



Particle observers for state estimation and adaptation in deterministic systems with random piecewise constant inputs

Silvère Bonnabel, Jean-Jacques Slotine

► To cite this version:

Silvère Bonnabel, Jean-Jacques Slotine. Particle observers for state estimation and adaptation in deterministic systems with random piecewise constant inputs. To be published in the proceedings of IEEE Conference on Decision and Control 2018. 2018. <hal-01893032>

HAL Id: hal-01893032

<https://hal-mines-paristech.archives-ouvertes.fr/hal-01893032>

Submitted on 11 Oct 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Particle observers for state estimation and adaptation in deterministic systems with random piecewise constant inputs

Silvère Bonnabel and Jean-Jacques Slotine

Abstract—In this paper we develop a methodology for state estimation of partially contracting systems inspired by the particle filter (PF). If a system is partially contracting, in the sense that once some state variables are assumed known and viewed as inputs, the remaining reduced system is contracting, then only those variables need to be estimated. Indeed, once those variables are known, the remaining variables are asymptotically recovered owing to the contraction properties. To estimate the state of such a system, we thus suggest to use a PF approach to estimate only the non-contracting part, whereas the other part is naturally recovered. The methodology is applied to a nontrivial tracking problem. The trajectories of objects one seeks to track, such as aircraft and marine vehicles, typically consist of smooth sections with large, but infrequent, unpredictable changes. As a result, their motion is well modeled by nonlinear deterministic ordinary differential equations driven by piecewise constant random control inputs. Following our methodology, we combine a type of PF called the variable rate particle filter (VRPF) with a bank of observers assumed to have convergence properties. Using the notion of virtual system, we devise a hybrid variable rate particle observer, that uses the particles to detect changes, and the observers for state estimation in between changes.

I. INTRODUCTION

Particle filters have attracted a lot of attention over the past two decades, with the advent of powerful computers, see e.g., [1]. Nowadays, they are able to execute at many hundreds of Hertz for realistic problems, as object tracking (with early implementations dating back to [2]), automotive state estimation, and parameter learning. Virtually all practical implementations of the particle filter resort to the Rao-Blackwellized particle filter (RBPF): the state consists of two groups variables (r, z) , such that given r the distribution of variable z may be estimated analytically, typically via a Kalman filter. It thus suffices to use particles to estimate r , and, using Bayes' theorem the whole distribution of (r, z) is then estimated. This reduces the number of particles and improves estimation accuracy [3].

Building upon this idea, we propose a methodology of particle observers (PO) for partially contracting systems, i.e., systems with state (r, z) , and such that if $r(t)$ is known and viewed as an input, the reduced system governing $z(t)$ is deterministic and is contracting, see [4]. Along the lines of the RBPF methodology, the rationale is to use particles to estimate the statistical distribution of variable r , whereas z may be asymptotically recovered owing to contraction

properties [5]. This leads to a parsimonious representation of the state, and makes the particle filtering problem more tractable and the filter possibly more accurate.

Intuitively, the main difference is as follows. In the RBPF, $r(\omega, t)$ and $z(\omega, t)$ are stochastic processes, and if $r(\cdot, t)$ is known, then the *distribution* of $z(\cdot, t)$ is supposed to be analytically computable. In our approach, $z(\cdot, t)$ is deterministically related to $r(\cdot, t)$ owing to the conditionally deterministic assumption. Thus, given r we need to estimate a *value* for $z(t)$. This is where contraction properties allow us to derive an approximation $\hat{z}(t)$ of this “true” value $z(t)$ instead of an entire distribution, making the use of an observer more suited to estimating the z part of the state.

We then tackle a tracking application. Trajectories of aircraft and marine vehicles are in fact smooth and well described by continuous time ordinary differential equations (ODE) driven by constant inputs, but subject to abrupt changes at each maneuver. Based on this approach to modeling the tracking task, S. Godsill with various co-authors proposed the variable rate particle filter [6], [7], [8], a particle filter well suited to piecewise deterministic models.

On the one hand, it is well known that particle filters have trouble converging in the presence of low process noise, and are subject to divergence when dealing with (even partially) deterministic models, see e.g., [9]. In particular, to the authors best knowledge, the stability of the variable rate particle filter has never been studied. On the other hand, the field of nonlinear observers, and notably the theory of contraction [5], provide powerful tools from dynamical systems theory to achieve stable estimation for noise-free models, but are not suited to tasks such as change detection. This prompts an application of our methodology of particle observers for partially contracting systems. We consider deterministic systems with inputs jumping randomly at unknown times, and run many observers in parallel (the particles), each assumed to be contracting in the absence of jumps. This way, we obtain a hybrid filter which is able to detect the jumps, and to reconverge to the true state between jumps. The theory is made rigorous using the notion of virtual systems introduced in [4], see also [10]. Our approach may be related to the flow particle filter [11], and the feedback particle filter [12]. Some preliminary ideas can be found in our recent paper [13].

