Bayesian Quickest Detection with Observation-Changepoint Feedback

Michael Ludkovski
Department of Statistics and Applied Probability, University of California
Santa Barbara, CA, USA 93106-3110 ludkovski@pstat.ucsb.edu

***Abstract*— We study Bayesian quickest detection problems where the observations and the underlying change-point are coupled. This setup supersedes classical models that assume independence of the two. We develop several continuous-time formulations of this problem for the cases of Poissonian and Brownian sensors. Our approach to detection uses methods of nonlinear filtering and optimal stopping and lends itself to an efficient numerical scheme that combines particle filtering with Monte Carlo dynamic programming. The developed models and algorithms are illustrated with numerical examples.**

Keywords: Bayesian Quickest Detection; Hawkes Process; Particle Filtering; Monte Carlo Dynamic Programming.

I. INTRODUCTION

Quickest detection of signal changes is an important problem in a variety of disciplines ranging from biosurveillance to network communications. The typical formulation is that at unknown instant θ, the statistical profile of an observation stream (Y_t) undergoes a disruption. The task of the controller is to fuse the sequentially collected information in order to detect the change-point as quickly as possible while controlling for probability of false alarms.

Models of quickest detection may be classified according to their assumptions on θ. Two main formulations of change-point detections are min-max and Bayesian [20]. The min-max approach takes θ as an unknown *constant* and aims to minimize worst-case performance, i.e. the optimization criterion involves supremum over all $\theta \in \mathbb{R}_+$. Consideration of worst-case allows for robustness, but is often too pessimistic in practice [17]. In the Bayesian setup, the disorder time is treated as a random variable with a specified prior distribution, $\theta \sim p(\theta)$ that captures accumulated knowledge about possible occurrences of the change-point. The quickest detection problem is then formulated in terms of the posterior probability of the disorder already having taken place. Models of increasing complexity, including a variety of observation schemes and multiple channels have been considered in [4], [7], [14], [15], [21], [22].

Fixing a model for θ, both methods then postulate given dynamics for observations (Y_t) and the respective impact of the disorder. A fundamental criticism of this approach is that it rules out any (explicit) dependence between θ and (Y_t). Indeed, in the min-max approach, θ is a deterministic constant; in the Bayesian approach all classical models assume autonomous description of θ, not involving (Y_t). In other words, prior to disorder, the conditional distribution $\theta | \theta > t$ is taken to be independent of $(Y_s)_{s \le t}$. For example, in the early seminal work of Shiryaev [21], θ has an exponential distribution, i.e. a constant hazard rate. While the optimality properties of detection rules such as CUSUM or Shiryaev-Roberts have been thoroughly studied in a variety of models [7], [16], [19], to our knowledge, no such proofs (and in fact only minimal analysis and numerical studies) are available for the case of observation-changepoint interaction.

In this paper, we extend the current literature by removing this independence assumption within the Bayesian paradigm. The main idea of our modeling is to view θ through its hazard rate μ_t; we then treat (μ_t, Y_t) as a coupled stochastic process. Detection of θ then reduces to a *stochastic filtering* problem for (μ_t). Our framework relies on the optimal stopping paradigm with the aim to approximate the exact optimal solution, rather than seeking other (asymptotic) notions of optimality. Given the complexity of our model, we focus on a flexible numerical approximation method, extending our earlier tools in [12], [13], [15]. This also allows us to consider a generic Bayesian formulation without requiring existence of low-dimensional sufficient statistics or making other restrictive assumptions. Moreover, as our approach is fully simulation-based, it can be easily adjusted depending on model specification, requiring no intermediate analytical computations.

As our main setup we consider *centralized Bayesian detection* with continuous-time observations modeled alternately as a point process or an Îto diffusion. Use of continuous-time is not only convenient analytically but is also more faithful for asynchronous systems without a well-defined time scale. Real-time information is bursty and consists of discrete events that punctuate periods of inactivity. We model such observations via a marked point process, linking to the theory of Poisson disorder problems [18]. Because the eventual numerical implementation is in discrete time, our methods are in fact directly applicable (with obvious modifications) also in discrete-time models.

A. Applications

Below we present three application area motivating the models we consider.

Biosurveillance of emerging pathogens: The aim of bio-surveillance is to monitor a range of prediagnostic and diagnostic data for the purpose of enhancing the ability of the public health infrastructure to detect, investigate, and respond to disease outbreaks. An important case is monitoring for emerging infectious diseases that may trigger a new pandemic, such as the avian H5N1 virus. Currently, H5N1 is primarily confined to birds and only rarely infects humans, with very few known cases of human-to-human transmission. As a result, despite its high mortality rate, the basic reproductive ratio of the pathogen is low in humans, preventing epidemics. However, each time a human is infected, further virus adaptation may result, enabling direct human-to-human transmission and triggering an epidemic onset [3]. Letting (Y_t) be the count of observed infections for an endemic pathogen, and θ the (unknown) instant when the disease goes epidemic we have a positive feedback between (Y_t) and hazard rate (μ_t) of θ.

