# Stability Enforced Bandit Algorithms for Channel Selection in Remote State Estimation of Gauss-Markov Processes

Alex S. Leong, *Member, IEEE*, Daniel E. Quevedo, *Fellow, IEEE*, and Wanchun Liu, *Member, IEEE*

*Abstract*—In this paper we consider the problem of remote state estimation of a Gauss-Markov process, where a sensor can, at each discrete time instant, transmit on one out of $M$ different communication channels. A key difficulty of the situation at hand is that the channel statistics are unknown. We study the case where both learning of the channel reception probabilities and state estimation is carried out simultaneously. Methods for choosing the channels based on techniques for multi-armed bandits are presented, and shown to provide stability. Furthermore, we define the performance notion of estimation regret, and derive bounds on how it scales with time for the considered algorithms.

*Index Terms*—Learning, multi-armed bandits, regret, stability, state estimation

## I. INTRODUCTION

The use of machine learning techniques for estimation and control has gained increasing attention in recent years [1]–[3]. Learning based estimation and control holds the promise of enabling the solution of problems which are difficult or even intractable using traditional control design techniques.

In this paper we focus on the problem of remote state estimation of a Gauss-Markov process, where the channel statistics are constant but unknown. As a motivating example, consider the use of unmanned aerial vehicles (UAVs) for providing wireless communication capabilities [4], [5], with the UAV acting as a mobile aerial base station collecting information from wireless sensors. The use of such systems can allow for monitoring of processes which are difficult to access otherwise, e.g., in a hostile or contested environment. The UAV would transmit collected information to a remote estimator via one of $M$ different channels, e.g., using $M$ different frequency bands. Characteristics of the individual channels can vary substantially at different locations/environments [4], [6], which makes it difficult to have accurate knowledge of the channel statistics whenever the UAV moves to a different location. In this paper we will assume that, at any given location, the channel statistics are constant but unknown.

Learning of transmission schedules for a single system over a channel with unknown packet reception probability has been studied in [7], while learning of power allocations for multiple control systems connected over an unknown non-stationary channel is considered in [8]. With multiple channels, channel allocation for multiple processes using deep reinforcement learning have been considered in [9]–[11]. While the use of deep reinforcement learning techniques can find schedules which perform well, in general stabilizing properties of the learned policies are not guaranteed. Furthermore, many training samples may be needed during learning. In this paper, we consider a remote state estimation where, at every time instant, a sensor or controller can choose one of $M$ unknown channels for transmission. The aim is to learn the best channel while guaranteeing stability of the system and being *sample efficient* (via maintaining *low regret*).

A. S. Leong is with Defence Science and Technology Group, Melbourne, Australia. E-mail: `alex.leong@defence.gov.au`. D. E. Quevedo is with the School of Electrical Engineering and Robotics, Queensland University of Technology, Brisbane, Australia. E-mail: `dquevedo@ieee.org`. W. Liu is with the School of Electrical and Information Engineering, University of Sydney, Sydney, Australia. Email: `wanchun.liu@sydney.edu.au`.

We will view the channel selection task as a multi-armed bandit type problem [12], [13]. The multi-armed bandit problem is a simple example of a reinforcement learning problem which captures the exploration vs. exploitation tradeoff, between basing decisions only on knowledge currently obtained about a system ("exploit") vs. trying out alternative decisions which may potentially be better ("explore"). Among many applications in a wide variety of domains, multi-armed bandit techniques have been used to address scheduling problems in wireless communications [6], [14]–[16].

In contrast to classical multi-armed bandit problems, here the underlying processes are dynamical systems, which raises additional challenges on stability and performance. In this paper, we will study the use of the $\varepsilon$-greedy algorithm [17], as well as several algorithms based on the Thompson Sampling technique [18], to carry out channel selection in the context of remote state estimation. Thompson Sampling is a sampling based algorithm for the multi-armed bandit problem, which has been shown to be optimal with respect to the accumulated regret in certain problems. Interestingly, Thompson sampling has been previously used to address problems in stochastic [19], [20] and linear quadratic control [21].

*Summary of contributions*: The main contributions of the current work are as follows:

- We propose the use of bandit algorithms for channel selection in the context of remote state estimation, including a new stability-aware Bayesian sampling algorithm.
- We prove that the considered algorithms can guarantee stability of the remote estimation process.
- We study performance of these algorithms by introducing the notion of estimation regret, and prove that the $\varepsilon$-greedy algorithm has linear estimation regret, while the sampling based algorithms can achieve logarithmic estimation regret.

The remainder of the paper is organized as follows: The system model is provided in Section II. A selection of bandit algorithms for channel selection are presented in Section III. Stability and performance bounds of these algorithms are studied in Sections IV and V respectively. Numerical studies are given in Section VI. Section VII draws conclusions.

*Notation*: Given a matrix $X$, we say that $X \geq 0$ if $X$ is positive semi-definite. For matrices $X$ and $Y$, we say that $X \leq Y$ if $Y - X \geq 0$. The spectral radius of a matrix $X$ is denoted by $\rho(X)$. The beta distribution with parameters $\alpha$ and $\beta$ is denoted by $\text{Beta}(\alpha, \beta)$, with expected value $\alpha/(\alpha + \beta)$. $\mathbb{1}(\cdot)$ is the indicator function that returns 1 if its argument is true and 0 otherwise.

## II. SYSTEM MODEL

We consider a discrete-time plant described by a Gauss-Markov process:

$$x_{k+1} = Ax_k + w_k, \tag{1}$$

with sensor measurements

$$y_k = Cx_k + v_k,$$

where $w_k \sim \mathcal{N}(0, Q)$, $v_k \sim \mathcal{N}(0, R)$, see Fig. 1. The sensor runs a local Kalman filter, with local state estimates $\hat{x}_k^s \triangleq$
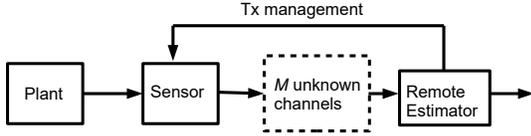
Fig. 1. Remote state estimation over unknown channels

$\mathbb{E}[x_k|y_0,\ldots,y_k]$ and estimation error covariances $P_k^s \triangleq \mathbb{E}[(x_k - \hat{x}_k^s)(x_k - \hat{x}_k^s)^T|y_0,\ldots,y_k]$. We assume that the pair $(A,C)$ is observable and the pair $(A,Q^{1/2})$ is controllable, which implies that $P_k^s$ converges to some $\overline{P}$ as $k \to \infty$ [22]. We will assume steady state, so that $P_k^s = \overline{P}, \forall k$.

As depicted in Fig. 1, the local state estimates are transmitted over lossy channels to a remote estimator. There are $M$ such channels, each i.i.d. Bernoulli with packet reception probabilities $\theta_m, m = 1,\ldots,M$. For notational simplicity, we will assume that all the packet reception probabilities are different, i.e., $\theta_i \neq \theta_j, \forall i \neq j$.

At every time step $k$, we can choose one of the $M$ channels to transmit over. This decision is determined by the remote estimator, which notifies the sensor via a transmission (Tx) management channel. As the information transmitted on the transmission management channel is normally only a few bits in length, we assume that these transmissions are error-free, see also [9].

A key issue is that none of the packet reception probabilities are known. Ideally, we would like to always use the best channel (i.e., the channel with the highest packet reception probability). But since the packet reception probabilities are unknown, this requires us to find/learn the best channel by trying out different channels and observing their outcomes. The situation resembles a multi-armed bandit type problem [12], [13], with the channels acting as the different arms. The difference between the situation considered in this paper and classical multi-armed bandit problems is that the underlying estimation process has dynamics. This raises additional issues such as stability, as investigated in the current work.

Before proceeding, we recall (see, e.g., [23]) that the remote estimator has estimation error covariance

$$P_k = \begin{cases} \overline{P}, & \gamma_k = 1 \\ h(P_{k-1}), & \gamma_k = 0 \end{cases} \tag{2}$$

where

$$h(X) \triangleq AXA^\top + Q. \tag{3}$$

In the above, $\gamma_k = 1$ means that the transmission at time $k$ was successful, while $\gamma_k = 0$ means that the transmission was lost. We assume that $P_0 = \overline{P}$.

