Measurement-Efficient Dynamics Change Detection in On-Off Models for Dynamic Spectrum Access

Simon Lindståhl*†, Alexandre Proutiere*§, Andreas Johnsson†‡
*KTH Royal Institute of Technology, Division of Decision and Control Systems, Sweden
† Ericsson Research, Research Area Artificial Intelligence, Sweden
‡ Uppsala University, Department of Information Technology, Sweden
§Digital Futures, Sweden

***Abstract*—To maintain efficient scheduling for dynamic spectrum access problems, it is crucial to promptly detect changes in the statistical properties of spectrum occupancy. Compared to traditional change detection problems, this is complicated by the fact that measurements are not independent through time and can instead have Markovian dependencies. Moreover, classical change detection methods neglect the cost associated with measurements and do not consider the potential benefits of adapting the measurement schedule based on the observed state and the perceived likelihood of a change. This may result in high measurement overhead. In this paper, we study measurement-efficient change detection in Markovian models and demonstrate its applicability for spectrum access problems. In particular, we study problems with two states corresponding to spectrum occupancy, so called *on-off* models, and show important properties of these problems. For these problems, we establish fundamental limits that are imposed when the detection agent must maintain a sufficiently small false alarm rate. We also propose two classes of algorithms designed to adapt to different aspects of the problem. We analyze the behavior of these algorithms and evaluate them, using both synthetic data as well as real Wi-Fi spectrum data.**

***Index Terms*—Change point detection, sequential analysis, non-i.i.d. data, measurements, cognitive radio, dynamic spectrum access.**

I. Introduction

Due to the scarcity of radio spectrum, it is vital to efficiently utilize radio spectrum without causing conflicts in transmission. To this end, Dynamic Spectrum Access (DSA) is useful, thanks to its ability to compact secondary users into the gaps left by primary users, and it is a promising technology for spectrum management [1]. However, any DSA algorithm will make assumptions about the dynamics of the spectrum, and these dynamics can sometimes change rapidly. Such changes are caused by specific events such as batch user arrivals or departures, or mass-handovers to/from neighboring cells. While the spectrum can be monitored to detect such changes, this monitoring induces costly overhead such as power consumption in the sensor, as well as costs of data storage and management.

In this paper, we study a DSA scenario where a change detection agent wishes to detect a change as soon as possible without raising too many false alarms, and without taking more measurements per time unit than necessary. This is illustrated in Figure 1. The detection agent observes the spectrum by communicating with a DSA agent in an access point or base station, and raises an alarm to the DSA agent when a change is detected. However, spectrum occupancy is typically independent of earlier occupancy given the current occupancy, which as such can be modeled as a Markov chain. Because of this Markov dependence, the information gained per measurement diminishes as the interval between measurements grows smaller. In addition, we are interested in devising *data-efficient* detection schemes, i.e., limiting the number of taken measurements. The problem of data-efficient quickest change detection has been studied before [2], [3], [4], but without the Markovian dependence (in scenarios where the results of measurements are independent and identically distributed). In particular, existing approaches all tend to measure very often when they believe a change has occurred, allowing for less measurements otherwise. These approaches would fail in case of Markovian dependence, as this would lead to vastly inefficient measurements being taken, even when no change has occurred.

This paper addresses this challenge and presents the following contributions. (1) We formalize a general framework for measurement-efficient dynamics change detection in finite state Markov chain models with known transition probabilities, tailored for DSA. We also analyze fundamental limits that any algorithm on this problem are bound to, in terms of detection

ISBN 978-3-903176-72-0 ©2025 IFIP

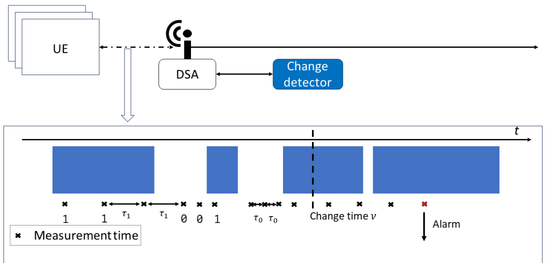

Fig. 1. The change detection agent communicating with a DSA algorithm. Also illustrated: typical spectrum behavior and measurement sequence, where the blue blocks represent the busy periods of the spectrum. The static measurement policy defined in Section V is also shown.

delay, false alarm rate and measurement volume. (2) We present two simple spectrum monitoring algorithms which adapt to this problem, one which focuses on adapting to potential asymmetries within the Markov chain and another which, inspired by related work, changes its measurement frequency based on the belief that a change has already occurred. (3) We evaluate these algorithms on real Wi-Fi spectrum data as well as synthetic data generated by finite-state Markov chain models.

II. Related work

Dynamic Spectrum Access is a well-researched topic with roots in opportunistic spectrum access [5]. It deals with utilizing spectrum frequency as efficiently as possible over time, while ensuring that some primary users can transmit effectively with minimal conflicts. Generally, the problem is treated either through modelling the spectrum [6], [7], [8], [9], learning and predicting user behavior [10], [11] or learning optimal policies directly [12]. Another framework, which we do not consider in this paper, is one where the secondary users themselves attempt to sense and access the spectrum in a distributed manner [13], [14], [15]. None of these methods, however, directly treat the rapid changes in spectrum behavior that can be caused by events within the network, leaving an interesting gap for our work.

The framework of statistical change detection in a frequentist setting was designed by Lorden [16], adapting a cumulative sum statistic of a pre-change and post-change distribution and showing its asymptotic optimality. Lai extended this framework to work for non-i.i.d. settings in which a clear concept of average information per measurement could be established, showing finite state Markov chains to be such case [17]. Instead, Banerjee and Veeravalli extended this problem by allowing the agent to choose not to take a measurement in each slot [2] while Lindståhl et al. used a formulation with continous time to show more general results for any measurement-efficient agent, as well as analyzed the suboptimality of periodic measurement agents [3], [4]. However, neither of these frameworks were able to handle the non-i.i.d. measurements suggested by Lai.

