RoME-QCD: Robust and Measurement Efficient Quickest Change Detection in 5G Networks

Simon Lindståhl*†, Alexandre Proutiere*§, Andreas Johnsson†‡
*KTH Royal Institute of Technology
†Ericsson Research
‡Uppsala University
§Digital Futures

***Abstract*—To effectively monitor a network and verify its performance, it is essential to quickly detect sudden changes in its state, even when the form of such a change is initially unknown. While classical quickest change detection methods are potentially useful, they rely on probing the network state periodically, which in turn, may induce high measurement costs. In this paper, we extend existing frameworks in quickest change detection to allow for both adaptive measurement periods and unknown post-change measurement distribution. In our extended framework, the agent decides both when to raise an alarm and when to take the next measurement (if any), maintaining a trade-off between detection delay, false alarm rate and measurement costs. We evaluate both classical methods with periodic measurements as well as our adaptive scheme, called RoME-QCD (Robust and Measurement Efficient Quickest Change Detection). We demonstrate the latter's superiority analytically and verify this observation via numerical experiments, both using one-way delay data from a 5G testbed and synthetically generated data.**

I. Introduction

With the rapid global deployment of 5G networks, a great number of emerging applications have been enabled. Some of these applications include real-time communication, industrial control and autonomous vehicles. To ensure their continuous successful operation, it is critical that certain metrics such as round trip time (RTT) or one-way delay (OWD) stay within service level agreements. However, these metrics, for any given user equipment (UE) are dependent on the wireless channel, exogeneous processes within the UE, base station configuration and UE hardware and configuration. Several of these are subject to change at any point, and there is a non-trivial relationship between this network state and the performance metrics. It is therefore vital to probe the performance metrics and respond to changes which cause them to deviate from pre-set parameters.

In this paper, we study a network where a probe attempts to detect performance changes, having a negative impact on the service operations, as soon as possible. The probe has plenty of data of the normal state in day-to-day operations, but must be able to respond to a large variety of changes. The probe's job is further complicated by potentially expensive network measurements, since these measurements are taken by (for instance) injecting probe packets into the network using TWAMP [1], costing bandwidth and other resources. As such, the probe must adaptively measure less often when there is no reason to suspect a change has occurred, while still being ready to respond whenever a change does occur [2], [3]. The scenario under consideration is modelled in Figure 1.

978-3-903176-64-5

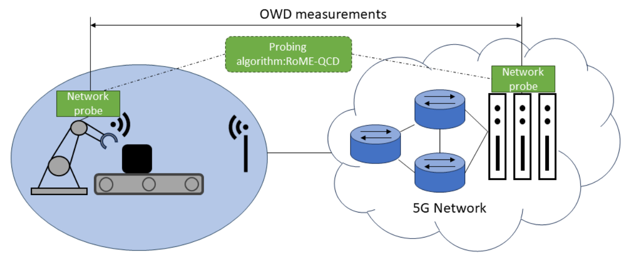

Fig. 1. An illustrative industrial 5G network, where a measurement agents attempts to detect One-Way Delay (OWD) changes through active probing.

The problem of determining a change in the network performance appears a good match for the well-known problem of *quickest change detection* (QCD) [4]. In this framework, the probe monitors a sequence of random variables, which could for instance be the uplink or downlink OWD of a UE in the network. Whenever these random variables change in distribution from one well-known distribution to another, the probe raises an alarm. However, any such network probe will face two major challenges. First, while the distribution before the change may be well-known, its change could be caused by any one of many possible changes in the network. These could be a node failure, batch traffic arrival or a hand-over. As such, it is difficult to a priori predict what the distribution will look like after the change has occurred. Then, network measurements can be costly especially when taken in bulk, as they consume power, bandwidth and CPU time. Hence, the network probe should both be *robust* to a large set of potential changes, as well as use *adaptive* measurement schedules to minimize network overhead of probing. While both of these aspects have been explored separately within the framework of QCD [5], [6], it remains an open question how to combine the two and what impact this will have on network management.

This paper addresses this challenge and presents the following contributions. (1) We propose RoME-QCD (Robust and Measurement Efficient Quickest Change Detection), an adaptive probing scheme that monitors the network performance efficiently, even in the presence of unknown changes. (2) We analyze this scheme and prove that it achieves an efficient trade-off between measurement cost, false alarm rate and detection delay. (3) We evaluate RoME-QCD on data from a 5G testbed and show its utility for realistic network scenarios.

II. Related work

Adaptive network monitoring is far from a new topic in literature. The challenges with measurement overhead in software defined networking (SDN) were tackled by Wang and Su [7] and Adrichem et al [8] by proposing mechanisms and methods for adaptive polling and sampling. Further, Adaptive monitoring for analysis and orchestration in 5G systems were discussed and targeted by Xie et al. [2]. Their framework was used in the context of service assurance, adapting monitoring to analytical results in order to balance monitoring cost to network performance. Further, an adaptive probing system with the purpose of detecting outages in edge networks was designed by [3]. This system probes the network regularly in order to guarantee freshness, increasing probing frequency to resolve uncertainty whenever necessary. Finally, Steinert and Gillblad proposed a heuristic method for adapting measurement frequency of RTT measurements over a network path [9]. Their methods rely on computing the information distance between empirically estimated RTT distributions.

Another method to detect anomalies in a network is to study and detect anomalous flows. Phan et al. [10] do this by studying flows at different granularities, and Tan et al. [11] follow up on this in their lightweight application FLOWSPOTTER.

Initially, the framework of detecting a change without prior knowledge on the change time was formulated by Lorden [4]. He showed the value of a scheme considering only the maximal cumulative sum of log-likelihood ratios, called the CUSUM statistic, by showing an asymptotic lower bound on the number of measurements post-change and that an agent tracking only the CUSUM statistic could meet such a bound asymptotically. Later, Moustakides showed that an algorithm tracking only the CUSUM statistic is exactly optimal [12].

Lorden simulateously extended the framework to the case where the post-change parameter belonged to a one-parameter exponential family, and showed how one could asymptotically optimize the algorithm even when the post-change parameter was initially unknown [4]. Lai [13] extended this case to pre- and post-change distributions which had some history dependence. It was, however Unnikrishnan and Veeravalli who formulated the framework for robust change point detection that we use today [5]. They considered the pre- and post-change distributions to lie within some uncertainty sets, and demanded that the algorithm used should always minimize the worst-case post-change delay, taken over the entire set.

Later, Liang and Veeravalli relaxed the stochastic ordering assumption into the notion of weakly stochastically bounded sets, which, informally, only requires that some pair of distributions is closer to each other than any other pair within the uncertainty sets [14]. They also showed that a general set of post-change distributions, only defined by bounds on mean and variance, was weakly stochastically bounded with respect to a singleton set of any pre-change distribution.

Banerjee and Veeravalli in 2013 extended this scenario by adding a component of data-efficiency [15]. They allowed the agent to "skip" measurements and as such, control the measurement frequency in a slotted time formulation. Their method, which they proved was asymptotically optimal when the constraint on false alarm rate grows stricter, utilizes the undershoot of the CUSUM statistic, normally set to 0, and set the number of skipped slots as a linear function of this undershoot.

Lindståhl et al. [6] introduced a new framework, where the agent is allowed to set the next measurement time at any point after the current time. They set the physical time detection delay and false alarm rate as constraint, and attempted to optimize the measurement costs pre- and post-change subject to these constraint. They showed that the post-change costs is invariant to the measurement schedule given the stopping rule, and that a simple adaptive scheme could qualitatively outperform classical periodic measurements. They did not, however, show an optimal measurement scheme in terms of pre-change measurement frequency.