When $A$ is unstable, we know that if the same channel $m$ is chosen at every time step, then the remote estimator has bounded expected estimation error covariance for all $k$ if and only if its packet reception probability satisfies $\theta_m > \theta_c$ [24], [25], where

$$\theta_c \triangleq 1 - \frac{1}{\rho(A)^2}. \tag{4}$$

Accordingly, in this paper we make the assumption that at least one of the channels has packet reception probability larger than $\theta_c$, i.e.,

$$\theta^* \triangleq \max_m \theta_m > \theta_c. \tag{5}$$

*A. Preliminaries*

For later reference, we mention some properties of the operator $h(.)$ defined in (3).

**Lemma 1.** *The operator $h(.)$ defined in (3) satisfies:*
*(i) $h(X) \leq h(Y)$ if $X \leq Y$*

---

**Algorithm 1** $\varepsilon$-greedy
1: **Initialize**: $\alpha_{m,1} = 1, \beta_{m,1} = 1, \quad m = 1,\ldots,M$
2: **for** $k = 1,2,\ldots$ **do**
3:     Set $\hat{\theta}_{m,k}^\varepsilon = \frac{\alpha_{m,k}}{\alpha_{m,k}+\beta_{m,k}}, \quad m = 1,\ldots,M$
4:     With probability $\varepsilon > 0$, choose $m_k^\varepsilon$ uniformly from $\{1,\ldots,M\}$, else choose[1] $m_k^\varepsilon = \arg\max_m \hat{\theta}_{m,k}^\varepsilon$
5:     Observe $\gamma_k$
6:     Update sample means via

$$\begin{aligned} \alpha_{m,k+1} &= \alpha_{m,k} + \gamma_k, & \text{if } m = m_k^\varepsilon \\ \beta_{m,k+1} &= \beta_{m,k} + 1 - \gamma_k, & \text{if } m = m_k^\varepsilon \\ \alpha_{m,k+1} &= \alpha_{m,k}, \beta_{m,k+1} = \beta_{m,k}, & \text{if } m \neq m_k^\varepsilon \end{aligned}$$

---

*(ii) $h(\overline{P}) \geq \overline{P}$*
*(iii) $\mathrm{tr}(h(\overline{P})) > \mathrm{tr}(\overline{P})$.*

*Proof.* The proof of (i) follows easily from the definition (3). Proofs of (ii) and (iii) can be found in [26]. $\square$

Note that Lemma 1 implies that $P_k \geq \overline{P}, \forall k$, for $P_k$ evolving according to (2).

The following recursion (6) for the expected error covariances will also be used later. Let $\theta(k)$ be the reception probability at time $k$. From (2) we have

$$\begin{aligned} \mathbb{E}[P_{k+1}|P_k] &= \mathbb{E}[P_{k+1}|P_k,\gamma_k = 1]\mathbb{P}(\gamma_k = 1) \\ &\quad + \mathbb{E}[P_{k+1}|P_k,\gamma_k = 0]\mathbb{P}(\gamma_k = 0) \\ &= \theta(k)\overline{P} + (1-\theta(k))(AP_kA^\top + Q). \end{aligned}$$

Then, since $h(.)$ is linear, we can obtain

$$\mathbb{E}[P_{k+1}] = \mathbb{E}[\mathbb{E}[P_{k+1}|P_k]] = \theta(k)\overline{P} + (1-\theta(k))h(\mathbb{E}[P_k]). \tag{6}$$

III. BANDIT ALGORITHMS FOR CHANNEL SELECTION

In this section we present a selection of multi-armed bandit algorithms for carrying out channel selection, mainly based on the $\varepsilon$-greedy and Thompson sampling approaches [18]. Thompson sampling is computationally efficient and can be easily implemented, requiring very little tuning apart from the choice of a prior distribution. There exist other bandit algorithms such as those based on the upper confidence bound (UCB) approach, which are not considered in the current work. Numerical studies however have shown that in many situations the performance of Thompson sampling is highly competitive against even well-tuned UCB algorithms [27]–[29].

*A. $\varepsilon$-greedy*

We first consider the well-known $\varepsilon$-greedy algorithm [17], which with probability $\varepsilon > 0$ chooses a channel uniformly at random ("explore"), and with probability $1 - \varepsilon$ chooses the channel with the maximum estimated (via a sample mean) packet reception probability ("exploit"). After transmitting over this channel and observing $\gamma_k$, i.e., whether the transmission was successful, we update the sample means of the packet reception probabilities. The procedure is stated in Algorithm 1, where we have expressed the estimated packet reception probabilities in terms of the quantities $\alpha_{m,k}$ and $\beta_{m,k}$. These can be regarded as pseudo-counts of the number of successful and non-successful transmissions respectively on channel $m$ [18], and are updated using the formulas $\alpha_{m,k+1} = \alpha_{m,k} + \gamma_k$ and $\beta_{m,k+1} = \beta_{m,k} + 1 - \gamma_k$. This allows us to emphasize similarities in the implementation with the other algorithms considered below.

---
[1]If multiple channels have the same maximum estimated packet reception probability, then a channel is chosen uniformly at random from these (maximizing) channels.

---

**Algorithm 2** Thompson Sampling (TS)

1: **Initialize**: $\alpha_{m,1} = 1, \beta_{m,1} = 1, \quad m = 1, \dots, M$
2: **for** $k = 1, 2, \dots$ **do**
3:      Sample $\hat{\theta}_{m,k}^{\mathrm{TS}} \sim \mathrm{Beta}(\alpha_{m,k}, \beta_{m,k}), \quad m = 1, \dots, M$
4:      Choose $m_k^{\mathrm{TS}} = \arg\max_m \hat{\theta}_{m,k}^{\mathrm{TS}}$ and observe $\gamma_k$
5:      Update posteriors to $\mathrm{Beta}(\alpha_{m,k+1}, \beta_{m,k+1})$ via

$$\begin{aligned}
\alpha_{m,k+1} &= \alpha_{m,k} + \gamma_k, & &\text{if } m = m_k^{\mathrm{TS}} \\
\beta_{m,k+1} &= \beta_{m,k} + 1 - \gamma_k, & &\text{if } m = m_k^{\mathrm{TS}} \\
\alpha_{m,k+1} &= \alpha_{m,k}, \ \beta_{m,k+1} = \beta_{m,k}, & &\text{if } m \neq m_k^{\mathrm{TS}}
\end{aligned}$$

---

**Algorithm 3** Optimistic Bayesian Sampling (OBS)

1: **Initialize**: $\alpha_{m,1} = 1, \beta_{m,1} = 1, \quad m = 1, \dots, M$
2: **for** $k = 1, 2, \dots$ **do**
3:      Sample $\hat{\theta}_{m,k}^{\mathrm{TS}} \sim \mathrm{Beta}(\alpha_{m,k}, \beta_{m,k}), \quad m = 1, \dots, M$
4:      Set $\hat{\theta}_{m,k}^{\mathrm{OBS}} = \max\left( \hat{\theta}_{m,k}^{\mathrm{TS}}, \frac{\alpha_{m,k}}{\alpha_{m,k}+\beta_{m,k}} \right)$
5:      Choose $m_k^{\mathrm{OBS}} = \arg\max_m \hat{\theta}_{m,k}^{\mathrm{OBS}}$ and observe $\gamma_k$
6:      Update posteriors to $\mathrm{Beta}(\alpha_{m,k+1}, \beta_{m,k+1})$ via

$$\begin{aligned}
\alpha_{m,k+1} &= \alpha_{m,k} + \gamma_k, & &\text{if } m = m_k^{\mathrm{OBS}} \\
\beta_{m,k+1} &= \beta_{m,k} + 1 - \gamma_k, & &\text{if } m = m_k^{\mathrm{OBS}} \\
\alpha_{m,k+1} &= \alpha_{m,k}, \ \beta_{m,k+1} = \beta_{m,k}, & &\text{if } m \neq m_k^{\mathrm{OBS}}
\end{aligned}$$

---

*B. Thompson Sampling*

Thompson sampling is a sampling based approach which draws samples $\hat{\theta}_{m,k}^{\mathrm{TS}}, m = 1, \dots, M$ from the posterior distribution of the unknown packet reception probability of each arm/channel. The posterior distribution can be shown to have a beta distribution if the prior distribution is beta distributed [18], as we will assume here. In particular, the prior beta distribution with $\alpha_{m,1} = 1, \beta_{m,1} = 1, m = 1, \dots, M$ is uniformly distributed in $[0,1]$. The channel that is selected is then given by

$$m_k^{\mathrm{TS}} = \arg\max_m \hat{\theta}_{m,k}^{\mathrm{TS}}.$$

After transmitting over this channel and observing $\gamma_k$, the posterior is then updated to a $\mathrm{Beta}(\alpha_{m,k+1}, \beta_{m,k+1})$ distribution, using the same formulas for updating $\alpha_{m,k}$ and $\beta_{m,k}$ as in Section III-A. The procedure is summarized in Algorithm 2.

