Covert Sequential Hypothesis Testing

Meng-Che Chang and Matthieu R. Bloch
School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332
Email: {mchang301,matthieu}@gatech.edu

2021 IEEE Information Theory Workshop (ITW) | 978-1-6654-0312-2/21/$31.00 ©2021 IEEE | DOI: 10.1109/ITW48936.2021.9611391

***Abstract*—We consider the problem of covert sequential testing, in which a legitimate party attempts to run a sequential test while escaping detection from an adversary. Specifically, the legitimate party's decisions should meet prescribed risk constraints and, simultaneously, the adversary's observations induced by the test should remain indistinguishable from the observations obtained in the absence of a test. Our main result is the characterization of the risk exponent γ_θ, which captures the asymptotic exponential decrease of the risk with the square-root of the averaged stopping time in the limit of low risk. An example is provided to illustrate how the covertness constraint influences the design of the sequential test.**

I. Introduction

The problem of ensuring covertness, or low probability of detection, has attracted significant attention in the context of communication systems [1]–[5]. A square root low has been established, which requires the transmitted codewords to have weight scaling as the square root of the blocklength to maintain covertness. A similar law has also been established for *covert sensing* [6], [7], which refers to a scenario in which the estimation of parameters of interest requires the use of probing signals that emit energy and are therefore detectable. Specifically, [6], [7] have considered the problem of estimating an unknown phase while keeping the sensing undetectable by a passive quantum adversary. This operation is made possible by the presence of thermal noise, which allows one to hide the useful sensing signal in the background thermal noise and results in a mean-square phase estimation error scaling as $O(\frac{1}{\sqrt{n}})$ if n is the number of modes.

In contrast to conventional hypothesis testing, active hypothesis testing, also known as *controlled sensing*, adaptively selects the kernel through which unknown parameters are observed. The estimation can adapt to the observations, resulting in a potentially faster strategy than non-adaptive ones. Motivated by recent progress in this area [8]–[10], the problem of *covert active sensing* has been analyzed [5], [11], characterizing the exponent of the probability of detection error subject to the covertness constraint with non-sequential and non-adaptive tests and showing the superiority of adaptive (but still non-sequential) strategies

In the present work, we formulate the problem of *covert sequential testing*, in which the objective is to carry out a sequential test that meets certain risk constraints while keeping the existence of the test undetectable by the adversary at any time before the sequential test stops. This problem formulation is inspired by applications in which the adversary has the ability to hinder the estimation if the existence of a test is detected. In contrast to [5], our work allows the strategy to be adaptive *and* sequential. Our main contribution is to characterize the ratio of the risk exponent to the square root of the expected number of tests in the limit of low risk. Conceptually related to the present work, [12] analyzes the covert communication problem under the setting in which the adversary does not know when the communication starts and uses specific sequential tests to determine the existence of the communication in real-time. An optimization problem is then formulated and a transmission scheme that maximizes the total amount of information under the covertness constraint is proposed. This approach differs from ours in that the sequential test is a constraint placed on the type of detectors deployed by adversary, while our model allows the adversary to deploy its optimal testing strategy.

The rest of the paper is organized as follows. After defining some notation in Section II, we formulate the problem and state the main results in Section III. We detail the proofs of the main results in Section IV. Finally, we provide a numerical example in Section V.

II. Notations

For two distributions P, Q on some common alphabet $\mathcal{X}$, $D(P\|Q) \triangleq \sum_x P(x)\log\frac{P(x)}{Q(x)}$ is the Kullback-Leibler (KL) divergence between P and Q. We say P is absolutely continuous with respect to (w.r.t.) Q, denoted by $P \ll Q$, if for all $x \in \mathcal{X}$ $P(x) = 0$ if $Q(x) = 0$. We denote $P^{\otimes n}$ the product distribution $\prod_{\ell=1}^n P$ on $\mathcal{X}^n$. We also define $\chi_2(P\|Q) \triangleq \sum_x \frac{(P(x)-Q(x))^2}{Q(x)}$. The set of all distributions on $\mathcal{X}$ is defined as $\mathcal{P}_\mathcal{X}$. For any sequence $\mathbf{x}$, we define $\widehat{p}_\mathbf{x}$ as the type of the sequence $\mathbf{x}$.

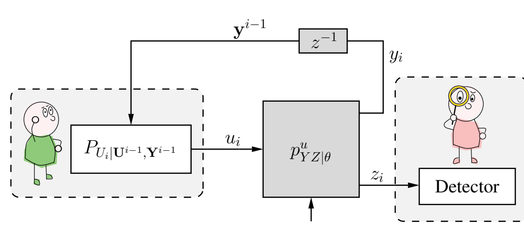

Fig. 1. Model for covert detection

III. Problem Formulation and Main Result

The problem of covert sequential hypothesis testing is illustrated in Fig 1 and defined as follows. A legitimate party

This work was supported in part by NSF grant 1955401.