Both of the above works assume knowledge of the pre- and post-change measurement distribution. Our work is novel in that we remove this assumption while taking measurement frequency into account, and show the application of such schemes in network performance verification.

III. Monitoring changes through Quickest change detection

In this section, we describe how measurements are taken, and how the framework of QCD is useful for detecting performance changes in the network.

A. Quickest Change Detection (QCD) framework

A measurement probe (synonymously called a measurement agent) observes a sequence $\{X_i\}_{i=1}^N$ of OWD measurements, which are modelled as random variables. These measurements are independent of each other and, before some time $\nu \in \mathbb{R}$, they are all identically distributed according to some known distribution[1] p_∞. At time ν, some event occurs in the network, causing the measurements to change and instead be distributed according to some unknown distribution p_0. The network probe is then tasked with raising an alarm as soon as possible.

While the agent cannot know p_0, it may know some (potentially large) *uncertainty set* $\mathcal{P}_0$ such that $p_0 \in \mathcal{P}_0$. If the agent knows $\mathcal{P}_0$, that is because it may know some general properties of the change, or just that the change will cause some increase or decrease a certain performance indicator. In the sequel, we assume that $\mathcal{P}_0$ is stochastically bounded with respect to p_∞ . That is, there is some distribution $\overline{p}_0$ which

[1]In state-of-the-art robust QCD, p_∞ is often unknown, but in this paper, we assume that it is known. This is partly because the key results would not change even if it was unknown, partly for simplicity, and partly because having detailed information about the pre-change distribution is far more realistic than having detailed information about the post-change distribution.

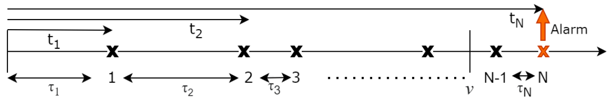

Fig. 2. Quickest detection with adaptive measurement schedules.

is closer to p_∞ than any other distribution in p_0[2]. In Section IV-D we show how to compute $\overline{p}_0$ for a variety of uncertainty sets.

Classical QCD algorithms treat network measurements as an optimal stopping problem, with the aim of minimizing the number of measurements after the change has occurred, as well as the rate of incurring false alarms. Implementing these algorithms in practice, one can only assume that measurements are taken periodically. However, this leaves us unable to model scenarios where the network probe may change its measurement frequency in real time. Indeed, the probe may aim to balance delay and false alarm rate to the measurement frequency pre-change. This frequency represents the operation cost of the network probe in scenarios where change events are relatively uncommon. In this paper, we adapt the framework introduced in [6] where the agent adaptively schedules measurements based on past observations. More precisely, the network probe makes the following decisions.

1) *Measurement strategy:* The first measurement is taken at time τ_1. For any $i \geq 1$, after the i-th measurement, the probe makes a decision on a measurement interval $\tau_{i+1} \geq 0$ based on the previous measurement intervals and outcomes $(\tau_k, X_k)_{k=1}^{i}$. The next potential measurement is taken at time $t_{i+1} := \sum_{k=1}^{i+1} \tau_k$.
2) *Alarm strategy:* For any $i \geq 1$, the probe may raise an alarm after the i-th measurement, based on the previous measurement intervals and outcomes $(\tau_k, X_k)_{k=1}^{i}$.

We denote by n_ν the number of measurements taken before the change time ν, and by N the number of measurements taken before the probe raises an alarm. We let $\mathbb{P}_{\nu,p_0}$ be the probability measure when the change time is ν and the post-change distribution is p_0, $\mathbb{P}_\infty$ be the probability measure when there is no change (so $\nu = \infty$) and $\mathbb{E}_{\nu,p_0}$ and $\mathbb{E}_\infty$ be the corresponding expectations. The detection problem with adaptive measurement schedules is visualized in Figure 2.

B. Performance metrics and trade-offs

In the QCD framework, especially when using adaptive measurement schedules, the network probe needs to balance several performance metrics. For instance, detecting a change faster can be done by taking measurements more frequently, increasing the operation cost, or by setting a lower bar for alarms, increasing false alarm rate. These metrics are formalized below.

False alarm rate: To keep the operation cost down, the network probe wants to avoid an abundance of false alarms. The false alarm rate is captured by the physical Average Run Length (ARL) to false alarm, which is defined as $\mathbb{E}_\infty[t_N]$. That is, we let the probe keep up operations on a network where no change occurs, and see how long it takes for it to raise an alarm.

Worst case detection delay: It is business critical that changes in the network are detected fast, no matter what the change is or when it occurs. The network probe thus needs to ensure a sufficiently small worst case average detection delay, here defined as $\overline{\mathbb{E}}[t_N] := \sup_{\nu \geq 0, p_0 \in \mathcal{P}_0} \mathbb{E}_{\nu,p_0}[t_N - \nu | t_N \geq \nu]$.

Measurement costs: In addition to these metrics, we are also interested in minimizing the number of measurements, i.e. the number of probe packets sent over the network. We distinguish two types of measurement costs, those induced before the change and those after the change.

Pre-change measurement frequency: Since it is unknown when a change will occur, one cannot measure the exact number of measurements pre-change and must instead study the measurement frequency. Letting $(\tau_k)_{k\geq 1}$ denote the lengths of the measurement intervals, we define $\mathbb{E}_\infty[\hat{\tau}] = \liminf_{n\to\infty} \mathbb{E}_\infty\left[\frac{1}{n}\sum_{k=1}^{n} \tau_k\right]$ as the stationary measurement period. $\mathbb{E}_\infty[\hat{\tau}]$ is a long term average of inter-measurement times.

Post-change measurement cost: We define the post-change measurement cost as the greatest possible average number of measurements taken after the change has occurred. Similarly to the detection delay, this is formally defined as

$$\overline{\mathbb{E}}[N] := \sup_{0 \leq \nu < \infty, p_0 \in \mathcal{P}_0} \left(\text{ess sup } \mathbb{E}_{\nu,p_0}\left[(N - n_\nu)^+ | t_N \geq \nu\right]\right).$$

It is well-known (see [4], [14], or [6]) that this quantity is asymptotically lower bounded as

$$\overline{\mathbb{E}}[N] \geq \left(\overline{I}^{-1} + o(1)\right) \log\left(\frac{\gamma}{\beta}\right) \tag{1}$$

when $\gamma \to \infty$. It remains to study whether this lower bound is tight, which we establish that it is in Section IV.

Trade-offs and objectives: The network probe is interested in maintaining a sufficiently small false alarm rate while keeping the delay $\overline{\mathbb{E}}[t_N]$ small (defined for the worst possible post-change distribution $p_0 \in \mathcal{P}_0$). It is always possible to increase the ARL while maintaining the same detection delay, simply by increasing the measurement frequency, but this increases measurement costs. In the interest of guaranteeing a certain level of performance in terms of detection delay and false alarm rate, we treat both of these objectives as constraints. While keeping those objectives within pre-defined limits, we aim to minimize the measurement cost, particularly before the change (i.e., maximizing $\mathbb{E}_\infty[\hat{\tau}]$). As such, we define a set of network probes, we call them $(\gamma, \beta, \mathcal{P}_0)$-compliant probes, which meet the relevant constraints, and attempt to optimize the measurement frequency over this set of agents.