*C. Optimistic Bayesian Sampling*

In [28] a modification of Thompson sampling called "Optimistic Bayesian sampling" (OBS) was proposed and studied. Recall that in Thompson sampling, at time $k$ one would draw samples $\hat{\theta}_{m,k}^{\mathrm{TS}}, m = 1, \dots, M$ and then select $m_k^{\mathrm{TS}} = \arg\max_m \hat{\theta}_{m,k}^{\mathrm{TS}}$. In OBS we would also draw samples $\hat{\theta}_{m,k}^{\mathrm{TS}}, m = 1, \dots, M$, but now select

$$m_k^{\mathrm{OBS}} = \arg\max_m \hat{\theta}_{m,k}^{\mathrm{OBS}},$$

where

$$\hat{\theta}_{m,k}^{\mathrm{OBS}} \triangleq \max\left( \hat{\theta}_{m,k}^{\mathrm{TS}}, \mathbb{E}[\hat{\theta}_{m,k}^{\mathrm{TS}}] \right) = \max\left( \hat{\theta}_{m,k}^{\mathrm{TS}}, \frac{\alpha_{m,k}}{\alpha_{m,k}+\beta_{m,k}} \right).$$

The procedure is summarized in Algorithm 3. The motivation for OBS is to result in increased selection probabilities for arms which are uncertain [28]. For classical multi-armed bandit problems, simulations have shown that OBS can outperform Thompson sampling in certain situations [28], [30].

---

**Algorithm 4** Stability-aware Bayesian Sampling (SBS)

1: **Initialize**: $\alpha_{m,1} = 1, \beta_{m,1} = 1, \quad m = 1, \dots, M$
2: **for** $k = 1, 2, \dots$ **do**
3:      Sample $\hat{\theta}_{m,k}^{\mathrm{TS}} \sim \mathrm{Beta}(\alpha_{m,k}, \beta_{m,k}), \quad m = 1, \dots, M$
4:      Set $\hat{\theta}_{m,k}^{\mathrm{SBS}} = \eta_{m,k}\hat{\theta}_{m,k}^{\mathrm{OBS}} + (1 - \eta_{m,k})\hat{\theta}_{m,k}^{\mathrm{TS}}$, where $\hat{\theta}_{m,k}^{\mathrm{OBS}} = \max\left( \hat{\theta}_{m,k}^{\mathrm{TS}}, \frac{\alpha_{m,k}}{\alpha_{m,k}+\beta_{m,k}} \right)$ and $\eta_{m,k} = \mathbb{1}(\alpha_{m,k}/(\alpha_{m,k}+\beta_{m,k}) > \theta_c)$
5:      Choose $m_k^{\mathrm{SBS}} = \arg\max_m \hat{\theta}_{m,k}^{\mathrm{SBS}}$ and observe $\gamma_k$
6:      Update posteriors to $\mathrm{Beta}(\alpha_{m,k+1}, \beta_{m,k+1})$ via

$$\begin{aligned}
\alpha_{m,k+1} &= \alpha_{m,k} + \gamma_k, & &\text{if } m = m_k^{\mathrm{SBS}} \\
\beta_{m,k+1} &= \beta_{m,k} + 1 - \gamma_k, & &\text{if } m = m_k^{\mathrm{SBS}} \\
\alpha_{m,k+1} &= \alpha_{m,k}, \ \beta_{m,k+1} = \beta_{m,k}, & &\text{if } m \neq m_k^{\mathrm{SBS}}
\end{aligned}$$

---

*D. Stability-aware Bayesian Sampling*

In the context of selecting channels for the purpose of remote state estimation of dynamical systems as in (1), it is important to take into account the long term estimation performance. In particular, the fundamental bound (4) suggests that, unlike when using OBS, one should not artificially stimulate the selection probabilities for channels which seem to be poor. More specifically, if the current mean-estimate $\alpha_{m,k}/(\alpha_{m,k} + \beta_{m,k})$ of the reception probability is less than the critical threshold $\theta_c$, then it is questionable whether this channel should be chosen with a probability higher than that prescribed by Thompson sampling.

Motivated by the above, in this paper, we propose a *stability-aware modification* of OBS which uses the OBS sample $\hat{\theta}_{m,k}^{\mathrm{OBS}}$ only if $\alpha_{m,k}/(\alpha_{m,k} + \beta_{m,k}) > \theta_c$, while using the Thompson sample $\hat{\theta}_{m,k}^{\mathrm{TS}}$ otherwise. The algorithm, here named *Stability-aware Bayesian Sampling* (SBS), can thus be characterized by:

$$\hat{\theta}_{m,k}^{\mathrm{SBS}} = \eta_{m,k}\hat{\theta}_{m,k}^{\mathrm{OBS}} + (1 - \eta_{m,k})\hat{\theta}_{m,k}^{\mathrm{TS}}, \quad m = 1, \dots, M \qquad (7)$$

and

$$m_k^{\mathrm{SBS}} = \arg\max_m \hat{\theta}_{m,k}^{\mathrm{SBS}},$$

where

$$\eta_{m,k} = \begin{cases} 1 & \text{if } \alpha_{m,k}/(\alpha_{m,k} + \beta_{m,k}) > \theta_c \\ 0 & \text{otherwise.} \end{cases} \qquad (8)$$

The method is summarized in Algorithm 4.

## IV. STABILITY OF REMOTE STATE ESTIMATION WITH CHANNEL SELECTION

In this section, we show that the $\varepsilon$-greedy algorithm will ensure stability of the remote estimator, provided the exploration rate $\varepsilon$ is not too high, while all of the sampling based channel selection schemes presented in the previous section (Thompson sampling, OBS, and SBS) will ensure stability. It is important to emphasize that Theorems 1 and 2, stated below, cover the challenging case where channel qualities are unknown, and where channel selection and learning are done in real-time.

Let $m_k$ be the channel chosen at time $k$, and define

$$m^* \triangleq \operatorname{argmax}_m \theta_m.$$

**Theorem 1.** *Suppose condition* (5) *holds. Then under the $\varepsilon$-greedy algorithm, the expected estimation error covariance $\mathbb{E}[P_k]$ is bounded for all $k$ if and only if*

$$0 < \varepsilon < \frac{\theta^* - \theta_c}{\theta^* - \frac{1}{M}\sum_{m=1}^{M} \theta_m}. \qquad (9)$$

*Proof.* Because of the exploration phase, where with probability $\varepsilon > 0$ a channel is chosen randomly, it is easy to see that each channel will be accessed infinitely often. Hence the estimates of the packet reception probabilities for each of the channels will converge to their true values as $k \to \infty$. Thus, during the exploitation phase we have

$$\mathbb{P}(m_k = m^* | \text{exploit}) \to 1 \text{ as } k \to \infty,$$

while during the exploration phase we have

$$\mathbb{P}(m_k = m^* | \text{explore}) = \frac{1}{M}.$$

Then, given any $\delta > 0$, there exists a $K(\delta) < \infty$ such that

$$\mathbb{P}(\text{reception at time } k) \geq (1 - \varepsilon)\theta^*(1 - \delta) + \frac{\varepsilon}{M} \sum_{m=1}^{M} \theta_m$$

for all $k > K(\delta)$. Now let

$$\theta' \triangleq (1 - \varepsilon)\theta^*(1 - \delta) + \frac{\varepsilon}{M} \sum_{m=1}^{M} \theta_m.$$

For $k > K(\delta)$, we have by Lemma 1 and a similar derivation to (6)

$$\mathbb{E}[P_{k+1}] \leq \theta'\overline{P} + (1 - \theta')h(\mathbb{E}[P_k]).$$

By induction, it follows that $\{\mathbb{E}[P_k]\}$ will be upper bounded by a sequence $\{V_k\}$ defined by

$$V_k = \mathbb{E}[P_k], \quad k \leq K(\delta)$$
$$V_{k+1} = \theta'\overline{P} + (1 - \theta')h(V_k), \quad k \geq K(\delta).$$

The sequence $\{V_k\}$ converges if and only if $\theta' > \theta_c$. If

$$(1 - \varepsilon)\theta^* + \frac{\varepsilon}{M} \sum_{m=1}^{M} \theta_m > \theta_c, \quad (10)$$

then one can always find a sufficiently small $\delta > 0$ to satisfy $\theta' > \theta_c$. As (10) is equivalent to the condition (9), this proves the "if" direction of the theorem.