A related, well-studied problem to ours is to apply the change detection framework to detect incoming and exiting primary users, this has been done by Badawy et al. [18], Shakhov and Koo [19] and Haug et al. [20]. Our problem differs in that we are less interested in the entering and exiting of individual primary users and are instead interested in the system dynamics (or primary user behavior) changing from one pair of entering/exiting rate to another. This problem has been treated by Khan et al. [21], which took an application-based, less analytical approach as well as Zou and Wang [22] who studied offline change detection of autoregressive time series (rather than the more commonly used online stopping problem). Furthermore, our setup differs in that we consider physical time as well as the number of measurements, introducing an additional degree of freedom and complexity for the agent. The payoff for this additional complexity is the ability to reduce day-to-day measurement costs while maintaining good detection performance.

III. Preliminaries

Problem formulation: When sensing a single channel, the measurement agent can typically only infer if the channel is idle (0) or busy (1). Furthermore, if the channel is idle, it is reasonable to assume that the time to the next busy period is an exponentially distributed random variable. If one also makes the same assumption about the length of the busy period, this defines a two-state continuous time Markov chain with state space $\mathcal{S} = \{0, 1\}$. This process is memoryless, and the transition rate from state 0 to state 1 is denoted λ_0^∞, the converse transition rate is denoted λ_1^∞. This is known as the *on-off* model, and is common in DSA [6], [5]. At some unknown time ν, this system changes to instead having transition rates $\{\lambda_i^0\}_{i=0,1}$. Our task, then, is to detect this change as quickly as possible on average over the worst possible change time as well as pre-change measurements, while maintaining a low false alarm rate as well as low measurement overhead. In this endeavor, a measurement agent can observe the current state of the process at any time, and we denote the n-th such measurement as X_n. The probability measure with respect to the pre-change process is denoted as $\mathbb{P}_\infty$ while the probability measure with respect to the post-change process is denoted as $\mathbb{P}_0$. If the change occurs at time $0 < \nu < \infty$, the probability measure for the entire system is denoted as $\mathbb{P}_\nu$ (expectations are notated analogously).

We model the ability of observing at any time as a sequential sampling strategy, with the window between the $n-1$-th and n-th measurement denoted as τ_n, and the n-th measurement taking place at time $t_n = \sum_{k=1}^{n} \tau_k$. We denote by $\mathcal{F}_n$ the sigma-algebra $\mathcal{F}_n := \sigma(\tau_1, X_1, ..., \tau_n, X_n)$. The agent is tasked with making a decision to stop and raise an alarm whenever it believes the change has occurred, which is modeled as a stopping time N (with respect to the filtration $\{\mathcal{F}_n\}_{n=1}^{\infty}$).

We wish to minimize the average detection delay in the worst case, defined as $\overline{\mathbb{E}}[t_N] := \sup_{\nu \geq 0} \text{ess sup } \mathbb{E}_\nu[t_N - \nu | \mathcal{F}_{n_\nu}]$[1], where n_ν is the index of the last measurement before the change. Similarly, we define the average worst case post-change measurement volume as $\overline{\mathbb{E}}[N] := \sup_{\nu \geq 0} \text{ess sup } \mathbb{E}_\nu[N - n_\nu | \mathcal{F}_{n_\nu}]$. In the i.i.d. case, it is possible to make $\overline{\mathbb{E}}[t_N]$ arbitrarily small by increasing the measurement frequency. This is, however, not possible for on-off systems.

Given a sufficiently low detection delay, we wish to minimize the pre-change measurement frequency. Define $n(T)$ as the greatest index such that $t_n < T$, this is of course a random variable. We wish to minimize the average frequency, defined as $\mathbb{E}_\infty[f] := \limsup_{T\to\infty} \mathbb{E}_\infty[n(T)/T]$, defined only when no stopping occurs. We also demand that the average run length to false alarm is sufficiently great, modeled as the constraint $\mathbb{E}_\infty[t_N] \geq \gamma$.

Properties of on-off models: For on-off models, the probability of observing a transition after waiting for period τ can be calculated exactly, and is (pre-change)

$$p_{01}^{\infty}(\tau) = \frac{\lambda_0^\infty}{\lambda_0^\infty + \lambda_1^\infty}(1 - \exp(-(\lambda_0^\infty + \lambda_1^\infty)\tau)) \quad (1)$$

when transitioning from state 0 to state 1 (the reverse transition is calculated analogously). Similar quantities for the post-change distribution can be defined analogously. Some important properties of these probabilities are that $\lim_{\tau\to\infty} p_{ij}^\infty(\tau) = \pi_j^\infty$, the stationary probability of the Markov chain, and that $\lim_{\tau\to 0} \frac{p_{ij}^\infty(\tau)}{\tau} = \lambda_i^\infty$ by definition of rate of transition.

We define by $\boldsymbol{\tau} := (\tau_0, \tau_1)$ a static policy, which always waits for time τ_0 after observing state 0 and time τ_1 after observing state 1. Such a policy defines a discrete time Markov chain when applied to the continuous time Markov chain of the original problem. Furthermore, the stationary distribution of the corresponding Markov chain is defined (pre-change) by $\pi_0^\infty(\boldsymbol{\tau}) = \frac{p_{10}^\infty(\tau_1)}{p_{10}^\infty(\tau_1) + p_{01}^\infty(\tau_0)}$.

An important question is how to discern between the pre-change and post-change processes. This is done by observing the outcome distributions from each possible previously observed state. Having observed state 0, the average information gained after waiting for time τ and then measuring is

$$D(p_{01}^0(\tau)||p_{01}^\infty(\tau)) = \sum_{j=0}^{1} p_{0j}^0(\tau) \log\left(\frac{p_{0j}^0(\tau)}{p_{0j}^\infty(\tau)}\right).$$

The average information per measurement of a static policy is denoted $I(\boldsymbol{\tau})$ and is calculated as

$$I(\boldsymbol{\tau}) = \sum_{i=0}^{1} \pi_i^0(\boldsymbol{\tau}) D(p_i^0(\tau_i)||p_i^\infty(\tau_i)).$$

Finally, the average waiting time of a static policy $\boldsymbol{\tau}$ depends on whether the change has occurred, but under the post-change distribution it is $\overline{\tau}^0(\boldsymbol{\tau}) := \pi_0^0(\boldsymbol{\tau})\tau_0 + \pi_1^0(\boldsymbol{\tau})\tau_1$.