Definition 1. *A network probe $(\{\tau_n\}_{n\geq 1}, N)$ is said to be $(\gamma, \beta, \mathcal{P}_0)$-compliant if it fulfills $\mathbb{E}_\infty[t_N] \geq \gamma$ and $\overline{\mathbb{E}}[t_N] \leq \beta$*

While our framework allows for any γ and β, note that the problem is trivial unless $\gamma > \beta$, as the probe could otherwise

[2]More formally, if $D(p||q)$ denotes the Kullback-Leibler divergence from distribution p to distribution q, we assume that there exists some distribution $\overline{p}_0$ such that, for any $p_0 \in \mathcal{P}_0$, $D(\overline{p}_0||p_\infty) \leq D(p_0||p_\infty) - D(p_0||\overline{p}_0)$.

simply wait for exactly γ time units and then stop immediately, measuring only once.

IV. Analysis and Algorithms

In this section, we present QCD schemes and analyze their performance. These schemes are based on the CUSUM statistic, commonly used in QCD. We first analyze the performance of probes with periodic measurement schedules. We then introduce RoME-QCD, a strategy with adaptive schedules, and establish their superiority against schemes with periodic schedules.

A. Cumulative sum statistics

While many schemes for network probes can be proposed, some statistics have shown to be more useful than others. In particular, the case with fixed measurement intervals suggests that it is optimal to study a probing scheme which depends only on the *cumulative sum* (CUSUM) statistic. We establish that, even with adaptive measurement schedules in the robust setting, these probes are efficient.

When p_0 is known, the CUSUM statistic, taken after n measurements, is defined as $S_n := \max_{1\leq i\leq n}\sum_{k=i}^{n}\log\left(\frac{p_0(X_k)}{p_\infty(X_k)}\right)$. Obviously the CUSUM statistic is a random variable, and is interpreted as the sum of log-likelihood ratios maximized over starting indices, ending at the current measurement index n. When p_0 is unknown, this statistic is no longer available to the probe. Instead, we use an alternatively statistic when $\mathcal{P}_0$ is stochastically bounded with respect to $\{p_\infty\}$, still called the CUSUM statistic, defined as $\overline{S}_n := \max_{1\leq i\leq n}\sum_{k=i}^{n}\log\left(\frac{\overline{p}_0(X_k)}{p_\infty(X_k)}\right)$. Here $\overline{p}_0$ is the least favorable distribution of $\mathcal{P}_0$[3]. Using the next result, we are able to introduce a useful class of probes, depending only on this statistic.

Lemma 1. *Let a probe depend only on CUSUM-statistics $\overline{S}_n$, that is $\forall n\geq 0, \tau_{n+1} = g(\overline{S}_n)$ for some non-increasing function g and $N = \min\{n\geq 1 : \overline{S}_n > S^{(0)}\}$. Then we have:*
(i) $\mathbb{E}_{\nu,p_0}[(N-n_\nu)^+|t_N > \nu]$ is independent of the post-change measurement intervals $\{\tau_n\}_{n=n_\nu+1}^{N}$.
(ii) For any $n\geq 1$, $\mathbb{E}_\infty[\tau_n|N > n-1] \geq \mathbb{E}_\infty[\tau_n] \geq \mathbb{E}_\infty[\hat{\tau}]$.

In the above lemma, (i) allows us to put less effort into minimizing the number of measurements post-change, as it depends only on the stopping rule and not on the sampling strategy. Further, (ii) justifies using $\mathbb{E}_\infty[\hat{\tau}]$ as a metric to capture the pre-change measurement frequency. Lemma 1 is proven in Appendix A.

B. Periodic measurement schedules

A natural measurement strategy consists in scheduling measurements periodically, i.e., for some fixed $\tau > 0$ and for all n, $\tau_n = \tau$. We investigate how τ and the stopping rule should be chosen in order to create $(\gamma,\beta,\mathcal{P}_0)$-compliant probes with periodic measurement schedules.

[3]This statistic is indeed known to the probe.

First, we note that, under some mild assumptions on $\mathcal{P}_0$, it is possible to create $(\gamma,\beta,\mathcal{P}_0)$-compliant probes by solving the system of equations

$$\tau = \frac{\beta\overline{I}}{S^{(0)}+\xi} \quad \text{and} \quad S^{(0)} = \log\left(\frac{\gamma}{\tau}\right). \tag{2}$$

It has previously been shown [6] that such periodic measurement schemes meet the requirements on ARL and delay when the true distribution is $\overline{p}_0$, and intuitively this distribution is harder to detect than any other post-change distribution $p_0\in\mathcal{P}_0$. Here, $\xi := \max_r \mathbb{E}_{0,\overline{p}_0}[Z-r|Z\geq r]$ is the maximal overshoot and is mainly a technical quantity. When $\gamma > \beta$, the above system of equations has two solutions, and we pick that with the greatest inter-measurement time. However, it is notable that since $S^{(0)}$ is an increasing function of γ, the measurement period τ is a decreasing function of γ. This is a rather undesirable property, as we would prefer that the average measurement frequency only depends on detection delay requirement β. As such, fixing the measurement schedules to be periodic comes with a rather steep cost in terms of measurement frequency. This is stated clearly in the following proposition, which is a straightforward consequence of results in [6]. We show that there exist classes of probes with adaptive measurement schedules which are $(\gamma,\beta,\mathcal{P}_0)$-compliant for any value of γ but which all have the same measurement frequency, $\mathbb{E}_\infty[\hat{\tau}]$. This means that indeed, periodic schedules come with a high measurement cost.

Proposition 1. *Fix the delay guarantee $\beta > 0$. For any period $\tau > 0$, there exists $\bar{\gamma}(\tau)$ such that for any false alarm rate guarantee $\gamma > \bar{\gamma}(\tau)$, there exists no periodic probe with measurement period τ that is $(\gamma,\beta,\mathcal{P}_0)$-compliant.*

C. The RoME-QCD algorithm

With periodic measurement schedules, decreasing the false alarm rate comes at the cost of greater measurement frequency before the change occurs. A natural question is whether it is possible to devise an intelligent measurement schedule under which this trade-off can be avoided. As was shown in [6], this can be done whenever the post-change distribution is known. We now show that this is also possible whenever the post-change distribution lies within the uncertainty set $\mathcal{P}_0$. Using the least favorable distribution $\overline{p}_0\in\mathcal{P}_0$, we can use the associated CUSUM statistic $\overline{S}_n$ to introduce the algorithm RoME-QCD. This algorithm, just like periodic measurement schedule probes, aims to stop whenever $\overline{S}_n$ exceeds some threshold $S^{(0)}$. Furthermore, it also aims to measure periodically most of the time. However, it also uses some crisis threshold $S^{(1)}$ such that whenever $S^{(1)} < \overline{S}_n \leq S^{(0)}$, it measures aggressively. The idea behind this algorithm is that this crisis region will rarely be entered when $t_n < \nu$, but will constitute the majority of the measurements while $t_n > \nu$. The pseudo-code of probes with RoME-QCD is presented in Algorithm 1.

We wish to verify that such probes have desired properties, and that the aggressive measurements in the crisis region has a minimal impact on the pre-change measurement frequency.