For the "only if" direction, first define

$$\tilde{\theta} \triangleq \theta^*(1 - \varepsilon) + \frac{\varepsilon}{M} \sum_{m=1}^{M} \theta_m, \quad (11)$$

and note that, for a given $\varepsilon$,

$$\mathbb{P}(\text{reception at time } k) \leq \tilde{\theta}$$

holds, because $\tilde{\theta}$ is the reception probability assuming the optimal $m^*$ is always chosen during the exploitation phase. If $\varepsilon \geq (\theta^* - \theta_c)/(\theta^* - \frac{1}{M}\sum_{m=1}^{M}\theta_m)$, then $\mathbb{P}(\text{reception at time } k) \leq \tilde{\theta} \leq \theta_c$, and thus

$$\mathbb{E}[P_{k+1}] \geq \theta_c\overline{P} + (1 - \theta_c)h(\mathbb{E}[P_k]).$$

Define a sequence $\{\tilde{V}_k\}$ by

$$\tilde{V}_{k+1} = \theta_c\overline{P} + (1 - \theta_c)h(\tilde{V}_k),$$

which from the definition of $\theta_c$ is a divergent sequence. An induction argument shows that $\mathbb{E}[P_k] \geq \tilde{V}_k, \forall k$. This implies that $\mathbb{E}[P_k]$ also diverges and completes the proof. $\square$

**Remark 1.** *If $(\theta^* - \theta_c)/(\theta^* - \frac{1}{M}\sum_{m=1}^{M}\theta_m) > 1$, then Theorem 1 establishes that the $\varepsilon$-greedy algorithm provides estimation stability for all $\varepsilon > 0$.*

We next consider stability of the sampling based algorithms. We first give a preliminary result.

**Lemma 2.** *Under i) Thompson sampling, ii) OBS, and iii) SBS, all channels will be used infinitely often, and*

$$\mathbb{P}(m_k = m^*) \to 1 \text{ as } k \to \infty. \quad (12)$$

*Proof.* For Thompson sampling and OBS, this is proved in [28] (see also [27] for the two armed bandit case under Thompson sampling). For SBS we note that, in view of (7), it holds that $\hat{\theta}_{m,k}^{\text{TS}} \leq \hat{\theta}_{m,k}^{\text{SBS}} \leq \hat{\theta}_{m,k}^{\text{OBS}}$. Thus, intuitively (12) should also hold, as SBS in a sense lies in between Thompson sampling and OBS (and convergence holds for both of these schemes). A rigorous proof of Lemma 2 for SBS can be given using similar arguments as in [28]. The details are omitted for brevity. $\square$

**Theorem 2.** *Suppose condition (5) holds. Then under i) Thompson sampling, ii) OBS, and iii) SBS, the expected error covariance $\mathbb{E}[P_k]$ is bounded for all $k$.*

*Proof.* From Lemma 2, we know that under all three sampling based schemes, given any $\delta > 0$, there exists a $K(\delta) < \infty$ such that

$$\mathbb{P}(\text{reception at time } k) \geq \theta^*(1 - \delta), \forall k > K(\delta).$$

Pick a sufficiently small $\delta$ such that $\theta' \triangleq \theta^*(1 - \delta) > \theta_c$ still holds. Then for $k > K(\delta)$, we have by Lemma 1 and a similar derivation to (6) that

$$\mathbb{E}[P_{k+1}] \leq \theta'\overline{P} + (1 - \theta')h(\mathbb{E}[P_k]).$$

Define a sequence $\{V_k\}$ by

$$V_k = \mathbb{E}[P_k], \quad k \leq K(\delta)$$
$$V_{k+1} = \theta'\overline{P} + (1 - \theta')h(V_k), \quad k \geq K(\delta).$$

Using an induction argument, we can show that $V_k$ upper bounds $\mathbb{E}[P_k]$ for all $k$. Since $\theta' > \theta_c$, the sequence $\{V_k\}$ converges, and in particular $V_k$ is bounded for all $k$. Hence $\mathbb{E}[P_k]$ is also bounded for all $k$. $\square$

**Remark 2.** *Suppose the system matrix $A$, and hence $\theta_c$, is not known exactly, but we believe $\theta_c$ to be equal to $\hat{\theta}_c$. Then in Line 4 of Algorithm 4 we would replace $\theta_c$ by $\hat{\theta}_c$. However, even with an incorrect value $\hat{\theta}_c$, we still have $\hat{\theta}_{m,k}^{\text{TS}} \leq \hat{\theta}_{m,k}^{\text{SBS}} \leq \hat{\theta}_{m,k}^{\text{OBS}}$ and Lemma 2 would still hold. Hence SBS guarantees stability even for an incorrect $\hat{\theta}_c$, provided condition (5) holds for the true $\theta_c$.*

## V. PERFORMANCE BOUNDS

In the multi-armed bandit literature, it is common to analyze the performance of an algorithm via the notion of *regret* [12], [13], which is defined as the difference between the expected cumulative reward from always playing the optimal arm and the expected cumulative reward using a particular bandit algorithm. The $\varepsilon$-greedy algorithm when applied to the standard multi-armed bandit problem is well known to achieve linear regret. More sophisticated algorithms which can achieve logarithmic regret (which, in a sense, is best possible [31]) include UCB [32], $\varepsilon$-greedy with a carefully chosen decaying exploration probability [32], and Thompson sampling [33], [34].

For the remote estimation problem at hand, it is convenient to regard the trace of the estimation error covariance $\text{tr}P_k$ as a *cost*, or alternatively regard $-\text{tr}P_k$ as a *reward*, see also [9]. Motivated by the notion of regret in standard multi-armed bandit problems, in this paper we define the *estimation regret* over a horizon $T$ for the remote state estimation problem as

$$\text{regret}^E(T) \triangleq \sum_{k=1}^{T} \left( -\text{tr}\mathbb{E}[P_k^*] - (-\text{tr}\mathbb{E}[P_k]) \right)$$
$$= \sum_{k=1}^{T} \text{tr}(\mathbb{E}[P_k] - \mathbb{E}[P_k^*]). \quad (13)$$

In (13), $\mathbb{E}[P_k^*]$ is the expected error covariance assuming that the optimal channel $m^*$ is always chosen, and $\mathbb{E}[P_k]$ is the expected error covariance when using a particular channel selection algorithm. Thus, (13) constitutes a measure of the degree of suboptimality incurred from using a particular algorithm. Stationary and transient performance of an algorithm can be quantified by inspecting how fast the regret grows in relation to the horizon length $T$.

To state our results, we first recall some order notation.

**Definition 1.** *Given functions $f(.)$ and $g(.)$, we say that $f(T) = O(g(T))$ if there exist constants $c$ and $T_0$, such that $|f(T)| \leq cg(T), \forall T \geq T_0$. We say that $f(T) = \Omega(g(T))$ if $g(T) = O(f(T))$, and that $f(T) = \Theta(g(T))$ if both $f(T) = O(g(T))$ and $f(T) = \Omega(g(T))$. We say that $f(T) = o(g(T))$ if $\lim_{T \to \infty} f(T)/g(T) = 0$.*

For $\varepsilon$-greedy, the estimation regret scales linearly with $T$.