Section II is devoted to the particle observer methodology. Section III considers the piecewise deterministic estimation problem, and the variable rate particle filter is introduced. In Section IV, assuming we have at our disposal an observer with some convergence properties, a hybrid strategy that

S. Bonnabel is with Mines ParisTech, PSL Research University, Centre for Robotics, 60 bd Saint-Michel, 75006 Paris, France silvere.bonnabel@mines-paristech.fr

J.J. Slotine is with the Department of Mechanical Engineering and the Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139 USA, jjs@mit.edu

combines it with a particle filter is proposed.

II. PARTICLE FILTERS AND PARTICLE OBSERVERS

A. The Rao-Blackwellized particle filter (RBPF)

Consider a discrete Markov process r_n having initial distribution $p(r_0)$ and transition kernel $p(r_n | r_{n-1})$. The variable r_n is supposed hidden, and assume we have observations in the form of a random variable y_n at time n correlated with r_n . Moreover, assume the observations to be conditionally independent given the hidden variables. The goal of discrete time filtering is to estimate the hidden variables at best from the observations, i.e., to ideally compute

$$p(r_n | y_{1:n}), \quad \text{where } y_{1:n} = \{y_1, \dots, y_n\}, \quad (1)$$

or more generally $p(r_{1:n} | y_{1:n})$, a task known as smoothing. Assume now, that observations involve another process z_n , such that $p(z_n | y_{1:n}, r_{1:n})$ may be analytically evaluated, typically using a Kalman filter.

If we may approximate the desired distribution using a set of weighted particles, i.e., $p(r_{1:n} | y_{1:n}) \approx \sum_{i=1}^N w_n^{(i)} \delta_{r_{1:n}}^{(i)}(r_{1:n})$, then the best estimate based on past observations $y_{1:n}$ of any quantity $F(z_n, r_n)$ in the sense of minimal mean squared error may be approximated as:

$$\mathbb{E}_{p(z_n, r_n | y_{1:n})} (F(z_n, r_n)) \approx \sum_{i=1}^N w_n^{(i)} \mathbb{E}_{p(z_n | y_{1:n}, r_{1:n}^{(i)})} (F(z_n, r_n^{(i)}))$$

where we used $p(z_n, r_{1:n} | y_{1:n}) = p(z_n | r_{1:n}, y_{1:n}) p(r_{1:n} | y_{1:n})$.

This gives rise to the celebrated Rao-Blackwellized Particle Filter (RBPF). Algorithm 1 provides its basic version.

Algorithm 1 RBPF with prior sampling

Draw N particles from the prior initial distribution $p(r_0)$

for $i = 1 \dots N$ **do**

Sample from the prior

$$r_n^{(i)} \sim p(r_n | r_{n-1}^{(i)}), \quad \text{and let } r_{1:n}^{(i)} = (r_n^{(i)}, r_{1:n-1}^{(i)})$$

Evaluate and update weights

$$w_n^{(i)} = p(y_n | y_{1:n-1}, r_{1:n}^{(i)}) w_{n-1}^{(i)}$$

Normalize weights $\tilde{w}_n^{(i)} = w_n^{(i)} / [\sum_j w_n^{(j)}]^{-1}$

Our estimate for value $F(z_n, r_n)$ given past observations is

$$\sum_i \tilde{w}_n^{(i)} \mathbb{E}_{p(z_n | y_{1:n}, r_{1:n}^{(i)})} (F(z_n, r_n^{(i)})).$$

Resample if necessary, i.e., multiply/delete particles to obtain N random samples with weights equal to $1/N$.

end for

B. Particle observers for partially contracting systems

In this paragraph, we summarize the results of [13]. Consider a noisy dynamical system of the form

$$\frac{d}{dt} z = f(z, r) \quad (2)$$

$$\frac{d}{dt} r = g(r) + w(t) \quad (3)$$

where f, g are smooth maps, and $w(t)$ process noise.

Definition 1: If equation (2) is a contraction when $r(t)$ is considered as a known input, system (2)-(3) is said to be a partial contraction [4].

This implies, see [5], that if $r(t)$ is known and viewed as an input of (2), then (2) has the property that for any two trajectories $z_1(t), z_2(t)$ we have:

$$d(z_1(t), z_2(t)) \leq d(z_1(0), z_2(0)) e^{-2\lambda t}, \quad (4)$$

with λ the contraction rate, and d a well-chosen distance.

Assume discrete time noisy measurements are available:

$$y_n = h(z(t_n), r(t_n)) + V_n \quad \text{at times } t_0 < t_1 < t_2 < \dots, \quad (5)$$

with V_n unknown independent identically distributed random variables with density l , i.e., $p(y | z, r) = l(y - h(z, r))$.