Algorithm 1 ROME-QCD

Input: Crisis threshold $S^{(1)}$, stopping threshold $S^{(0)}$, calm period τ
Initialization: $\overline{S}_0 = 0$
for $n = 1, \ldots$ **do**
  if $\overline{S}_{n-1} > S^{(0)}$ **then**
    Stop and raise alarm.
  else if $\overline{S}_{n-1} > S^{(1)}$ **then**
    $\tau_n := 0$
  else
    $\tau_n := \tau$
  end if
  Wait for time period τ_n.
  Obtain measurement outcome X_n.
  Update $\overline{S}_n := \log\left(\frac{\overline{p}_0(X_n)}{p_\infty(X_n)}\right) + \max(0, \overline{S}_{n-1})$
end for

As such, arbitrarily harsh false alarm rate constraints can be imposed without affecting measurement costs. To do so, we need the following assumption on the set $\mathcal{P}_0$, restricting its higher moments.

Assumption 1. *There exists a constant σ_Z such that for all $p_0 \in \mathcal{P}_0$, the random variable $Z := \log(\frac{\overline{p}_0(X)}{p_\infty(X)})$ is subgaussian under p_0 with parameter no greater than σ_Z. In other words, for any $z \in \mathbb{R}$ it holds that $\mathbb{P}_{X\sim\overline{p}_0}(|Z| \geq z) \leq 2\exp\left(-\frac{z^2}{\sigma_Z^2}\right)$.*

While Assumption 1 might not always hold, it is a milder assumption than it may seem. In practice, the metrics of interest are often bounded, e.g. OWD is truncated by some upper bound when a packet is discarded. Such metrics are then sub-gaussian by nature. In theory, it is usually enough to limit higher moments in other ways, often with just a bound on the variance, to guarantee a well-performing algorithm. It is, however, considerably more difficult to state explicitly the parameters of such an algorithm. We are now ready to state our main result.

Theorem 1. *Fix the delay requirement β, and assume Assumption 1. Select $(\tau, S^{(0)}, S^{(1)})$ satisfying:*

$$\tau = \beta\frac{(1-\exp(-\overline{I}^2/2\sigma_Z^2))}{\exp(S^{(1)}\overline{I}/\sigma_Z^2)} \tag{3}$$

$$S^{(0)} = \log\left(\frac{\gamma}{\tau(1-\exp(-S^{(1)}))}\right). \tag{4}$$

Then, for any $\gamma > 0$, ROME-QCD using parameters $(\tau, S^{(0)}, S^{(1)})$ is $(\gamma, \beta, \mathcal{P}_0)$-compliant, and its pre-change measurement cost satisfies $\mathbb{E}_\infty[\hat{\tau}] \geq \tau_c$, for some constant $0 < \tau_c < \infty$ that does not depend on γ. In addition, as $\gamma \to \infty$, it fulfills

$$\overline{\mathbb{E}}[N] \leq \left(\overline{I}^{-1} + o(1)\right)\log\left(\frac{\gamma}{\beta}\right).$$

Theorem 1, proven in Appendix A, contrasts with Proposition 1. It states that the same pre-change measurement frequency can be achieved by ROME-QCD for any value of $\gamma > 0$, and is constructive in that it specifies values of $(\tau, S^{(0)}, S^{(1)})$ that guarantee that the corresponding QCD network probe is $(\gamma, \beta, \mathcal{P}_0)$-compliant. We also establish that the resulting network probe has (asymptotically) minimal post-change measurement cost, which perhaps should not be surprising in light of Lemma 1.

D. Computing the least favorable distribution

While the above results are promising, the astute reader may have noted that their utility depends on being able to identify the least favorable distribution $\overline{p}_0$, and being able to evaluate the expression $\frac{\overline{p}_0(X)}{p_\infty(X)}$ for any X. For complicated models of p_∞ and $\mathcal{P}_0$, this may not be immediately obvious. However, relying on previous work [5], [14], we below show some commonly used uncertainty sets $\mathcal{P}_0$ where these quantities can indeed be computed.

One-parameter exponential families: First, we study the case of one-parameter exponential families. Letting h, T and Λ be real-valued functions, such distributions are defined by their probability density function $p(x;\theta) = h(x)\exp(\theta T(x) - \Lambda(\theta))$. Here x is a measurement outcome and θ is a parameter. We are interested in the case $p_\infty(x) = p(x;\theta_\infty)$ and where $\mathcal{P}_0 = \{p_0 : p_0(x) = p(x;\theta_0), \theta_0 \in [\theta_{\min}, \theta_{\max}]\}$. Whenever $\theta_\infty \notin [\theta_{\min}, \theta_{\max}]$, it is well known that the least favorable distribution is

$$\overline{p}_0(x) = \begin{cases} p(x;\theta_{\min}), & \theta_\infty < \theta_{\min} \\ p(x;\theta_{\max}), & \theta_\infty > \theta_{\max}. \end{cases} \tag{5}$$

In the sequel, we assume $\theta_\infty < \theta_{\min}$, the analysis is analogous for $\theta_\infty > \theta_{\max}$. We note that, under these families $Z_n = \log\left(\frac{\overline{p}_0(X_n)}{p_\infty(X_n)}\right) = (\theta_{\min} - \theta_\infty)T(X_n) - (\Lambda(\theta_{\min}) - \Lambda(\theta_\infty))$. As such, we can calculate $\overline{I} := \mathbb{E}_{0,\overline{p}_0}[Z] = (\theta_{\min} - \theta_\infty)\mathbb{E}_{0,\overline{p}_0}[T(X)] - (\Lambda(\theta_{\min}) - \Lambda(\theta_\infty))$, and we also find that Assumption 1 can be met if $T(X)$ is subgaussian under $p(X;\theta_{\min})$, since it will also be subgaussian for all $\theta \in [\theta_{\min}, \theta_{\max}]$.

Mean-shifted uncertainty sets: While one-parameter exponential families are interesting for the sake of analysis, it is often unrealistic that one would find a family containing both the pre-change distribution and the entire set of post-change distributions. Another approach, rooted in the idea that whenever an important change occurs, some key metric will increase, is the approach of using *mean-shifted* sets. In this case, we deal with sets of distributions which have mean greater than some threshold, as well as bounded second moments. Formally, we assume that we know the pre-change distribution p_∞ and that for some real numbers μ_0 and η we have $\mathbb{E}_\infty[X] := \mu_\infty < \eta$. The mean-shifted set is defined as

$$\mathcal{P}_0 = \{p_0 : \mathbb{E}_{0,p_0}[X] \geq \eta, V_{0,p_0}(X) < \mathcal{V}\}$$

where $\mathcal{V}$ is some (potentially large) upper bound on the variance of distributions within $\mathcal{P}_0$. These sets are very general, and analyzing them only relies on the knowledge that X will in expectation shift from μ_0 to some value greater than η. Considering these sets is realistic. Indeed, η can be be chosen such that if the shift in X is smaller than η in expectation, one can disregard the change as it would have negligible impact on the network performance.

A key feature of these sets is that we can identify the least favorable distribution of these sets and that it takes a rather comfortable form for the sake of QCD. This is a result first shown in [14] and stated below.

Proposition 2. *Define by* $\kappa_0(\lambda) = \log\left(\mathbb{E}_\infty[\exp(\lambda X)]\right)$ *the cumulant generative function for* p_∞*. Then, if* $\mathcal{V}$ *is sufficiently large,* $\mathcal{P}_0$ *is stochastically bounded with respect to* p_∞*, and there exists a value* $\lambda^\star$ *such that the least favorable distribution is given by*

$$\overline{p}_0(x) = p_\infty(x)\exp(\lambda^\star x - \kappa_0(\lambda^\star)) \tag{6}$$

Furthermore, $\lambda^\star$ *is the solution to the fixed point equation in* λ

$$\mathbb{E}_\infty[X\exp(\lambda X)] = \eta\mathbb{E}_\infty[\exp(\lambda X)]. \tag{7}$$

Equation (7) often does not exhibit an explicit solution, but it can be solved numerically. A consequence of Proposition 2 is that the log-likelihood ratio of mean-shifted uncertainty sets can be expressed as $Z_n = \lambda^\star X_n - \kappa_0(\lambda^\star)$, which allows us to immediately see that $\overline{I} = \lambda^\star\eta - \kappa_0(\lambda^\star)$.

