

A Fundamental Limit of Distributed Hypothesis Testing Under Memoryless Quantization

Yunus İnan, Mert Kayaalp, Ali H. Sayed, Emre Telatar
EPFL, Lausanne, Switzerland

Abstract—We consider a distributed binary hypothesis testing setup where multiple nodes send quantized information to a central processor, which is oblivious to the nodes’ statistics. We study the regime where the missed detection (type-II error) probability decays exponentially and the false alarm (type-I error) probability vanishes. For memoryless quantization, we characterize a tradeoff curve that yields a lower bound for the feasible region of type-II error exponents and the average number of bits sent under the null hypothesis. Moreover, we show that the tradeoff curve is approached at high rates with lattice quantization.

I. INTRODUCTION

Modern designs allow the construction of complex systems consisting of many smaller devices and/or sensors. For instance, consider an autonomous vehicle with multiple sensors sharing information with a central processor. The latter assists with critical decisions such as collision avoidance. This is one example of a distributed detection or distributed hypothesis testing problem. In such scenarios, the sensors usually cannot communicate at arbitrarily high rates; this may be due to the medium they communicate over, or due to the processing capabilities of the devices. Hence, it is reasonable to impose communication constraints on a distributed hypothesis testing problem.

Building upon the formulation in [1], communication constrained hypothesis testing has been studied from an information theoretic perspective by several authors — a useful survey appears in [2]. In its simple form, the problem proposed in [1] is as follows. We have independent and identically distributed (i.i.d.) random variables (X_i, Y_i) with $X^n := (X_1, \dots, X_n)$ observed at a remote sensor. An nR -bit description $f_n(X^n)$ is provided to a decision maker who also possesses Y^n , and performs a binary hypothesis test about the joint distribution P_{XY} . The aim in [1] was to characterize the best type-II error exponent with vanishing type-I error. Extensions of this problem were studied in later work, e.g., including the compression of Y^n [3], operating at zero-rate [3], [4], or communicating over noisy channels [5], [6]. Note that the communication constraint imposed in [1], [3] is ‘hard’, i.e., for any n , f_n can take at most 2^{nR} possible values. Recently, there has been some work centered around relaxed communication constraints, e.g., restricting the average number of bits sent for the task of independence testing [7], [8].

The signal detection community approaches the distributed hypothesis testing setup under communication constraints from a different perspective. They focus on scalar quantization of signals at peripheral nodes, i.e., quantization of X instead of X^n , and assume that a central processor (usually called a fusion center) combines the data conveyed from peripheral nodes, without side information. Deviating from block quantization procedures allows memory-efficiency and low-latency

at the peripheral nodes. This perspective requires the individual quantization procedures to be well-performing, if not optimal. For instance, in [9], the authors give an iterative procedure to find locally optimal quantizers for a class of performance metrics. The quantization scheme they assume is also ‘hard’ in the sense that the signal is quantized into N bins. Subsequent work also approached the problem by trying to find optimal truncation of log-likelihood ratios (LLRs) together with an optimal fusion rule [10], or by studying the point densities of optimal quantizers at high rates [11]. One can also find examples of studies on sparsifying the communication [12], [13] and including quantization to tackle noisy transmissions to the fusion center [14].

This work contains flavors from both approaches: (i) By adopting an information-theoretic perspective, we study the tradeoff between communication constraints and type-II exponents by formulating a problem amenable to rate-distortion methods; and (ii) we show that the limits can be approached closely (up to 1.047 bits of difference) at high communication rates by scalar lattice quantization, for which efficient encoding and decoding algorithms exist [15]. Unlike previous work, we provide impossibility results for memoryless quantization under an average communication constrained scheme. The detailed problem formulation and assumptions are given in Section III.

Notation: Random variables are denoted with uppercase letters whereas their realizations are denoted with lowercase letters (e.g., X_t and x_t). Sets and events are denoted with script-style letters (e.g., \mathcal{A}). $|\mathcal{A}|$ denotes the cardinality of set \mathcal{A} . Vectors and sequences are denoted by boldface letters (e.g., \mathbf{x}_t). $\mathcal{B}(\mathbb{R})$ denotes the Borel σ -algebra of \mathbb{R} . All the logarithms are assumed natural logarithms unless we explicitly state its base.