**Theorem 3.** *Suppose conditions (5) and (9) hold. Then under the $\varepsilon$-greedy algorithm, we have*

$$\text{regret}^E(T) = \Theta(T).$$

*Proof.* We first show that $\text{regret}^E(T) = O(T)$. By Theorem 1, there exists some bounded $\mathcal{P}^\varepsilon$ such that

$$\mathbb{E}[P_k] \leq \mathcal{P}^\varepsilon, \quad \forall k.$$

We have

$$\begin{aligned}
\text{regret}^E(T) &= \sum_{k=1}^{T} \text{tr}(\mathbb{E}[P_k] - \mathbb{E}[P_k^*]) \\
&\leq \sum_{k=1}^{T} \text{tr}\mathbb{E}[P_k] \\
&\leq \text{tr}(\mathcal{P}^\varepsilon)T = O(T).
\end{aligned}$$

We now show that $\text{regret}^E(T) = \Omega(T)$. Recall the definition $\tilde{\theta}$ from (11). Note that we have

$$\mathbb{E}[P_k] \geq \tilde{\theta}\overline{P} + (1 - \tilde{\theta})h(\mathbb{E}[P_{k-1}]) \geq \tilde{\theta}\overline{P} + (1 - \tilde{\theta})h(\mathbb{E}[P_{k-1}^*]),$$
$$\mathbb{E}[P_k^*] = \theta^*\overline{P} + (1 - \theta^*)h(\mathbb{E}[P_{k-1}^*]).$$

Then

$$\begin{aligned}
\text{tr}(\mathbb{E}[P_k] - \mathbb{E}[P_k^*]) &\geq \text{tr}\big((\theta^* - \tilde{\theta})(h(\mathbb{E}[P_{k-1}^*]) - \overline{P})\big) \\
&\geq (\theta^* - \tilde{\theta}) \min_{k \geq 1} \text{tr}\big(h(\mathbb{E}[P_{k-1}^*]) - \overline{P}\big) \\
&\triangleq r. \quad (14)
\end{aligned}$$

By the definition (11) and Lemma 1, we have $r > 0$. Hence

$$\begin{aligned}
\text{regret}^E(T) &= \sum_{k=1}^{T} \text{tr}(\mathbb{E}[P_k] - \mathbb{E}[P_k^*]) \\
&\geq rT = \Omega(T).
\end{aligned}$$
$\square$

For the Bernoulli bandit problem with per stage rewards in $\{0, 1\}$, it is known that the regret scales logarithmically with $T$ under Thompson sampling [33], [34]. An "age-of-information regret" measure was recently considered in [29], where the per stage cost is unbounded, and can increase by (at most) one at every time step. It is shown that this age-of-information regret also scales logarithmically with $T$. In the current work, it follows from (3) that the per stage cost $\mathbb{E}[P_k]$ (or reward $-\mathbb{E}[P_k]$) may also become unbounded, and furthermore can increase at an exponential rate. Interestingly, for the estimation regret introduced in (13), a bound that is logarithmic in $T$ still holds.
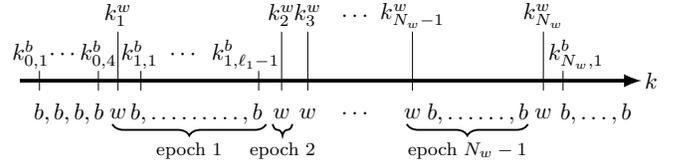


Fig. 2. Channel selection based on schedule $\mathcal{A}$, where 'w' and 'b' refer to the worst and the best channel selection, respectively.

Let us refer to sub-optimal channels as those whose packet reception probabilities are not equal to the optimal $\theta^*$. We begin with a preliminary result.

**Lemma 3.** *Let $N_{\text{sub}}(T)$ denote the number of uses of sub-optimal channels over horizon $T$. Then, under i) Thompson sampling, ii) OBS, and iii) SBS, we have*

$$\mathbb{E}[N_{\text{sub}}(T)] = \Theta(\log T).$$

*Proof.* The property that $\mathbb{E}[N_{\text{sub}}(T)] = O(\log T)$ is shown for Thompson sampling in [33], [34]. The corresponding result for OBS is proved in [30], based on similar arguments as in [33]. For SBS, $\mathbb{E}[N_{\text{sub}}(T)] = O(\log T)$ can be shown by making use of the relation $\hat{\theta}_{m,k}^{\text{TS}} \leq \hat{\theta}_{m,k}^{\text{SBS}} \leq \hat{\theta}_{m,k}^{\text{OBS}}$ and the arguments of [30]. The (rather lengthy) details are omitted for brevity.

We next recall a result from [31], namely that any policy with $\mathcal{R}(T) = o(T^a)$ for every $a > 0$ satisfies $\mathbb{E}[N_{\text{sub}}(T)] = \Omega(\log T)$, where $\mathcal{R}(T) \triangleq \mathbb{E}\big[\sum_{k=1}^{T}(\theta^* - \theta(k))\big]$ is the classical notion of regret for the Bernoulli bandit problem, see e.g. [33], [34]. Since for Thompson sampling, OBS, and SBS, we have

$$\mathcal{R}(T) = O\big(\mathbb{E}[N_{\text{sub}}(T)]\big) = O(\log T) = o(T^a), \forall a > 0,$$

where the first equality follows from standard arguments [33], [34] and the second equality is what we have just shown, the property $\mathbb{E}[N_{\text{sub}}(T)] = \Omega(\log T)$ therefore holds for all three schemes. $\square$

**Theorem 4.** *Suppose that condition (5) holds. Then under i) Thompson sampling, ii) OBS, and iii) SBS, we have*

$$\text{regret}^E(T) = \Theta(\log T).$$

Before giving the formal proof of this result, we first provide a sketch of the proof strategy for the upper bound $\text{regret}^E(T) = O(\log T)$. Following the idea of [29], which showed a logarithmic bound for the age-of-information regret, we will upper bound the estimation regret with the regret that is achieved using an alternative schedule, say $\mathcal{A}$, that replaces all uses of sub-optimal channels with the worst channel. In [29], this schedule $\mathcal{A}$ is then further upper bounded by a "worst-case" schedule $\mathcal{B}$ that groups all the uses of the worst channel together. However, for the problem at hand, as the expected error covariance can increase exponentially fast if the worst channel does not stabilize the remote estimator, this argument will give the desired logarithmic regret bound only in problem instances where the worst channel (and hence all channels) is stabilizing. To cover more interesting and challenging situations, to establish Theorem 4, we will consider a division of the time interval into "epochs" separated by successive uses of the worst channel, see Fig. 2. Through a careful analysis, we will show that the accumulated regret within each epoch is bounded, while the expected number of epochs is logarithmic in $T$. As a consequence the required logarithmic regret upper bound is established. The detailed proof of Theorem 4 now follows.

*Proof.* We will first prove that $\text{regret}^E(T) = O(\log T)$. Define the reception probability of the worst channel as

$$\theta_w \triangleq \min_m \theta_m.$$

Given any schedule of channel selections, let schedule $\mathcal{A}$ denote an alternative schedule which replaces all uses of sub-optimal channels with the worst channel. Let $\mathbb{E}[P_k^{\mathcal{A}}]$ be the expected error covariance under this replacement procedure. Then from (6) and Lemma 1, we can easily show that $\mathbb{E}[P_k] \leq \mathbb{E}[P_k^{\mathcal{A}}], \forall k$. In particular,

$$\sum_{k=1}^{T} \text{tr}\mathbb{E}[P_k] \leq \sum_{k=1}^{T} \text{tr}\mathbb{E}[P_k^{\mathcal{A}}]. \tag{15}$$

Then, we have

$$\text{regret}^E(T) \leq \text{regret}^{\mathcal{A}}(T) \triangleq \sum_{k=1}^{T} \text{tr}\left(\mathbb{E}[P_k^{\mathcal{A}}] - \mathbb{E}[P_k^*]\right). \tag{16}$$

Thus, we only need to prove that $\text{regret}^{\mathcal{A}}(T) = O(\log T)$.

Let $k_i^w \in \mathbb{N}, i \geq 1$ denote the time index of the $i$th usage of the worst channel based on schedule $\mathcal{A}$. The time horizon is divided by $\{k_i^w\}$ into epochs, as illustrated in Fig. 2. The $i$th epoch starts from $k_i^w$ and ends before $k_{i+1}^w$, and is of epoch length $\ell_i \triangleq k_{i+1}^w - k_i^w \geq 1$. Note that $\ell_i$ is a random process with a time-varying distribution due to the bandit algorithms. We define

$$\Delta_i^w \triangleq \mathbb{E}[P_{k_i^w}^{\mathcal{A}}] - \mathbb{E}[P_{k_i^w}^*] \tag{17}$$

as the gap between the expected estimation error covariances achieved by schedule $\mathcal{A}$ and the persistent schedule of the best channel.