978-1-6654-0312-2/21/$31.00 ©2021 IEEE

(Alice) and an adversary (Willie) are engaged in a sequential hypothesis testing problem to estimate an unknown parameter θ, which belongs to a known parameter set Θ. Let $\mathcal{Y}$, $\mathcal{Z}$ be the alphabets for Alice's and Willie's observations, respectively, and $\mathcal{U}$ be the set of control inputs that Alice can use to adapt her estimation. At any time $k > 0$, the observations of Alice and Willie depend on the true hypothesis θ and the control input $u_k \in \mathcal{U}$ chosen by Alice. Specifically, $y_k \in \mathcal{Y}$ and $z_k \in \mathcal{Z}$ are generated from the kernels $p_\theta^{u_k}$ and $q_\theta^{u_k}$, respectively, and the action u_k is generated from a closed loop control policy $P_{U_k|Y^{k-1},U^{k-1}}$. Therefore, for any $n > 0$, the joint distribution $P_{U^nY^nZ^n|\theta}$ of (y^n, z^n, u^n) under the hypothesis θ is

$$P_{U^nY^nZ^n|\theta}(y^n, z^n, u^n) = \prod_{k=1}^{n} p_\theta^{u_k}(y_k) q_\theta^{u_k}(z_k) P_{U_k|Y^{k-1},U^{k-1}}(u_k|y^{k-1}, u^{k-1}). \quad (1)$$

We assume that i) the sets $\mathcal{Y}$, $\mathcal{Z}$, and $\mathcal{U}$ are finite; ii) there exists a prior $\{\pi_i\}_{i\in\Theta}$ on the parameter θ known by both Alice and Willie; iii) the kernel sets $\{p_\theta^u\}_{u\in\mathcal{U}}$ and $\{q_\theta^u\}_{u\in\mathcal{U}}$ are known to both Alice and Willie; iv) the control sequence u^τ is unknown by Willie; and v) $\forall\theta \neq \theta'$ and $\forall u \in \mathcal{U} \setminus \{0\}$, it holds that $0 < D(p_\theta^u \| p_{\theta'}^u) < \infty$, and $D(p_\theta^0 || p_{\theta'}^0) = 0$ for all $\theta' \neq \theta$.

Remark 1. *We assume that $D(p_\theta^u \| p_{\theta'}^u)$ is finite for $u \in \mathcal{U} \setminus \{0\}$ to avoid the trivial case in which the hypothesis can be determined by a single action. In our problem setting, the null action 0 refers to the situation in which Alice does not take any effective action for estimating the hypothesis. Therefore, to make the action 0 uninformative from the perspective of identifying θ, we also assume that $D(p_\theta^0 || p_{\theta'}^0)$ for all $\theta' \neq \theta$.*

The sequential test $\sigma = (\phi, \tau, \delta)$ consists of the following components: i) a *control policy* $\phi \triangleq \left\{P_{U_k|Y^{k-1},U^{k-1}}\right\}_{k\geqslant 1}$ that allows Alice to generate action u_k at time k according to the distribution $P_{U_k|Y^{k-1},U^{k-1}}$; ii) a *stopping rule* τ that allows Alice to stop the sequential estimation iii) a *decision rule* δ that allows Alice and Bob to decide on a hypothesis when the test stops. Let $\mathbb{P}_\theta$ be the probability measure under the hypothesis θ. We require that the tests satisfy the risk constraints

$$\sum_{\theta' \neq \theta} \pi_{\theta'} \mathbb{P}_{\theta'}\{\delta(y^\tau) = \theta\} < R_\theta \quad (2)$$

for all $\theta \in \Theta$, where $\{R_\theta\}_{\theta\in\Theta}$ are predefined values. Alice's goal is to perform a sequential test to meet the risk constraints in (2) while maintaining covertness with respect to Willie. The covertness is measured in terms of divergence. Specifically, we require that for all $\theta \in \Theta$ and $0 < n < \tau$,

$$\lim_{R_{\max}\to 0} \mathbb{P}_\theta \left\{D(P_{Z^n|\theta}||(q_\theta^0)^{\otimes n}) \leqslant \eta\right\} = 1 \quad (3)$$

for some predetermined value $\eta > 0$, where $R_{\max} = \max_{\theta\in\Theta} R_\theta$. We say that the control policy converges almost surely to P_U^* if for any $u \in \mathcal{U}$ the sequence $\{P_{U_k|Y^{k-1}U^{k-1}}(u)\}_{k\geqslant 1}$ converge almost surely to $P_U^*(u)$. We define $\Sigma(\mathbf{R}, \eta)$ as the set of sequential tests that have almost surely convergent control policies and satisfy the constraint in (2) and (3), where $\mathbf{R} = (R_0, ..., R_{|\Theta|-1})$ is a vector of risk constraints. We are interested in the exponent γ_θ which is related to the test σ and is defined as

$$\gamma_\theta(\sigma) = \liminf_{R_{\max}\to 0} \frac{-\log R_\theta}{\sqrt{\mathbb{E}_\theta(\tau)}}. \quad (4)$$

Definition 1. *We call the exponent γ_θ achievable if there exists a test $\sigma \in \Sigma(\mathbf{R}, \eta)$ such that*

$$\gamma_\theta(\sigma) = \liminf_{R_{max}\to 0} \frac{-\log R_\theta}{\sqrt{\mathbb{E}_\theta(\tau)}} \geqslant \gamma_\theta. \quad (5)$$

The optimal exponent γ_θ^* is the supremum of all achievable exponent γ_θ, i.e.,

$$\gamma_\theta^* = \sup_{\sigma\in\Sigma(\mathbf{R},\eta)} \gamma_\theta(\sigma). \quad (6)$$

Our main result characterizes the value of γ_θ^*.