II. PROBLEM FORMULATION

Consider a binary hypothesis testing problem where m nodes communicate with a fusion center. At time t , under the null hypothesis \mathcal{H}_0 , node i observes data coming from distribution $P^{(i)}$ and under the hypothesis \mathcal{H}_1 observes data coming from distribution $Q^{(i)}$, with both distributions being time invariant. Assume $P^{(i)}$ and $Q^{(i)}$ admit density functions $p^{(i)}$, $q^{(i)}$. Furthermore, assume $P^{(i)}$ is absolutely continuous with respect to $Q^{(i)}$. Unlike the previous information theoretic approaches, we assume that each node is only aware of its own set of distributions, i.e., node i only knows $P^{(i)}$ and $Q^{(i)}$. Similarly, the fusion center neither has any knowledge about the nodes’ distributions nor a helper information. The fusion center performs a test based on the information conveyed by the nodes and outputs a prediction $\hat{\mathcal{H}}$.

We assume independent observations across the nodes and across time under each hypothesis. Namely, at time t and under \mathcal{H}_0 , the joint density of the observations is given by $p(x_t^{(1)}, \dots, x_t^{(m)}) = \prod_{i=1}^m p^{(i)}(x_t^{(i)})$, where $x_t^{(i)}$ is the observation of node i . Since the samples are independent across time, $p(\mathbf{x}_1, \dots, \mathbf{x}_t) = \prod_{\tau=1}^t p(x_\tau^{(1)}, \dots, x_\tau^{(m)})$, where $\mathbf{x}_t := (x_t^{(1)}, \dots, x_t^{(m)})$ is a length- m vector. All of the above holds under \mathcal{H}_1 by replacing p 's with q 's as well.

Being oblivious to nodes' statistics, the fusion center assumes each node sends a "score". For example, each node may send its LLR, $L_t^{(i)} := \log \frac{p^{(i)}(X_t^{(i)})}{q^{(i)}(X_t^{(i)})}$, and an optimal test can be performed at the fusion center using the sufficient statistic $\sum_{\tau=1}^t \sum_{i=1}^m L_\tau^{(i)}$. It is therefore reasonable to assume that the central processor trusts the nodes and sums the scores it receives. Formally, we denote the score sent by node i at time t by $S_t^{(i)}$. The center then performs a test based on $\bar{S}_t := \frac{1}{t} \sum_{\tau=1}^t \sum_{i=1}^m S_\tau^{(i)}$ and outputs

$$\hat{\mathcal{H}} = \begin{cases} \mathcal{H}_0, & \bar{S}_t \geq \eta_t \\ \mathcal{H}_1, & \text{else} \end{cases} \quad (1)$$

for a threshold η_t . The nodes take into account that the central processor performs its test based on \bar{S}_t . In other words, they know *a priori* that the scores sent by them will be averaged in the end. As we mentioned before, had there not been communication constraints, the nodes would have sent their LLRs; the central processor would have received a sufficient statistic and performed an optimal test. However, since $P^{(i)}$, $Q^{(i)}$ are continuous, the LLRs are continuous random variables in general and must be subject to some quantization or compression before being conveyed. Moreover, the communication rate between the nodes and the processor might be limited. Among many possible ways of restricting communication, we limit the average number of bits sent by each node under \mathcal{H}_0 .

Remark 1. The communication constraint is not symmetric, i.e., there is no constraint under \mathcal{H}_1 . This fits in with many real-world scenarios when \mathcal{H}_1 represents a high-risk situation in which the system is allowed to violate communication constraints in order to identify the risk — responding to an emergency takes priority over communication constraints — recall the collision avoidance example at the beginning of this manuscript. This view of \mathcal{H}_1 also implies that the type-II error must be very rare. In fact, in many hypothesis testing problems, it is desired that the type-II error decays exponentially. This is indeed the approach we will follow for the rest of this work.

Let us focus on one of the nodes. We drop the node's superscript (i) for brevity. At time t , after observing the data X_t and calculating its LLR L_t , the node sends its compressed score $S_t := f_t(L_t)$ with a simple function $f_t: \mathbb{R} \rightarrow \mathbb{R}$, i.e., f_t is measurable and takes finitely many values [16, Chapter 1]. Taking into account that S_t 's are discrete random variables, it is possible to generate a lossless code whose average length ℓ

in terms of bits is bounded as [17], [18]

$$H_P(S_t) \log_2 e - \log_2(H_P(S_t) \log_2 e + 1) - \log_2 e \leq \ell \leq H_P(S_t) \log_2 e \quad (2)$$

where $H_P(S_t)$ is the entropy of S_t under P .¹ Thus, we define the communication constraint as

$$\frac{1}{t} \sum_{\tau=1}^t H_P(S_\tau) \leq R \quad (3)$$

which ensures that S_1, \dots, S_t can be conveyed losslessly with an average number of bits less than R under \mathcal{H}_0 .