Let $k_{i,j}^b \in \mathbb{N}, i, j \geq 1$ denote the time index based on schedule $\mathcal{A}$ of the $j$th usage of the best channel in epoch $i$, when the epoch length $\ell_i > 1$ (see Fig. 2). Let $k_{0,j}^b \in \mathbb{N}, j \geq 1$ denote the time index of the $j$th usage of the best channel before the first epoch and $\ell_0$ denote the total number of time slots before the first epoch. Then, we define

$$\Delta_{i,j}^b \triangleq \mathbb{E}[P_{k_{i,j}^b}^{\mathcal{A}}] - \mathbb{E}[P_{k_{i,j}^b}^*], i \geq 0, j \geq 1, \text{ if } \ell_i > 1. \tag{18}$$

In particular, we have $\Delta_{i,j}^b = 0$ when $i = 0$, since schedule $\mathcal{A}$ is identical to the persistent schedule of the best channel before the first epoch.

We use random variable $N_w$ to denote the number of uses of the worst channel within a time interval of length $T$. Then, $\text{regret}^{\mathcal{A}}(T)$ in (16) can be upper bounded as

$$
\begin{aligned}
\text{regret}^{\mathcal{A}}(T) &\leq \mathbb{E}\left[\sum_{j=1}^{\ell_0 - 1} \text{tr}\Delta_{0,j}^b + \sum_{i=1}^{N_w}\left(\text{tr}\Delta_i^w + \sum_{j=1}^{\ell_i - 1} \text{tr}\Delta_{i,j}^b\right)\right] \\
&= \mathbb{E}\left[\sum_{i=1}^{N_w}\left(\text{tr}\Delta_i^w + \sum_{j=1}^{\ell_i - 1} \text{tr}\Delta_{i,j}^b\right)\right],
\end{aligned} \tag{19}
$$

where the expectation is over $N_w$ and $\{\ell_i\}$. To prove the desired result, we will analyze $\left(\text{tr}\Delta_i^w + \sum_{j=1}^{\ell_i - 1} \text{tr}\Delta_{i,j}^b\right)$ first.

*Analysis of per epoch regret*: By applying the estimation error covariance updating rule (2) into (18), we have

$$
\begin{aligned}
\Delta_{i,j}^b &= (1 - \theta^*)h(\mathbb{E}[P_{k_{i,j}^b - 1}^{\mathcal{A}}]) + \theta^*\overline{P} \\
&\quad - \left((1 - \theta^*)h(\mathbb{E}[P_{k_{i,j}^b - 1}^*]) + \theta^*\overline{P}\right) \\
&= (1 - \theta^*)A\left(\mathbb{E}[P_{k_{i,j}^b - 1}^{\mathcal{A}}] - \mathbb{E}[P_{k_{i,j}^b - 1}^*]\right)A^\top.
\end{aligned} \tag{20}
$$

For the time slots in epoch $i$ when $\ell_i > 1$, it directly follows that

$$\Delta_{i,j}^b = (1 - \theta^*)^j A^j \Delta_i^w (A^j)^\top, 1 \leq j \leq \ell_i - 1. \tag{21}$$

Applying the trace operator to inequality (21), it can be obtained that

$$\text{tr}\Delta_{i,j}^b = (1 - \theta^*)^j \text{tr}\left(A^j \Delta_i^w (A^j)^\top\right) \tag{22}$$

$$\leq (1 - \theta^*)^j \text{tr}((A^j)^\top A^j)\text{tr}\Delta_i^w \tag{23}$$

$$\leq \kappa_1(\epsilon)(1 - \theta^*)^j (\rho(A) + \epsilon)^{2j} \text{tr}\Delta_i^w \tag{24}$$

where $\epsilon > 0$ is arbitrary and $\kappa_1(\epsilon) > 0$. Inequality (23) is due to the property that $\text{tr}(UV) \leq \text{tr}(U)\text{tr}(V)$ for any positive semidefinite matrices $U$ and $V$, see e.g. [35, p.445]. Inequality (24) is due to the fact that $\text{tr}((A^j)^\top A^j)$ is the sum of squares of all elements of $A^j$, and is obtained based on the lemma below about the element-wise bounds of matrix powers [36, Lemma 2]:

**Lemma 4.** *Consider a $z$-by-$z$ matrix $Z$, and let $[Z]_{i,i'}$ denote the element at the $i$th row and $i'$th column of $Z$.*

*(i) For any $\epsilon > 0$, there exists a $\kappa_1(\epsilon) > 0$ such that $\left|[Z^j]_{i,i'}\right|^2 \leq \kappa_1(\epsilon)(\rho(Z) + \epsilon)^{2j}, \forall i, i' \in \{1, \ldots, z\}$.*

*(ii) There exist $\kappa_2 > 0$ and $L \in \mathbb{N}$ such that $\left|[Z^j]_{i,i'}\right|^2$ is periodically lower bounded by $\kappa_2 \rho(Z)^{2j}$ with period $L$.*

From assumption (5), we can find a sufficiently small $\epsilon$ such that $(1 - \theta^*)(\rho(A) + \epsilon)^2 < 1$. Using the inequality (24), we obtain

$$\sum_{j=1}^{\ell_i - 1} \text{tr}\Delta_{i,j}^b \leq \frac{\kappa_1(\epsilon)\text{tr}\Delta_i^w}{1 - (1 - \theta^*)(\rho(A) + \epsilon)^2}. \tag{25}$$

Therefore, the regret of epoch $i$ is bounded as

$$\text{tr}\Delta_i^w + \sum_{j=1}^{\ell_i - 1} \text{tr}\Delta_{i,j}^b \leq \kappa(\epsilon)\text{tr}\Delta_i^w, \tag{26}$$

where $\kappa(\epsilon) \triangleq 1 + \frac{\kappa_1(\epsilon)}{1 - (1 - \theta^*)(\rho(A) + \epsilon)^2} > 0$.

*Analysis of* $\text{tr}\Delta_i^w$: By applying the estimation error covariance updating rule (2) into (17), we have

$$
\begin{aligned}
\Delta_i^w &= (1 - \theta_w)h(\mathbb{E}[P_{k_i^w - 1}^{\mathcal{A}}]) + \theta_w\overline{P} \\
&\quad - \left((1 - \theta^*)h(\mathbb{E}[P_{k_i^w - 1}^*]) + \theta^*\overline{P}\right) \\
&\leq (1 - \theta_w)h(\mathbb{E}[P_{k_i^w - 1}^{\mathcal{A}}]) + \theta_w\overline{P} \\
&= (1 - \theta_w)h(\mathbb{E}[P_{k_i^w - 1}^{\mathcal{A}}] - \mathbb{E}[P_{k_i^w - 1}^*]) \\
&\quad + (1 - \theta_w)h(\mathbb{E}[P_{k_i^w - 1}^*]) + \theta_w\overline{P} \\
&\leq (1 - \theta_w)A(\mathbb{E}[P_{k_i^w - 1}^{\mathcal{A}}] - \mathbb{E}[P_{k_i^w - 1}^*])A^\top + G
\end{aligned} \tag{27}
$$

where $G \triangleq (1 - \theta_w)\left(A\mathcal{P}A^\top + Q\right) + \theta_w\overline{P}$, and $\mathcal{P}$ is a bound on $\mathbb{E}[P_k^*]$ which exists due to the stability condition (5).