Security intrusions: In counter-terrorism and other security intrusion applications, θ is used to denote the start of an attack by the adversary. In that context, it is clear that the assumption of independence between observations and θ is unlikely, since the enemy agent would also monitor the situation and try to strategically select a good "opportunity window" [17]. In particular, it might be expected that the attack would occur after a period of relative "quiet" when the monitor is lulled into complacency, inducing a *negative* dependence between (Y_t) and θ.

Quality control: In quality control applications [1], θ is the instant the system goes out-of-control. A common reason for break-down is accumulated damage, i.e. the previous minor degradations, captured by (Y_t), build up to a major failure (e.g. engineering structure failure after a series of small shocks, such as earthquakes over the years). Hence, observed shocks cause a positive feedback with the hazard rate of the change-point θ, and should be taken into account by the detection algorithm.

The rest of the paper is organized as follows. Section II sets up our stochastic model and formalizes the detection objectives. Section III presents our Monte Carlo based solution method. Section IV discusses further extensions of the approach which is illustrated in Section V.

II. STOCHASTIC MODEL

Fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $(Y_t)_{t\in\mathbb{R}_+}$ be the data stream observed by the controller. The statistical properties of the information received undergo a transition when the signal is present. The instant of disruption, denoted θ, is henceforth called a change-point. Crucially, θ is unobserved. Define $\mathbb{F} = \{\mathcal{F}_t\}_{t\geq 0}$ to be the right-continuous augmentation of the natural filtration $\mathcal{F}_t = \sigma(Y_s : 0 \leq s \leq t)$ of (Y_t), as well as the extended filtration $\mathcal{G}_t \triangleq \mathcal{F}_t \vee \sigma\{\theta\}, \ t \geq 0$.

The dynamic system state is described by the disorder indicator process $X_t \in \{0,1\}$ which encodes the present state of the signal,

$$X_t := 1_{\{\theta \leq t\}}. \tag{1}$$

We view (X_t) as a point process with a single arrival date θ. Denote by μ_t the hazard rate of θ. Formally, this means that

$$M_t = 1_{\{\theta\leq t\}} - \int_0^{t\wedge\theta} \mu_s \, ds$$

is a $(\mathbb{P}, \mathbb{F})$-martingale and (X_t) admits the decomposition

$$dX_t = \mu_t(1 - X_t)\, dt + dM_t.$$

In the special case where θ is a mixture of the point mass at zero and an exponential $Exp(\lambda)$ distribution, $\mu_t \equiv \mu$ is constant, making (X_t) a time-homogenous continuous-time Markov chain.

Once the change-point occurs, no further disorders can take place and we set $\mu_t = \Delta$, the 'cemetery state'. The state-space of μ_t is denoted as $\mathfrak{R} := \mathbb{R}_+ \cup \{\Delta\}$.

A. Observations Model

A variety of formulations are possible for the dynamics of observations (Y_t). As our main example, we consider monitoring of asynchronous "lumpy" events, whereby we model (Y_t) via a doubly stochastic Poisson process driven by the signal. Precisely, we assume that (Y_t) is a point process with intensity $\Lambda(X_t)$. In other words, we have that

$$dY_t = 1_{\{t\leq\theta\}} dN^0(t) + 1_{\{t>\theta\}} dN^1(t), \tag{2}$$

where $(N^j(t))$ are conditionally independent Poisson processes with intensity $\Lambda(j)$, $j = 0, 1$ (to be specified). The arrivals of (Y_t) are denoted as $0 = \sigma_0 < \ldots < \sigma_k < \sigma_{k+1}$.

In Section IV-A below we consider the alternative case where (Y_t) is specified as an Îto diffusion.

B. Feedback Effect

As a first approach for linking the hazard rate (μ_t) of θ and the observations (Y_t), we consider a bivariate Hawkes process specification [10]. Namely, (μ_t) has a shot-process feedback effect from arrivals in (Y_t),

$$\begin{aligned} \mu_t &= \bar{\mu} + (\mu_0 - \bar{\mu})e^{-\beta t} + \int_0^t a e^{-\beta(t-s)} dY_s \\ &= \bar{\mu} + (\mu_0 - \bar{\mu})e^{-\beta t} + \sum_{\ell:\sigma_\ell\leq t} a e^{-\beta(t-\sigma_\ell)}, \end{aligned} \tag{3}$$

for some constants $a, \beta, \bar{\mu}$. Thus, the transition rate of (X_t) increases by a after each arrival σ_k; this effect dissipates exponentially at rate β. Hence, if $a > 0$ then the change-point θ is likely to be "triggered" by a cluster of observed events, creating a correlation between (Y_t) and (X_t).

The model parameters are $\Xi := (a, \beta, \bar{\mu}, \Lambda(0), \Lambda(1))$, at least some of which are likely to be unknown. We assume a fully Bayesian specification with a given hyper-prior $\Xi \sim p(\Xi)$ on some domain $\Xi \in \mathcal{D}_\Xi$.

C. Bayes Risk Performance Criterion

The controller aims to detect signal presence by raising the alarm at decision time $\tau \leq T$, where T is a given horizon. Since θ is not directly observed, we require that the decision is based on available information, namely $\tau \in \mathcal{S}_T$, where $\mathcal{S}_T$ denotes the set of all $\mathbb{F}$-adapted stopping times bounded by T. Given Bayesian priors, the basic objective is to achieve quickest detection while maintaining a bound on false alarm frequency. We assume that the decision criteria are based on the detection delay $(\tau - \theta)^+$, where $a^+ := \max(a, 0)$, and the probability of the false alarm $\{\tau < \theta\}$.