Remark 2. The memoryless quantization procedures we consider are practically appealing since the peripheral devices can be designed in a memory-efficient manner. Moreover, the assumption that each node only knows their own P 's and Q 's allows independent design of the peripheral nodes, as opposed to the joint design of all sensors which may be impractical. Note that without independence across the nodes, joint design might be necessary. We assume that the network subject to this study is designed such that the peripheral nodes have a spatial configuration that yields, or at least approximates, independence across nodes. ■

III. PERFORMANCE TRADEOFFS UNDER MEMORYLESS QUANTIZATION

We continue to focus on one of the nodes, omitting its superscript (i) for notational simplicity. In the previous section, we defined the scores as $S_t = f_t(L_t)$. We analyze the performance of this scheme. Let the type-I and type-II errors be

$$\begin{aligned} \alpha_t &:= P(\bar{S}_t < \eta_t) \\ \beta_t &:= Q(\bar{S}_t \geq \eta_t) \end{aligned} \quad (4)$$

respectively. For the rest of the analysis we set the decision threshold to

$$\eta_t = \frac{1}{t} \sum_{\tau=1}^t E_P[S_\tau] - \epsilon \quad (5)$$

with $\epsilon > 0$, and where $E_P[\cdot]$ denotes expectation under P . Although we previously assumed that the center does not know the statistics of the node, this threshold adjustment is without loss of generality. The node sends $S_t - E_P[S_t]$ and by setting $\eta_t = -\epsilon$, the performance remains the same and the test does not require knowledge of P or Q . Moreover, since centering of S_t does not change $H_P(S_t)$, the communication constraints are not violated. Therefore, for simplicity let us allow the center to set its threshold as in (5).

In light of Remark 1, we aim to keep α_t vanishing while guaranteeing exponential decay of β_t .

Definition 1. Given P and Q , (R, θ) is an achievable pair if there exists a sequence $\{f_t\}$ of simple functions and thresholds $\{\eta_t\}$ such that

- (a) $\frac{1}{t} \sum_{\tau=1}^t H_P(S_\tau) \leq R$ for all t
- (b) $\lim_{t \rightarrow \infty} \alpha_t = 0$
- (c) $\liminf_{t \rightarrow \infty} \frac{1}{t} \log \frac{1}{\beta_t} \geq \theta$

¹We adopt the definition of entropy from [19], which makes it necessary that S is discrete for $H_P(S)$ to be finite.

where S_t , α_t , β_t are defined in the beginning of this section.

We characterize the decay of β_t in the following theorem given in two parts.

Theorem 1. Let $\theta^*(R) := \sup\{\theta : (R, \theta) \text{ achievable}\}$ and define

$$\theta_t(R) := \sup_{\{f_1, \dots, f_t\} \in \mathcal{F}_t(R)} \frac{1}{t} \sum_{\tau=1}^t \left(E_P[S_\tau] - \log E_P[e^{S_\tau - L_\tau}] \right) \quad (6)$$

where $\mathcal{F}_t(R)$ is the set of all simple real-valued functions f_1, \dots, f_t on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ such that $\frac{1}{t} \sum_{\tau=1}^t H_P(S_\tau) \leq R$. Then,

(i) $\lim_{t \rightarrow \infty} \theta_t(R)$ equals to the upper concave envelope $\theta_1(R)$ of

$$\theta_1(R) = \sup_{f_1 \in \mathcal{F}_1(R)} E_P[S_1] - \log E_P[e^{S_1 - L_1}] \quad (7)$$

(ii) $\theta^*(R) = \lim_{t \rightarrow \infty} \theta_t(R) = \check{\theta}_1(R)$.

In words, for any scheme under the communication constraint in (3), and with vanishing type-I error; the worst-case type-II error decay θ — defined in Definition 1 (c) — cannot be greater than $\theta^*(R)$.

Proof: See Appendix A in [20]. ■

Even though (7) is a single letter expression, the optimization domain is non-convex. Consequently, computation of $\theta^*(R)$ does not appear to be easy. We now propose the following single-letter optimization problem, which turns out to be (i) tractable, (ii) yield an upper bound to $\theta^*(R)$, (iii) be a good approximation to $\theta^*(R)$ for high R :

$$\begin{aligned} \theta_U(R) := \sup_{p_{V|U}} & E_P[V] - \log E_P[\exp(V - U)] \\ \text{s.t.} & I_P(U; V) \leq R \end{aligned} \quad (8)$$

where U has the same distribution as $L_t = \log \frac{p(X_t)}{q(X_t)}$, $I_P(U; V)$ is the mutual information under P , and the transition kernel $p_{V|U} : \mathcal{B}(\mathbb{R}) \times \mathbb{R} \rightarrow [0, 1]$ relates to the (possibly randomized) quantization procedure $V = f(U)$. This suggests that the tradeoff may be studied with methods of rate-distortion theory. Observe that (i) the set of feasible $p_{V|U}$'s is larger than the set of simple functions, and (ii) the communication constraint is relaxed as $I_P(V; U) \leq H_P(V)$. Hence, $\theta_U(R) \geq \theta_1(R)$.