Next, using a similar argument as in the proof of Theorem 2, we note that the expected error covariances $\{\mathbb{E}[P_k^{\mathcal{A}}]\}$ are bounded by a positive semidefinite matrix $\mathcal{P}'$ for all $k$. From (27), it directly follows that $\Delta_i^w \leq (1 - \theta_w)A\mathcal{P}'A^\top + G$. Then, we can find a positive constant $\xi$ such that

$$\text{tr}\Delta_i^w \leq \xi, \forall i. \tag{28}$$

By taking (26) and (28) into (19) and using Lemma 3,

$$\text{regret}^E(T) \leq \text{regret}^{\mathcal{A}}(T)$$

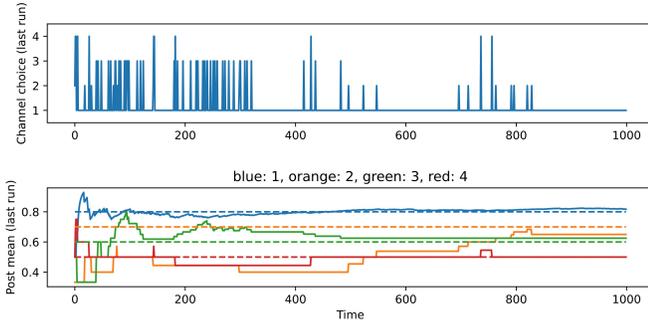$$\leq \mathbb{E}\left[\sum_{i=1}^{N_w} \kappa(\epsilon)\xi\right] = \kappa(\epsilon)\xi\mathbb{E}[N_w] = O(\log T).$$

Fig. 3. Learning using Stability-aware Bayesian Sampling

To conclude, we now prove that $\text{regret}^E(T) = \Omega(\log T)$. We have

$$\text{regret}^E(T) = \sum_{k=1}^{T} \text{tr}(\mathbb{E}[P_k] - \mathbb{E}[P_k^*])$$

$$= \mathbb{E}\left[ \sum_{k=1}^{T} \text{tr}(\mathbb{E}[P_k] - \mathbb{E}[P_k^*]) \mathbb{1}(\theta(k) = \theta^*) \right.$$

$$\left. + \sum_{k=1}^{T} \text{tr}(\mathbb{E}[P_k] - \mathbb{E}[P_k^*]) \mathbb{1}(\theta(k) \neq \theta^*) \right]$$

$$\geq \mathbb{E}\left[ \sum_{k=1}^{T} \text{tr}(\mathbb{E}[P_k] - \mathbb{E}[P_k^*]) \mathbb{1}(\theta(k) \neq \theta^*) \right].$$

Let $\theta^\circ$ denote the second largest packet reception probability. When $\theta(k) \neq \theta^*$, we can show similar to (14) that

$$\text{tr}(\mathbb{E}[P_k] - \mathbb{E}[P_k^*]) \geq \text{tr}\big((\theta^* - \theta^\circ)(h(\mathbb{E}[P_{k-1}^*]) - \overline{P})\big)$$

$$\geq (\theta^* - \theta^\circ) \min_{k \geq 1} \text{tr}\big(h(\mathbb{E}[P_{k-1}^*]) - \overline{P}\big)$$

$$\triangleq r^\circ > 0.$$

Thus

$$\text{regret}^E(T) \geq r^\circ \mathbb{E}\left[ \sum_{k=1}^{T} \mathbb{1}(\theta(k) \neq \theta^*) \right]$$

$$= r^\circ \mathbb{E}[N_{\text{sub}}(T)]$$

$$= \Omega(\log T),$$

where the last line follows from Lemma 3. This completes the proof that $\text{regret}^E(T) = \Omega(\log T)$, and hence establishes Theorem 4. □

**Remark 3.** *We note that some works on learning based control have made assumptions that the underlying system is stable [37], or can be stabilized for all possible controls [21], when carrying out a regret analysis. This stands in contrast to Theorem 4, stated above, which merely requires that at least one of the $M$ channels is stabilizing for the remote estimator. Our result thus considers scenarios where some of the sub-optimal channels can cause the expected error covariance to diverge if they are used too often.*

## VI. NUMERICAL STUDIES

We consider a situation with $M = 4$ channels. Figure 3 illustrates how Stability-aware Bayesian sampling (SBS) learns, and increasingly uses, the optimal channel, for a situation where

$$A = \begin{bmatrix} 1.5 & 0.2 \\ 0.3 & 0.9 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 1 \end{bmatrix}, \quad Q = I, \quad R = 1$$

and the (unknown) channel success probabilities are given by

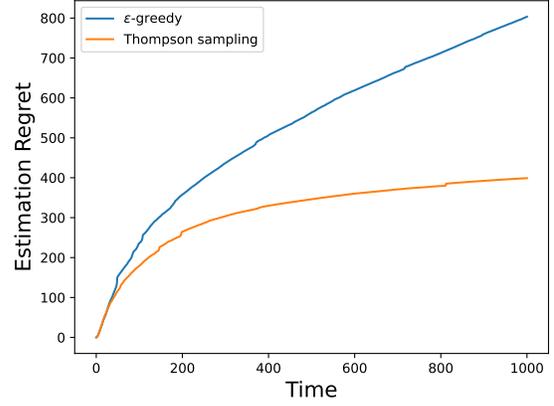$$\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3, \theta_4) = (0.8, \ 0.7, \ 0.6, \ 0.5).$$



Fig. 4. Estimation regret over time

Table I illustrates the empirical estimation regret for a number of different channel probabilities $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3, \theta_4)$ and system matrices $A$, $C$, $Q$, $R$. All bandit algorithms studied in this work, namely $\varepsilon$-greedy, Thompson sampling (TS), OBS, and SBS, are considered. The regret is computed over horizon $T = 1000$, and averaged over 100000 runs. For $\varepsilon$-greedy, we considered different values of $\varepsilon$ in steps of 0.02, with the value giving the smallest simulated regret presented in Table I.

As expected, the $\varepsilon$-greedy algorithm is outperformed by the sampling based algorithms. In all the considered scenarios OBS and SBS can achieve at least some improvement in performance over Thompson sampling.[2] In many scenarios the performance of OBS and SBS are quite close to each other. However in more challenging situations, where there are only few stabilizing channels (just one or two), SBS seems to be able to outperform OBS. This could be due to the fact that in these cases using non-stabilizing channels more than necessary can significantly degrade performance.

In Fig. 4 we plot the estimation regret over time for the $\varepsilon$-greedy and Thompson sampling algorithms. Plots for OBS and SBS are qualitatively similar to Thompson sampling and are omitted. The parameters used are the same as those mentioned at the start of Section VI. We see a linear scaling (at larger times) for $\epsilon$-greedy and a logarithmic scaling in the case of Thompson sampling, in agreement with Theorems 3 and 4.

## VII. CONCLUSION

We have studied a remote state estimation problem where a sensor can choose between a number of channels with unknown channel statistics, with which to transmit over. We have made use of bandit algorithms to learn the statistics of the channels, while simultaneously carrying out the estimation procedure. Stability of the estimator using these algorithms has been shown, and bounds on the estimation regret have been derived. Future work will include the use of bandit algorithms in the allocation of channels to multiple processes.

## REFERENCES

[1] L. Buşoniu, T. de Bruin, D. Tolić, J. Kober, and I. Palunko, "Reinforcement learning for control: Performance, stability, and deep approximators," *Annual Reviews in Control*, vol. 46, pp. 8–28, 2018.

[2] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger, "Learning-based model predictive control: Toward safe learning in control," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 3, pp. 269–296, 2020.

---

[2]Note that when looking at individual simulation runs, we cannot say that any one scheme will outperform any other.

TABLE I
ESTIMATION REGRET OF DIFFERENT ALGORITHMS

| $\theta$ | $A$ | | $C$ | $Q$ | $R$ | $\theta_c$ | $\varepsilon$ | $\varepsilon$-greedy | TS | OBS | SBS |
|---|---|---|---|---|---|---|---|---|---|---|---|
| $(0.8, 0.75, 0.55, 0.5)$ | 1.45 | | 1 | 1 | 1 | 0.524 | 0.12 | 263 | 143 | 136 | **135** |
| $(0.7, 0.6, 0.5, 0.4)$ | 1.45 | | 1 | 1 | 1 | 0.524 | 0.18 | 996 | 679 | **596** | 603 |
| $(0.6, 0.5, 0.4, 0.3)$ | 1.45 | | 1 | 1 | 1 | 0.524 | 0.22 | 10319 | 8253 | 6823 | **6276** |
| $(0.7, 0.6, 0.4, 0.3)$ | $\begin{bmatrix} 1.2 & 0.1 \\ 0.2 & 1.1 \end{bmatrix}$ | | $[1 \ 1]$ | $I$ | 1 | 0.408 | 0.10 | 582 | 295 | 274 | **273** |
| $(0.65, 0.55, 0.45, 0.35)$ | $\begin{bmatrix} 1.2 & 0.1 \\ 0.2 & 1.1 \end{bmatrix}$ | | $[1 \ 1]$ | $I$ | 1 | 0.408 | 0.12 | 886 | 517 | **473** | 483 |
| $(0.55, 0.45, 0.35, 0.25)$ | $\begin{bmatrix} 1.2 & 0.1 \\ 0.2 & 1.1 \end{bmatrix}$ | | $[1 \ 1]$ | $I$ | 1 | 0.408 | 0.18 | 4723 | 3424 | 2727 | **2430** |
| $(0.9, 0.8, 0.7, 0.5)$ | $\begin{bmatrix} 1.5 & 0.2 \\ 0.3 & 0.9 \end{bmatrix}$ | | $[1 \ 1]$ | $I$ | 1 | 0.603 | 0.14 | 290 | 104 | 98 | **97** |
| $(0.8, 0.7, 0.6, 0.5)$ | $\begin{bmatrix} 1.5 & 0.2 \\ 0.3 & 0.9 \end{bmatrix}$ | | $[1 \ 1]$ | $I$ | 1 | 0.603 | 0.18 | 803 | 403 | **354** | 363 |
| $(0.7, 0.6, 0.5, 0.4)$ | $\begin{bmatrix} 1.5 & 0.2 \\ 0.3 & 0.9 \end{bmatrix}$ | | $[1 \ 1]$ | $I$ | 1 | 0.603 | 0.22 | 7185 | 6178 | 3903 | **3106** |

[3] G. Cherubini, M. Guay, and S. Tarbouriech, Eds., *Special Issue on Learning and Control*. IEEE Control Systems Letters, Jul. 2020, vol. 4, no. 3.

