On exact filters for continuous signals with discrete observations

MICHAEL A. KOURITZIN^{*} INSTITUTE FOR MATHEMATICS AND ITS APPLICATIONS 514 VINCENT HALL UNIVERSITY OF MINNESOTA MINNEAPOLIS, MINNESOTA 55455 (612) 624-1824, FAX: (612) 626-7370 KOURITZI@IMA.UMN.EDU.

May 28, 1996

ABSTRACT. Many filtering applications are characterized by continuous state dynamics $X_t = \int_0^t m(X_s) ds + \sigma W_t + \rho$, discrete observations $Y_k = Y_{t_k}$, and observation noise that is non-additive or non-Gaussian. In most such instances, neither exact finite-dimensional filters nor known on-line/off-line splitting methods apply. Thus, there is a pressing issue to determine how best to calculate the conditional density $\Pr\{X_t \in dz | Y_k, 1 \le k \le l\}$. Ideally, one would like an answer which avoids solving partial differential equations on-line. In this note, we show that a combination of convolution, scaling, and substitutions efficiently solves this problem under certain conditions. The most noteable aspects about our method are that it is extremely easy to use and that it assumes nothing about the observations other than the ability to construct $p_{Y_k|X_{t_k}}$, the conditional density of the kth observation given the current state.

1. INTRODUCTION

Traditionally, non-linear filtering theory for continuous-time state variables consisted of elegant mathematical solutions which were too abstract and too computationally expensive for use in real tracking and prediction problems. How should one implement continuous-state filters with continuous or discrete observations in real time for practical problems with either non-linear equations where the Kalman filter does not supply the optimal state estimate or with non-Gaussian noise where the Kalman filter does not provide enough information to reconstruct conditional densities? Until recently, this imposing question invariably led system engineers to mathematically unjustified adaptations of the Kalman filter such as the extended Kalman filter. The

^{*}The author is grateful to Lockheed Martin Tactical Defense Systems-Eagan and the National Science Foundation for their support.

resulting filters are often unsatisfactory and our participation in the current comprehensive effort to develop real-time implementable non-linear filters is well justified.

Since many process measurements are not made continuously but rather at discrete instances and computers are digitals instruments it is often natural to work with discrete observations. Suppose measurement Y_j is taken at time t_j , j = 1, 2, 3, ...Then, the basic problem of continuous/discrete non-linear filtering is to obtain the distribution of state X_t conditioned on the past and current observations $\{Y_j, 0 \le t_j \le t\}$. The state is modeled by a \Re^d -valued Itô equation

$$dX_t = m_t(X_t) dt + \sigma_t(X_t) dW_t, \quad X_0 = \rho$$
(1)

driven by some \Re^e -valued standard Brownian motion W_t . The observations are described by a conditional density

$$p_{Y_k|X_k}(\cdot|\cdot)$$
 defined by $\int_A p_{Y_k|X_k}(z|X_k)dz = \Pr\left\{Y_k \in A \mid \sigma(X_k)\right\}$ a.s. (2)

and usually calculated from a mapping $Y_k = h_{t_k}(X_{t_k}, v_k)$, where each h_{t_k} is a function from $\Re^d \times \Re^l \to \Re^n$ and $\{v_k, k = 1, 2, 3, ...\}$ is a sequence of independent \Re^l -valued random variables independent of W and ρ . Under certain regularity conditions the conditional distribution for X_t given $\{Y_j, 0 \leq t_j \leq t\}$ admits a density which can be determined iteratively by solving Kolmogorov's forward equation (also called the Fokker-Planck equation) for (1) with random initial conditions between samples and using a Bayes' formula update with (2) at sample times (see Ch. 6 of Jazwinski [10]).