Theorem 2. *Let Θ be the set of parameters that are indistinguishable from another parameter by choosing the null action 0, i.e. $p_\theta^0 = p_{\theta'}^0$ for all $\theta \neq \theta'$. For all $\theta \in \Theta$, we assume that no distribution $\bar{P}$ over $\mathcal{U} \setminus \{0\}$ is such that $\sum_{u\neq 0} \bar{P}(u) q_\theta^u = q_\theta^0$. Then, we have*

$$\gamma_\theta^* = \sqrt{2\eta} \max_{\bar{P}} \frac{\min_{\theta'\neq\theta} \sum_{u\neq 0} \bar{P}(u) D(p_\theta^u||p_{\theta'}^u)}{\sqrt{\chi_2(\sum_{u\neq 0} \bar{P}(u) q_\theta^u||q_\theta^0)}}. \quad (7)$$

Remark 2. *The assumption that $\sum_{u\neq 0} \bar{P}(u) q_\theta^u \neq q_\theta^0$ for all $\bar{P} \in \mathcal{P}_{\mathcal{U}\setminus\{0\}}$ is made to circumvent the special case in which Alice can simulate the null action output distribution at Willie by any distribution over $\mathcal{U} \setminus \{0\}$.*

Remark 3. *We limit ourselves to analyze the best exponent among all sequential tests that have almost surely convergent strategies. This assumption allows us to characterize the time averaged distribution of actions as well as the type of the action sequence. Furthermore, this assumption is reasonable because a good strategy is usually a function of the estimated hypothesis [10], and Alice's estimation of the true hypothesis is increasingly accurate as the number of tests increases.*

IV. Proof of Main result

A. Achievability

We first specify the test σ we choose in the achievability proof. For any $\bar{P} \in \mathcal{P}_{\mathcal{U}\setminus\{0\}}$, $\theta \in \Theta$ and any fixed $\eta > 0$, we define the function $f(\bar{P}, \theta)$ as

$$f(\bar{P}, \theta) \triangleq \min_{\theta'\neq\theta} \sqrt{2\eta} \frac{\sum_{u\neq 0} \bar{P}(u) D(p_\theta^u||p_{\theta'}^u)}{\sqrt{\chi_2(\sum_{u\neq 0} \bar{P}(u) q_\theta^u||q_\theta^0)}}. \quad (8)$$

Let $\widehat{\theta}_k$ be the maximum likelihood estimate of the true hypothesis θ at time k, and define $\bar{P}_k$ as

$$\bar{P}_k = \operatorname*{argmax}_{\bar{P}\in\mathcal{P}_{\mathcal{U}\setminus\{0\}}} f(\bar{P}, \widehat{\theta}_k), \quad (9)$$

i.e., $\bar{P}_k$ is the maximizing distribution of the function f with the second parameter replaced by $\widehat{\theta}_k$. Then, the control policy $P_{U_k|X^{k-1},U^{k-1}}$ at time k is defined as

$$P_{U_k|X^{k-1},U^{k-1}}(u) \triangleq P_k(u) = \begin{cases} 1-\alpha_k & \text{if } u=0 \\ \alpha_k \bar{P}_k(u) & \text{if } u \neq 0, \end{cases}$$

where

$$\alpha_k = \frac{2\eta}{\chi_2(\sum_{u\neq 0} \bar{P}_k(u) q^u_{\widehat{\theta}_k} || q^0_{\widehat{\theta}_k})} \times \frac{\min_{\theta' \neq \widehat{\theta}_k} \sum_{u\neq 0} \bar{P}_k(u) D(p^u_{\widehat{\theta}_k} || p^u_{\theta'})}{|\log R_{\widehat{\theta}_k}|}.$$

We define the generalized likelihood ratio $L_\theta(n)$ as

$$L_\theta(n) = \frac{\prod_{k=1}^n p_\theta^{U_k}(Y_k)}{\max_{\theta'\neq\theta} \prod_{k=1}^n p_{\theta'}^{U_k}(Y_k)}, \tag{10}$$

and the stopping time τ as

$$\tau = \min_{\theta\in\Theta} \inf\{n : L_\theta(n) \geqslant \exp(b_\theta)\}, \tag{11}$$

where

$$b_\theta = \log(1/R_\theta) + \log(|\Theta|-1) + \max_{\theta\in\Theta} \pi_\theta. \tag{12}$$

Finally, the decision rule is defined as

$$\delta(y^\tau) = \theta \quad \text{if} \quad L_\theta(\tau) = \max_{\theta'\in\Theta} L_{\theta'}(\tau). \tag{13}$$

To evaluate the expected stopping time, we use Theorem 4.2 in [13], which states that if the averaged log likelihood ratio (LLR) $\frac{1}{n} Z_{\theta,\theta'}(n)$ with $Z_{\theta,\theta'}(n)$ defined as

$$Z_{\theta,\theta'}(n) = \sum_{k=1}^n \log \frac{p_\theta^{U_k}(Y_k)}{p_{\theta'}^{U_k}(Y_k)} \tag{14}$$

converges 1-quickly [13] to a number $D_{\theta,\theta'}$ under the measure $\mathbb{P}_\theta$, then $\mathbb{E}_\theta(\tau) = \frac{|\log R_\theta|}{\min_{\theta'\neq\theta} D_{\theta,\theta'}}(1+o(1))$. Therefore, in the following, we would like to show that $\frac{1}{n} Z_{\theta,\theta'}$ converge 1-quickly to $\sum_{u\in\mathcal{U}} P_U^*(u) D(p_\theta^u || p_{\theta'}^u)$, where

$$P_U^* = \begin{cases} 1-\alpha^* & \text{if } u=0 \\ \alpha^* \bar{P}_U^*(u) & \text{if } u\neq 0, \end{cases} \tag{15}$$