It is immediate that $\theta_U(R)$ is non-decreasing. It turns out that it is also concave and thus continuous, as we will see shortly. Considering the extremes, (i) if $R \rightarrow \infty$, one can set $V = U$ and $\theta_U(R) \rightarrow D(P||Q)$ and (ii) if $R = 0$, one does no better than choosing V as a constant and $\theta_U(R) = 0$.

A. Concavity of $\theta_U(R)$

The following lemma gives another characterization of $\theta_U(R)$, from which we conclude that it is concave.

Lemma 1. Let

$$\begin{aligned} \tilde{\theta}_U(R) := \sup_{p_{V|U}} & E_P[V] - E_P[\exp(V - U)] + 1 \\ \text{s.t.} & I_P(U; V) \leq R. \end{aligned} \quad (9)$$

Then, $\theta_U(R) = \tilde{\theta}_U(R)$.

Proof: See Appendix B in [20]. ■

Note the equivalence between $-\tilde{\theta}_U(R)$ and the distortion-rate curve $D(R)$ for the distortion function $d(u, v) = -v + e^{v-u} - 1$ [21, Chapter 10]. Since the distortion-rate curve is known to be convex, $\tilde{\theta}_U(R) = \theta_U(R)$ is concave.

Corollary 1. $\theta_U(R) \geq \check{\theta}_1(R) = \theta^*(R)$.

Proof: $\theta_U(R)$ is a concave function that dominates $\theta_1(R)$. Hence it also dominates $\check{\theta}_1(R)$. ■

Although the formulations (8) and (9) are the same, for any candidate $p_{V|U}$, (8) gives a tighter bound to $\theta^*(U)$. Hence, we primarily study the formulation (8) for the rest of this work.

B. The gap function $\delta_U(R)$

Starting from this section, we assume all expectations (including the mutual information) are taken under P , which we omit from subscripts. Recall that $\theta_U(R)$ can at most be $D(P||Q) = E[U]$. Hence, the tradeoff can also be studied by the gap to the optimal; $\delta_U(R) := D(P||Q) - \theta_U(R)$. Thus,

$$\begin{aligned} \delta_U(R) = \inf_{p_{V|U}} & \log E[\exp(V - U)] - E[V - U] \\ \text{s.t.} & I(U; V) \leq R \end{aligned} \quad (10)$$

which is also equivalent, by defining $Z := V - U$, to

$$\begin{aligned} \delta_U(R) = \inf_{p_{Z|U}} & \log E[\exp(Z)] - E[Z] \\ \text{s.t.} & I(U; U + Z) \leq R. \end{aligned} \quad (11)$$

Note that $\delta_U(R)$ is convex and non-increasing for $R > 0$, therefore its generalized inverse $R_U(\delta)$ exhibits the same properties for $\delta > 0$ as well.

IV. CALCULATION OF $\delta_U(R)$

Recall the definition of $\delta_U(R)$ given in (11). First, we obtain a simple upper bound for $\delta_U(R)$ as follows. Choose a Gaussian Z with variance σ^2 and independent of U . Then,

$$\delta = \log E_P[e^Z] - E_P[Z] = \sigma^2/2 \quad (12)$$

and

$$\begin{aligned} R = I(U; U + Z) &= h(U + Z) - h(Z) \\ &\leq \frac{1}{2} \log(1 + \text{Var}(U)/\sigma^2). \end{aligned} \quad (13)$$

where $h(\cdot)$ denotes the differential entropy. Thus,

$$R_U(\delta) \leq \frac{1}{2} \log \left(1 + \frac{\text{Var}(U)}{2\delta} \right). \quad (14)$$

Observe that (14) is not tight for low rates, i.e., it approaches 0 only when $\delta \rightarrow \infty$, whereas the actual $R_U(\delta)$ curve is known to be 0 at $\delta = D(P||Q)$. Since any (δ, R) pair is achievable by time-sharing, any convex combination of $(D(P||Q), 0)$ with the upper bound curve is achievable. Taking such combinations, one can tighten the upper bound, see Figure 1 for an example.

If we attempt to calculate the exact curve, at first glance one might argue that the problem is a concave minimization, hence it might not be amenable to convex optimization methods. However, we exploit the property of the mutual information

$I(X; X+Z)$ being invariant under any shift of Z . Thus, adding the constraint $E[Z] = 0$ does not change the feasible region. After adding the constraint, we formulate the problem as

$$\begin{aligned} \delta_U(R) &= \inf_{p_{Z|U}} \log E[\exp(Z)] \\ \text{s.t. } & I(U; U+Z) \leq R \\ & E[Z] = 0 \end{aligned} \quad (15)$$

and any infimizer $p_{Z|U}$ of the above problem infimizes the same problem with objective function $E[\exp(Z)]$ as the logarithm is a strictly increasing function.