Efforts to implement continuous-state non-linear filters with either continuous or discrete observations have been focused in three directions: (i) Obtaining exact "finite-dimensional filters" which can be implemented on-line as a finite number of ordinary differential equations (e.g. Benesh [1], Ocone et. al. [16], and Daum [6], [7]), (ii) decomposing the non-linear filtering problem into time-consuming off-line and computationally efficient on-line components (e.g. Bensoussan et. al. [2], Budhiraja and Kallianpur [3], Mikulevicius and Rozovskii [14], and Lototsky, Mikulevicius and Rozovskii [13]), or (iii) solving Kolmogorov's second order partial differential equation on-line using multi-grid or other numerically-sophisticated methods (Cai et. al. [4], Mirković [15]. Whereas the first method has only been proven for a limited number of signal models

$$dX_t = m_t(X_t) dt + \sigma_t dW_t \tag{3}$$

where σ does not depend on X_t and m satisfies certain conditions, the second method invariably imposes its own restrictive conditions, involves approximations, and entails considerable apriori off-line work. Moreover, all of these works from the first two methods require limiting conditions on the observation noise (e.g. Gaussian and additive) which is a severe limitation in certain problems. Indeed, the very method of looking for finite dimensional filters forces restrictive conditions on the observation model. The third method is general but requires long, sophisticated computer algorithms which are well-tuned for a particular filtering problem. Moreover, to obtain similar accuracy as is possible with the former two methods, one must often utilize a large number of grid points and/or high order difference approximations for the derivatives making this method quite slow even when multi-grid or other efficient algorithms are used.

Our present work is in the direction of implementing exact (infinite dimensional) continuous/discrete filters under more general conditions on the observations. As will be shown in later sections, our method is easy to implement, does not require any apriori off-line work, and avoids the usual restrictive conditions on the observations by not requiring that there be finite-dimensional sufficient statistics.

To motivate our method consider the case where $X_t \equiv W_t + \rho$ is just a \Re^d valued standard Brownian motion added to some (possibly non-Gaussian) random variable ρ which is independent of W. Moreover, suppose that $p_{Y_k|X_k}(\cdot|\cdot)$ denotes the conditional density of Y_k given X_{t_k} and $p_{X_j|\mathcal{Y}_k}(\cdot)$ is the conditional density of X_{t_j} given $\mathcal{Y}_k \doteq \sigma\{Y_i, i \leq k\}$. Then, a fundamental solution to Kolmogorov's forward equation (which gives the density for X's transition probability function) is just the heat kernel

$$\Phi_{t-\tau}(x-\xi) = \Phi(x,t;\xi,\tau) = \frac{\exp\left\{-|x-\xi|^2/2(t-\tau)\right\}}{(2\pi(t-\tau))^{d/2}}$$
(4)

and the continuous/discrete filtering problem is solved as: (I) $p_{X_0|y_0} = p_{\rho}$, the density of ρ ; and then for all k = 0, 1, ...

(II)
$$p_{X_{k+1}|\mathcal{Y}_k}(x) = \int_{\Re^d} \Phi(x, t_{k+1}; \xi, t_k) p_{X_k|\mathcal{Y}_k}(\xi) d\xi = \Phi_{\delta t_k} * p_{X_k|\mathcal{Y}_k}(x)$$
 (5)

where $\delta t_k \doteq t_{k+1} - t_k$; and

(III)
$$p_{X_{k+1}|\mathcal{Y}_{k+1}}(x) = \frac{p_{Y_{k+1}|X_{k+1}}(Y_{k+1}|x)p_{X_{k+1}|\mathcal{Y}_{k}}(x)}{\int_{\Re^{d}} p_{Y_{k+1}|X_{k+1}}(Y_{k+1}|\xi)p_{X_{k+1}|\mathcal{Y}_{k}}(\xi)\,d\xi}.$$
 (6)