Note that, by denoting $(r_n, z_n) := (r(t_n), z(t_n))$ one may want to apply the RBPF methodology to dynamics (2)-(3) with observations (5). However, the usual assumption underlying the RBPF that $p(z_n | y_{1:n}, r_{1:n})$ may be analytically evaluated does not hold. It may yet be approximated under the partial contraction assumption introduced in Definition 1. Indeed, if we call $\hat{z}(t)$ a solution of (2) associated to some given trajectory $\{r(t)\}_{t \geq 0}$, then asymptotically we have $p(z(t) | \{r(s)\}_{s \leq t}) \approx \delta(z(t) - \hat{z}(t))$, since (4) implies exponential convergence of $\hat{z}(t)$ to $z(t)$. Thus, letting $(\hat{z}(t)^{(i)}, r(t)^{(i)})$ be a solution to the stochastic differential equations (SDE) (2)-(3), we have the approximations:

$$p(z(t) | y_{1:n}, \{r(s)^{(i)}\}_{s \leq t}) \approx \delta(z_t - \hat{z}(t)^{(i)}), \quad t \geq t_n \quad (6)$$

$$\mathbb{E}_{p(z(t) | y_{1:n}, \{r(s)^{(i)}\}_{s \leq t})} (F(z(t), r(t)^{(i)})) \approx F(z(t)^{(i)}, r(t)^{(i)}),$$

$$p(y_n | y_{1:n-1}, \{r(s)^{(i)}\}_{s \leq t_n}) \approx l(y_n - h(\hat{z}(t_n)^{(i)}, r(t_n)^{(i)})). \quad (7)$$

An application of the RBPF methodology to (2)-(3) with observations (5) under the partial contraction assumption and the subsequent approximations (6)-(7) thus yields the particle observer (PO) Algorithm 2.

Algorithm 2 The PO with prior sampling

Draw N particles from the prior initial distribution $p(z_0, r_0)$

for $i = 1 \dots N$ **do**

Sample $(z(t_n)^{(i)}, r(t_n)^{(i)})$ by numerically integrating the stochastic differential equations (2)-(3) from t_{n-1} to t_n .

Evaluate and update weights

$$w_n^{(i)} = l(y_n - h(z(t_n)^{(i)}, r(t_n)^{(i)})) w_{n-1}^{(i)}.$$

Normalize weights $\tilde{w}_n^{(i)} = w_n^{(i)} / [\sum_j w_n^{(j)}]^{-1}$.

Our estimate for value $F(z, r)$ given past observations is

$$\sum_i \tilde{w}_n^{(i)} F(r(t)^{(i)}, z(t)^{(i)}).$$

In particular the state estimate is $\sum_i \tilde{w}_n^{(i)} (r(t)^{(i)}, z(t)^{(i)})$.

Resample if necessary.

end for

III. STATE ESTIMATION FOR DETERMINISTIC SYSTEMS WITH RANDOM JUMPS

The reader might wonder to what extent it is realistic to consider systems of the form (2)-(3) where noise affects only (3). First, it may serve as a meaningful approximation, if for instance (2) is subject to small noise. Then, and above all, it may readily be applied to various realistic systems. One example consists of nonlinear piecewise deterministic systems subject to abrupt unpredictable jumps. These are relevant to tracking applications, and in the remainder of this paper we apply the ideas of Section II to such systems. Various other perspectives are offered in Section V.

A. Original considered model

In this section we review the conditionally deterministic inference problem addressed by the variable rate particle filter [6], see also [8] for a somehow more concise presentation. We consider a deterministic dynamical model with continuous-time state $x(t) \in \mathbb{R}^n$, and parameter (or control input) $u \in \mathbb{R}^m$. We assume that changepoints occur at random times $\{\tau_0 = 0, \tau_1, \dots, \tau_k, \dots\}$, each with associated parameters $\{u_0, u_1, \dots, u_k, \dots\}$. The pairs (τ_k, u_k) form elements of a marked point process. Consider some discrete measurement times $t_0, t_1, \dots, t_n, \dots$, along with associated measurements $y_0, y_1, \dots, y_n, \dots$ of the form

$$y_n = h(x(t_n)) + V_n \quad (8)$$

with V_n a random variable with known distribution. By letting l denote the probability density function of V_n we have

$$p(y_n | x(t_n)) = l(y_n - h(x(t_n))). \quad (9)$$

The objective is to estimate from partial and noisy measurements of $x(t)$, the changepoints, the associated parameters, and the state $x(t)$ in real time.

Let θ_n represent the past and present variables

$$\theta_n = \{\tau_1, \dots, \tau_j, u_1, \dots, u_j : 0 \leq \tau_j < t_n \leq \tau_{j+1}\}.$$

Let us define the following quantities

$$K(t) = \max\{k : \tau_k < t\}, \quad K_n = K(t_n)$$

Note that, we could alternatively write $K_n = \tilde{K}(\theta_n)$ as

$$K_n = \max\{i : \tau_i \in \theta_n\} \quad (10)$$

Processes τ_k and u_k are assumed to be Markov, and the transitions to be given by

$$p(\tau_k | \tau_{k-1}) = p_\tau(\tau_k - \tau_{k-1}), \quad p_u(u_k | u_{k-1})$$

with p_τ a probability density such that $p_\tau(\tau) = 0$ for $\tau < 0$.