IV. Fundamental Limits

Following conventional wisdom about change point detection, it stands to reason that for any detection agent to have sufficiently low false alarm rate, one can only stop once a sufficient amount of information has been acquired. As such, the amount of measurement required, $\overline{\mathbb{E}}[N]$, multiplied by the information gained per measurement, $I(\boldsymbol{\tau})$, must exceed some threshold $S^{(0)}$. The number of measurements required would then be $\overline{\mathbb{E}}[N] \approx \frac{S^{(0)}}{I(\boldsymbol{\tau})}$ and the worst case detection delay would proportional to $\frac{\tau}{I(\boldsymbol{\tau})}$. We then wish to examine the limit of this expression as $\tau \to 0$. As such, we examine the function $d(\tau) := \frac{\tau}{I(\tau,\tau)}$ for $\tau > 0$.

When τ is small, the probability of transitioning exactly once from state 0 to state 1 is, under λ^0, approximately $\lambda_0^0 \tau \exp(-\lambda_0^0 \tau)$, and the information gain within one slot is then approximately

$$\begin{aligned} D(p_{01}^0(\tau)||p_{01}^\infty(\tau)) &\approx \\ &\approx \tau\lambda_0^0 \log\left(\frac{\lambda_0^0}{\lambda_0^\infty}\right) + \log\left(\frac{1-\lambda_0^0\tau}{1-\lambda_0^\infty\tau}\right) \\ &\approx \tau\left(\lambda_0^0 \log\left(\frac{\lambda_0^0}{\lambda_0^\infty}\right) + \lambda_0^\infty - \lambda_0^0\right) \end{aligned}$$

where, in both approximations, we have preserved only the first order terms in τ. We now let $\tau \to 0$ and notice that

$$d(0) := \lim_{\tau\to 0} d(\tau) \quad (2)$$

$$= \frac{1}{\sum_{i=0}^{1} \pi_i^0 \left(\lambda_i^0 \log\left(\frac{\lambda_i^0}{\lambda_i^\infty}\right) + \lambda_i^\infty - \lambda_i^0\right)}. \quad (3)$$

We call $d(0)$ the *characteristic delay* of the on-off model, and it provides a fundamental limit of the delay of any change point algorithm which seeks to avoid unnecessarily great false alarm rates, regardless of measurement frequency.

To show the relevance of this characteristic delay, we will consider a limiting case of an algorithm which

[1]The essential supremum, "ess sup", is the supremum over all elements $\mathcal{E} \in \mathcal{F}_{n_\nu}$ such that $\mathbb{P}_\nu(\mathcal{E}) > 0$.

at every time has knowledge of the current state. As such, it would be notified immediately whenever a state transition occurs and would otherwise be secure in knowledge that no transition has occurred. We denote these as *perfect knowledge agents*. For such algorithms, we have the following result.

Proposition 1. *Any perfect knowledge agent with* $\mathbb{E}_\infty[t_N] \geq \gamma$ *must also fulfill*

$$\overline{\mathbb{E}}[t_N] \geq (d(0) + o(1))C(\gamma) \tag{4}$$

as $\gamma \to \infty$, *where* $C(\gamma) = \log(\gamma \sum_i \pi_i^\infty \lambda_i^\infty)$.

The proof of Proposition 1, as well as other theoretical results, can be found in Appendix A. It stands to reason that no agent with imperfect information can outperform an agent with perfect information, and so, the above lower bound also holds for any change point detection agent with imperfect information. This makes the characteristic delay $d(0)$ interesting to study.

Since it is impossible to detect changes with a delay smaller than $C(\gamma)d(0)$, we can study agents which detect with a delay only slightly greater than this, say $\beta C(\gamma)$ with $\beta \geq d(0)$. This, in turn, gives a lower bound on the required number of post-change measurements, and this bound can be met by a static measurement policy. This is outlined by the following result.

Proposition 2. *Any agent with* $\mathbb{E}_\infty[t_N] \geq \gamma$ *and* $\overline{\mathbb{E}}[t_N] \leq \beta C(\gamma)$ *has its expected post-change measurements lower-bounded as*

$$\overline{\mathbb{E}}[N] \geq (I_{\max}(\beta)^{-1} + o(1))C(\gamma) \tag{5}$$

as $\gamma \to \infty$, *where* $I_{\max}(\beta)$ *is the value of the following optimization problem*

$$\max_{\boldsymbol{\tau} \geq 0} \; I(\boldsymbol{\tau}) \quad \text{s.t. } \frac{\overline{\tau}^0(\boldsymbol{\tau})}{I(\boldsymbol{\tau})} \leq \beta. \tag{6}$$

Notably, the statement of Proposition 2 also outlines how to find a static policy which asymptotically meets the lower bound, and that this bound is concerned with post-change measurement volume rather than the physical time delay of Proposition 1.

V. Approaches to Quickest Change Detection in On-Off Models

While measurement policies are able to use the entire history of measurements and sampling times, past occupancies of the spectrum should not affect future occupancy, conditioned on current occupancy. As such, the Markov nature of the problem allows us to only study two aspects of the state. One of them is the last observed spectrum state X_n and the other is the Cumulative Sum (CUSUM) statistic of the change detection problem, defined at measurement point n as

$$S_n = \max_{1 \leq i \leq n-1} \sum_{k=i}^{n-1} \log\left(\frac{p^0_{X_k X_{k+1}}(\tau_k)}{p^\infty_{X_k X_{k+1}}(\tau_k)}\right). \tag{7}$$

The CUSUM statistic is the sum of log-likelihood ratios with a moving starting index, and is very common in change detection literature [16], [17]. While an optimal agent would likely use both of these aspects, it is of particular interest to study agents which use only one aspect each, which we do in the sequel. We'll initially analyze policies using only X_n and show lower and upper bounds on their performance, leading to optimization over such policies. Then, we'll examine policies using only S_n, finding the necessary partial results for empirical evaluation.