The Bayesian quickest detection problem is to compute

$$V(p_0) := \inf_{\tau\in\mathcal{S}_T} \mathbb{E}\left\{(\tau-\theta)^+ + c1_{\{\tau<\theta\}}|X_0\sim p_0\right\} \quad (4)$$
$$= \inf_{\tau\in\mathcal{S}_T} \mathbb{E}\left\{\int_0^\tau 1_{\{X_s\neq\mathbf{0}\}}ds + c1_{\{X_\tau=\mathbf{0}\}}|X_0\sim p_0\right\}.$$

The first term on the right-hand-side in (4) is the expected detection delay (EDD), while the second term is the probability of a false alarm (PFA) given that alarm is raised at τ. The parameter c is the tunable penalty for false alarms; small c will induce aggressive detection, while $c\to\infty$ is the case where false alarms are not tolerated.

If (X_t) was observed (i.e. $\mathbb{F}$-measurable), then the optimal detection rule would simply be $\tau=\theta$; the crux of the problem is therefore to construct a good approximation to θ using information flow $(\mathcal{F}_t)$ only. From a control perspective, this means that to minimize Bayes risk requires now to solve a partially observable optimal stopping problem. Indeed, the costs in (4) are not measurable with respect to the decision variable τ. Accordingly, the solution approach [18] is to first perform *filtering* of the latent state (X_t) by computing the posterior distribution

$$\tilde{\Pi}_t := \mathbb{P}\left\{X_t = 1|\mathcal{F}_t\right\}. \quad (5)$$

However, due to the interaction between (Y_t) and (μ_t), $(\tilde{\Pi}_t)$ is not Markov, presenting multiple challenges of using this approach. We therefore consider the larger filtering problem for the hazard rate (μ_t) and parameters Ξ,

$$\Pi_t(A) := \mathbb{P}\left\{(\mu_t,\Xi)\in A|\mathcal{F}_t\right\}, \qquad A\in\mathcal{B}(\Re\times\mathcal{D}_\Xi). \quad (6)$$

We may identify Π_t as an element of $\mathcal{M}(\Re\times\mathcal{D}_\Xi)$, the set of all probability measures on $\Re\times\mathcal{D}_\Xi$. Thus, (Π_t) is a diffuse-measure-valued process, but trivially possesses the Markov property. In other words, to gain the Markov property, one must lift from the restricted filter $(\tilde{\Pi}_t)$ to the full filter (Π_t).

In particular, the probability that the change-point already took place is simply $\Pi_t(1_\Delta)$ and the Bayesian performance functional in (4) is equivalent to

$$J(\tau;\pi_0) := \mathbb{E}_{\pi_0}\left\{\int_0^\tau H^1(\Pi_s)\,ds + H^2(\Pi_\tau)\right\}, \quad (7)$$

where $H^1(\pi) := \pi(1_{\{\Delta\}\times\mathcal{D}_\Xi})$, $H^2(\pi) := c\pi(1_{\{\Re\setminus\Delta\}\times\mathcal{D}_\Xi})$ and $\mathbb{E}_{\pi_0}\{\cdot\}$ denotes expectation under the initial condition $p(\mu_0,\Xi|\mathcal{F}_0)\sim\pi_0$.

A Bayesian optimal detection rule τ^* is the optimizer in (7) and can be viewed as the functional mapping histories to decisions $\tau^*:\mathcal{F}_t\to\{stop, continue\}$. Since the enlarged state variable is Markov, the detection rule is simplified to a function of the current Π_t. This point of view highlights the key challenges in working with (7), namely the need to (i) characterize and solve the evolution equations of the filter process (Π_t) and (ii) overcome the curse of dimensionality associated with optimizing over the state space of Π_t. In fact, without further assumptions Π_t is infinite-dimensional making V a functional on the nonlocally compact space $\mathcal{M}(\Re\times\mathcal{D}_\Xi)$. Thus, a complete Bayesian solution requires consideration of non-Markov optimal stopping (if working with $(\tilde{\Pi}_t)$) or infinite-dimensional Markov optimal stopping problems. The resulting complexity has earned this approach the stigma of analytical and computational intractability.

D. Solution Outline

Because there are no finite sufficient statistics for the measure-valued process (Π_t), the key to its characterization are stochastic filtering techniques [2], [24]. Namely, (Π_t) satisfies a variant of the Kushner-Stratonovich nonlinear filtering equation. In general this equation is not analytically tractable and requires a numerical approximation. An efficient and flexible approach to computing Π_t is via Sequential Monte Carlo methods (SMC), also known as particle filters [5]. The main mechanism of SMC consists of a mutation-selection procedure applied to an interacting particle system.