The state evolution is assumed to be governed by an ODE which depends on the changepoints and the corresponding parameters as follows:

$$\frac{d}{dt}x(t) = f(x(t), t, u_{K(t)}). \quad (11)$$

This means that the state is governed by an ODE driven by parameters that remain constant over random time intervals. Let us denote ϕ the flow of the ODE $\dot{x} = f(x, t, u)$, that is,

$\phi(x_1, u, t_1, t_2)$ is the solution at t_2 to the latter ODE with fixed parameter u and initial condition x_1 at t_1 . We have:

$$x(t) = \phi(x(\tau_{K_n}), u_{K_n}, \tau_{K_n}, t), \quad \tau_{K_n} < t \leq \tau_{K_n+1}. \quad (12)$$

By choosing $t = \tau_{K_n+1}$, this equation specifies the state at the next changepoint, and by reiterating to the following changepoints, we get the whole future trajectory. This proves that the knowledge of the process θ_n and the initial state $x(0)$ uniquely specifies the state $x(t)$. For simplicity, we will systematically assume $x(0)$ to be known. This is not a very restrictive assumption, as the aim of tracking is to keep track of $x(t)$ despite the random changepoints, and after a few jumps this knowledge about the initial state is forgotten.

Recalling (1), the goal of the variable rate particle filter is to compute an approximation to the posterior

$$p(\theta_n | y_{1:n}) \quad (13)$$

and $x(t)$ as a byproduct, through (12), i.e. by integrating (11).

B. The variable rate particle filter

The bootstrap variable rate particle filter (VRPF) is a particle filter that computes an approximation

$$\hat{p}(\theta_n | y_{1:n}) = \sum_{j=1}^N w_n^{(j)} \delta_{\theta_n^{(j)}}(\theta_n) \quad (14)$$

to the posterior (13). In its simplest bootstrap form, the piecewise deterministic VRPF is implemented along the lines of Algorithm 1, see [8], yielding Algorithm 3.

Algorithm 3 Variable rate particle filter

For all $1 \leq j \leq N$, initialize changepoint sequence as $\theta_0^{(j)} = (\tau_0^{(j)}, u_0^{(j)})$ with $\tau_0^{(j)} = 0$, $u_0^{(j)} \sim p_u(u_0)$, and x on $x(0)$.

for measurement time t_n **do**

for $j = 1 \dots N$ **do**

 Sample $\theta_n^{(j)} \sim p(\theta_n | \theta_{n-1}^{(j)})$ for each particle j .

 Using equation (10) to get $K_n^{(j)}$, we calculate the associated $x(t)$ at current measurement time t_n :

$$x_n^{(j)} = \phi(x(\tau_{K_n^{(j)}}), u_{K_n^{(j)}}, \tau_{K_n^{(j)}}, t_n). \quad (15)$$

 Evaluate and update weight $w_n^{(j)} = p(y_n | x_n^{(j)})w_{n-1}^{(j)}$ using (9), and normalize weights to get $\tilde{w}_n^{(j)}$.

 Resample if necessary.

end for

The seeked distribution is approximated by (14). This yields $\sum_j \tilde{w}_n^{(j)} x_n^{(j)}$ as an estimate for the state $x(t_n)$.

end for

Let us now comment on some of the steps. Regarding the sampling step, assume that we want to sample τ_{K_n} given $\tau_{K_{n-1}^{(j)}}$. This means there has been exactly $K_{n-1}^{(j)}$ changepoints before t_{n-1} . We can sample then the next changepoint time as $\tau_{K_{n-1}^{(j)}+1} \sim p_\tau(\tau - \tau_{K_{n-1}^{(j)}})$. Either $t_n \leq \tau_{K_{n-1}^{(j)}+1}$, then we let $\bar{\tau}$ denote the latter time and we let $K_n^{(j)} = K_{n-1}^{(j)}$. θ does not

need to be sampled anymore until the first measurement time t_k , $k > n$ to become larger than $\bar{\tau}$. Or $\tau_{K_{n-1}^{(j)}+1} < t_n$, then we integrate the flow (12) to obtain $x^{(j)}(\tau_{K_{n-1}^{(j)}+1})$, then sample $u_{K_{n-1}^{(j)}+1} \sim p_u(u | u_{K_{n-1}^{(j)}})$, and reiterate until the changepoint time becomes larger than t_n .

Regarding the weight update, particles are sampled using the prior and the weights are classically defined as $w_n^{(j)} = p(y_n | \theta_n^{(j)}, y_{1:n-1}) w_{n-1}^{(j)}$, see Algorithm 1. As $x(0)$ is known (and thus considered as a fixed parameter), given the sequence $\theta_n^{(j)}$ of changepoint times and parameters, the state corresponds to the integration of system (11), that is $x_n^{(j)}$ as defined by (15). Using that $x_n^{(j)}$ is a function $\mathcal{F}(\theta_n^{(j)})$ of the past changepoints we have $p(y_n | \theta_n^{(j)}, y_{1:n-1}) = p(y_n | \theta_n^{(j)}, \mathcal{F}(\theta_n^{(j)}), y_{1:n-1}) = p(y_n | x_n^{(j)})$ indeed, as y_n is independent of all the other variables conditionally on x_n .