A. Static measurement policies

A lesson learned from i.i.d. adaptive measurement change point detection is that finding the optimal change point detection policy is far from easy, as it requires understanding the average measurement period value $\mathbb{E}_\infty[\tau]$ which in general depends on the behavior of a continuous state Markov chain (that is, the evolution of the CUSUM statistic S_n under $\mathbb{P}_\infty$). One way to approach this is to optimize over the set of *static* policies $\boldsymbol{\tau}$, that is, the measurement time τ_n is simply defined as

$$\tau_n(S_n, X_n) = \sum_{i=0}^{1} \tau_i \mathbb{1}(X_n = i)$$

If $\tau_i = \tau$ everywhere, this reduces to a periodic measurement policy. Of course, τ_n is still a random variable, but does not depend on S_n (and is in that sense static). The static policy is illustrated in Figure 1. The following proposition is useful for understanding how to go about minimizing the delay for such policies, at least asymptotically. It serves as an improved detection delay lower bound when compared to Proposition 1.

Proposition 3. *For any* γ*-compliant algorithm with static measurement policy* $\boldsymbol{\tau}$, *the worst case average detection delay is asymptotically lower bounded as*

$$\overline{\mathbb{E}}[t_N] \geq (I(\boldsymbol{\tau})^{-1}\overline{\tau}^0(\boldsymbol{\tau}) + o(1))\log(\gamma) \tag{8}$$

as $\gamma \to \infty$.

Proposition 3 tells us that minimizing the delay is asymptotically the same as maximizing the per-time information gain $I(\boldsymbol{\tau})/\overline{\tau}^0(\boldsymbol{\tau})$. As such, we should choose measurement times $\boldsymbol{\tau}$ in order to maximize this gain while maintaining a sufficiently small pre-change measurement frequency, and after this choose

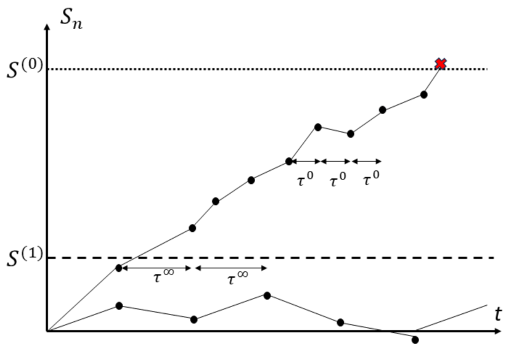

Fig. 2. Evolution of the CUSUM statistics pre-change and post-change, when measurements are taken according to a CUSUM-threshold policy.

a stopping rule to achieve the desired false alarm rate. We thus set up the optimization problem

$$\max_{\boldsymbol{\tau}\geq 0} \frac{I(\boldsymbol{\tau})}{\overline{\tau}^0(\boldsymbol{\tau})} \quad \text{s.t. } \overline{\tau}^\infty(\boldsymbol{\tau}) \geq \frac{1}{f} \tag{9}$$

where f is some desired measurement frequency. This optimization problem cannot be solved analytically and we instead resort to (somewhat involved) numerical methods. The next proposition shows that by doing so, it is indeed possible to construct an optimal static policy.

Proposition 4. *Let an agent sample according to static policy $\boldsymbol{\tau}$ and let it stop according to $N = \min\{n : S_n > C(\gamma)\}$. Then, this policy fulfills $\mathbb{E}_\infty[t_N] \geq \gamma$ and, as $\gamma \to \infty$,*

$$\overline{\mathbb{E}}[t_N] \leq \left(I(\boldsymbol{\tau})^{-1}\overline{\tau}^0(\boldsymbol{\tau}) + o(1)\right)\log(\gamma) \tag{10}$$

B. CUSUM-threshold measurement policies

Next, we study the opposite type of measurement policies. Inspired by the similar problem with i.i.d. random variables, we explore policies on the form

$$\tau_n(S_n, X_n) = \begin{cases} \tau^\infty & \text{if } S_n < S^{(1)} \\ \tau^0 & \text{otherwise.} \end{cases} \tag{11}$$

We call τ^∞ the calm period and τ^0 the crisis mode period of the measurement policy. Thus, the policy has no direct dependence on X_n (only through its update on S_n) and is therefore less exploitative of the particular problem structure. However, it instead exploits the nature of the change point detection problem by attempting to take measurements only when they are needed. The form of this policy, illustrated by the evolution of the CUSUM statistic S_n, is shown in Figure 2.

In order to compare the delay of CUSUM-threshold policies to other benchmarks, we must first obtain an estimate of the pre-change measurement frequency. Of course, this is a weighted average between $1/\tau^0$ and $1/\tau^\infty$, and we must thus find the weights between the two. To investigate this, we will use semi-fluid analysis. Imagine a related measurement agent which measures with a short, fixed measurement period τ_s. For this fixed measurement agent, we can use existing analysis of Lai [17] combined with Lindståhl et al. [3] to determine that when $\nu = \infty$, it will spend a maximum of a fraction of $\exp(-S^{(1)})$ of time in the regime $S_n > S^{(1)}$. Thus, we can conclude that over a long time T, at most $\exp(-S^{(1)})T$ will be spent in the regime $S_n > S^{(1)}$. For a CUSUM-threshold agent, the number of measurements in this regime can be at most $\exp(-S^{(1)})T/\tau^0 + o(T)$, and the total number of measurements taken can then at most be $\exp(-S^{(1)})T/\tau^0 + (1 - \exp(-S^{(1)}))T/\tau^\infty + o(1)$. Since this is true for any large T, the stationary measurement frequency is then upper bounded by

$$\mathbb{E}_\infty[f] \leq 1/\tau^\infty + \exp(-S^{(1)})(1/\tau^0 - 1/\tau^\infty) \tag{12}$$