$\bar{P}_U^* = \text{argmax}_{\bar{P}\in\mathcal{P}_{\mathcal{U}\setminus\{0\}}} f(\bar{P},\theta)$, and

$$\alpha^* = \frac{2\eta}{\chi_2(\sum_{u\neq 0} \bar{P}_U^*(u) q^u_\theta || q^0_\theta)} \times \frac{\min_{\theta'\neq\theta} \sum_{u\neq 0} \bar{P}_U^*(u) D(p^u_\theta || p^u_{\theta'})}{|\log R_\theta|}.$$

To show 1-quickly convergence, we need to show that for any $\epsilon > 0$, $\mathbb{E}_\theta[T(\epsilon)]$ is bounded, where

$$T(\epsilon) = \sup\left\{ n : \left| Z_{\theta,\theta'}(n) - \sum_{u\in\mathcal{U}} P_U^*(u) D(p_\theta^u || p_{\theta'}^u) \right| \geqslant n\epsilon \right\}.$$

We can further define $T_1(\epsilon_1)$ and $T_2(\epsilon_2)$ as follows

$$T_1(\epsilon) = \sup\left\{ n : \left| Z_{\theta,\theta'}(n) - \sum_{k=1}^n \sum_{u\in\mathcal{U}} P_k(u) D(p_\theta^u || p_{\theta'}^u) \right| \geqslant n\epsilon_1 \right\}$$

$$T_2(\epsilon) = \sup\left\{ n : \left| \sum_{k=1}^n \sum_{u\in\mathcal{U}} (P_k(u) - P_U^*(u)) D(p_\theta^u || p_{\theta'}^u) \right| \geqslant n\epsilon_2 \right\}.$$

By the triangle inequality, it can be shown that $T(\epsilon) \leqslant \max\{T_1(\epsilon/2), T_2(\epsilon/2)\}$. Defining N as the time at which the estimate of the hypothesis θ is correct for all $n > N$, then it is known from [10, Eq. (63)] that $\mathbb{P}(N \geqslant n) \leqslant O_n(n^{-c})$ for any $c > 0$. Therefore, for some $K > 0$, we have

$$\begin{aligned} &\mathbb{P}_\theta \left\{ \left| \sum_{k=1}^n \sum_{u\in\mathcal{U}} (P_k(u) - P_U^*(u)) D(p_\theta^u || p_{\theta'}^u) \right| \geqslant n\epsilon/2 \right\} \\ &\leqslant \mathbb{P}_\theta\{NK \geqslant n\epsilon/2\} \\ &\leqslant O(n^{-c}). \end{aligned} \tag{16}$$

Similarly, one can reuse the argument in [14, Appendix] to show that

$$\mathbb{P}_\theta \left\{ \left| Z_{\theta,\theta'}(n) - \sum_{k=1}^n \sum_{u\in\mathcal{U}} P_k(u) D(p_\theta^u || p_{\theta'}^u) \right| \geqslant n\epsilon/2 \right\} \leqslant O(\gamma_1^n) \tag{17}$$

for some $\gamma_1 < 1$. Note that both $\sum_{n=1}^\infty O(n^{-c})$ and $\sum_{n=1}^\infty O(\gamma_1^n)$ are finite when we choose $c \geqslant 2$, so $\mathbb{E}_\theta\{T(\epsilon)\}$ is bounded, and we obtain from Theorem 4.2 in [13] that

$$\begin{aligned} \mathbb{E}_\theta(\tau) &= \frac{|\log R_\theta|}{\alpha^* \min_{\theta'\neq\theta} \sum_{u\neq 0} \bar{P}_U^*(u) D(p_\theta^u || p_{\theta'}^u)}(1+o(1)) \\ &= \left(\frac{|\log R_\theta|}{\min_{\theta'\neq\theta} \sum_{u\neq 0} \bar{P}_U^*(u) D(p_\theta^u || p_{\theta'}^u)} \right)^2 \\ &\quad \times \frac{\chi_2(\sum_{u\neq 0} \bar{P}_U^*(u) q_\theta^u || q_\theta^0)}{2\eta}(1+o_\tau(1)). \end{aligned}$$

Finally, we can lower bound the value of $\gamma_\theta(\sigma)$ by

$$\begin{aligned} \gamma_\theta(\sigma) &\geqslant \sqrt{2\eta} \times \frac{\min_{\theta'\neq\theta} \sum_{u\neq 0} \bar{P}_U^*(u) D(p_\theta^u || p_{\theta'}^u)}{\sqrt{\chi_2(\sum_{u\neq 0} \bar{P}_U^*(u) q_\theta^u || q_\theta^0)}} \\ &= \sqrt{2\eta} \max_{\bar{P}\in\mathcal{P}_{\mathcal{U}\setminus\{0\}}} \frac{\min_{\theta'\neq\theta} \sum_{u\neq 0} \bar{P}(u) D(p_\theta^u || p_{\theta'}^u)}{\sqrt{\chi_2(\sum_{u\neq 0} \bar{P}(u) q_\theta^u || q_\theta^0)}}. \end{aligned}$$