C. Possible caveats in a partially deterministic context

It is well known that particle filters are not efficient for systems governed by deterministic dynamics, typically for fixed parameter inference. Actually, they are prone to divergence, see e.g., [9]. In the context of the VRPF, though, the state space is continually explored thanks to the jumps. But asymptotic convergence may necessitate a large number of particles to get full coverage. Another possible limitation of PF is in the presence of low measurement noise. Since $p(y_n | x_n)$ is then peaked, virtually all weights drop to zero after one measurement, leading in turn to estimation failure.

IV. A PARTICLE OBSERVER BASED APPROACH

In the present section, we consider the model of Section III-A, and propose a particle observer, that addresses some of the issues raised in Section III-C.

Up to a few subtleties related to marked point processes in place of discrete-time Markov processes, Algorithm 3 is close in essence to the standard Algorithm 1, since $\theta = (\tau, u)$ plays the role of hidden variable r , and x plays the role of auxiliary variable z , see notably the discussion after Algorithm 3. As a result, one may want to pursue the methodology of particle observers described at Section II-B for this application. Our idea is that, indeed, if the sequence of jump times τ_1, τ_2, \dots were known, then an asymptotic observer might be devised to estimate the remaining variables, that is, $x(t)$ and input $u(t)$. And under contraction properties of the observer's dynamics, approximations akin to (6)-(7) may then be used. To formalize this approach, we will need to split the observations in two. One type of observation will be used by the observer, and the other type by the particle filter to assess the weights. We will assume to have discrete the noisy measurements (8), but also some auxiliary continuous-time noise-free measurements of the form:

$$z(t) = \tilde{h}(x(t)). \quad (16)$$

Note that the two sources of information are assumed independent, and h and \tilde{h} may bear a quite different structure. For instance, $z(t)$, albeit noise free, may be of very small dimension compared to $x(t)$.

A. Considered problem and further assumptions

To recap, the considered system is

$$\frac{d}{dt}x(t) = f(x(t), t, u_K(t)) \quad (17)$$

$$y_n = h(x(t_n)) + V_n, \quad t_1 < t_2 < \dots \quad (18)$$

$$z(t) = \tilde{h}(x(t)) \quad (19)$$

where $u_K(t)$ are parameters jumping at random times. Besides, let us consider the following alternative system

$$\frac{d}{dt}x(t) = f(x(t), t, u(t)) \quad (20)$$

$$\frac{d}{dt}u(t) = 0 \quad (21)$$

$$z(t) = \tilde{h}(x(t)) \quad (22)$$

which corresponds to the noise free continuous-time system (17)-(19) between two jumps (as the parameters are fixed between jumps). Furthermore, assume we have at our disposal an asymptotic observer for the latter system of the form:

$$\frac{d}{dt}\hat{x} = f(\hat{x}, t, \hat{u}) + L_1(t, \hat{x}, \hat{u})(z - \tilde{h}(\hat{x})) \quad (23)$$

$$\frac{d}{dt}\hat{u} = L_2(t, \hat{x}, \hat{u})(z - \tilde{h}(\hat{x})) \quad (24)$$

having asymptotic convergence properties for arbitrary initial condition $(\hat{x}(0), \hat{u}(0))$. Typically, a desirable property would be exponential contraction [5], leading to

$$d(\hat{x}(t), x(t)) \leq d(\hat{x}(0), x(0))e^{-2\lambda t}. \quad (25)$$

On the one hand, this observer can readily be used on the original problem (17)-(19). This will lead to a solution that is unlikely to diverge (provided average time between consecutive jumps is larger than the convergence time of the observer), but in the transitory phases that follow changepoints, it will yield inaccurate estimates until $(\hat{x}(t), \hat{u}(t))$ has re-converged to $(x(t), u_K(t))$.

On the other hand, the variable rate particle filter cannot readily account for the information provided by $z(t)$ of eq. (19), since noise free measurements (as they appear in the theory of observers), or very small measurement noise, lead to failure of particle filters as explained in Section III-C. Besides, owing to the fact the dynamics are piecewise deterministic, the variable rate particle filter based solely on the information brought by y_n of eq. (18) will never be able to guarantee asymptotic convergence to the true parameters, even if the time between jumps is large (as the particles $u^{(j)}$ remain fixed between two jumps). For those reasons, we propose a hybrid strategy, combining both approaches.

B. Proposed hybrid strategy

Our idea is to devise a variable rate particle filter for system (17)-(18), and to use the side partial information (19) to help the sampled parameters $u^{(j)}$ (and thus the state) move towards the ‘‘true’’ parameters and state between consecutive changepoints. Crucially, this can be done by resorting to the notion of virtual systems [4], [10], i.e., by applying the variable rate methodology to an auxiliary virtual system.