In the general case where X solves (1) or (3) steps (I) and (III) remain unchanged but Φ must be replaced with an alternate fundamental solution Γ and convolution in (II) is often not possible. Therefore, to solve (II) directly one would have to store Γ for all $x \in \Re^d$ (or subdomain where X lives) and ξ "close to x". However, when convolution is possible one only has to store a function of $x - \xi$ which, being a fundamental solution of a parabolic equation, has rapidly decreasing tails. Our goal is to determine conditions under which this convolution method holds. It is often the case, for example in position velocity models, that only some of the states are randomly excited. Hence, we assume that $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P, X, W)$ with $X_t \in \Re^d$ and $W_t \in \Re^e$ $(e \leq d)$ is the weakly unique weak solution to

$$\begin{bmatrix} dX_t^1 \\ dX_t^2 \end{bmatrix} = \begin{bmatrix} m(X_t^1) dt + \sigma dW_t \\ (\alpha + \mu X_t^1) dt \end{bmatrix} \doteq \widetilde{m}(X_t) dt + \widetilde{\sigma} dW_t, \quad X_0 = \rho;$$
(7)

4

and show under conditions on m similar to Benesh's [1] that the convolution method applies and, moreover, one can always (at every iteration) perform the convolution with respect to the standard normal density. This means that (i) the fundamental solution in (II) does not have to be stored at all and (ii) standard (fixed-grid) Fast Fourier transform (FFT) routines, which can be implemented on a computer with about $2n \log_2 n$ operations, are ideally suitable. It is possible to devise O(n) adaptations of our methods. However, our only sources of error are in approximating Fourier transforms by FFTs and it is felt that the loss of accuracy incurred by additional approximation would usually far outweigh the very limited asymptotic time advantage gained by such an adaptation. Finally, it is well known that slight modifications to filtering algorithms like ours provide prediction and smoothing estimates for X_t .

In the following section we give the remainder of our notation, state the conditions under which this "standard normal convolution" method is known to hold, state a supportive mathematical proposition, and outline an algorithm for this method. In Section 3, we prove the proposition of Section 2, in Section 4 we discuss a simple example which has been implemented and tested in both Mathematica and C, and in Section 5 we briefly mention the more technical details in implementing our algorithm.

2. **Result and Algorithm**

Throughout the remainder of this note we let $|\cdot|$ denote Euclidean distance, assume that X and W are as defined in (7), let $\sigma \in \Re^{e \times e}$, $\alpha \in \Re^{d-e}$, and $\mu \in \Re^{(d-e) \times e}$ be constants with σ non-singular and μ having full row rank, and fix S > 0 large enough that our time between observations $t_{k+1} - t_k \leq S$ for all k. For Proposition 1 (to follow), we also assume:

C0 For each $x \in \Re^e$, $y \in \Re^{d-e}$ there exists a weakly unique, weak solution on [0, S] to the Itô equation

$$Y_t^{x,y} = \begin{bmatrix} x \\ y \end{bmatrix} - \int_0^t \widetilde{m}(Y_s^{x,y}) \, ds + \widetilde{\sigma} W_t. \tag{8}$$

C1 *m* is continuously differentiable and satisfies a linear growth condition i.e. $\partial_i m \doteq \frac{\partial_i m}{\partial x_i}$ is a continuous function on \Re^e for i = 1, 2, ..., e and $|m(x)|^2 \leq K \cdot (1 + |x|^2) \quad \forall x \in \Re^e$.

For the next two conditions we let $l(x) \doteq \sigma^{-1}m(\sigma x) \ \forall x \in \Re^e$ and some K > 0.

- **C2** There are constants k, r, q and Q with $Q = Q^T \ge 0$ such that $\frac{|l(\sigma^{-1}x)|^2 \sum_{j=1}^e \partial_j l_j(\sigma^{-1}x)}{2} + \sum_{i=1}^d \partial_i \widetilde{m}_i(x) = x^T Q x + r^T x + k \text{ and } 2Qq = -r.$
- **C3** $\frac{\partial l_j}{\partial x_k} = \frac{\partial l_k}{\partial x_j}$ holds for all $1 \leq j < k \leq e, x \in \Re^e$ which is necessary and sufficient for existence of a function $L : \Re^e \to \Re$, unique to within an additive constant, such that $l_j = \partial L / \partial x_j \ \forall 1 \leq j \leq e$.
- C4 For every $g \in C_0(\Re^d)$, the space of continuous functions with compact support, the Cauchy initial data problem on $[0, S] \times \Re^d$ for the Kolmogorov forward equation corresponding to (7)