V. Testbed description

To evaluate RoME-QCD on a realistic network scenario, we use data from a research 5G testbed instrumented with measurement and analysis capabilities [16], [17]. It corresponds to a commercial 5G mmWave NSA system where the control is performed through a 4G LTE eNB and user-plane through a 5G gNB. We configure the 4G LTE eNB to operate on band B3 (1800 MHz, 5 MHz bandwidth), and 5G gNB on band n257 (28 GHz, 100 MHz bandwidth). The spectrum is then shared between uplink and downlink using a 4:1 Time-Division Duplex (TDD) pattern [18]. Finally, the transmission power is set to 2 W.

In the testbed, there are, for the purpose of experiments considered in this paper, two user equipment (UEs). One UE is for performance measurements and the other one is for generating traffic load scenarios. Change points are defined in terms of change in load or UE movement pattern.

Data-plane network performance in this testbed is measured in terms of One-Way Delay (OWD) between a OWD probe source and a probe packet reflector every 50 ms using 1400 byte TWAMP packets [1] as illustrated in Figure 3. Precise measurements of OWD is enabled by time stamping all TWAMP packets, when sent and reflected, in the OWD probe source by using a repeater. Using the repeater we capture 4 time stamps, namely (T1) sending the probe packet from the probe source, (T2) receiving the probe packet at the reflector, (T3) returning the packet from the reflector, and (T4) receiving

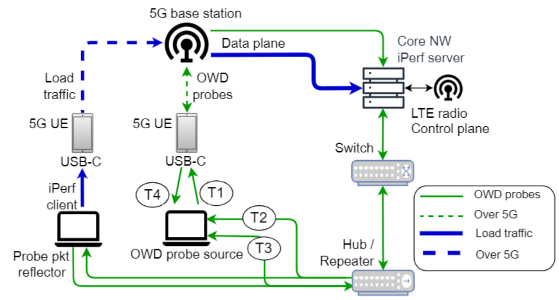

Fig. 3. An illustration of the 5G testbed setup [17]. Green lines show the path of TWAMP packets, and blue lines correspond to traffic load. A dashed line indicate transmission over the air.

the packet at the probe source. From these time stamps we then calculate OWD in the uplink (i.e., T2 - T1).

The testbed is located in an industrial facility with approximate dimensions of 30 x 7 meters. The placement and direction of the antennas are such that a UE moving in the facility may be both in and out-of-coverage.

VI. Evaluation

A. Methodology

Periodic schemes and RoME-QCD: We evaluated QCD agents using both periodic measurement schedules and RoME-QCD. We used the parameters in equations (2) and (4) in order to guarantee $(\gamma, \beta, \mathcal{P}_0)$-compliance. Recall here that γ is the constraint on the average run length to false alarm, and β is the constraint on the average detection delay.
Performance metrics: We wished to evaluate the following performance metrics:

- **Pre-change measurement frequency:** We evaluated pre-change measurement frequency by running the agent on the pre-change distribution with no changes. We estimated the frequency by $\mathbb{E}_\infty[N/t_N] \approx 1/\mathbb{E}_\infty[\hat{\tau}]$.
- **Average run length to false alarm:** We also estimated the average run length to false alarm $\mathbb{E}_\infty[t_N]$. We capped run lengths to $t_N < 5\gamma$ and then forced a stop. This biased the false alarm rate negatively, but all results still showed that the ARL is significantly greater than γ.
- **Worst case average detection delay:** It is well-known in QCD literature that whenever the stopping rule is a threshold on the CUSUM statistic, the worst case detection delay (in terms of both post-change measurement costs and physical time delay) occurs whenever $\nu = 0$. As such, we imposed a change immediately and evaluated both $\overline{\mathbb{E}}[t_N] = \mathbb{E}_{0,p_0}[t_N]$ and $\overline{\mathbb{E}}[N] = \mathbb{E}_{0,p_0}[N]$.

B. Testbed experiments

We first evaluated our QCD schemes using the tesbed described in Section V.
Data processing: To create proper measurement scenarios, we first performed seven experiments where the measurement UE was stationary, and the load was introduced through the second UE using Iperf with an intensity ranging from 0 to 60 Mbps. After this, another experiment was performed where no load

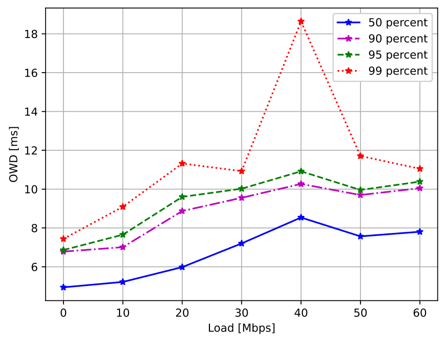

Fig. 4. Percentiles of the OWD distributions for different loads.

was introduced but the measurement UE moved back and forth in the curved corridor. For each experiment, the uplink OWD values were measured for 250 seconds, with 50 ms intervals, resulting in a total of 5 000 measurements each. Through these measurements, we could estimate p_∞ (which the agent has access to) as well as the different distributions p_0 (which it does not) in the different measurement scenarios.

Armed with the data from the experiments, we estimated p_∞ and p_0 by using Gaussian kernel density estimation (KDE). A choice of bandwidth is needed, and we heuristically set it to $\sigma_B = 0.4$ ms as this showed a good balance of resistance to measurement noise without information loss. Some key percentiles of the distributions are presented in Figure 4.

Change scenarios: Any given change scenario is defined by (1) the true post-change distribution p_0, (2) the network probe's knowledge of the post-change distribution, characterized by the uncertainty set $\mathcal{P}_0$ and (3) the true pre-change distribution p_∞. For simplicity, we let the pre-change distribution be the same everywhere, that is, the KDE generated by the stationary UE without load. Thus, we have $\mathbb{E}_\infty[X_n] = 5.11$ ms in all change scenarios. The uncertainty set $\mathcal{P}_0$ was set as a mean-shifted uncertainty set as described in Section IV-D, where the mean threshold η varies between the different change scenarios, and with the variance bound $\mathcal{V} = 100$ $(\text{ms})^2$ everywhere, just to be sufficiently large. To ensure that Assumption 1 is met, we imposed an additional constraint on $\mathcal{P}_0$ that the distributions within are subgaussian. Assuming that the distributions within $\mathcal{P}_0$ have a similar shape to p_∞, we set $\sigma_Z = \lambda^\star \sigma_B$. Here, $\lambda^\star$ was calculated as in equation (7), which can be written explicitly when p_∞ is a Gaussian KDE.