From now on, consider the following linear optimization on a convex domain:

$$\begin{aligned} \Delta_U(R) &:= \inf_{p_{Z|U}} E[\exp(Z)] \\ \text{s.t. } & I(U; U+Z) \leq R \\ & E[Z] = 0 \end{aligned} \quad (16)$$

Observe that $\Delta_U(R)$ is non-increasing, convex, hence continuous and $\lim_{R \rightarrow \infty} \Delta_U(R) = 1$. Suppose U is a continuous random variable. Since 1 is only attained when Z is chosen identically equal to 0, and since this cannot happen at a finite R , this shows that the function is also strictly decreasing. Also note that $\log \Delta_U(R) = \delta_U(R)$, which shows that $\delta_U(R)$ is also strictly decreasing. The boundary to the feasible (Δ, R) pairs can also be found by the curve

$$\begin{aligned} R_U(\Delta) &:= \inf_{p_{Z|U}} I(U; U+Z) \\ \text{s.t. } & E[\exp(Z)] \leq \Delta \\ & E[Z] = 0. \end{aligned} \quad (17)$$

This formulation is useful to find a lower bound for feasible (Δ, R) pairs. We follow a method that is similar to the one that yields a lower bound for the rate distortion problem under mean square distortion.

Lemma 2. *Define the parametric curve*

$$\begin{aligned} \underline{R}_U(\alpha) &= h(U) - \log(\Gamma(\alpha)) + \alpha\psi(\alpha) - \alpha, \\ \delta(\alpha) &= \log(\alpha) - \psi(\alpha), \quad \alpha > 0, \end{aligned} \quad (18)$$

where $\Gamma(\cdot)$, $\psi(\cdot)$ are gamma and digamma functions respectively. Then $\underline{R}_U(\delta) \leq R_U(\delta)$.

Proof: See Appendix C in [20]. ■

Recall that we have $R_U(\delta) \leq \frac{1}{2} \log(1 + \frac{\text{Var}(U)}{2\delta})$ from (14). Now we obtain another upper bound for U with a differentiable probability density function $p_U(u)$ to show that the lower bound is approached at high rates with additive Gaussian noise. Note that $h(U+Z)$ in (13) is a function of σ . We expand $h(U+Z)$ around $\sigma = 0$. From De Bruijn's identity [21] and Taylor's theorem it is known that

$$h(U+Z) \leq h(U) + \frac{\sigma}{2} J(U) \quad (19)$$

where $J(U) := E\left[\left(\frac{d}{du} \log p_U(U)\right)^2\right]$ is the Fisher information of U , [21]. The bound above makes sense only for finite $J(U)$. Assuming finite $J(U)$, we obtain

$$R_U(\delta) \leq h(U) + \sqrt{\frac{\delta}{2}} J(U) - \frac{1}{2} \log(4\pi e \delta). \quad (20)$$




Fig. 1. The solid curve \underline{R}_U corresponds to the lower bound (18), the dashed curve corresponds to the upper bound in (14) of a Gaussian LLR U and the dotted line segment (TS) corresponds to boundary of the pairs achievable with time sharing. No point in the shaded region is achievable. The mean of U is $E[U] = D(P||Q) = 10$, $\text{Var}(U) = 2E[U] = 20$ and $h(U) \approx 2.9$.

Now we approximate (18) for large α . From [22], we use the inequalities $\log(\alpha) - \frac{1}{2\alpha} \geq \psi(\alpha) \geq \log(\alpha) - \frac{1}{2\alpha} - \frac{1}{12\alpha^2}$ and Stirling's approximation $\log(\Gamma(\alpha)) \leq \alpha \log(\alpha) - \alpha - \frac{1}{2} \log(\alpha) + \frac{1}{2} \log(2\pi) + O(1/\alpha)$ to obtain

$$\underline{R}_U(\delta) \geq h(U) - \frac{1}{2} \log(4\pi e \delta) - O(2\delta), \quad (21)$$

which asymptotically matches the upper bound given in (20). This is consistent with the behavior of the curves in high-rate regime in Figure 1.

V. HIGH-RATE COMPRESSION REGIME

In this section, we evaluate the performance of lattice quantization and show that at high rates the lower bound (18) can be approached within a small gap. We refer to [23] for a detailed definition of lattice quantization procedures. Again, assume U has the same distribution as the LLR L_t , which admits a v -regular probability density.