In terms of the control step, since the state variable is Π_t, analytic representations, through, e.g. quasi-variational inequalities, of the resulting value function $V(\pi)$ are difficult to come by. Instead we recall the probabilistic characterization of V through its dynamic programming equations. Precisely, for any stopping time σ (in particular the first arrival time of (Y_t)) define the monotone operator $\mathcal{J}$ acting on a measurable test function $v:[0,T]\times\mathcal{M}(\Re\times\mathcal{D}_\Xi)\to\mathbb{R}$ via

$$\mathcal{J}v(t,\pi) = \inf_{\tau\in\mathcal{S}_T}\mathbb{E}_\pi\Big\{\int_t^{\tau\wedge\sigma} H^1(\Pi_s)\,ds + 1_{\{\tau\leq\sigma\}}H^2(\Pi_\tau) + 1_{\{\tau>\sigma\}}v(\sigma,\Pi_\sigma)\Big\}. \quad (8)$$

Then using the Bellman optimality principle,

Lemma 2.1: $V(\pi)$ is the largest fixed point of $\mathcal{J}$ smaller than $H^2(\pi)$ and one can approximate V as

$$V = \lim_{n\to\infty} V_n, \qquad \text{where } V_n = \mathcal{J}V_{n-1}, \qquad \text{with } V_0 = H^2.$$

Moreover, the optimal stopping rule is given by

$$\tau^* = \inf\{t : V(t,\Pi_t)\geq H^2(\Pi_t)\}\wedge T, \quad (9)$$

and can be approximated through $\tau^n = \inf\{t : V_n(t,\Pi_t))\geq H^2(\Pi_t)\}\wedge T$.

The ensuing representation of the value function as the Snell envelope corresponding to the reward functional $J(\tau;\cdot)$ in (7) allows to resolve the curse of dimensionality through the use of Monte Carlo dynamic programming (MCDP) methods [6], [12], [13], [15]. MCDP approximately solves the dynamic programming equations of optimal stopping by applying stochastic simulation/regression framework. This allows a fully simulation-based solution of the Bayesian formulation (4), seamlessly merging SMC inference and MCDP for the optimization step. The resulting Monte Carlo algorithm first uses particle filtering to obtain a high-dimensional approximation $(\hat{\Pi}_t^{(N)})$ to the true (Π_t) with arbitrarily small errors as $N\to\infty$, and then applies MCDP to solve the optimal stopping problem for $(\hat{\Pi}_t^{(N)})$.

Remark 2.1: When $a>0$, it follows that the posterior likelihood of a change-point decreases on each inter-arrival interval $[\sigma_k,\sigma_{k+1})$ and stopping is only optimal upon arrival $\tau^*\in\{\sigma_k : k=1,\ldots\}$. While convenient for the controller, this property is hard to exploit within the numerical solution.

Remark 2.2: If all the parameters Ξ of the model are known, one can use Bayes rule to analytically evaluate the posterior probability of the change-point. Indeed, conditioning on $\theta = s$ and observations (Y_t), one can compute the full trajectory (μ_s), $s \le t$, and therefore evaluate the resulting likelihood $p(Y_{[0,t)}|\mu_{[0,t)})$. Integrating out θ yields the posterior probability of $\{\mu_t = \Delta\} = \{X_t = 1\}$ up to an evaluation of a 1-dim. integral.

III. NUMERICAL ALGORITHM

A. Particle Filtering

We utilize sequential Monte Carlo approach to approximate $\Pi_t \simeq \hat{\Pi}_t^{(N)}$, where the discrete measure $\hat{\Pi}_t^{(N)}$ consists of N particles,

$$\hat{\Pi}_t^{(N)} := \frac{1}{W(t)} \sum_{n=1}^{N} w^n(t) \delta_{m^n(t),\xi^n}(\cdot). \tag{10}$$

Above $w^n(t) \in \mathbb{R}_+$ are the particle weights, $W(t)$ is a normalizing constant, and $m^n(t), \xi^n$ are the particle versions of μ_t and Ξ, respectively. Hence, any posterior probability of an event $A \subseteq \mathfrak{R} \times \mathcal{D}_\Xi$ is approximated via

$$\mathbb{P}\{(\mu_t, \Xi) \in A | \mathcal{F}_t\} \simeq \frac{1}{W(t)} \sum_{n:(m^n(t),\xi^n)\in A} w^n(t).$$

The SMC algorithm is now specified through the recursive evolution of the particles $(w^n(t), m^n(t), \xi^n)_{n=1}^N$, allowing for a sequential (online) update of the particle filter as new information is collected. This evolution is given by the genetic mutation-selection steps. In general, the particles are supposed to mimic μ_t, i.e. follow (3). The parameters ξ^n are therefore static. Given m_t^n, ξ^n it is trivial to back out the dynamic disorder indicator $x_t^n := 1_{\{m^n(t)=\Delta\}}$ and the corresponding arrival intensity $\Lambda(x_t^n; \xi^n)$. The weights $w^n(t)$ then correspond to the likelihood of observations $(Y_s)_{s\le t}$ given the particle history $(m^n(s))_{s\le t}$, or iteratively

$$w^n(t) = w^n(s) \cdot \exp\left(-\int_s^t \Lambda(x_u^n; \xi^n)\, du \right) \cdot \prod_{s \le \sigma_k \le t} \Lambda(x_{\sigma_k}^n). \tag{11}$$