To evaluate and discuss different aspects of the robust QCD algorithms, we studied the following four change scenarios:

- **Small increase (SI):** The systems has a small increase in load, which the agent must detect. Here, we interpreted this as the system going from no load at all, to a load of 10 Mbps. Thus, p_0 is a KDE, with mean $\mathbb{E}_{0,p_0}[X_n] = 5.49$ ms and we set $\eta = 5.4$ ms.
- **Great increase (GI):** The system has a great change in load, which the agent is prepared to detect. Here, we interpreted this as the system going from no load at all, to a load of 50 Mbps. The average post-change OWD was $\mathbb{E}_{0,p_0}[X_n] = 7.66$ ms and we set $\eta = 7.0$ ms.
- **Over-increase (OI):** The system has a great change in load, but the agent is prepared to also detect a small change in load. Combined with the two above cases, this gives a good idea of the prize of robustness. The load then changed from 0 Mbps to 50 Mbps, but the mean threshold η was set to be the same as for the scenario Small increase, that is $\mathbb{E}_{0,p_0}[X_n] = 7.66$ ms and $\eta = 5.4$ ms.
- **Moving UE (MUE):** For the final scenario, we kept the load constant at 0 Mbps, but the UE started moving from a stationary position to one where it moves back and forth in a corridor. Here, we find $\mathbb{E}_{0,p_0}[X_n] = 14.9$ ms, and we set $\eta = 7.0$ ms, the same as in the scenario GI.

Network probes and evaluation: We set $(\tau, S^{(0)})$ for the periodic schedule network probes as in equation (2) and $(\tau, S^{(0)}, S^{(1)})$ for the RoME-QCD network probes as in equation (4). We did not see a natural way of setting the parameter ξ, but found that $\xi = 15$ was useful in similar experiments before.

For the experiment parameters, we set the delay constraint $\beta = 70$ s and ARL constraint $\gamma = 200\beta = 14\ 000$ s. Rather than evaluating the algorithms directly on the data, we evaluated them online with data from the estimated distributions, allowing us to freely control elements of the experiments such as pre- and post-change distribution or change time.

Presenting our results, we begin by illustrating how the network probe with periodic schedules operate contrasted with RoME-QCD. To do so, we have run an example episode for both in the scenario "Small increase". We set $\nu = 20$ s for both probes and illustrate when measurement were taken, what their outcome was and how the CUSUM statistic evolved over time. The example episodes are shown in Figure 5.

To see the sensitivity of the ARL constraint γ, we performed experiments on the scenario SI where γ varied logarithmically from β to 500β (β being the delay constraint). The average measurement frequency and detection delay is found in Figure 6. The ARL is omitted to save space, but for all agents it is close to the time limit 5γ in all experiments.

For each scenario in the subsection "Change scenarios", we evaluated both pre-change and post-change distribution, and the robust method (with $\lambda^\star$ generated by solving equation (7)) was evaluated with both periodic measurement schedules and RoME- QCD. The evaluation results are found in Table I.

Furthermore, to get a more nuanced view of the impact of the post-change load, we evaluated the worst case average detection delay for a number of different loads, increasing from 10 Mbps to 60 Mbps in increments of 10 Mbps. We evaluated these delays using both $\eta = 5.4$ ms and $\eta = 7.0$ ms. Notably, the probe could not effectively detect changes when the post-change load was 10 Mbps and the delay threshold was $\eta = 7.0$ ms, as the post-change load was in this case much smaller than η. We have thus excluded this data point from the result for the sake of presentation. These delays are found in Figure 7.

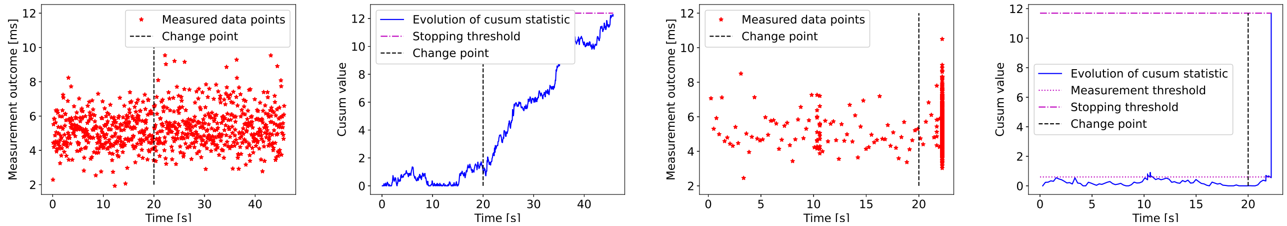

(a) Measurements, periodic (b) CUSUM, periodic (c) Measurements, ROME-QCD (d) CUSUM, ROME-QCD

Fig. 5. Example episode for network probe with periodic measurement schedules and RoME-QCD in the scenario SI, showing measurement outcomes and CUSUM statistic as a function of time.

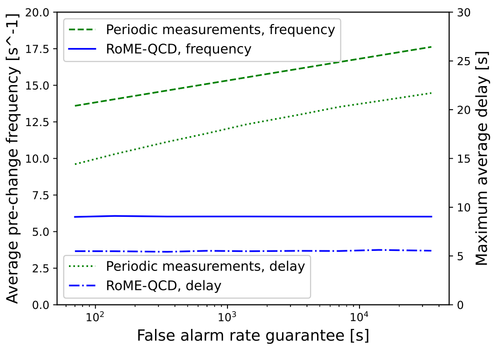

Fig. 6. Measurement frequency and detection delay for robust detection agents with periodic measurement schedules or ROME-QCD on testbed data for different values of γ, with the scenario SI.

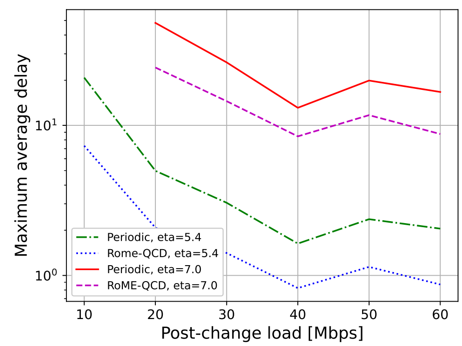

Fig. 7. Worst case average detection delay as a function of the post-change load. Note the logarithmic scale on the y-axis, and that a data point is missing in the upper two curves, which is explained in text.

TABLE I
PERFORMANCE UNDER PERIODIC VS. ROME-QCD SCHEDULES.

Scenario	Schedules	Meas. frequency	Delay	ARL
SI	Periodic	$17.0\ s^{-1}$	$20.8\ s$	$68\ 800\ s$
SI	RoME-QCD	$\mathbf{3.68}\ s^{-1}$	$\mathbf{7.3}\ s$	$69\ 300\ s$
OI	Periodic	$17.0\ s^{-1}$	$2.36\ s$	$68\ 500\ s$
OI	RoME-QCD	$\mathbf{3.68}\ s^{-1}$	$\mathbf{1.14}\ s$	$69\ 300\ s$
GI	Periodic	$1.08\ s^{-1}$	$20.0\ s$	$67\ 000\ s$
GI	RoME-QCD	$\mathbf{0.338}\ s^{-1}$	$\mathbf{6.74}\ s$	$68\ 900\ s$
MUE	Periodic	$1.08\ s^{-1}$	$\mathbf{4.96}\ s$	$67\ 100\ s$
MUE	RoME-QCD	$\mathbf{0.339}\ s^{-1}$	$9.53\ s$	$68\ 900\ s$

C. Synthetic data - Gamma distributions

In order to evaluate the cost of robustness in a controlled environment, we performed simulations on synthetic data. Since our goal was to apply our methods to delay monitoring, we investigated data generated by gamma distributions [19]. Gamma distributions, denoted by $\Gamma(\alpha, \beta)$ are defined by the probability density function

$$p(x) = \frac{\beta^\alpha x^{\alpha-1}}{\Gamma(\alpha)} \exp(-\beta x) \tag{8}$$

on the set $[0, \infty)$. Here α is the shape parameter and β is the rate parameter, and $\Gamma(\alpha)$ is the gamma function identical to $(\alpha - 1)!$ whenever α is a positive integer. The mean of the gamma distribution is $\mathbb{E}[X] = \frac{\alpha}{\beta}$ and its variance is $V(X) = \frac{\alpha}{\beta^2}$. For the pre-change distribution, we used $p_\infty \sim \Gamma(2, 2)$. The post-change distribution set was set to $\mathcal{P}_0 = \{p_0 : p_0 \sim \Gamma(2, \beta), \beta \in [3, 10]\}$. The true distribution generating post-change measurements, unknown to the agent, used $\beta = 4$, while the least favorable distribution of $\mathcal{P}_0$ is $\overline{p}_0 \sim \Gamma(2, 3)$.