Definition 2 (v -regular density, [23]). *Given $v : \mathbb{R} \rightarrow \mathbb{R}$, a continuous and differentiable density function p is called v -regular if $|\frac{d}{du} p(u)| \leq v(u)p(u)$.*

It is shown in [23] that when U has a v -regular density, the entropy of the lattice-quantized U — denoted by V to be consistent with the previous section — is upper bounded by

$$H(V) \leq h(U) - \log r + r C_U(r) \quad (22)$$

where r is the length of quantization intervals and $C_U(r)$ is a function of r depending on the density of U and $v(u)$. Moreover, if $v(u)$ is Lipschitz-continuous almost everywhere and if $E[v(U)]$ is finite, then $C_U(r)$ can be shown to be bounded for bounded r , see Appendix-D in [20].

Given a lattice quantization scheme, we want to relate δ and r in order to upper bound (22) in terms of δ . Observe that

$$\begin{aligned}\delta &= \log(E[e^{V-U}]) - E[V-U] \\ &\leq E[e^{V-U}] - E[V-U] - 1 \\ &= E[e^Z - Z - 1]\end{aligned}\quad (23)$$

and since $|Z| \leq r/2$ surely, $e^Z - Z - 1 \leq e^{r/2} - r/2 - 1$. Consequently, $\delta \leq e^{r/2} - r/2 - 1$. Now, we can obtain an achievability bound in parametric form based on the lattice quantizer described above. Using (22), we have

$$\begin{aligned}\overline{R}_U(r) &:= h(U) - \log r + rC_U \\ \overline{\delta}(r) &:= e^{r/2} - r/2 - 1\end{aligned}\quad (24)$$

and the $(\overline{\delta}, \overline{R}_U)$ pairs lie above the curve $R_U(\delta)$. Now observe that for low δ , i.e. for small radius r , $\overline{\delta}(r) = \frac{r^2}{8} + O(r^3)$. In fact, $\overline{\delta}(r) \leq \frac{r^2}{4}(k^2 e^{1/k} - k - k^2)$ for $r < 2/k$, which allows us to conclude that all $(\overline{\delta}, \overline{R}_U)$ pairs lie under

$$\begin{aligned}\overline{\overline{R}}_U(\delta) &:= h(U) - \frac{1}{2} \log(4\pi e\delta) \\ &\quad + \frac{1}{2} \log(\pi e(k^2 e^{1/k} - k - k^2)) + o(1)\end{aligned}\quad (25)$$

which for instance at $k = 10$ gives at most ≈ 0.743 nats ≈ 1.072 bits of difference from the lower bound $\underline{R}_U(\delta)$ in (21). As $k \rightarrow \infty$, the lower bound in (21) is approached with a difference of $\frac{1}{2} \log(\pi e/2) \approx 0.726$ nats ≈ 1.047 bits; this is the gap we mentioned in Section I. Recall that all achievable (δ, R) pairs must lie above the \underline{R}_U curve described in (21). Equation (25) shows the existence of (δ, R) pairs that live close to the \underline{R}_U curve, and these pairs are achieved with a lattice quantization scheme.




Fig. 2. Plot of the required number of bits (R) versus δ in logarithmic scale at high rate regime for the same Gaussian U and with same $\underline{R}_U(\delta)$ given in Fig. 1. The dashed curve is the upper bound $\overline{R}_U(\delta)$ in (24) obtained with a lattice quantization scheme. No point in the shaded region is achievable. The difference of 1.047 bits is visible for low δ .

Remark 3. Although we have not imposed any communication constraints under \mathcal{H}_1 , the average number of bits sent is

not arbitrarily large under certain circumstances. Suppose we generate a standard binary entropy code on S , which assigns a length $\lceil \log_2 \frac{1}{P_S(s)} \rceil$ to symbol $s \in \mathcal{S}$. Then the expected length of the code under Q will be less than $(H_Q(S) + D(Q_S||P_S)) \log_2 e + 1 \leq (H_Q(S) + D(Q||P)) \log_2 e + 1$ bits per symbol, with Q_S and P_S denoting the distribution of S under \mathcal{H}_1 and \mathcal{H}_0 respectively. The last inequality is a result of data processing inequality. We therefore conclude that if $H_Q(S)$ and $D(Q||P)$ are finite, the expected number of bits can be kept finite under \mathcal{H}_1 as well. ■

VI. MULTIPLE-NODE CASE

The study until this point can be extended to multiple nodes. Suppose m nodes communicate with the center, where at time t each node i observes $X_t^{(i)}$, computes the LLR $L_t^{(i)}$; and is subject to a rate constraint R_i . We give an extension of Definition 1 to multiple nodes.