As information is collected, most particles will diverge from observations and their weights will collapse $w^n(t) \to 0$. To avoid the resulting particle degeneracy, the SMC approach applies sequential resampling to multiply "good" particles, and cull poor ones, ensuring particle diversity. Thus, we introduce re-sampling instances R_k, $k = 1, \ldots$, at which we draw (with replacement) N times from the atomic measure $\hat{\Pi}_{R_k-}^{(N)}$ according to the weights $w^n(R_k-)$, and then reset the particle weights to $w^n(R_k) = 1$. We use the Effective Sample Size (ESS) measure of particle diversity, $ESS(t) = \{\sum_{n=1}^N (w^n(t))^2\}^{-1}$, to resample whenever ESS drops below a threshold $R_k = \inf\{t \ge R_{k-1} : ESS(t) \le \overline{ess}\}$. If the parameters Ξ are unknown, then further steps (such as the Liu-West particle rejuvenation [11]) are needed to avoid particle degeneracy in those dimensions. We also remark that advanced resampling schemes (residual, stratified, etc.) should be used to minimize Monte Carlo variance.

To simulate m_t^n, we let

$$\theta_k^n := \sigma_k + I^{-1}(E_{n,k}; m_{\sigma_k}^n), \tag{12}$$

where $I^{-1}(y;m)$ is the inverse of the cumulative hazard map $t \mapsto \int_0^t \bar{\mu} + e^{-\beta s}(m - \bar{\mu})\, ds$ and $E_{n,k}$ are independent $Exp(1)$ random variables. Thus, θ_k^n denotes the particle-specific disorder date at the k-th stage. Note that to maintain maximal particle diversity, we keep re-setting θ_k^n as long as $m_{\sigma_k}^n \neq \Delta$. Finally, for $\sigma_k \le t < \sigma_{k+1}$ we set

$$m_t^n = \begin{cases} \Delta & \text{if } \theta_k^n < t; \\ \bar{\mu}^n + e^{-\beta^n(t-\sigma_k)}(m_{\sigma_k}^n - \bar{\mu}^n) & \text{otherwise,} \end{cases} \tag{13}$$

and $m_{\sigma_k}^n = m_{\sigma_k -}^n + a^n$. Algorithm 1 below summarizes particle filtering in the Hawkes model. For simplicity it assumes that resampling takes place at arrival dates $R_k = \sigma_k$.

Algorithm 1 Particle Filtering in a Hawkes model of signal-observation feedback

```
Sample m^n ~ π, ξ^n ~ p(Ξ) n = 1,...,N
Set w^n(0) ← 1, n = 1,...,N
for k = 1,... do
   for each particle n = 1,...,N do
      Compute (x_t^n) on the interval t ∈ (σ_k, σ_{k+1}]
      Calculate weights w^n(σ_{k+1}) using (11)
   end for
   if ESS(σ_{k+1}) < ess then
      Re-sample n' ∝ w^n(σ_{k+1}) for n' = 1,...,N
      Update m^n_{σ_{k+1}} ← m^{(n')}_{σ_{k+1}}
      Re-set weights w^n(σ_{k+1}) ← 1
   end if
end for
```

B. Monte Carlo Dynamic Programming

Equipped with the filter of (μ_t), Bayesian sequential detection reduces to solving the optimal stopping problem (7) with the Markovian state variable Π_t. Since (Π_t) is measure-valued, instead of computationally intractable analytic approaches we use a simulation-based method. Recall that for a discrete-time problem,

$$V^\Delta(0, \pi) = \inf_{\tau \in \mathcal{S}_T^\Delta} \mathbb{E}_\pi \left\{ \sum_{s=0}^{\tau-1} H^1(\Pi_s) + H^2(\Pi_\tau) \right\}$$

where $\mathcal{S}_T^\Delta = \{\tau \in \mathcal{S} : \tau \in \{0, \Delta, 2\Delta, \ldots, (T/\Delta)\Delta\}\}$, Bellman's optimality principle implies that

$$V^\Delta(t, \Pi_t) = \mathbb{E}\Big\{ \sum_{s=t}^{\tau^*(t)-1} H^1(\Pi_s)\Delta t + H^2(\Pi_{\tau^*(t)}) \big| \mathcal{F}_t \Big\};$$
$$\tau^*(t) = t 1_{\{\mathbf{S}_t\}} + \tau^*(t+\Delta t) 1_{\{\mathbf{S}_t^c\}}, \tag{14}$$

where $\tau^* = \tau^{*,\Delta}(t)$ is the optimal stopping time conditioned on not stopping so far and $\mathbf{S}_t^c$ is the complement of the set

$$\mathbf{S}_t := \Big\{ H^2(\Pi_t) < H^1(\Pi_t)\Delta t + \mathbb{E}\left\{ V^\Delta(t+\Delta t, \Pi_{t+\Delta t}) | \mathcal{F}_t \right\} \Big\}. \tag{15}$$