Next, we analyze the divergence between $P_{Z^\tau|\theta}$ and $(q_\theta^0)^{\otimes\tau}$ under the test σ. By Jensen's inequality and convexity of KL divergence, we have

$$D(P_{Z^\tau|\theta} || (q_\theta^0)^{\otimes\tau}) \leqslant \mathbb{E}\left[D(P_{Z^\tau|\theta,N} || (q_\theta^0)^{\otimes\tau}) \right], \tag{18}$$

where the expectation is taken over N. For each N, the term $D(P_{Z^\tau|\theta,N} || (q_\theta^0)^{\otimes\tau})$ can be expressed as

$$\begin{aligned} &D(P_{Z^\tau|\theta,N} || (q_\theta^0)^{\otimes\tau}) \\ &= \sum_{k=1}^N D(\sum_u P_k(u) q_\theta^u || q_\theta^0) + \sum_{k=N+1}^\tau D(\sum_u P^*(u) q_\theta^u || q_\theta^0) \\ &\leqslant N\lambda + (\tau - N) D(\sum_u P^*(u) q_\theta^u || q_\theta^0), \end{aligned}$$

for some $\lambda > 0$. Then,

$$\mathbb{E}\left[D(P_{Z^\tau|\theta} || (q_\theta^0)^{\otimes\tau}) \right]$$

$$\leqslant \tau \left(\frac{(\alpha^*)^2}{2} \chi_2(\sum_u P^*(u) q_\theta^u || q_\theta^0) + o_\tau((\alpha^*)^2) \right) + o_\tau(1)$$

by [2, Eq. (13)] and the fact that $\mathbb{E}_\theta[N] = o_\tau(1)$. Note that it can be shown that τ concentrates around $\mathbb{E}_\theta[\tau]$ with high probability because the stopping rule only depends on the likelihood ratio, and (16) and (17) imply that $\frac{1}{n} Z_{\theta,\theta'}(n)$ converges in probability to $P_U^*(u) D(p_\theta^u || p_{\theta'}^u)$ for any $\theta' \neq \theta$. Therefore, for any $\xi > 0$, with probability at least $1 - \xi$ it holds that

$$\tau \leqslant (1+\xi) \left(\frac{|\log R_\theta|}{\min_{\theta' \neq \theta} \sum_{u \neq 0} \bar{P}^*(u) D(p_\theta^u || p_{\theta'}^u)} \right)^2 \times \frac{\chi_2(\sum_{u \neq 0} \bar{P}^*(u) q_\theta^u || q_\theta^0)}{2\eta} \quad (19)$$

when $|\log R_\theta|$ is large enough. By plugging in the definition of α^* and upper bounding the value of τ by (19), it holds with probability at least $1 - \xi$ that

$$\mathbb{E}\left[D(P_{Z^\tau|\theta} || (q_\theta^0)^{\otimes \tau}) \right] \leqslant (1+\xi)\eta(1 + o_\tau(1)). \quad (20)$$

Since $\xi > 0$ can be arbitrary small, we have shown that

$$D(P_{Z^\tau|\theta} || (q_\theta^0)^\tau) \leqslant \eta \quad (21)$$

with probability 1 when $R_\theta \to 0$. Furthermore, it is known that risk constraints in (2) are satisfied when the stopping rule is defined as (11) by Lemma 3 of [14]. Therefore, $\sigma \in \Sigma(\mathbf{R}, \eta)$, and

$$\gamma_\theta^* \geqslant \sqrt{2\eta} \max_{\bar{P} \in \mathcal{P}_{\mathcal{U} \setminus \{0\}}} \frac{\min_{\theta' \neq \theta} \sum_{u \neq 0} \bar{P}(u) D(p_\theta^u || p_{\theta'}^u)}{\sqrt{\chi_2(\sum_{u \neq 0} \bar{P}(u) q_\theta^u || q_\theta^0)}}. \quad (22)$$

B. Converse

Let the action U_k be generated from some policy $P_{U_k|X^{k-1}U^{k-1}}$. From the assumption, we know that for any $u \in \mathcal{U}$, $P_{U_k|X^{k-1}U^{k-1}}(u|X^{k-1}, U^{k-1})$ converge almost surely to some $P_U^*(u)$. Therefore, the number of indices k such that $P_{U_k|X^{k-1}U^{k-1}}(u|X^{k-1}, U^{k-1}) \neq P_U^*(u)$ for some $u \in \mathcal{U}$ is finite. Define the set $\mathcal{K}$ as

$$\mathcal{K} \triangleq \{k \in \mathbb{N} : P_{U_k|X^{k-1}U^{k-1}}(u|X^{k-1}, U^{k-1}) = P_U^*(u) \ \forall u \in \mathcal{U}\}.$$

With probability 1, it holds that $|\mathbb{N} \setminus \mathcal{K}|$ is finite. Then, for all $k \in \mathcal{K}$, U_k is generated identically and independently from P_U^*. Defining the log likelihood ratio $Z_{\theta,\theta'}$ as

$$Z_{\theta,\theta'}(n) = \sum_{k=1}^n \log \frac{p_\theta^{U_k}(Y_k)}{p_{\theta'}^{U_k}(Y_k)}, \quad (23)$$

then the following lemma from [10] relates the log likelihood ratio $Z_{\theta,\theta'}$ and the risk constraints $\{R_\theta\}_{\theta \in \Theta}$.

Lemma 3. *For all sequential tests $\sigma \in \Sigma(\mathbf{R}, \eta)$, we have*

$$\lim_{R_{max} \to 0} \inf_{\sigma \in \Sigma(\mathbf{R},\eta)} \mathbb{P}_\theta \{ Z_{\theta,\theta'}(\tau) \geqslant \rho |\log R_\theta| \} = 1 \quad (24)$$

for all $\theta' \neq \theta$ and $0 < \rho < 1$.