Notably, we see that this expression blows up whenever $\tau^0 \to 0$, yet another explanation as to why we cannot measure infinitely often. That said, if τ^0 is moderate and $S^{(1)}$ is large, this expression approaches $1/\tau^\infty$. Furthermore, we are now ready to bound

$$\mathbb{E}_\infty[t_N] \approx \mathbb{E}_\infty[N]/\mathbb{E}_\infty[f] \geq$$
$$\exp(S^{(0)})\left(1/\tau^\infty + \exp(-S^{(1)})(1/\tau^0 - 1/\tau^\infty)\right)^{-1}.$$

Due to the varying information per measurement depending on past measurement outcomes, it is very difficult to accurately bound the detection delay. However, if $\tau^0 < \tau^\infty$, using classical change detection results, we can trivially bound $S^{(0)}d(\tau^0) \leq \overline{\mathbb{E}}[t_N] \leq S^{(0)}d(\tau^\infty)$. To get a better estimate of $\overline{\mathbb{E}}[t_N]$, we evaluate the detection delay empirically.

VI. Numerical Evaluation

A. Synthetic data

To evaluate the performance of our algorithms compared to a benchmark of fixed-interval measurements on their intended model, we performed simulations on synthetic data. Here, the data was generated by an on-off model, the parameters and transition probabilities of which were known to the agent. We evaluated two different scenarios:

- **Rate change:** This scenario corresponds to when the stationary probability remains the same but the rate of the system changes. Here we expected that measurement frequencies will have a great impact on detection performance. In this scenario, we set $\lambda_1 = 5\lambda_0$ both before and after the change, with $\lambda_0^\infty = 0.2$ and $\lambda_0^0 = 0.6$. This scenario

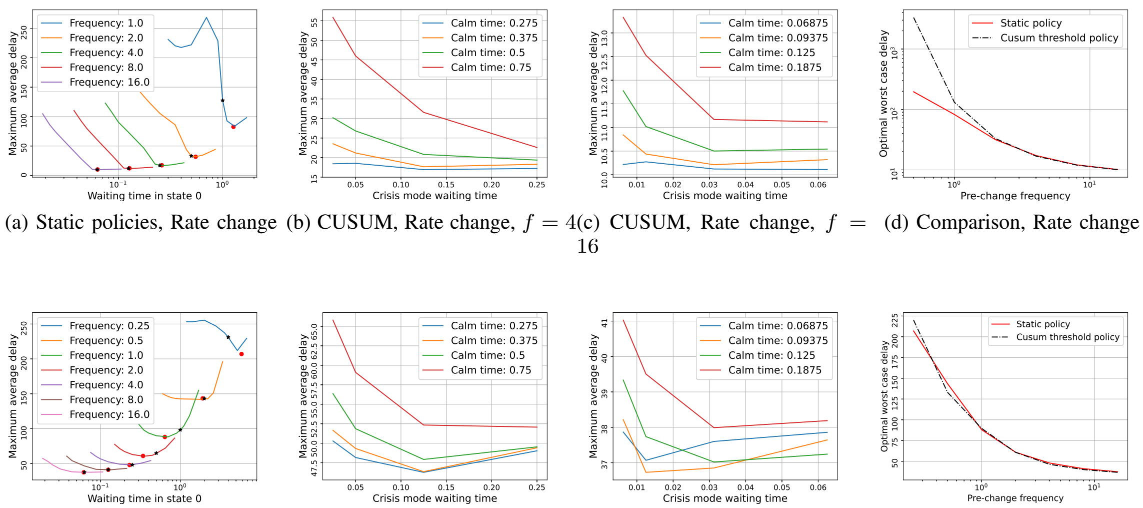

(a) Static policies, Rate change (b) CUSUM, Rate change, $f = 4$ (c) CUSUM, Rate change, $f = 16$ (d) Comparison, Rate change

(e) Static policies, Add user (f) CUSUM, Add user, $f = 4$ (g) CUSUM, Add user, $f = 16$ (h) Comparison, Add user

Fig. 3. Comparison of CUSUM-threshold policies for two pre-change frequencies, for the mirror change scenarios. Different values of τ_∞ are represented by different curves, while different values of τ_0 are represented by the x-axis.

can be interesting in order to model behavior changes from high-throughput services to low-latency services.

- **Add user:** Also, we looked at a scenario where the number of users in the system doubles in number. Thus we set $\lambda_0^\infty = \lambda_1^\infty = \lambda_1^0 = 1$, $\lambda_0^0 = 2$. Modeling traffic increase is generally useful to understand detection performance.

We required $\gamma = 200$ everywhere, and varied the measurement frequencies to see the impact on detection delay and volume of post-change measurements. We evaluated delay only at $\nu = 0$ since it is known that this causes maximal delay whenever the stopping rule is a CUSUM-threshold rule. We compared static sampling policies and CUSUM-threshold sampling policies to a benchmark with fixed measurement periods, and we varied the frequencies logarithmically from $1/2$ to 16 (doubling the frequency at each step). The stopping rule is $N = \min\{n : S_n > C(\gamma)\}$ for all algorithms.

For static policies, we evaluated the delay for different values of τ_0 between $0.3/f$ and $1.7/f$, where f is the frequency. The value for τ_1 is calculated from equation (9). Then, we also solved the optimization problem (9) for (τ_0, τ_1) and compare this with the different delays above, to see if the optimized values are indeed optimal in practice. Furthermore, by setting $\tau_0 = \tau_1 = 1/f$ we obtained a fixed period sampling policy to be used as a benchmark.