By the Markov property, the conditional expectation $\mathbb{E}\left\{V^{\Delta}(t+\Delta t, \Pi_{t+\Delta t})|\mathcal{F}_t\right\} =: \hat{E}(t;\Pi_t)$ is a function of the measure-valued Π_t for some functional $\hat{E} : [0,T] \times \mathcal{M}(\Re \times \mathcal{D}_\Xi) \to \mathbb{R}$. The MCDP method first replaces $V(t+\Delta t, \Pi_{t+\Delta t})$ in the last term of (15) with an *empirical pathwise continuation value* $v_{t+\Delta t}$ (computed according to (14)). It then implements (15) by replacing the conditional expectation operator $\mathbb{E}[\cdot|\mathcal{F}_t]$ (characterized as the L^2-minimizer) with an L^2-projection onto the $span(B_i(\Pi_t) : i = 1, \ldots, r)$,

$$\mathbb{E}\left\{V^{\Delta}(t+\Delta t, \Pi_{t+\Delta t})|\mathcal{F}_t\right\} \simeq \sum_{i=1}^{r} \alpha^i(t) B_i(\Pi_t), \qquad (16)$$

where $(B_i(\pi))_{i=1}^r$ are the basis functions and $\alpha^i(t)$ the corresponding regression coefficients. This is implemented through a *cross-sectional regression* of a Monte Carlo collection $(v^m_{t+\Delta t})_{m=1}^M$ to find (α^i). Comparing the prediction $\sum_i \alpha^i(t) B_i(\Pi_t)$ and the immediate payoff $H^2(\Pi_t)$ we then construct the approximate stopping region $\mathbf{S}_t$ for (15).

Finally, since we do not have access to (Π_t), we instead work with the approximate filter $\hat{\Pi}^{(N)}$. Thus, we simulate M realizations $(Y_t^{(m)})$ of (Y_t) (along with the "shadow" simulations of $(\mu_t^{(m)})$), and generate $(\hat{\Pi}_t^{(N),m})$ along each Monte Carlo path using the particle filter above. Simulation of the Hawkes process $(Y_t^{(m)}, \mu_t^{(m)})$ is done using the Poisson thinning algorithm of [10] which is similar to (12), except we also simultaneously generate the event times σ_k. We then approximate (i) $B_i(\Pi_t^m) \simeq B_i(\hat{\Pi}_t^{(N),m})$ and using backward recursion implement (15) by (ii) regressing the empirical $(v^m_{t+\Delta t})$ against the simulated $\{B_i(\Pi_t^{(N),m})\}_{m=1}^M$ to (iii) obtain the empirical regression coefficients from the simulation of size M, $\alpha^{(M),\cdot}(t)$, and the approximate value function $V^{\Delta}(0,\pi; M, N, r, \Delta t)$, see [13]. General theory [12] implies that $V^{\Delta}(0,\pi; M, N, r, \Delta t) \to V$ as $r \to \infty$, $N \to \infty$, $M \to \infty$, $\Delta t \to 0$.

The Bayesian detection rule is a map between Π_t and the stopping decision. This suggests that to obtain good tests, it is important first and foremost to identify the key features in Π_t. In the MCDP method, this translates into appropriately *parametrizing* candidate rules in terms of the basis functions. For example, the posterior probability of a disorder, $\Pi_t(1_{\{\Delta\}})$ drives the immediate payoff H^2 and is certainly a relevant quantity. If we therefore take $B_1(t,\pi) = \pi(\{X_t = 1\})$, and $r = 1$ the resulting detection test consists of declaring alarms based *solely* on $\mathbb{P}\{\theta \leq t|\mathcal{F}_t\}$. This cannot be optimal, since it would imply ignoring all other information in (Π_t), but provides a good starting point. Such a procedure can then be iteratively used to pick the basis functions B_i while employing natural stopping rules and remaining faithful to the true non-Markovian system dynamics. For instance, knowing the posterior likelihood that μ_t is "large" might provide additional insight about whether the change-point is likely to have occurred or not.

IV. EXTENSIONS

Our framework of applying stochastic filtering and optimal stopping techniques on quickest detection problems is highly flexible and can handle a variety of modified models. Below we briefly discuss several potential generalizations.

A. Diffusion Observations

In applications where the sensors collect high-frequency information, a model based on observations under white noise may be more appropriate. The resulting diffusion-based model of quickest detection (see e.g. [18]) takes (Y_t) to be a Brownian motion with drift driven by (X_t). To relate (Y_t) and the X-hazard rate (μ_t) we assume that the latter is also a diffusion, correlated with (Y_t). For simplicity, we assume constant volatilities of the background noises,

$$dY_t = \Lambda(X_t)\,dt + \sigma\,dW_t; \qquad (17)$$

$$d\mu_t = A(\mu_t)\,dt + b\,dZ_t, \qquad (18)$$

where the Wiener processes (W_t) and (Z_t) have correlation $\mathbb{E}\{Z_t W_s\} = \rho(t \wedge s)$. In a typical example, (Y_t) is drift-less $\Lambda(0) = 0$ before the change-point, and acquires positive drift $\Lambda(1) > 0$ thereafter. Classical techniques [2, Ch 3] imply that in the case of (17) the filter (Π_t) satisfies the Zakai equation of nonlinear filtering. A corresponding particle filter can then be obtained as in our previous work [12]; no changes are needed in the MCDP step.

