance. Therefore,

$$R_M \geq 4\delta^2 \inf_{\varphi_1^m, \varphi_2^m, \psi} \inf \mathbb{E}[\ell_H(W, \widehat{W})], \quad (11)$$

where the second infimum is now over all remote estimators of $W \in \{-1, +1\}^d$. Let ρ be a nonnegative integer. Then

$\mathcal{L}(W, \rho) = \sum_{\tau=0}^{\rho} \binom{d}{\tau} / 2^d$, and $\log(1/\mathcal{L}(W, \rho)) \geq d/6$ for $\rho \leq d/6$ and $d \geq 12$. Thus from Theorem 1,

$$\inf_{\varphi_1^m, \varphi_2^m, \psi} \inf \mathbb{E}[\ell_H(W, \widehat{W})] \geq \rho \left(1 - \frac{I(W; V^{m \times T}) + \log 2}{\log(1/\mathcal{L}(W, \rho))} \right) \\ \geq \frac{d}{6} \left(1 - \frac{I(W; V^{m \times T}) + \log 2}{d/6} \right).$$

From Corollary 2 and Theorem 3, we have $I(W; V^{m \times T}) \leq m\delta^2 \eta_T \min\{d, b\}$. Thus from (11),

$$R_M \geq \frac{2d\delta^2}{3} \left(1 - \frac{m\delta^2 \eta_T (d \wedge b) + \log 2}{d/6} \right).$$

With $\delta^2 = \min\{1, d/(24m\eta_T(d \wedge b))\}$, the quantity in the parentheses is $\geq 1/4$, and we obtain the desired result. ■

In the noiseless channel case, Corollary 7 recovers and improves the lower bound in Proposition 2 of [1], which can be achieved within a constant factor using a method described there. When the processors communicate to the estimator via noisy channels, the effect of the noise on the minimax risk is captured by η_T in the denominator of the lower bound.

ACKNOWLEDGMENT

The authors would like to thank Y. Polyanskiy and Y. Wu for helpful discussions and for making Ref. [13] available.

REFERENCES

- [1] J. Duchi, M. Jordan, M. Wainwright, and Y. Zhang, "Optimality guarantees for distributed statistical estimation," *arXiv preprint*, 2014. [Online]. Available: <http://arxiv.org/abs/1405.0782>
- [2] O. Shamir, "Fundamental limits of online and distributed algorithms for statistical learning and estimation," in *28th Annual Conference on Neural Information Processing Systems (NIPS)*, 2014.
- [3] R. Gallager, "Finding parity in a simple broadcast network," *IEEE Trans. Inform. Theory*, vol. 34, no. 2, pp. 176–180, Mar 1988.
- [4] T. Han and S. Amari, "Statistical inference under multiterminal data compression," *IEEE Trans. Inform. Theory*, vol. 44, no. 6, Sep. 1998.
- [5] A. Xu and M. Raginsky, "A new information-theoretic lower bound on distributed function computation," in *Proc. Int. IEEE Symp. on Inform. Theory*, June 2014, pp. 2227–2231.
- [6] J. E. Cohen, Y. Iwasa, G. Rautu, M. B. Ruskai, E. Seneta, and G. Zbaganu, "Relative entropy under mappings by stochastic matrices," *Lin. Algebra Appl.*, vol. 179, pp. 211–235, 1993.
- [7] V. Anantharam, A. Gohari, S. Kamath, and C. Nair, "On maximal correlation, hypercontractivity, and the data processing inequality studied by Erkip and Cover," *arXiv preprint*, 2013. [Online]. Available: <http://arxiv.org/abs/1304.6133>
- [8] M. Raginsky, "Strong data processing inequalities and Φ -Sobolev inequalities for discrete channels," *arXiv preprint*, 2014. [Online]. Available: <http://arxiv.org/abs/1411.3575>
- [9] X. Chen, A. Guntuboyina, and Y. Zhang, "On Bayes Risk Lower Bounds," *arXiv preprint*, 2014. [Online]. Available: <http://arxiv.org/abs/1410.0503>
- [10] I. Csiszár, "A class of measures of informativity of observation channels," *Periodica Math. Hungar.*, vol. 2, no. 1–4, pp. 191–213, 1972.
- [11] R. Ahlswede and P. Gács, "Spreading of sets in product spaces and hypercontraction of the Markov operator," *Ann. Probab.*, vol. 4, no. 6, pp. 925–939, 1976.
- [12] W. Evans and L. Schulman, "Signal propagation and noisy circuits," *IEEE Trans. Inform. Theory*, vol. 45, no. 7, pp. 2367–2373, 1999.
- [13] Y. Polyanskiy and Y. Wu, "Strong data processing in networks: note on a method of Evans and Schulman," 2014, unpublished manuscript.
- [14] R. Gallager, *Information theory and reliable communication*. New York: Wiley, 1968.
- [15] B. S. Clarke and A. R. Barron, "Information-theoretic asymptotics of Bayes methods," *IEEE Trans. Inform. Theory*, vol. 36, no. 3, pp. 453–471, May 1990.