$$\partial_t q_t = \sum_{i,j=1}^e A_{i,j} \partial_{i,j}^2 q_t - \sum_{i=1}^d \widetilde{m}_i \, \partial_i q_t - \sum_{i=1}^d \partial_i \widetilde{m}_i \, q_t, \ q_0 = g; \ A \doteq \frac{\sigma \sigma^T}{2} \tag{9}$$

has a solution q_t which satisfies the polynomial growth condition

$$\max_{0 \le t \le S} |q(t, x)| \le M_g (1 + |x|^{\mu_g}), \tag{10}$$

where $M_g > 0$ and $\mu_g \ge 1$ are constants which depend on g.

Remark 1. It is well known (see Benes [1] and Daum [5]) that there are interesting examples of filters for non-linear m satisfying (C2).

Remark 2. Suppose *m* is affine i.e. $\widetilde{m}(x) \equiv \widetilde{m}_0 + \widetilde{m}_1 x$. Then, under Condition (C4) there is a Fundamental solution to (9) which has the form

$$\Gamma(x, y, t; \xi, \theta, \tau) = \exp\left\{-tr(\widetilde{m}_1)(t-\tau)\right\} \Phi_{\Sigma_{t-\tau}}(\nu_{t-\tau}(x, y) - (\xi, \theta)),$$
(11)

where Φ_{Σ} is as in (17), and $\nu_{t-\tau}(x, y)$ and $\Sigma_{t-\tau}$ are the mean and covariance of

$$Y_t^{x,y} = \begin{pmatrix} x \\ y \end{pmatrix} - \int_{\tau}^t \widetilde{m}(Y_s^{x,y}) \, ds + \widetilde{\sigma} \begin{pmatrix} W_{t-\tau} \\ 0 \end{pmatrix}. \tag{12}$$

(This is easily derived and is an immediate consequence of (22), (25), and (26) of the proof in Section 3.) Hence, our convolution method and the algorithm given below

apply even if Condition (C3) is not satisfied! This affine case was first presented at Lockheed Martin Tactical Defense Systems-Eagan (formerly LORAL Defense-Eagan) by the author on November 9, 1995. A project to consider its use in altitude tracking for air traffic control was immediately thereafter initiated. The results have been encouraging and are being reported in [12].

Remark 3. If L is not known in closed form it may of course be solved for at any x by evaluating the line integral $L(x) = \int_{\gamma} \sum_{i=1}^{e} l_i dx_i$ along any path γ connecting the origin and x.

Remark 4. Condition (C4) is a mild assumption; there are known constraints on m that ensure (C4) will hold (see e.g. Friedman [9] p. 147, [8] p. 140, and Sonin [17]).

Now, under the above conditions, the classical filtering theory applies and the density for X_{t_k} given $\sigma\{Y_i, i \leq k\}$ is given by (I-III) of the introduction. The only remaining question is how are we going to ensure that we can factor its fundamental solution in such a manner that convolution in (II) will take place. For notational convenience, we define $R \doteq$ the symmetric, positive semi-definite square root of $\sigma^T Q \sigma$, $R' \doteq \sigma R \sigma^{-1}$, $\overline{R} \doteq (\sigma^{-1})^T R \sigma^{-1}$, $M(x) \doteq L(\sigma^{-1}x)$, $k' \doteq k + |q|^2$, and semi-group

$$U(s) \doteq \exp\left\{s \begin{bmatrix} -2R' & 0\\ -\mu & 0 \end{bmatrix}\right\}.$$
(13)