Definition 3. Given $\{P^{(i)}\}_{i=1}^m$ and $\{Q^{(i)}\}_{i=1}^m$, $(R_1, \dots, R_m, \theta)$ is an achievable pair if there exists m sequences $\{f_t^{(1)}\}, \dots, \{f_t^{(m)}\}$ of simple functions and thresholds $\{\eta_t\}$ such that

- $\frac{1}{t} \sum_{\tau=1}^t H_P(S_\tau^{(i)}) \leq R_i$ for all t and for all i
- $\lim_{t \rightarrow \infty} \alpha_t = 0$
- $\liminf_{t \rightarrow \infty} \frac{1}{t} \log \frac{1}{\beta_t} \geq \theta$

where $S_t^{(i)} = f_t^{(i)}(L_t^{(i)})$, and α_t, β_t are type-I and type-II errors respectively, defined in the beginning of Section III.

Define $\theta_t(R_1, \dots, R_m) := \sum_{i=1}^m \theta_t^{(i)}(R_i)$, where $\theta_t^{(i)}(R_i)$ is defined similarly as in Theorem 1, but with respect to node i . Since we assume independence across nodes and if we follow the same steps we did in the proof of Theorem 1, we get

$$\begin{aligned}\theta^*(R_1, \dots, R_m) &:= \sup\{\theta : (R_1, \dots, R_m, \theta) \text{ achievable}\} \\ &= \lim_{t \rightarrow \infty} \theta_t(R_1, \dots, R_m) \\ &= \sum_{i=1}^m \check{\theta}_1^{(i)}(R_i) \\ &\leq \sum_{i=1}^m D(P^{(i)}||Q^{(i)}) - \delta_{L^{(i)}}(R_i),\end{aligned}\quad (26)$$

where $\delta_{L^{(i)}}(R_i)$ are defined as in Section III-B. This is intuitive because there is independence across the nodes and the decay rate should be the sum of individual decay rates. Also the upper bound can be reached closely with lattice quantization followed by lossless coding at each node, independently, as we have shown in Section V.

VII. CONCLUSION AND DISCUSSION

In this work, we have studied a distributed binary hypothesis testing problem where the agents quantize and forward each observation to a central processor. We have proved that for all schemes that send R bits on average under the null hypothesis and have a vanishing type-I error, the exponential decay rate of type-II error, which is defined as θ in Definition 1, cannot be greater than $\theta^*(R)$. It turns out that $\theta^*(R)$ is difficult to calculate. Therefore, we derived the lower bounds (18), (21) in closed form for the achievable (δ, R) pairs, where

$\delta := D(P||Q) - \theta$ quantifies the gap to the optimal decay rate $D(P||Q)$. Lastly, we have shown that if each node (i) calculates the log-likelihood ratio of its observation and (ii) performs lattice quantization on the log-likelihood ratio, the lower bound in (21) is approached closely up to a 1.047 bits of difference.

One may ask what would the curves look like if we allowed quantization of multiple data, i.e., $f_{kt}(x_1, \dots, x_k)$ for a finite k . We have preliminary results suggesting that the curves may be as follows. Let $R_U^{(k)}(\delta)$ be the curve defined similarly for data consisting of k elements. Then $R_U^{(k)}(\delta) = \frac{1}{k} R_{U_1 + \dots + U_k}(k\delta)$. Proceeding in a similar way to obtain (6), we observe that

$$R_U^{(k)}(\delta) \leq \frac{1}{2k} \log \left(1 + \frac{\text{Var}(U)}{2\delta} \right). \quad (27)$$

This suggests that the average number of bits sent under \mathcal{H}_0 can be made arbitrarily small for large k .

The argument above may therefore recover the case that at zero rate one may not do worse than the optimal test. This is easily seen for the 1-node case. Suppose the node performs its own Neyman-Pearson test and sends its decision at time t , represented by one bit. Therefore the rate is $1/t \rightarrow 0$ and the center performs optimally. Observe that this scheme ensures zero-rate communication under both \mathcal{H}_0 and \mathcal{H}_1 . Also for multiple-node case, there is a simple fusion rule that attains the optimal exponent which is based on each node performing their own tests: Decide \mathcal{H}_1 if at least one node decides \mathcal{H}_1 . Whether to give such autonomy to every single node is a question of design, e.g., a faulty node with full autonomy may significantly increase the type-I error.