Consider indeed the following *virtual system* associated to the observer

$$\frac{d}{dt}\chi = f(\chi, t, \mathbf{v}) + L_1(t, \chi, \mathbf{v})(z - \tilde{h}(\chi)) \quad (26)$$

$$\frac{d}{dt}\mathbf{v} = L_2(t, \chi, \mathbf{v})(z - \tilde{h}(\chi)) \quad (27)$$

and rewrite those dynamics as

$$\frac{d}{dt}(\chi, \mathbf{v}) = \hat{f}(\chi, \mathbf{v}, t, z) \quad (28)$$

Building upon the latter systems, we let $z(t) = \tilde{h}(x(t))$ as before, and we define auxiliary dynamics (recall that $\tau_1, \tau_2, \dots, u_1, u_2, \dots$ are in fact unknown) as

$$\begin{aligned} \tau_{k-1} < t < \tau_k &\Rightarrow \frac{d}{dt}(\chi, \mathbf{v}) = \hat{f}(\chi, \mathbf{v}, t, z), \\ t = \tau_k &\Rightarrow \mathbf{v}(t) = u_k, \end{aligned} \quad (29)$$

that is, we integrate (28) between two consecutive change-points, and, at each changepoint at time τ_k , the variable \mathbf{v} jumps to the corresponding parameter value u_k .

Letting $\tilde{\phi}$ denote the flow of (28), i.e., $\tilde{\phi}(\chi_1, \mathbf{v}_1, t_1, z^{(t_1, t_2)}, t_2)$ is the solution at t_2 to (28) being initialized on (χ_1, \mathbf{v}_1) at t_1 , and where we let $z^{(t_1, t_2)} = (\tilde{h}(x(s)))_{t_1 \leq s \leq t_2}$. Our auxiliary virtual system can be alternatively defined by

$$\begin{aligned} (\chi(t), \mathbf{v}(t)) &= \tilde{\phi}(\chi(\tau_{K_n}), \mathbf{v}(\tau_{K_n}), \tau_{K_n}, z^{(\tau_{K_n}, t)}, \tau_{K_n} < t \leq \tau_{K_{n+1}}) \\ y_n &= h(\chi(t_n)) + V_n, \quad t_1 < t_2 < \dots \end{aligned} \quad (30)$$

that is, we recover an auxiliary system to which the variable rate methodology of Section III-B can be applied. This yields the proposed Algorithm 4.

Algorithm 4 Hybrid variable rate particle observer

For all $1 \leq j \leq N$, initialize changepoint sequence as $\theta_0^{(j)} = (\tau_0^{(j)}, \mathbf{v}_0^{(j)})$ with $\tau_0^{(j)} = 0$, $\mathbf{v}_0^{(j)} \sim p_u(\mathbf{v}_0)$, and $\chi^{(j)}(0) = x(0)$.
for measurement time t_n **do**

for $j = 1 \dots N$ **do**

Sample $\theta_n^{(j)} \sim p(\theta_n | \theta_{n-1}^{(j)})$ for each particle j .

Using equation (10) to get $K_n^{(j)}$, we calculate the associated $\chi(t), \mathbf{v}(t)$ at current measurement time t_n :

$$(\chi_n^{(j)}, \mathbf{v}_n^{(j)}) = \tilde{\phi}(\chi^{(j)}(\tau_{K_n^{(j)}}), \mathbf{v}^{(j)}(\tau_{K_n^{(j)}}), \tau_{K_n^{(j)}}^j, z^{(\tau_{K_n^{(j)}}^j, t)}, t).$$

Evaluate weight $w_n^{(j)} = p(y_n | \chi_n^{(j)})w_{n-1}^{(j)} = l(y_n - j(\chi_n^{(j)}))w_{n-1}^{(j)}$ and normalize weights to get $\tilde{w}_n^{(j)}$.

Resample if necessary.

end for

$\sum_j w_n^{(j)} \chi_n^{(j)}$ is an estimate for $x(t_n)$.

end for

Using the flow of auxiliary system (26)-(27) in place of the flow of original system (17) to estimate the state of the original system through a VRPF, may be justified through the notion of virtual system, see Section IV-C below.

Remark 1: Particle filters have drawn criticism regarding their execution time. In this respect, our method is particularly efficient, owing to the use of observers as building blocks, each being very cheap to implement, as for each particle we need to integrate an ODE and weight computation is also cheap since $p(y | \chi) = l(y - h(\chi))$, see (9). This is in contrast to the standard RBPF where Kalman filters are commonly used to estimate the z variable, see Section II.