For CUSUM-threshold sampling policies, we varied both the calm period τ^∞ between $1/f$ and $2/f$ as well as the crisis mode period τ^0 between $0.1/f$ and $1/f$. We determined $S^{(1)}$ through equation (12). Then, we took the minimal delay of all these points and compared them to the optimal static policy delay as found above.

We present the result of our evaluation in Figure 3. The rows correspond to the two different scenarios. The first column shows the detection of static policies for different frequencies and different values of τ_0. Here, the policies with $\tau_0 = \tau_1$ are marked in black and the policies obtained by solving the optimization problem (9) are marked in red. The second and third column shows the detection delay of CUSUM-threshold polices for two frequencies and different parameters (τ^0, τ^∞). The fourth column shows a comparison between the best performing static and CUSUM policies for different frequencies. Confidence intervals are on an order of ± 0.1 time units and are omitted to increase visibility.

In general, the optimizer seems to work well in identifying the best static policy. That said, at high frequencies (when the delay approaches the characteristic delay) the delay curve around the optimal policy is very flat, so in these scenarios identifying the optimal policy becomes less important. For CUSUM-threshold policies, it seems like the configuration of τ^0 and τ^∞ instead become more important at higher frequencies, as the delay curves become flatter at lower frequencies. In the mirror flip scenario, it seems to be optimal to keep τ^∞ close to $1/f$ while this ceases to be true for the other scenarios at high frequencies - but remains true at low frequencies.

Overall, static policies seem to outperform CUSUM-threshold policies at low frequencies while CUSUM-threshold policies become competitive at higher frequencies, but the difference between the two at higher

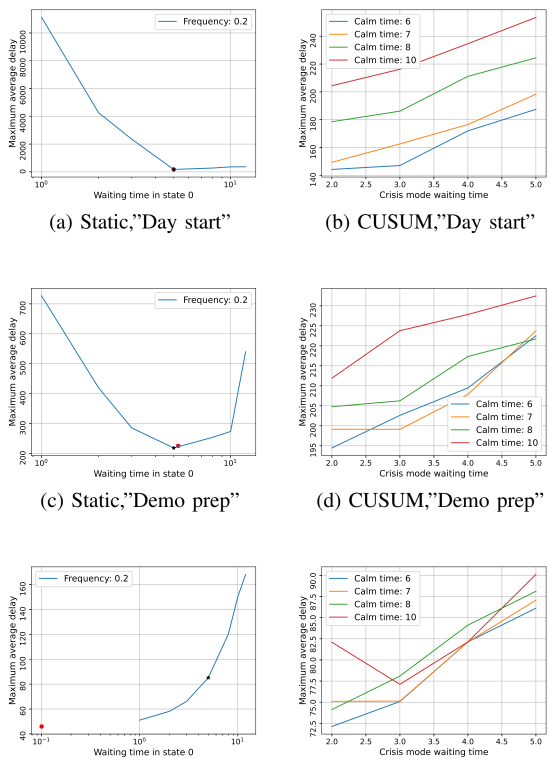

(a) Static,"Day start" (b) CUSUM,"Day start"

(c) Static,"Demo prep" (d) CUSUM,"Demo prep"

(e) Static,"Day end" (f) CUSUM,"Day end"

Fig. 4. Worst case average delays for different static policies and CUSUM-threshold policies on the three different change scenarios. Red dots mark optimal parameters by equation (9) and black stars mark periodic measurement schedules.

frequencies is too small to be significant. In particular, for the rate change scenario both policies seem to perform about the same everywhere, suggesting that fixed period policies may be optimal for this scenario. Indeed, Figure 3a suggests that this may be the case for all but very low frequencies. In particular, we note that the delay increases rapidly around $f = 0.5$ and below, and that it becomes next to impossible to detect changes around $f = 0.3$ without a highly optimized static policy. However, it is notable that for static policies, the shape of the boundary becomes flatter as the frequency increases. We attribute this to the delay getting closer to the characteristic delay, in whose regime the exact values of (τ_0, τ_1) become less important (notice the cancellation of time units in equation (3)). For lower frequencies, the asymmetries in the model matter more, which also explain why the shape of the low-frequency curves becomes steeper.

B. Spectrum data

Using synthetic data is helpful as it allows us to understand how the heuristic policies work on their intended model. It does not, however, allow us to understand how these policies and our modeling framework apply on a real spectrum access scenario. To do so, we studied data generated by a Wi-Fi access point programmed to sense the usage of a Wi-Fi channel. This access point was mounted in a test factory environment and as such, we expected that the usage of the channel would change throughout the day. In particular, we studied the usage of a channel on a work day where (1) at first, very few users were in the factory and as such, the spectrum was mostly idle (about 5% utilization), (2) as workers entered the factory, the utilization of the spectrum increases dramatically (to about 25%), (3) preparation for an upcoming demo began and the utilization increased to about (36%) and (4) the day ended and the utilization drops to $< 0.1\%$. We denote the change (1)-(2) as **Day start**, the change (2)-(3) as **Demo prep** and the change (3)-(4) as **Day end**, and we evaluated our algorithms on all of these scenarios.

Since the access point measured the channel usage with a frequency of 1 ms^{-1}, we required that the algorithms always sample with period $\tau_n \geq 1$ ms. This is somewhat limiting in terms of potential sampling frequencies, as most blocks of transmissions lasted shorter than 10 ms. As such, we only evaluated the pre-change measurement frequency $f = 0.2$ ms^{-1}. We required $\gamma = 2000$ ms everywhere.

For static policies, we evaluated $\tau_0 \in [1, 2, 3, 5, 8, 10, 12]$ ms and decided on τ_1 by using equation (9). For CUSUM-threshold policies, we evaluated the calm period in the range $[6, 7, 8, 10]$ ms and the crisis mode period in the range $[2, 3, 4, 5]$ ms - the threshold $S^{(1)}$ between them was set through equation (12). The results for both policies evaluated on all three scenarios is found in Figure 4. The confidence intervals are on the order ± 1 ms and have again been omitted for visibility.