We evaluated both periodic schedules and ROME-QCD, using the robust methods of Section IV. We compared them both to each other as well as a benchmark method which knows the true post-change distribution, in order to see the cost of using a robust set. To show the qualitative results as a function of γ, we fixed $\beta = 1$ and let γ vary from 1 to 500 in logarithmic steps. The pre-change measurement frequency and delay of both measurement schedules is shown in Figure 8.

D. Discussion

The main conclusion drawn from the results, particularly Table I, is that ROME-QCD is able to significantly reduce measurement overhead for network probes, compared to previous robust approaches with periodic measurements. It also does so while maintaining as good or better detection delay and false alarm rate. As expected, Figures 6 and 8 show that whenever the false alarm rate constraint γ grows, ROME-QCD significantly outperform periodic measurement schemes. However, in the scenario MUE, the delay of ROME-QCD is somewhat worse compared to periodic measurements. We attribute this to the pre- and post-change distributions having

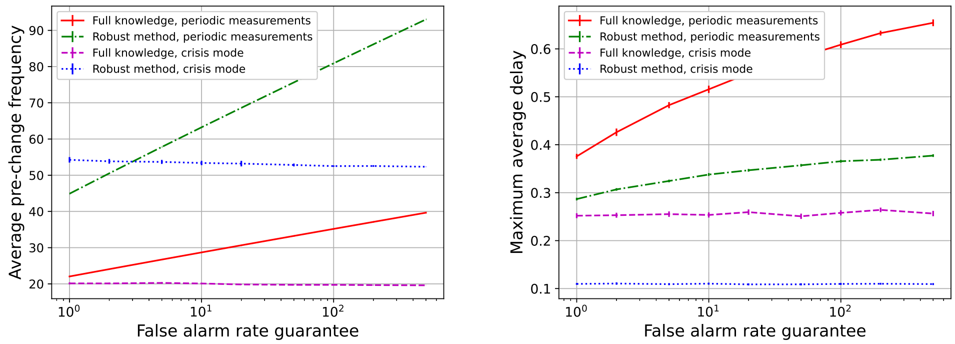

(a) Pre-change measurement frequencies (b) Maximum average delays

Fig. 8. Measurement frequency and detection delay for robust detection agents with periodic measurement schedules or ROME-QCD.

wildly different shapes, making it difficult to leave the slow measurement regime $\overline{S}_n \leq S^{(1)}$.

In Figure 5 we can get an idea of why crisis mode schedules are so effective. In terms of time, they can measure considerably less frequently before the change. After the change, they let the CUSUM statistic jump nearly vertically once suspicions of a change having occurred become large enough.

The different load distributions in Figure 7 show that the delay falls of quickly in all cases from 10 Mbps and 20 Mbps. The marginal decrease afterwards is diminishing. This suggests that even a small increase in the mean threshold could potentially have a great effect on measurement costs and monitoring performance. Furthermore, ROME-QCD consistently outperforms periodic measurement schedules in terms of delay, on all loads. Note that the results from the setup 40 Mbps are less reliable, as Figure 4 suggests that this setup is contaminated. This does not, however, affect any of our conclusions.

As Figure 8 shows, along with comparing the scenarios Over-increase and Great increase in Table I, there is a significant performance drop when using a least favorable distribution considerably different from the real distribution. This is well-known as the price of robustness and can be quite high. Currently, this issue can only be resolved by reducing the size of the uncertainty set.

While only OWD measurements were considered in these experiments, ROME-QCD extends very well to any metric which correlates positively or negatively with performance.

VII. Conclusion

In 5G networks, it is crucial for network probes to efficiently monitor and detect unknown changes in performance. In this paper, we have presented a statistical framework for approaching this monitoring task. Combining the frameworks of measurement-efficient quickest change detection and robust change point detection, we have shown how these interact and allow for more efficient and realistic network probes compared to previous approaches.

Within this framework, we have analyzed the adaptive measurement schedule called ROME-QCD and shown its superiority compared to classical periodic measurement schedules. We have proven our results and verified them, both using synthetic data and using data from a 5G testbed, on scenarios of load increases and changes in the user's behavioral patterns.

For future work, one interesting open question is what happens to measurement schedules as the assumption of independent measurements is relaxed. Indeed, in this case the magnitude of the interdependence between measurements may depend on the interval between them. We also aim to investigate the performance of ROME-QCD in a real-world deployment.

Acknowledgements

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation. The research by Simon Lindståhl and Andreas Johnsson was also supported by the Swedish Governmental Agency for Innovation Systems (VINNOVA) through the project Celtic IMMINENCE (C2020/2-2).

References

[1] A. Morton and L. Ciavattone, "Two-way active measurement protocol (twamp) reflect octets and symmetrical size features," tech. rep., AT&T Labs, 2010.

[2] M. Xie, Q. Zhang, A. J. Gonzalez, P. Grønsund, P. Palacharla, and T. Ikeuchi, "Service assurance in 5g networks: A study of joint monitoring and analytics," in *2019 IEEE 30th Annual International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, (Istanbul, Turkey), pp. 1–7, IEEE, 2019.

[3] L. Quan, J. Heidemann, and Y. Pradkin, "Trinocular: Understanding internet reliability through adaptive probing," *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4, pp. 255–266, 2013.

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

[5] J. Unnikrishnan, V. V. Veeravalli, and S. P. Meyn, "Minimax robust quickest change detection," *IEEE Transactions on Information Theory*, vol. 57, no. 3, pp. 1604–1614, 2011.

[6] 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.