C. Discussion and formal results

System (17)-(18)-(19) is stochastic, since it is driven by the process $\omega \mapsto (\tau_k(\omega), u_k(\omega))$, and $\omega \mapsto V_n(\omega)$ is a random variable. Notably, $z(t)$ is a random process $\omega \mapsto z(t, \omega)$. Based on the fact that in the VRPF methodology the initial condition $x(0)$ is assumed to be known, the following result is noteworthy:

Lemma 1: Consider system (17)-(19) on the one hand, and system (29) on the other. Then, if $\chi(0) = x(0)$, we have $\chi(t) = x(t)$ at all times $t \geq 0$.

Proof: It suffices to note that, since $z(t) = \tilde{h}(x(t))$, we have $z(t) - \tilde{h}(x(t)) = 0$, and the pair $(x(t), u_{K(t)})$ is thus a solution to the stochastic differential equation (29). ■

This property lies at the heart of the use of virtual systems for observer design [4], [10]. In terms of our particle filter approach, this means Algorithm 4 may be viewed as a (bootstrap variable rate) particle filter for the original problem, given that initial conditions coincide $\chi(0) = x(0)$. Indeed, in this case both dynamical systems coincide as explained in the proof of Lemma 1. Normally, Algorithm 4 should use observations related to its own latent variables, i.e., of the form $\tilde{y}_n = h(\chi(t_n)) + V_n$, whereas it evaluates the weights based on the measurements $y_n = h(x(t_n)) + V_n$ instead. However, it is legitimate to use y_n as measurements for the auxiliary system since we know from Lemma 1 that at all times $\chi(t) = x(t)$ and thus $\tilde{y}_n = y_n$ indeed.

However, Lemma 1 does not hold under an arbitrary small error in measurements described by (19), or for $\chi(0) \neq x(0)$. But those are precisely the cases where our observer-based approach applies. Indeed, by assuming observer (23)-(24) to have convergence properties, *each* simulated particle $(\chi^{(j)}(t), \mathbf{v}^{(j)}(t))$ obtained through $\tilde{\phi}$ in Algorithm 4 tends to move closer to $(x(t), u_{K(t)})$ between jumps. As a result, particles tend to be less dispersed around $x(t)$ and to yield more precise and robust estimates, see Fig. 1 for an illustration.

D. Generalizations

Note that Eq. (17) being quite general, our approach can be used to enhance existing adaptive observers to track more rapidly variations in some unknown parameters u subject to abrupt changes.

Another possibility is to let u denote a model index, and then τ denotes the unknown times of switching.

Besides, the assumption (21) of constant $u(t)$ between jumps is not necessary. Between jumps, generic dynamical behavior for $u(t)$ may in fact be assumed as long as it is contracting or may be tracked by an observer. In this case $u_{K(t)}$ in (17) shall be replaced with $u_{K(t)}(t)$ and (21) by the ad hoc differential equation $\frac{d}{dt}u(t) = g(u(t), t)$.




Fig. 1. Graphical illustration of Algorithm 4 with $x(t) \in \mathbb{R}^2$. Plain stars indicate changing points at unknown times τ_1, τ_2, \dots , each being associated with unknown parameter u_i . The particle observer of Algorithm 4 keeps exploring the space of jump times τ and parameter u by constantly creating tentative couples (τ^j, v^j) with associated variable $\chi^j(t)$ referred to as particles, and makes use of a great number N of such particles. Owing to the assumed convergence properties of observer (23)-(24) for each particle $1 \leq j \leq N$, variable $\chi^j(t)$ tends to move towards true $x(t)$ even if the corresponding tentative couple (τ^j, v^j) is misidentified. The philosophy of particle filtering, that has been compared to Darwin’s “survival of the fittest”, consists in evaluating weights w^j for each particle indicating its fit to measured data. Particles with low weights are discarded, and particles with high weight are multiplied to populate the likely regions of the space, yielding quicker convergence. Here particle i corresponds to a tentative (τ^i, v^i) being close to actual (τ_{n+1}, u_{n+1}) and thus $\chi^i(t)$ rapidly converges to the true $x(t)$. Particle i is thus likely to be selected and multiplied, whereas particle j having less adequate (τ^j, v^j) might be suppressed.

V. PERSPECTIVES

Building upon the notion of virtual systems [4], we proposed a novel method which bridges the gap between nonlinear observer theory and particle filters for a broad class of practical applications.

We anticipate the method might prove useful in various contexts. One such instance is target tracking using one particular model of the nonlinear deterministic models with random piecewise constant inputs proposed in [7]. For a 3D target model inspired by [7], an invariant extended Kalman filter based on the work [14] was shown in [15] to yield good tracking performance, and might be used as a building block in a particle observer approach.

The method can also be applied to adaptive estimation of deterministic systems parameterized by an unknown vector selected from a finite set of known vectors. The switch is randomly connected to one of the vectors and may either remain there throughout the duration of the process, or change at random times associated to a prior temporal distribution. A typical example application would be Simultaneous Localization and Mapping (SLAM) with a vector accounting for the multi-hypotheses that stem from perceptual ambiguity, as e.g. in [16]. To address such questions one could combine the contracting observer of [17], based on the construction of simple synthetic measurements for SLAM, with the particle approach developed in this article. Another possible application is contact detection in robotics, see also Section IV-D. In the recent paper [18] for instance, leg state switching is based on a periodic phase-based state schedule, which provides a prior on times of change of a contact parameter s between stance ($s = 0$) and swing ($s = 1$).