Proposition 1. Under Conditions (C0-C4), there is a fundamental solution to (9) which is given by

$$\Gamma(x, y, t; \xi, \theta, \tau) = \Delta(t - \tau) \Upsilon(x) \Phi_{\Sigma_{t-\tau}}(\nu_{t-\tau}(x, y) - (\xi, \theta)) \Upsilon^{-1}(\xi)$$
(14)

for all $(x, y), (\xi, \theta) \in \Re^d$ and $0 \le t - \tau \le S$, where

$$\Delta(s) \doteq \exp\left\{-(k'+tr(R))s\right\}, \quad \Upsilon(z) = \exp\left\{M(z) - (q+z)^T \overline{R}(q+z)\right\}, \quad (15)$$

$$\nu_s(x,y) \doteq U(s) \begin{bmatrix} x+q\\ y \end{bmatrix} - \begin{bmatrix} q\\ \alpha s+\mu qs \end{bmatrix},\tag{16}$$

and Φ_{Σ} is the zero mean Gaussian kernel

$$\Phi_{\Sigma}(z) = \frac{1}{((2\pi)^d \det(\Sigma))^{1/2}} \exp\left\{-1/2 \, z^T \Sigma^{-1} z\right\}$$
(17)

with covariance

$$\Sigma_s \doteq \int_0^s U(u) \tilde{\sigma} \tilde{\sigma}^T U^T(u) du.$$
(18)

The decomposition (14-18) may at first seem complicated and difficult to use. However, it is already convenient for computer implementation and in practise it will simplify further. For instance, when we use Proposition 1 in our algorithm $t - \tau$ will normally just be a fixed time step $t_{k+1} - t_k$ so $\Delta(t - \tau)$, $\Sigma_{t-\tau}$, and $U(t - \tau)$ become fixed matrices.

Remark 5. In the linear case mentioned in Remark 2, Σ_s and $\nu_s(x, y)$ will have a similar form and the following algorithm will still apply.

2.1. ALGORITHM. In this subsection, we expand the algorithm (I-III) given in the introduction to include systems (7) having fundamental solutions to Kolmogorov's forward equation which have the form given in Proposition 1. We define P_s to be the symmetric, positive definite matrix

$$P_s \doteq \Sigma_s^{-1/2}$$
 and let $\mathcal{N}(z) \doteq \exp\left\{-|z|^2/2\right\}/(2\pi)^{d/2}$. (19)

Moreover, to ease the notation we make the following definitions

$$\delta t_j \doteq t_{j+1} - t_j, \ p_j \doteq p_{X_j|\mathcal{Y}_j}, \ p_{j+1|j} \doteq p_{X_{j+1}|\mathcal{Y}_j}, \ \Phi_{\delta t_j} \doteq \Phi_{\Sigma_{\delta t_j}}.$$
 (20)

Then, the basic algorithm is as follows

(a) $p_0(x,y) = p_\rho(x,y)$, the density of ρ Do j = 0, 1, 2, ...

(b)
$$r_j(x,y) \doteq \Upsilon^{-1}(x) p_j(x,y)$$

(c) $\lambda_j = \int_{\Re^d} {x \brack y} r_j(x,y) dx dy / \int_{\Re^d} r_j(x,y) dx dy$
(d) $\beta_j({w \brack y}) \doteq \Phi_{\delta t_j} * r_j({w \brack y}) = \int_{R^d} \mathcal{N}(P_{\delta t_j}({w \brack y} + \lambda_j) - z) r_j(P_{\delta t_j}^{-1}z - \lambda_j) dz$
(e) $p_{j+1|j}(x,y) = \Delta(\delta t_j) \Upsilon(x) \beta_j(U(\delta t_j) {x+q \brack y} - {q \brack (\alpha + \mu q) \delta t_j})$
(f) Evaluate "prediction" statistics $E[\phi(X_{t_{j+1}})|\mathcal{Y}_j] = \int \phi(x,y) p_{j+1|j}(x,y) dx dy$
(g) Wait for next observation $Y_{j+1}(\omega)$
(h) $q_{j+1}(x,y) \doteq p_{Y_{j+1}|X_{j+1}}(Y_{j+1}|x,y) p_{j+1|j}(x,y)$
(i) $p_{j+1}(x,y) = q_{j+1}(x,y) / \int_{\Re^d} q_{j+1}(\xi,\theta) d\xi d\theta$

(j) Evaluate "tracking" statistics $E[\phi(X_{t_{j+1}})|\mathcal{Y}_{j+1}] = \int \phi(x,y)p_{j+1}(x,y) dxdy$.