Interestingly, both "Day start" and "Demo prep" exhibit similar situations despite being quite qualitatively different. For static policies, in this case, the fixed period policy seems to be optimal. However, it is instead possible to outperform the fixed policy with a CUSUM-threshold policy with about 20% of the detection delay without taking more measurements pre-change. Instead, on the "Day end" scenario the optimized static policy is quite far away from the fixed period policy and also outperforms CUSUM-threshold policy. While this policy does violate the bound $\tau_n \geq 1$ ms, the same properties are true if one constrains the policy to staying within these bounds. These results imply that an increase in load could be better handled by a CUSUM-threshold policy, while a decrease in load are better handled by a static policy, but this must be investigated further to be conclusively established.

VII. Conclusion

For many spectrum access problems, it is crucial to detect when a change has occurred in order to take corrective action. This is complicated by the typically Markovian time dependency on the spectrum occupancy. In this paper, we have presented a statistical framework specifically designed for this purpose, by adapting the classical framework on change point detection to on-off models while also requiring measurement efficiency.

Within this framework, we have presented novel fundamental limits for any change point detection algorithm with sufficiently low false alarm rate, quantified with a constraint on the average run length to false alarm, and presented two approaches designed to meet these limits. We have analyzed each approach and compared them to each other as well as non-measurement-efficient approaches in evaluation, both on synthetic data as well as on Wi-Fi spectrum sensing data. Our findings demonstrate that the CUSUM-threshold policy is often more efficient compared to static measurement policies, but may fail when the scenario is heavily asymmetric.

For future work, we would like to study cases where the pre- and post-change rates are not known a priori, necessitating robust methods. We would also like to create algorithms which can monitor multiple channels simultaneously.

Acknowledgements

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

References

[1] F. Hu, B. Chen, and K. Zhu, "Full spectrum sharing in cognitive radio networks toward 5g: A survey," *IEEE Access*, vol. 6, pp. 15754–15776, 2018.

[2] T. Banerjee and V. V. Veeravalli, "Data-efficient quickest change detection in minimax settings," *IEEE Transactions on Information Theory*, vol. 59, no. 10, pp. 6917–6931, 2013.

[3] S. Lindståhl, A. Proutiere, and A. Johnsson, "Change point detection with adaptive measurement schedules for network performance verification," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 7, no. 3, pp. 1–30, 2023.

[4] S. Lindståhl, A. Proutiere, and A. Johnsson, "Rome-qcd: Robust and measurement efficient quickest change detection in 5g networks," in *2024 8th Network Traffic Measurement and Analysis Conference (TMA)*, pp. 1–11, IEEE, 2024.

[5] Y.-C. Liang, *Opportunistic Spectrum Access*, pp. 19–40. Singapore: Springer Singapore, 2020.

[6] C. E. C. Bastidas, J. A. Stine, A. Rennier, M. Sherman, A. Lackpour, M. M. Kokar, and R. Schrage, "Ieee 1900.5. 2: Standard method for modeling spectrum consumption: Introduction and use cases," *IEEE Communications Standards Magazine*, vol. 2, no. 4, pp. 49–55, 2019.

[7] H. Mosavat-Jahromi, Y. Li, L. Cai, and J. Pan, "Prediction and modeling of spectrum occupancy for dynamic spectrum access systems," *IEEE Transactions on Cognitive Communications and Networking*, vol. 7, no. 3, pp. 715–728, 2021.

[8] L. Wang, J. Hu, D. Jiang, C. Zhang, R. Jiang, and Z. Chen, "Deep learning models for spectrum prediction: A review," *IEEE Sensors Journal*, 2024.

[9] M. Wellens, J. Riihijarvi, and P. Mahonen, "Modelling primary system activity in dynamic spectrum access networks by aggregated on/off-processes," in *2009 6th IEEE Annual Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks Workshops*, pp. 1–6, IEEE, 2009.

[10] A. Agarwal, S. Dubey, M. A. Khan, R. Gangopadhyay, and S. Debnath, "Learning based primary user activity prediction in cognitive radio networks for efficient dynamic spectrum access," in *2016 international conference on Signal Processing and Communications (SPCOM)*, pp. 1–5, IEEE, 2016.

[11] H. Eltom, S. Kandeepan, R. J. Evans, Y. C. Liang, and B. Ristic, "Statistical spectrum occupancy prediction for dynamic spectrum access: a classification," *EURASIP Journal on Wireless Communications and Networking*, vol. 2018, pp. 1–17, 2018.

[12] H. Song, L. Liu, J. Ashdown, and Y. Yi, "A deep reinforcement learning framework for spectrum management in dynamic spectrum access," *IEEE Internet of Things Journal*, vol. 8, no. 14, pp. 11208–11218, 2021.

[13] Y. Lin, C. Wang, J. Wang, and Z. Dou, "A novel dynamic spectrum access framework based on reinforcement learning for cognitive radio sensor networks," *Sensors*, vol. 16, no. 10, p. 1675, 2016.

[14] O. Naparstek and K. Cohen, "Deep multi-user reinforcement learning for distributed dynamic spectrum access," *IEEE transactions on wireless communications*, vol. 18, no. 1, pp. 310–323, 2018.

[15] S. Jang, J. Kim, J. Byun, and Y. Shin, "Game theory based dynamic spectrum allocation for secondary users in the cell edge of cognitive radio networks," *KSII Transactions on Internet and Information Systems (TIIS)*, vol. 8, no. 7, pp. 2231–2245, 2014.

[16] G. Lorden, "Procedures for reacting to a change in distribution," *The annals of mathematical statistics*, vol. 42, no. 6, pp. 1897–1908, 1971.

[17] T. L. Lai, "Information bounds and quick detection of parameter changes in stochastic systems," *IEEE Transactions on Information theory*, vol. 44, no. 7, pp. 2917–2929, 1998.