By Azuma's inequality, we have for any $\epsilon > 0$ and any $\sigma \in \Sigma(\mathbf{R}, \eta)$ that

$$\lim_{R_{\max} \to 0} \mathbb{P}_\theta \left\{ Z_{\theta,\theta'}(\tau) \geqslant \sum_{k=1}^\tau D(p_\theta^{U_k} || p_{\theta'}^{U_k})(1+\epsilon) \right\} = 0. \quad (25)$$

Combining (25) and Lemma 3, we have

$$\lim_{R_{\max} \to 0} \inf_{\sigma \in \Sigma(\mathbf{R},\eta)} \mathbb{P}_\theta \left\{ \sum_{k=1}^\tau D(p_\theta^{U_k} || p_{\theta'}^{U_k})(1+\epsilon) \geqslant \rho |\log R_\theta| \right\} = 1$$

for all $\theta' \neq \theta$ and $0 < \rho < 1$. Defining $\rho' = \frac{\rho}{1+\epsilon}$, we have for all $\epsilon' > 0$

$$\begin{aligned}
&\lim_{R_{\max} \to 0} \mathbb{P}_\theta \left\{ \sum_{k=1}^\tau D(p_\theta^{U_k} || p_{\theta'}^{U_k}) \geqslant \rho' |\log R_\theta| \right\} \\
&= \lim_{R_{\max} \to 0} \mathbb{P}_\theta \Bigg\{ \sum_{k=1}^\tau D(p_\theta^{U_k} || p_{\theta'}^{U_k}) \geqslant \rho' |\log R_\theta|, \\
&\qquad |\widehat{p}_{\mathbf{u}}(u) - P_U^*(u)| \leqslant P_U^*(u)\epsilon' \text{ for all } u \in \mathcal{U} \Bigg\} \\
&+ \lim_{R_{\max} \to 0} \mathbb{P}_\theta \Bigg\{ \sum_{k=1}^\tau D(p_\theta^{U_k} || p_{\theta'}^{U_k}) \geqslant \rho' |\log R_\theta|, \\
&\qquad |\widehat{p}_{\mathbf{u}}(u) - P_U^*(u)| > P_U^*(u)\epsilon' \text{ for some } u \in \mathcal{U} \Bigg\}. \quad (26)
\end{aligned}$$

The first term of (26) can be upper bounded by

$$\begin{aligned}
&\lim_{R_{\max} \to 0} \mathbb{P}_\theta \Bigg\{ \tau \sum_u \widehat{p}_{\mathbf{u}}(u) D(p_\theta^u || p_{\theta'}^u) \geqslant \rho' |\log R_\theta|, \\
&\qquad |\widehat{p}_{\mathbf{u}}(u) - P_U^*(u)| \leqslant P_U^*(u)\epsilon' \text{ for all } u \in \mathcal{U} \Bigg\} \\
&\leqslant \lim_{R_{\max} \to 0} \mathbb{P}_\theta \left\{ \tau \sum_u P_U^*(u)(1+\epsilon') D(p_\theta^u || p_{\theta'}^u) \geqslant \rho' |\log R_\theta| \right\} \\
&= \lim_{R_{\max} \to 0} \mathbb{P}_\theta \left\{ \tau \sum_u P_U^*(u) D(p_\theta^u || p_{\theta'}^u) \geqslant \frac{\rho'}{1+\epsilon'} |\log R_\theta| \right\},
\end{aligned}$$

while the second term of (26) can be upper bounded by

$$\begin{aligned}
&\lim_{R_{\max} \to 0} \mathbb{P}_\theta \left\{ |\widehat{p}_{\mathbf{u}}(u) - P_U^*(u)| > P_U^*(u)\epsilon' \text{ for some } u \in \mathcal{U} \right\} \\
&\leqslant \lim_{R_{\max} \to 0} \sum_{u \in \mathcal{U}} \mathbb{P}_\theta \left\{ \left| \frac{1}{\tau} \sum_{k=1}^\tau 1(U_k = u) - P_U^*(u) \right| > P_U^*(u)\epsilon' \right\} \\
&\leqslant \lim_{R_{\max} \to 0} \sum_{u \in \mathcal{U}} \mathbb{P}_\theta \Bigg\{ \left| \frac{1}{\tau} \sum_{k \in \mathcal{K} \cap [\tau]} 1(U_k = u) - \frac{\tau - |\mathcal{K}^c \cap [\tau]|}{\tau} P_U^*(u) \right| \\
&\qquad > P_U^*(u)\epsilon' - \frac{|\mathcal{K}^c \cap [\tau]|}{\tau} P_U^*(u) - \frac{|\mathcal{K}^c \cap [\tau]|}{\tau} \Bigg\} \\
&= 0
\end{aligned}$$

by the Chernoff bound and the fact that $|\mathcal{K}^c|$ is finite with probability 1. Therefore, we have for any $\sigma \in \Sigma(\mathbf{R}, \eta)$, it holds that

$$\lim_{R_{\max} \to 0} \mathbb{P}_\theta \left\{ \tau \min_{\theta' \neq \theta} \sum_u P_U^*(u) D(p_\theta^u || p_{\theta'}^u) \geqslant \frac{\rho'}{1+\epsilon'} |\log R_\theta| \right\} = 1,$$