B. Extended Point Process Model

In (3), all the noise in the hazard rate (μ_t) came from the observations. We can consider a more general setup where (μ_t) possesses its own stochastic sources. Two possibilities are to model (μ_t) as a jump-diffusion with jump times coming from (Y_t), or as a general shot-noise process with arrival times (s_k). In the latter case, some of the s_k's correspond to observed arrivals σ_ℓ and some are unobserved, resulting in a censored model. Efficient simulation and filtering of such models was recently studied in [23], [8], [9].

V. NUMERICAL EXAMPLE

We take $\mathbb{P}(X_0 = 0) = 1$ and known parameters $\beta = 4, \bar{\mu} = 0.2, \Lambda(0) = 5$. We assume that the feedback strength a and the post-disorder intensity are unknown and have the joint log-normal prior $\mathbb{E}\{\Lambda(1)\} = 10, Var(\Lambda(1)) = 100(e^{0.01}-1), \mathbb{E}\{a\} = 0.05, Var(a) = 0.05^2(e^{0.04}-1)$ with correlation $\rho = 0.5$. Thus, we assume that stronger feedback effects are more likely to trigger more severe disorders. Figure 1 shows a sample path of the resulting particle filter $(\hat{\Pi}_t^{(N)})$. From our numerical experiments, $N = 2000$ particles produces a good approximation to the true change-point probabilities Π_t; each path of $(\hat{\Pi}_t^{(2000)})_{t\in[0,8]}$ takes less than half a second to generate on a desktop.

We proceed to study the stopping rule for the Bayesian risk minimization problem. We use $\Delta t = 0.02$, $T = 8$, $M = 32{,}000$ and $N = 2000$, with the regression bases ($r = 4$)

$$B_1(\pi) = \mathbb{P}\{\theta \leq t|\mathcal{F}_t\}, \quad B_2(\pi) = \mathbb{E}\{\mu_t|\mathcal{F}_t\},$$
$$B_3(\pi) = \mathbb{P}\{\theta \leq t|\mathcal{F}_t\}^2 \quad B_4(\pi) = \mathbb{E}\{\mu_t|\mathcal{F}_t\} \cdot \mathbb{P}\{\theta \leq t|\mathcal{F}_t\}.$$

With these parameters, $\mathbb{E}\{\theta\} \simeq 3.88$ and $\mathbb{P}\{\theta < T\} \simeq 0.875$, so the horizon constraint is not negligible. Table I presents

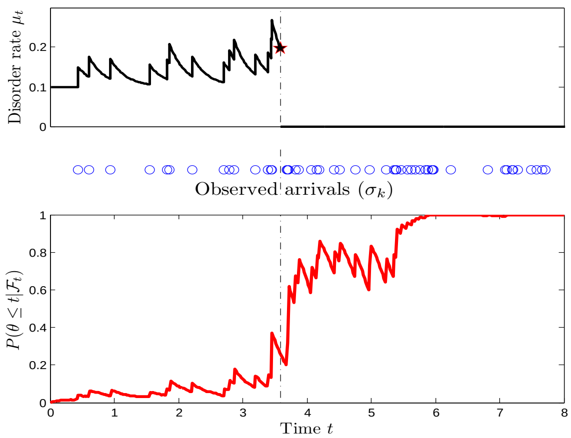

Fig. 1. Sample path of the observations (Y_t) (middle) and corresponding change-point intensity (μ_t) (top). The change-point θ is indicated with the vertical line. We also show the resulting posterior probability $\mathbb{P}\{\theta \le t|\mathcal{F}_t\}$ computed using $(\hat{\Pi}_t^{(N)})$ with $N = 2000$ (bottom). True parameters are $\Lambda(1) = 10, a = 0.05$.

TABLE I

SOLUTION OF THE BAYESIAN RISK MINIMIZATION PROBLEM (4) FOR A RANGE OF FALSE ALARM COSTS c.

c	$V(\mathbf{0})$	$\mathbb{E}_\mathbf{0}\{\tau^*\}$	PFA$_{\{\tau^*<T\}}$	EDD
5	1.96	4.12	0.135	1.00
10	2.85	4.57	0.069	1.30
20	4.36	4.97	0.037	1.63
50	8.42	5.41	0.011	2.04

some summary results as we vary the cost of false alarms c. We recall that the total Bayes risk can be decomposed into the probability of false alarm PFA and expected detection delay EDD,

$$V(\mathbf{0}) = \mathbb{E}_\mathbf{0}\{(\tau^* - \theta)^+\} + c\mathbb{P}_\mathbf{0}\{\tau^* \le \theta\} =: \text{EDD} + c \cdot \text{PFA}.$$

As expected, higher c reduces PFA and increases EDD, as well as the average time until first alarm. For example for $c = 10$, the PFA up to T, PFA$_{\{\tau^*<T\}} := \mathbb{P}_\mathbf{0}\{\tau^* < \theta|\tau^* < T\}$, is about 6.9% and the average detection delay is 1.30, while for stronger penalty $c = 20$ the probability of false alarm drops to 3.7% at the cost of increasing the detection delay to 1.63. Intuitively, the decision maker stops once the posterior probability of a change-point is "high enough". For example, for $c = 10$, the threshold is around $\mathbb{P}\{\theta \le t|\mathcal{F}_t\} \ge 0.91$—0.95, with the precise rule depending on the posterior mean of μ_t (stop sooner if μ_t is large, anticipating an imminent change-point which makes waiting costlier).