Disturbance observers [18] with convergence properties for known s might be run in parallel using our approach.

The methodology of particle observers might also prove useful in the context of differentially positive systems [19], where the state is known to converge to a Perron-Frobenius curve: the use of particles might allow us to identify accurately where the state is located on the curve.

VI. ACKNOWLEDGEMENTS

We thank Prof. Godsill for useful discussions, while the first author was visiting the University of Cambridge as a fellow of Sidney Sussex College. The authors would also like to thank Marion Pilté and Frédéric Barbaresco from the company Thales Air Systems for fruitful discussions.

REFERENCES

- [1] A. Doucet, S. Godsill, and C. Andrieu, “On sequential Monte Carlo sampling methods for bayesian filtering,” *Statistics and Computing*, vol. 10, no. 3, pp. 197–208, 2000.
- [2] P. D. Moral, J. Noyer, G. Rigal, and G. Salut, “Particle resolution and non-linear signal processing with radar/sonar application,” *Traitement du Signal*, vol. 12, no. 4, pp. 287–302, 1995.
- [3] A. Doucet, N. De Freitas, K. Murphy, and S. Russell, “ Rao-Blackwellised particle filtering for dynamic bayesian networks,” in *Proceedings of the Sixteenth conference on Uncertainty in artificial intelligence*. Morgan Kaufmann Publishers Inc., 2000, pp. 176–183.
- [4] W. Wang and J.-J. E. Slotine, “On partial contraction analysis for coupled nonlinear oscillators,” *Biological cybernetics*, vol. 92, no. 1, pp. 38–53, 2005.
- [5] W. Lohmiller and J.-J. E. Slotine, “On contraction analysis for non-linear systems,” *Automatica*, vol. 34, no. 6, pp. 683–696, 1998.
- [6] S. Godsill and J. Vermaak, “Variable rate particle filters for tracking applications,” in *IEEE/SP 13th Workshop on Statistical Signal Processing, 2005*. IEEE, 2005, pp. 1280–1285.
- [7] P. Bunch and S. Godsill, “Dynamical models for tracking with the variable rate particle filter,” in *Information Fusion (FUSION), 2012 15th International Conference on*. IEEE, 2012, pp. 1769–1775.
- [8] —, “Particle smoothing algorithms for variable rate models,” *IEEE Trans. on Signal Processing*, vol. 61, no. 7, pp. 1663–1675, 2013.
- [9] N. Kantas, A. Doucet, S. S. Singh, J. Maciejowski, N. Chopin *et al.*, “On particle methods for parameter estimation in state-space models,” *Statistical science*, vol. 30, no. 3, pp. 328–351, 2015.
- [10] J. Jouffroy and J.-J. Slotine, “Methodological remarks on contraction theory,” in *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, vol. 3. IEEE, 2004, pp. 2537–2543.
- [11] F. Daum, “Nonlinear filters: beyond the Kalman filter,” *IEEE Aerospace and Electronic Systems Magazine*, vol. 20, no. 8, pp. 57–69, 2005.
- [12] T. Yang, P. G. Mehta, and S. P. Meyn, “Feedback particle filter,” *IEEE Trans. on Automatic Control*, vol. 58, no. 10, pp. 2465–2480, 2013.
- [13] S. Bonnabel and J.-J. Slotine, “Particle observers for contracting dynamical systems,” in *International Conference on Geometric Science of Information*. Springer, 2017, pp. 310–317.
- [14] A. Barrau and S. Bonnabel, “The invariant extended Kalman filter as a stable observer,” *IEEE Transactions on Automatic Control*, vol. 62, no. 4, pp. 1797–1812, 2017.
- [15] M. Pilté, S. Bonnabel, and F. Barbaresco, “Tracking the Frenet-Serret frame associated to a highly maneuvering target in 3D,” in *Decision and Control (CDC), IEEE 56th Conference on*, 2017, pp. 1969–1974.
- [16] I. Kanitscheider and I. Fiete, “Training recurrent networks to generate hypotheses about how the brain solves hard navigation problems,” in *NIPS*, 2017, pp. 4529–4538.
- [17] F. Tan, W. Lohmiller, and J.-J. Slotine, “Analytical SLAM without linearization,” *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1554–1578, 2017.
- [18] G. Bledt, P. M. Wensing, S. Ingersoll, and S. Kim, “Contact model fusion for event-based locomotion in unstructured terrains,” in *International Conf. on Robotics and Automation (ICRA)*, 2018, pp. 1–8.
- [19] F. Forni and R. Sepulchre, “Differentially positive systems,” *IEEE Trans. on Automatic Control*, vol. 61, no. 2, pp. 346–359, 2016.