and

$$\lim_{R_{\max}\to 0} \mathbb{P}_\theta \left\{ \tau \geqslant \frac{\rho'}{1+\epsilon'} \frac{|\log R_\theta|}{\min_{\theta'\neq\theta} \sum_u P_U^*(u) D(p_\theta^u || p_{\theta'}^u)} \right\} = 1$$

for any $\epsilon' > 0$ and $0 < \rho' < 1$. For any $n > 0$, we define the time averaged distribution of actions as follows

$$\begin{aligned}
\bar{P}_n(u) &= \frac{1}{n}\sum_{k=1}^n P_{U_k|Y^{k-1}U^{k-1}}(u)\\
\alpha(\bar{P}_n) &= 1 - \bar{P}_n(0)\\
\tilde{P}_n(u) &= \begin{cases} \frac{\bar{P}_n(u)}{\alpha(\bar{P}_n)} & u \neq 0\\ 0 & u = 0.\end{cases}
\end{aligned}$$

Since the policy $P_{U_k|Y^{k-1}U^{k-1}}(u)$ converge almost surely to $P_U^*(u)$ for all u, we have

$$\lim_{n\to\infty} \bar{P}_n(u) = \lim_{n\to\infty} \frac{1}{n}\sum_{k=1}^n P_{U_k|Y^{k-1}U^{k-1}}(u) = P_U^*(u).$$

By [4], the covertness requirement implies that

$$\eta \geqslant D(P_{Z^n|\theta} || (q_\theta^0)^{\otimes n}) \geqslant n D(\bar{P}_{Z|\theta} || q_\theta^0)$$

for all $n \leqslant \tau$ with probability 1 when $\tau \to \infty$, where

$$\begin{aligned}
\bar{P}_{Z|\theta}(z) &= \frac{1}{n}\sum_{k=1}^n P_{Z_i|\theta}(z)\\
&= \frac{1}{n}\sum_{k=1}^n \sum_u P_{U_k|Y^{k-1}U^{k-1}}(u) q_\theta^u(z)\\
&= \sum_u \bar{P}_n(u) q_\theta^u(z).
\end{aligned}$$

Let $n = \frac{\rho'}{1+\epsilon'}\frac{|\log R_\theta|}{\min_{\theta'\neq\theta}\sum_u P_U^*(u)D(p_\theta^u||p_{\theta'}^u)}$, then for R_θ small enough we have

$$\begin{aligned}
\eta &\geqslant \frac{\rho'}{1+\epsilon'} \frac{|\log R_\theta|}{\min_{\theta'\neq\theta}\sum_u P_U^*(u) D(p_\theta^u||p_{\theta'}^u)} D(\sum_u \bar{P}_n(u) q_\theta^u || q_\theta^0)\\
&\geqslant \frac{\rho'}{1+\epsilon'} \frac{|\log R_\theta|}{\alpha(P_U^*) \min_{\theta'\neq\theta}\sum_{u\neq 0} \tilde{P}_U^*(u) D(p_\theta^u||p_{\theta'}^u)}\\
&\times \left(\frac{\alpha(P_U^*)^2}{2} \chi_2 \left(\sum_{u\neq 0} \tilde{P}_U^*(u) q_\theta^u \middle|\middle| q_\theta^0 \right) + o(\alpha(P_U^*)^2) \right).
\end{aligned}$$

Therefore,

$$\begin{aligned}
\alpha(\bar{P}_n) &\leqslant \frac{1+\epsilon'}{\rho'} \frac{\min_{\theta'\neq\theta}\sum_{u\neq 0}\tilde{P}_U^*(u) D(p_\theta^u||p_{\theta'}^u)}{|\log R_\theta|}\\
&\times \frac{2\eta}{\chi_2\left(\sum_{u\neq 0}\tilde{P}_U^*(u) q_\theta^u \middle|\middle| q_\theta^0\right)}(1+o(1)).
\end{aligned}$$

Then,

$$\begin{aligned}
\mathbb{E}_\theta[\tau] &\geqslant \left(\frac{|\log R_\theta|}{\min_{\theta'\neq\theta}\sum_{u\neq 0}\tilde{P}_U^*(u) D(p_\theta^u||p_{\theta'}^u)} \right)^2\\
&\times \frac{\chi_2\left(\sum_{u\neq 0}\tilde{P}_U^*(u) q_\theta^u \middle|\middle| q_\theta^0\right)}{2\eta}(1+o(1)).
\end{aligned}$$

Taking the infimum over all possible $\tilde{P}_U^*$, we have

$$\begin{aligned}
\mathbb{E}_\theta[\tau] &\geqslant \inf_{\tilde{P}_U^* \in \mathcal{P}_{\mathcal{U}\setminus\{0\}}} \left(\frac{|\log R_\theta|}{\min_{\theta'\neq\theta}\sum_{u\neq 0}\tilde{P}_U^*(u) D(p_\theta^u||p_{\theta'}^u)} \right)^2\\
&\times \frac{\chi_2\left(\sum_{u\neq 0}\tilde{P}_U^*(u) q_\theta^u \middle|\middle| q_\theta^0\right)}{2\eta}(1+o(1))
\end{aligned}$$