[7] B. Wang and J. Su, “Flexmonitor: A flexible monitoring framework in sdn,” *Symmetry*, vol. 10, no. 12, p. 713, 2018.
[8] N. L. Van Adrichem, C. Doerr, and F. A. Kuipers, “Opennetmon: Network monitoring in openflow software-defined networks,” in *2014 IEEE Network Operations and Management Symposium (NOMS)*, (Krakow, Poland), pp. 1–8, IEEE, 2014.
[9] R. Steinert and D. Gillblad, “Long-term adaptation and distributed detection of local network changes,” in *2010 IEEE Global Telecommunications Conference GLOBECOM 2010*, (Miami, FL, USA), pp. 1–5, IEEE, 2010.
[10] T. V. Phan, T. G. Nguyen, N.-N. Dao, T. T. Huong, N. H. Thanh, and T. Bauschert, “Deepguard: Efficient anomaly detection in sdn with fine-grained traffic flow monitoring,” *IEEE Transactions on Network and Service Management*, vol. 17, no. 3, pp. 1349–1362, 2020.
[11] S. Tan, D. He, S. Chan, and M. Guizani, “Flowspotter: Intelligent iot threat detection via imaging network flows,” *IEEE Network*, 2023.
[12] G. V. Moustakides, “Optimal stopping times for detecting changes in distributions,” *the Annals of Statistics*, vol. 14, no. 4, pp. 1379–1387, 1986.
[13] 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.
[14] Y. Liang and V. V. Veeravalli, “Non-parametric quickest mean-change detection,” *IEEE Transactions on Information Theory*, vol. 68, no. 12, pp. 8040–8052, 2022.
[15] 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.
[16] A. Rao, W. Tärneberg, E. Fitzgerald, L. Corneo, A. Zavodovski, O. Rai, S. Johansson, V. Berggren, H. Riaz, C. Kilinc, and A. Johnsson, “Prediction and exposure of delays from a base station perspective in 5G and beyond networks,” in *Proceedings of the ACM SIGCOMM Workshop on 5G and Beyond Network Measurements, Modeling, and Use Cases*, pp. 8–14, 2022.
[17] A. Rao, H. Riaz, A. Zavodovski, R. Mochaourab, V. Berggren, and A. Johnsson, “Generalizable one-way delay prediction models for heterogeneous ues in 5g networks,” in *NOMS 2024-2024 IEEE/IFIP Network Operations and Management Symposium*, pp. 1–9, IEEE, 2024.
[18] 3GPP, “NR; Physical layer procedures for control,” technical specification (ts) 38.213, 3rd Generation Partnership Project (3GPP), 2021. Version 16.5.0.
[19] H. Zhang, Y. Zhang, and Y. Liu, “Modeling internet link delay based on measurement,” in *2009 International Conference on Electronic Computer Technology*, pp. 420–424, IEEE, 2009.

APPENDIX A
PROOFS

A. Proof of Lemma 1

By definition of N and n_ν, $\mathbb{E}_{\nu,p_0}[(N-\nu)^+]$ depends only on measurement outcomes $\{X_i\}_{i\geq 1}$. But, conditioned on $\mathcal{G}_{n_\nu}$, it only depends on measurement outcomes $\{X_i\}_{i>n_\nu}$. Furthermore, these are distributed the same regardless of where they take place in time, making them independent on $\{\tau_i\}_{i>n_\nu}$. This proves statement (i).

For statement (ii), the proof logic of Lemma 1 in [6] applies exactly. □

B. Proof of Theorem 1

From the proof of Theorem 2 in [6], we know that the Theorem holds already for the least favorable distribution $\overline{p}_0$, that is, we can construct a 3-tuple $(\tau, S^{(0)}, S^{(1)})$ such that the corresponding crisis mode agent is (γ,β)-compliant and has $\mathbb{E}_\infty[\hat{\tau}] \geq \tau_c$ with τ_c independent of β. It thus suffices to show that, with the same choice of $(\tau, S^{(0)}, S^{(1)})$, such an agent is also $(\gamma,\beta,\mathcal{P}_0)$-compliant under Assumption 1, which in turn requires showing $\overline{\mathbb{E}}[t_N] \leq \beta$.

Recall that $Z_i = \log(\frac{\overline{p}_0(X_i)}{p_\infty(X_i)})$, and by this, it follows that $\mathbb{E}_{\nu,p_0}[Z_i] = D(p_0||p_\infty) - D(p_0||\overline{p}_0)$ for any $i > n_\nu$. Furthermore, by definition of stochastic boundedness, it follows that $\mathbb{E}_{\nu,p_0}[Z_i] \geq \overline{I}$. Note that, for any ν and $p_0 \in \mathcal{P}_0$, we can for a crisis mode agent rewrite $\mathbb{E}_{\nu,p_0}[(t_N-\nu)^+|\mathcal{F}_\nu] = \tau\sum_{n=n_\nu+1}^{\infty}\mathbb{P}_{\nu,p_0}(S_n \leq S^{(1)})\mathbb{P}_{\nu,p_0}(N > n)$. Introducing $\Delta_Z := Z - \mathbb{E}_{\nu,p_0}[Z]$ and $\Delta_n := n - n_\nu$, we then write

$$
\begin{aligned}
&\mathbb{E}_{\nu,p_0}\left[(t_N-\nu)^+|\mathcal{G}_{n_\nu}\right] \leq \tau\sum_{n=n_\nu+1}^{\infty}\mathbb{P}_{\nu,p_0}(S_n \leq S^{(1)})\\
&\leq \tau\sum_{n=n_\nu+1}^{\infty}\mathbb{P}_{\nu,p_0}\left(\sum_{i=n_\nu+1}^{n} Z_i \leq S^{(1)}\right)\\
&= \tau\sum_{n=n_\nu+1}^{\infty}\mathbb{P}_{\nu,p_0}\left(\sum_{i=n_\nu+1}^{n}\Delta_Z \leq S^{(1)} - \Delta_n\mathbb{E}_{\nu,p_0}[Z_i]\right)\\
&\leq \tau\sum_{n=n_\nu+1}^{\infty}\mathbb{P}_{\nu,p_0}\left(\sum_{i=n_\nu+1}^{n}\Delta_Z \leq S^{(1)} - \Delta_n\overline{I}\right)
\end{aligned}
$$

The last inequality follows from $\mathbb{E}_{\nu,p_0}[Z_i] \geq \overline{I}$, as discussed above. Now, by the proof of Theorem 2 in [6], it follows that by constructing an agent as if $\overline{p}_0$ was the true post-change distribution, it will also fulfill $\mathbb{E}_{p_0}\left[\overline{t}_N\right] \leq \beta$, by choosing

$$
\begin{aligned}
\tau &= \beta\frac{(1-\exp(-\overline{I}^2/2\sigma_Z^2))}{\exp(S^{(1)}\overline{I}/\sigma_Z^2)}\\
S^{(0)} &= \log\left(\frac{\gamma}{\tau(1-\exp(-S^{(1)}))}\right).
\end{aligned}
$$

Thus, we can construct a $(\gamma,\beta,\mathcal{P}_0)$-compliant agent with $\mathbb{E}_\infty[\hat{\tau}] \geq \tau_c$. It then remains to prove the upper bound on post-change measurement cost $\overline{\mathbb{E}}[N]$.

First, recall by Lemma 1 that we do not need to consider the measurement strategy, and can determine the post-change measurement cost by studying a periodic CUSUM agent with the same value of $S^{(0)}$. Note that $\overline{I}$ and σ_Z does not depend on γ or β, and so we can re-write

$$
S^{(0)} = \log\left(\frac{\gamma}{\beta}\right) + c
$$

where c does not depend on γ or β. Since the set $\mathcal{P}_0$ is stochastically bounded with respect to p_∞, the threshold rule $N = \min\{n : \overline{S}_n > S^{(0)}\}$ has a post-change measurement cost upper bounded as

$$
\overline{\mathbb{E}}[N] \leq (1+o(1))\frac{S^{(0)}}{\overline{I}}
$$

as $\gamma \to \infty$, by Lemma 2.1 in [14]. As such, fixing β, we see that

$$
\overline{\mathbb{E}}[N] \leq (1+o(1))\frac{\log(\gamma/\beta)+c}{\overline{I}} = (\overline{I}^{-1}+o(1))\log(\gamma/\beta)
$$

as $\gamma \to \infty$, which concludes the proof. □

Appendix B
Ethical considerations

The OWD data in this work is generated in an in-house testbed rather than from any real user data. Furthermore, the description and use of this testbed has been approved by its owner. As such, this work does not raise any ethical issues.

There was no use of Generative AI in the creation of this work.