The individual gap functions $\delta_{L^{(i)}}(R_i)$ for each node i can be used if the communication constraints are redefined such that subsets of nodes are restricted to communicate, e.g., node 1, 2 and 3 can together send at most $R_{(1,2,3)}$ bits in total. Then one can minimize the sum of gap functions of node 1, 2, and 3, i.e., $\delta_{L^{(1)}}(R_1) + \delta_{L^{(2)}}(R_2) + \delta_{L^{(3)}}(R_3)$, over the sum rate constraint $R_1 + R_2 + R_3 \leq R_{(1,2,3)}$ with an aim to find a lower bound to the achievable exponents. An extension of this work could be to study a similar setting in decentralized hypothesis testing problem, where a fusion center does not exist and the nodes rely on peer-to-peer communication.

REFERENCES

- [1] R. Ahlswede and I. Csiszar, "Hypothesis testing with communication constraints," *IEEE Transactions on Information Theory*, vol. 32, no. 4, pp. 533–542, 1986.
- [2] T. S. Han and S. Amari, "Statistical inference under multiterminal data compression," *IEEE Transactions on Information Theory*, vol. 44, no. 6, pp. 2300–2324, 1998.
- [3] T. Han, "Hypothesis testing with multiterminal data compression," *IEEE Transactions on Information Theory*, vol. 33, no. 6, pp. 759–772, 1987.
- [4] S. Salehkalaibar and V. Y. F. Tan, "Distributed sequential hypothesis testing with zero-rate compression," in *2021 IEEE Information Theory Workshop (ITW)*, 2021, pp. 1–5.
- [5] S. Sreekumar and D. Gündüz, "Distributed hypothesis testing over discrete memoryless channels," *IEEE Transactions on Information Theory*, vol. 66, no. 4, pp. 2044–2066, 2020.
- [6] S. Salehkalaibar and M. Wigger, "Distributed hypothesis testing over noisy broadcast channels," *Information*, vol. 12, no. 7, 2021. [Online]. Available: <https://www.mdpi.com/2078-2489/12/7/268>
- [7] —, "Distributed hypothesis testing with variable-length coding," in *Proc. International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOPT)*, 2020, pp. 1–5.
- [8] M. Hamad, M. Wigger, and M. Sarkiss, "Cooperative multi-sensor detection under variable-length coding," in *Proc. IEEE Information Theory Workshop (ITW)*, 2021, pp. 1–5.
- [9] M. Longo, T. Lookabaugh, and R. Gray, "Quantization for decentralized hypothesis testing under communication constraints," *IEEE Transactions on Information Theory*, vol. 36, no. 2, pp. 241–255, 1990.
- [10] C.-T. Yu and P. K. Varshney, "Paradigm for distributed detection under communication constraints," *Optical Engineering*, vol. 37, no. 2, pp. 417 – 426, 1998. [Online]. Available: <https://doi.org/10.1117/1.601627>
- [11] R. Gupta and A. Hero, "High-rate vector quantization for detection," *IEEE Transactions on Information Theory*, vol. 49, no. 8, pp. 1951–1969, 2003.
- [12] C. Rago, P. Willett, and Y. Bar-Shalom, "Censoring sensors: a low-communication-rate scheme for distributed detection," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 32, no. 2, pp. 554–568, 1996.
- [13] S. Appadwedula, V. V. Veeravalli, and D. L. Jones, "Decentralized detection with censoring sensors," *IEEE Transactions on Signal Processing*, vol. 56, no. 4, pp. 1362–1373, 2008.
- [14] L. Cao and R. Viswanathan, "Divergence-based soft decision for error resilient decentralized signal detection," *IEEE Transactions on Signal Processing*, vol. 62, no. 19, pp. 5095–5106, 2014.
- [15] J. Conway and N. Sloane, "Fast quantizing and decoding and algorithms for lattice quantizers and codes," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 227–232, 1982.
- [16] W. Rudin, *Real and Complex Analysis, 3rd Ed.* USA: McGraw-Hill, Inc., 1987.
- [17] A. Wyner, "An upper bound on the entropy series," *Information and Control*, vol. 20, no. 2, pp. 176–181, 1972. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0019995872903658>
- [18] N. Alon and A. Orlitsky, "A lower bound on the expected length of one-to-one codes," *IEEE Transactions on Information Theory*, vol. 40, no. 5, pp. 1670–1672, 1994.
- [19] M. S. Pinsker, *Information and Information Stability of Random Variables and Processes. Translated and edited by Amiel Feinstein.*, ser. Holden-Day series in time series analysis. San Francisco: Holden-Day, 1964.
- [20] Y. Inan, M. Kayaalp, A. H. Sayed, and E. Telatar, "A fundamental limit of distributed hypothesis testing under memoryless quantization." [Online]. Available: <http://infoscience.epfl.ch/record/290031>
- [21] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley-Interscience, July 2006.
- [22] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. New York: Dover, 1964.
- [23] V. Kostina, "Data compression with low distortion and finite block-length," in *Proc. Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2015, pp. 1127–1134.