Finally,

$$\gamma_\theta^* \leqslant \sqrt{2\eta} \sup_{\tilde{P}_U^* \in \mathcal{P}_{\mathcal{U}\setminus\{0\}}} \frac{\min_{\theta'\neq\theta}\sum_{u\neq 0}\tilde{P}_U^*(u) D(p_\theta^u||p_{\theta'}^u)}{\sqrt{\chi_2\left(\sum_{u\neq 0}\tilde{P}_U^*(u) q_\theta^u \middle|\middle| q_\theta^0\right)}}.$$

V. Numerical Example

In this section, we provide an example to illustrate how the covertness constraint can influence the policy for a fixed $\eta > 0$. This example is modified from [10] and [5]. Let $\Theta = \{0, 1, 2\}$, $\mathcal{U} = \{0, a, b, c\}$, and $\mathcal{Y} = \mathcal{Z} = \{0, 1\}$. We further assume that the true hypothesis is $\theta = 0$. The kernels $\{p_\theta^u\}$ are given in Table I, and the kernels $\{q_\theta^u\}$ for the first and second scenarios are given in Table II and Table III, respectively. In the first scenario, $\gamma_\theta^* = 0.0293\sqrt{\eta}$ and the optimal control policy is $P_U^*(u) = 1(u = a)$. Note that in the first scenario, choosing either action a, b, or c is the same from the perspective of covertness because $q_0^u = 0.4$ for all $u \in \mathcal{U}\setminus\{0\}$. Besides, there exists a $\theta' \neq \theta$ that is indistinguishable from θ when choosing the action b or c. Therefore, it is not surprise that the optimal strategy in the first scenario is $1(u = a)$. However, in the second scenario, $\gamma_\theta^* = 0.0146\sqrt{\eta}$, and the optimal control policy becomes $P_U^*(u) = 0.5 \times 1(u = b) + 0.5 \times 1(u = c)$. Note that, in this scenario, the difference between the kernels q_0^0 and q_0^a is much larger than the one between q_0^0 and q_0^b or q_0^c. Therefore, the covertness constraint compromises the benefit of choosing the action a, leading to a different optimal strategy.

TABLE I
$p_\theta^u(1)$ FOR ALL u AND θ

θ \ u	0	a	b	c
0	0.1	0.4	0.6	0.6
1	0.1	0.6	0.4	0.6
2	0.1	0.6	0.6	0.4

TABLE II
$q_\theta^u(1)$ FOR ALL u AND $\theta = 0$ IN SCENARIO 1

θ \ u	0	a	b	c
0	0.01	0.4	0.4	0.4

TABLE III
$q_\theta^u(1)$ FOR ALL u AND $\theta = 0$ IN SCENARIO 2

θ \ u	0	a	b	c
0	0.01	0.9	0.4	0.4

REFERENCES

[1] B. Bash, D. Goeckel, and D. Towsley, "Limits of reliable communication with low probability of detection on AWGN channels," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 9, pp. 1921–1930, September 2013.

[2] M. R. Bloch, "Covert communication over noisy channels: A resolvability perspective," *IEEE Transactions on Information Theory*, vol. 62, no. 5, pp. 2334–2354, 2016.

[3] P. H. Che, M. Bakshi, and S. Jaggi, "Reliable deniable communication: Hiding messages in noise," in *Proc. of IEEE International Symposium on Information Theory*, Istanbul, Turkey, July 2013, pp. 2945–2949.

[4] L. Wang, G. W. Wornell, and L. Zheng, "Fundamental limits of communication with low probability of detection," *IEEE Transactions on Information Theory*, vol. 62, no. 6, pp. 3493–3503, Jun. 2016.

[5] M. Tahmasbi and M. R. Bloch, "First and second order asymptotics in covert communication," *IEEE Transactions on Information Theory*, vol. 65, no. 4, pp. 2190 –2212, Apr. 2019.

[6] B. A. Bash, C. N. Gagatsos, A. Datta, and S. Guha, "Fundamental limits of quantum-secure covert optical sensing," in *2017 IEEE International Symposium on Information Theory (ISIT)*, 2017, pp. 3210–3214.

[7] C. N. Gagatsos, B. A. Bash, A. Datta, Z. Zhang, and S. Guha, "Covert sensing using floodlight illumination," in *2019 Conference on Lasers and Electro-Optics (CLEO)*, 2019, pp. 1–2.

[8] M. Naghshvar and T. Javidi, "Active sequential hypothesis testing," *The Annals of Statistics*, vol. 41, no. 6, pp. 2703 – 2738, 2013. [Online]. Available: https://doi.org/10.1214/13-AOS1144

[9] ——, "Active sequential hypothesis testing," *The Annals of Statistics*, vol. 41, no. 6, pp. 2703–2738, 12 2013.

[10] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Transactions on Automatic and Control*, vol. 58, no. 10, pp. 2451–2464, Oct. 2013.

[11] M. Tahmasbi and M. R. Bloch, "Active covert sensing," in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 840–845.

[12] K.-W. Huang, H.-M. Wang, D. Towsley, and H. V. Poor, "Lpd communication: A sequential change-point detection perspective," *IEEE Transactions on Communications*, vol. 68, no. 4, pp. 2474–2490, 2020.

[13] V. Draglia, A. Tartakovsky, and V. Veeravalli, "Multihypothesis sequential probability ratio tests .i. asymptotic optimality," *IEEE Transactions on Information Theory*, vol. 45, no. 7, pp. 2448–2461, 1999.

[14] G. K. Atia and V. V. Veeravalli, "Controlled sensing for sequential multihypothesis testing," in *2012 IEEE International Symposium on Information Theory Proceedings*, 2012, pp. 2196–2200.