[4] Y. Zeng, R. Zhang, and T. J. Lim, "Wireless communications with unmanned aerial vehicles: Opportunities and challenges," *IEEE Commun. Mag.*, vol. 54, no. 5, pp. 36–42, May 2016.

[5] N. Mohamed, J. Al-Jaroodi, I. Jawhar, A. Idries, and F. Mohammed, "Unmanned aerial vehicles applications in future smart cities," *Techn. Forecasting & Social Change*, vol. 153, no. 119293, Apr. 2020.

[6] F. Li, D. Yu, H. Yang, J. Yu, H. Karl, and X. Cheng, "Multi-armed-bandit-based spectrum sharing algorithms in wireless networks: A survey," *IEEE Wireless Commun.*, vol. 27, no. 1, pp. 24–30, Feb. 2020.

[7] S. Wu, X. Ren, Q.-S. Jia, K. H. Johansson, and L. Shi, "Learning optimal scheduling policy for remote state estimation under uncertain channel condition," *IEEE Trans. Control Netw. Syst.*, vol. 7, no. 2, pp. 579–591, Jun. 2020.

[8] M. Eisen, K. Gatsis, G. J. Pappas, and A. Ribeiro, "Learning in wireless control systems over nonstationary channels," *IEEE Trans. Signal Process.*, vol. 67, no. 5, pp. 1123–1137, Mar. 2019.

[9] A. S. Leong, A. Ramaswamy, D. E. Quevedo, H. Karl, and L. Shi, "Deep reinforcement learning for wireless sensor scheduling in cyber-physical systems," *Automatica*, vol. 113, pp. 476–486, Mar. 2020.

[10] A. Redder, A. Ramaswamy, and D. E. Quevedo, "Deep reinforcement learning for scheduling in large-scale networked control systems," in *Proc. IFAC NecSys*, Chicago, USA, Sep. 2019.

[11] W. Liu, K. Huang, D. E. Quevedo, B. Vucetic, and Y. Li, "Deep reinforcement learning for wireless scheduling in distributed networked control," 2021. [Online]. Available: https://arxiv.org/pdf/2109.12562.pdf

[12] A. Slivkins, *Introduction to Multi-Armed Bandit*. now Publishers Inc., 2019.

[13] T. Lattimore and C. Szepesvári, *Bandit Algorithms*. Cambridge University Press, 2020.

[14] Y. Gai, B. Krishnamachari, and R. Jain, "Learning multiuser channel allocations in cognitive radio networks: A combinatorial multi-armed bandit formulation," in *Proc. IEEE DySPAN*, Singapore, Apr. 2010.

[15] T. Stahlbuhk, B. Shrader, and E. Modiano, "Learning algorithms for scheduling in wireless networks with unknown channel statistics," *Ad Hoc Networks*, vol. 85, pp. 131–144, 2019.

[16] S. Takeuchi, M. Hasegawa, K. Kanno, A. Uchida, N. Chauvet, and M. Naruse, "Dynamic channel selection in wireless communications via a multi-armed bandit algorithm using laser chaos time series," *Scientific Reports*, vol. 10, no. 1574, 2020.

[17] R. S. Sutton and A. G. Barto, *Reinforcement Learning*, 2nd ed. Massachusetts: The MIT Press, 2018.

[18] D. J. Russo, B. Van Roy, A. Kazerouni, I. Osband, and Z. Wen, "A tutorial on Thompson sampling," *Foundations and Trends in Machine Learning*, vol. 11, no. 1, pp. 1–96, 2018.

[19] M. J. Kim, "Thompson sampling for stochastic control: The finite parameter case," *IEEE Trans. Autom. Control*, vol. 62, no. 12, pp. 6415–6422, Dec. 2017.

[20] D. Banjević and M. J. Kim, "Thompson sampling for stochastic control: The continuous parameter case," *IEEE Trans. Autom. Control*, vol. 64, no. 10, pp. 4137–4152, Oct. 2019.

[21] Y. Ouyang, M. Gagrani, and R. Jain, "Posterior sampling-based reinforcement learning for control of unknown linear systems," *IEEE Trans. Autom. Control*, vol. 65, no. 8, pp. 3600–3607, Aug. 2020.

[22] B. D. O. Anderson and J. B. Moore, *Optimal Filtering*. New Jersey: Prentice Hall, 1979.

[23] L. Shi, M. Epstein, and R. M. Murray, "Kalman filtering over a packet-dropping network: A probabilistic perspective," *IEEE Trans. Autom. Control*, vol. 55, no. 3, pp. 594–604, Mar. 2010.

[24] Y. Xu and J. P. Hespanha, "Estimation under uncontrolled and controlled communications in networked control systems," in *Proc. IEEE Conf. Decision and Control*, Seville, Spain, December 2005, pp. 842–847.

[25] L. Schenato, "Optimal estimation in networked control systems subject to random delay and packet drop," *IEEE Trans. Autom. Control*, vol. 53, no. 5, pp. 1311–1317, Jun. 2008.

[26] L. Shi and H. Zhang, "Scheduling two Gauss-Markov systems: An optimal solution for remote state estimation under bandwidth constraint," *IEEE Trans. Signal Process.*, vol. 60, no. 4, pp. 2038–2042, Apr. 2012.

[27] O.-C. Granmo, "Solving two-armed Bernoulli bandit problems using a Bayesian learning automaton," *Int. J. Intelligent Computing and Cybernetics*, vol. 3, no. 2, pp. 207–234, 2010.

[28] B. C. May, N. Korda, A. Lee, and D. S. Leslie, "Optimistic Bayesian sampling in contextual-bandit problems," *J. Machine Learning Research*, pp. 2069–2106, 2012.

[29] S. Fatale, K. Bhandani, U. Narula, S. Moharir, and M. K. Hanawal, "Regret of age-of-information bandits," *IEEE Trans. Commun.*, vol. 70, no. 1, pp. 87–100, Jan. 2022.

[30] J. Mellor, "Decision making using Thompson sampling," Ph.D. dissertation, University of Manchester, 2014.

[31] T. L. Lai and H. Robbins, "Asymptotically efficient adaptive allocation rules," *Adv. Appl. Math.*, vol. 6, no. 1, pp. 4–22, Mar. 1985.

[32] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine Learning*, vol. 47, pp. 235–256, 2002.

[33] S. Agrawal and N. Goyal, "Near-optimal regret bounds for Thompson sampling," *J. ACM*, vol. 64, no. 5, 2017.

[34] E. Kaufmann, N. Korda, and R. Munos, "Thompson sampling: An asymptotically optimal finite-time analysis," in *Proc. Int. Conf. Algorithmic Learning Theory*, Lyon, France, Oct. 2012, pp. 199–213.

[35] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed. Cambridge, UK: Cambridge University Press, 2013.

[36] W. Liu, D. E. Quevedo, Y. Li, K. H. Johansson, and B. Vucetic, "Remote state estimation with smart sensors over Markov fading channels," *IEEE Trans. Autom. Control*, vol. 67, no. 6, pp. 2743–2757, Jun. 2022.

[37] S. Lale, K. Azizzadenesheli, B. Hassibi, and A. Anandkumar, "Adaptive control and regret minimization in linear quadratic Gaussian (LQG) setting," in *Proc. Amer. Control Conf.*, New Orleans, USA, May 2021, pp. 2517–2522.