VI. CONCLUSION

Above we have developed a stochastic model for quickest detection explicitly accounting for changepoint-observations interaction. The key to our formulation is a Bayesian point of view which translates change-point detection into a nonlinear filtering step for the hazard rate of the change-point followed by an optimal stopping step. This approach allows a flexible specification of the effect of observations on the change-point which can be tailored to a variety of applications.

REFERENCES

[1] T. Aven and U. Jensen, *Stochastic models in reliability*. Springer Verlag, 1999, vol. 41.

[2] A. Bain and D. Crisan, *Fundamentals of stochastic filtering*, ser. Stoch. Model. Appl. Probab. New York: Springer, 2009, vol. 60.

[3] L. Bettencourt and R. Ribeiro, "Real time Bayesian estimation of the epidemic potential of emerging infectious diseases," *PLoS One*, vol. 3, no. 5, p. e2185, 2008.

[4] S. Dayanik, V. Poor, and S. Sezer, "Sequential multi-hypothetis testing for compound Poisson processes," *Stochastics*, vol. 80, no. 1, pp. 19–50, 2008.

[5] A. Doucet, N. de Freitas, and N. Gordon, Eds., *Sequential Monte Carlo methods in practice*, ser. Stat. Eng. Inform. Science. New York: Springer-Verlag, 2001.

[6] D. Egloff, "Monte Carlo algorithms for optimal stopping and statistical learning," *Ann. Appl. Probab.*, vol. 15, no. 2, pp. 1396–1432, 2005.

[7] E. Feinberg and A. Shiryaev, "Quickest detection of drift change for Brownian motion in generalized Bayesian and minimax settings," *Statist. Decisions*, vol. 24, no. 4, pp. 445–470, 2006.

[8] K. Giesecke, H. Kakavand, and M. Mousavi, "Exact simulation of point processes with stochastic intensities," *Oper. Res., forthcoming*, 2011.

[9] K. Giesecke and G. Schwenkler, "Filtered likelihood for point processes," *Ann. Stat., forthcoming*, 2011.

[10] P. Lewis and G. Shedler, "Simulation of nonhomogeneous Poisson processes by thinning," *Naval Res. Logist.*, vol. 26, no. 3, pp. 403–413, 1979.

[11] J. Liu and M. West, "Combined parameter and state estimation in simulation-based filtering," in *Sequential Monte Carlo methods in practice*, New York: Springer, 2001, pp. 197–223.

[12] M. Ludkovski, "A simulation approach to optimal stopping under partial observations," *Stoch. Proc. Appl.*, vol. 119, no. 12, pp. 4061–4087, 2009.

[13] ——, "Monte Carlo methods for adaptive disorder problems," in *Numerical Methods in Finance*, ser. Springer Proceedings in Mathematics, R. Carmona, P. D. Moral, P. Hu, and N. Oudjane, Eds. Springer, 2012, vol. 12, pp. 83–112.

[14] M. Ludkovski and S. Sezer, "Finite horizon decision timing with partially observable Poisson processes," *Stoch. Models*, vol. 28, no. 2, pp. 207–247, 2012.

[15] M. Ludkovski, "Bayesian quickest detection in sensor arrays," *Sequential Anal., To Appear*, 2012.

[16] G. Moustakides, "Optimality of the CUSUM procedure in continuous time," *Ann. Stat.*, vol. 32, no. 1, pp. 302–315, 2004.

[17] ——, "Sequential change detection revisited," *Ann. Stat.*, vol. 36, no. 2, pp. 787–807, 2008.

[18] G. Peskir and A. N. Shiryaev, *Optimal Stopping and Free-boundary problems*. Lectures in Mathematics, Zurich: Birkhauser-Verlag, 2006.

[19] M. Pollak and A. G. Tartakovsky, "Asymptotic exponentiality of the distribution of first exit times for a class of Markov processes with applications to quickest change detection," *Theory Probab. Appl.*, vol. 53, pp. 430–442, 2009.

[20] A. S. Polunchenko and A. G. Tartakovsky, "State-of-the-art in sequential change-point detection," *Meth. Comput. Appl. Probab., To Appear*, 2012.

[21] A. N. Shiryaev, "Optimal methods in quickest detection problems," *Teor. Verojatnost. i Primenen.*, vol. 8, pp. 26–51, 1963.

[22] A. G. Tartakovsky and V. V. Veeravalli, "General asymptotic Bayesian theory of quickest change detection," *Theory Probab. Appl.*, vol. 49, no. 3, pp. 458–497, 2005.

[23] E. Varini, "A Monte Carlo method for filtering a marked doubly stochastic Poisson process," *Stat. Methods Appl.*, vol. 17, pp. 183–193, 2008.

[24] M. H. Vellekoop and J. M. C. Clark, "A nonlinear filtering approach to changepoint detection problems: direct and differential-geometric methods," *SIAM J. Control Optim.*, vol. 42, no. 2, pp. 469–494, 2003.