The second equality in (d) is the key to making the algorithm practical and efficient. Defining $\tilde{r}_j(z) \doteq r_j(P_{\delta t_j}^{-1}z - \lambda_j)$ to be the centered and scaled version of r_j , one finds that this equation is efficiently evaluated as

$$\mathcal{F}^{-1}\left[\exp\left\{-\left|u\right|^{2}/2\right\}\exp\left\{-i\,u^{T}P_{\delta t_{j}}\lambda_{j}\right\}\,\mathcal{F}\left[\widetilde{r}_{j}\right]\left(u\right)\right]\left(P_{\delta t_{j}}\begin{bmatrix}w\\v\end{bmatrix}\right),\tag{21}$$

where \mathcal{F} denotes Fourier transform. On a computer FFTs would be used in place of the Fourier transforms. However, (21) suggests that: (i) the inverse Fourier transform is well approximated by an integral over a bounded domain centered at the origin, (ii) a fixed, uniform grid centered at the origin in the frequency domain is appropriate to evaluate the inverse FFT, and (iii) the scaling and size of grid can be determined by examining $\exp\left\{-|u|^2/2\right\}$. More technical details about implementing the algorithm are given in Section 5.

3. Proof of Proposition

In this section, we give the proof of Proposition 1 which was partially motivated by Benesh [1]. Throughout this proof we will let $\stackrel{D}{=}$ mean equal in distribution and P[Z] $(\tilde{P}[Z])$ denote the expectation of Z with respect to probability measure $P(\tilde{P})$.

Let us fix $0 \leq T \leq S$, $x \in \Re^e$, $y \in \Re^{d-e}$ and rewrite (9) subject to some continuous initial data g with compact support on \Re^d as

$$\partial_t q_t + cq_t = \sum_{i,j=1}^e A_{i,j} \partial_{i,j}^2 q_t - \sum_{i=1}^d \widetilde{m}_i \partial_i q_t, \quad q_0 = g; \quad c = \sum_{i=1}^d \partial_i \widetilde{m}_i.$$
(22)

Remark 6. Our linear growth assumption on m is necessary to ensure that the solution to (8) does not blow up. It is explicitly required for our use of Feynman-Kac's formula (see (5.7.2) of Karatzas and Shreve [11]) and Girsanov's theorem. Indeed, to show that Girsanov's theorem applies in (28) below, we use Fernique's lemma and Condition (C1) to find an $\varepsilon > 0$ respectively a $K'_x > 0$ such that

$$E[\exp\{\varepsilon |N|^2\}] < \infty \text{ and } |l(W_{s\wedge T}^x)|^2 \le K'_x(1+|W_{s\wedge T}|^2) \quad \forall \ s \ge 0,$$
(23)

where N is a standard normal random vector. Then, letting $t_n = \frac{2\epsilon}{K'_x T}n$ for n = 0, 1, 2, ... and using Jensen's inequality with measure $ds/(t_n - t_{n-1})$, we find

$$E[\exp\{\frac{1}{2}\int_{t_{n-1}}^{t_n} |l(W_{s\wedge T}^x)|^2 \, ds\}] \le \exp\{\frac{\varepsilon}{T}\} E[\exp\{\frac{t_n - t_{n-1}}{2}K_x'T \, |N|^2\}] < \infty$$
(24)

for all n = 1, 2, 3... and the Novikov condition (Corollary 3.5.14 of Karatzas and Shreve [11]) is satisfied so Girsanov's theorem applies.