[18] A. Badawy, A. El Shafie, and T. Khattab, "On the performance of quickest detection spectrum sensing: The case of cumulative sum," *IEEE Communications Letters*, vol. 24, no. 4, pp. 739–743, 2020.

[19] V. Shakhov and I. Koo, "Applying change point detection technique to dynamically support network security," in *2022 RIVF International Conference on Computing and Communication Technologies (RIVF)*, pp. 503–507, IEEE, 2022.

[20] S. Haug, A. Egbert, R. J. Marks, C. Baylis, and A. Martone, "Changepoint detection for real-time spectrum sharing radar," in *2023 IEEE Texas Symposium on Wireless and Microwave Circuits and Systems (WMCS)*, pp. 1–10, IEEE, 2023.

[21] M. S. Khan, M. Usman, and I. Koo, "Dynamic spectrum tracking through quickest detection techniques: a clustered approach," *Journal of Internet Technology*, vol. 19, no. 5, pp. 1363–1370, 2018.

[22] R. Zou and W. Wang, "Change detection based segmentation and modeling of lte spectrum tenancy," in *2019 IEEE Global Communications Conference (GLOBECOM)*, pp. 1–6, IEEE, 2019.

Appendix A
Proofs

Proof of Proposition 1. We consider an agent with perfect information, getting immediately notified whenever a change occurs. Any agent can only perform better than this agent by a small additive constant, vanishing in $\log(\gamma)$. This notified agent can consider the transition of states as a valued finite-state Markov chain, where the value of a transition is the time spent before the transition and the probability of transitioning from state 0 to state 1 (or reverse) is 1. Then the stationary probabilities of this Markov chain are $1/2$ for both states, both before and after the change. As such, the value of this Markov chain is simply the arithmetic average of the expected waiting time of each state, that is $\frac{(\lambda_0^\infty)^{-1}+(\lambda_1^\infty)^{-1}}{2}$. Thus, as $\gamma \to \infty$, in order to maintain $\mathbb{E}_\infty[t_n] \geq \gamma$ it is required that $\mathbb{E}_\infty[N] \geq \frac{2\gamma}{(\lambda_0^\infty)^{-1}+(\lambda_1^\infty)^{-1}}$. With some algebra, it can be seen that this expression is equivalent to $\exp(C(\gamma))$ To do so, Theorem 3 in Lai [17] states that it is required that $\overline{\mathbb{E}}[N] \geq (I^{-1}+o(1))C(\gamma)$, where I is the information per measurement of the finite state Markov chain. Since these are exponentially distributed random variables, it follows that

$$I = \frac{1}{2}\sum_{i=0}^{1}(\lambda_i^0 \log(\frac{\lambda_i^0}{\lambda_i^\infty}) + \lambda_i^\infty - \lambda_i^0)$$

Furthermore, the expected time per measurement post-change is $\overline{\tau}^0 = \frac{(\lambda_0^0)^{-1}+(\lambda_1^0)^{-1}}{2}$. Thus, since $\overline{\mathbb{E}}[t_N] \geq \overline{\mathbb{E}}[N]\overline{\tau}^0 + o(\log\gamma)$, we conclude that

$$\overline{\mathbb{E}}[t_N] \geq (I^{-1}\overline{\tau}^0 + o(1))C(\gamma) \geq (d(0) + o(1))c(\gamma)$$

which concludes the proof.

□

Proof of Proposition 2. It is not difficult to see that $I_{max}(\beta)$ maximizes per-measurement information of all static policies with sufficiently low delay. As such, among static policies, it is impossible to do better than the lower bound in terms of number of measurements, asymptotically as $\gamma \to \infty$. It thus remains to show that there is no non-static policy that asymptotically outperforms the optimal static policy.

Now, because the time per measurement and information per measurement are both functions only of the waiting time and the last measurement, any (stationary) non-static policy will act as a weighted average of static policies, with the weights being the probability of a certain tuple (τ_0, τ_1) being used. However, any such static policy cannot perform better than $I_{max}(\beta_i)$ for some value of β_i, and for the condition (2) to be met, it is necessary that the weighted average of $\beta_i \leq \beta$. However, because I_{max} is a concave function in β it is impossible to outperform $I_{max}(\beta)$ with a weighted average of $I_{max}(\beta_i)$, and it is thereby impossible to outperform the best static policy. This concludes the proof.

□

Proof of Proposition 3. By Lai (1995) [17], finite state Markov chains with known pre- and post-change transition matrices have lower bounded number of post-change measurements $\overline{\mathbb{E}}[N] \geq (I(\boldsymbol{\tau})^{-1} + o(1))\log(\gamma)$. We note in this that the right hand expression is unbounded as $\gamma \to \infty$. As such, when $\gamma \to \infty$, a vast majority of post-change measurements will be taken when the system is in a stationary regime, regardless of pre-change measurements. Thus, we conclude ess sup$\mathbb{E}_\nu[T-\nu|\mathcal{F}_\nu] \geq$ ess sup$\mathbb{E}_\nu[N-n_\nu|\mathcal{F}_\nu](\sum_i \pi_i^0(\boldsymbol{\tau})) + o(\log(\gamma))$, since the average time per measurement in the stationary regime is $\sum_i \pi_i^0(\boldsymbol{\tau})$. This concludes the proof.

□

Proof of Proposition 4. Since $\boldsymbol{\tau}$ is fixed, the time series X_n defines a finite-state, discrete time Markov chain. As such, it is well-known by Theorem 4 in Lai [17] that the stopping time N satisfies

$$\overline{\mathbb{E}}[N] \leq (I(\boldsymbol{\tau})^{-1} + o(1))\log(\gamma) \tag{13}$$

as $\gamma \to \infty$. Furthermore, by the law of large numbers, it holds that $\overline{\mathbb{E}}[t_N] \leq \overline{\mathbb{E}}[N]\overline{\tau}^0(\boldsymbol{\tau}) + o(\log\gamma)$. Combining this with equation (13) yields the upper bound in the Proposition.

□