Now, Condition (C4), an application of Feynman-Kac's formula (see e.g. p. 366 of Karatzas and Shreve [11] with a time reversal) and (7) yield

$$q(T, x, y) = P\left[g(Y_T^{x, y}) \exp\left\{-\int_0^T c\left(Y_s^{x, y}\right) ds\right\}\right]$$
(25)

where $B_t \doteq \sigma W_t$ and $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P, Y_t^{x,y}, W)$ is a weak solution on [0, T] to

$$Y_t^{x,y} = \begin{pmatrix} x \\ y \end{pmatrix} - \int_0^t \widetilde{m}(Y_s^{x,y}) \, ds + \begin{pmatrix} B_t \\ 0 \end{pmatrix}.$$
(26)

9

Now; we recall $l(x) \doteq \sigma^{-1}m(\sigma x)$; define $B_s^x \doteq x + B_s$,

$$W_s^x \doteq \sigma^{-1} B_s^x, \quad \frac{d\tilde{P}}{dP} \doteq \exp\left\{-\sum_{j=1}^e \int_0^T l_j(W_s^x) \, dW_s^i - \frac{1}{2} \int_0^T \sum_{j=1}^e |l_j(W_s^x)|^2 \, ds\right\}; \quad (27)$$

and note from Girsanov's theorem that

$$\{B_s^x, 0 \le s \le T\}_P \stackrel{D}{=} \left\{B_s^x + \int_0^s m(B_s^x) \, ds, \, 0 \le s \le T\right\}_{\widetilde{P}}.$$
(28)

Hence, by (7) and the weak uniqueness of (26) we have that

$$\{Y_s^{x,y}, 0 \le s \le T\}_P \stackrel{D}{=} \left\{ Z_s^{x,y} \doteq \begin{pmatrix} B_s^x \\ y - \alpha s - \mu \int_0^s B_u^x du \end{pmatrix}, 0 \le s \le T \right\}_{\widetilde{P}},$$
(29)

and from $M(x) \doteq L(\sigma^{-1}x)$, (27), Condition (C3) and Itô's formula that

$$M(B_T^x) - M(x) = \sum_{j=1}^e \int_0^T l_j(W_s^x) \, dW_s^j + \sum_{j=1}^e \int_0^T \frac{\partial_j l_j(W_s^x)}{2} ds.$$
(30)

Therefore, it follows by (25), (29), (22), Condition (C2), (27), and (30) that

$$q(T, x, y) = \widetilde{P}\left[g(Z_T^{x,y})\exp\left\{-\int_0^T c\left(Z_s^{x,y}\right)ds\right\}\frac{d\widetilde{P}}{dP}\frac{dP}{d\widetilde{P}}\right]$$

$$= \exp\left\{M(x)\right\}P[g(Z_T^{x,y})\exp\{-M(B_T^x) - \int_0^T B_s^{xT}QB_s^x + r^TB_s^x + kds\}]$$
(31)

$$= \exp \{M(x)\} P[g(Z_T^{x,y}) \exp\{-M(B_T^x) - \int_0^T B_s^{zT} Q B_s^z + k' ds\}],$$

where $z \doteq x + q$, 2Qq = -r, and $k' \doteq k + |q|^2$. Now; we let R be the symmetric, positive semi-definite square root of $\sigma^T Q \sigma$; define the processes

$$W_t^z = \sigma^{-1} B_t^z, \quad \zeta_t^z = -2 \int_0^t R \zeta_s^z ds + W_t^z,$$
 (32)

$$\zeta_t^x = -2\int_0^t R\zeta_s^x ds + W_t^x, \quad \xi_t^{x,y} \doteq \begin{pmatrix} \sigma\zeta_t^x \\ y - \alpha t - \mu \int_0^t \sigma\zeta_s^x ds \end{pmatrix}$$
(33)

for all $t \in [0, T]$; and redefine

$$\frac{d\tilde{P}}{dP} \doteq \exp\left\{-2\sum_{i=1}^{e} \int_{0}^{T} \left(RW_{s}^{z}\right)^{i} dW_{s}^{i} - \int_{0}^{T} (B_{s}^{z})^{T} QB_{s}^{z} ds\right\}.$$
(34)

Then, it follows from Itô's formula with $\overline{R} \doteq (\sigma^{-1})^T R \sigma^{-1}$ that

$$2\sum_{i=1}^{e} \int_{0}^{T} \left(RW_{s}^{z}\right)^{i} dW_{s}^{i} = \left(B_{T}^{z}\right)^{T} \overline{R} B_{T}^{z} - z^{T} \overline{R} z - T \ tr(R)$$

$$(35)$$

and from Girsanov's and weak uniqueness of the linear stochastic differential equation in (32) that $\{\sigma\zeta_s^z, 0 \le s \le T\}_P \stackrel{D}{=} \{B_s^z, 0 \le s \le T\}_{\widetilde{P}}$ so by (13)

$$\left\{\eta_s^{x,y} \doteq \xi_s^{x,y} + U(s) \begin{bmatrix} q\\ 0 \end{bmatrix} - \begin{bmatrix} q\\ \mu qs \end{bmatrix}, \ 0 \le s \le T \right\}_P \stackrel{D}{=} \{Z_s^{x,y}, \ 0 \le s \le T \}_{\widetilde{P}} . \tag{36}$$

Therefore, by (34), (31), and (35) it follows that

$$q(T, x, y) = \Psi(T, x) P\left[g(\eta_T^{x, y}) \exp\left\{\zeta_T^{z^T} \sigma^T \overline{R} \sigma \zeta_T^z - M(\sigma \zeta_T^x + \exp\left\{-2R'T\right\} q - q)\right\}\right]$$
(37)

where

$$\Psi(T,x) = \exp\left\{M(x) - (x+q)^T \overline{R}(x+q) - k'T - tr(R)T\right\}.$$
(38)

The proposition follows by the observation

$$\xi_t^{x,y} = \int_0^t \begin{bmatrix} -2R' & 0\\ -\mu & 0 \end{bmatrix} \xi_s^{x,y} + \begin{bmatrix} 0\\ -\alpha \end{bmatrix} ds + \begin{bmatrix} \sigma W_t\\ 0 \end{bmatrix} + \begin{bmatrix} x\\ y \end{bmatrix}.$$
(39)

4. EXAMPLE

To illustrate use of our algorithm on a simple, practical¹ yet unsolved problem (meaning the Kalman filter, other finite-dimensional exact filters, and on-line/offline splitting methods do not apply) it suffices to consider the following simple "position/velocity" example with non-linear observations corrupted by non-Gaussian noise. Let h_t denote the position of some object and suppose we would like to track h_t . Over

¹This problem is a simplified version of an altitude tracking model used in air traffic control. Indeed, our algorithm is also being used on the more realistic model (see [12]) but we do not wish to add technicalities to this exposition so we will stick with the idealized version here.

short time periods one often assumes that velocity $\frac{dh_t}{dt}$ is just a Brownian motion and that h_t itself is not randomly excited. The observations are rounded to the nearest quantization level (increments of 40 meters here) and subject to non-Gaussian noise. Thus, our model is given by

$$dX_t = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} X_t dt + \begin{bmatrix} \sigma \\ 0 \end{bmatrix} dW_t, \ X_0 = \rho$$

$$y_k = 40 * [h_{t_k}/40 + .5] + v_k, \ X_t = \begin{bmatrix} \frac{dh_t}{dt} \\ h_t \end{bmatrix},$$

$$(40)$$

with $\{v_k, k = 1, 2, ...\}$ being a zero-mean i.i.d. sequence with marginal distribution p_{v_k} and $[\cdot]$ denoting truncation to the nearest meter. In this case, it is easy to determine that

$$p_{Y_{k+1}|X_{k+1}}(Y_{k+1}|x,y) = p_{v_k}(Y_{k+1} - 40 * [y/40 + .5]).$$
(41)

For comparison purposes, we also constructed an extended Kalman filter by linearizing the observation function, which entails treating the rounding as a second, independent noise source, and then replacing $\{v_k, k = 1, 2, ...\}$ plus roundoff noise with a zero-mean, i.i.d. Gaussian sequence. In the case where the variance of v_k plus roundoff noise exists the Gaussian random variables are chosen by matching this variance. In the case where p_{v_k} is a Cauchy or other heavy-tailed distribution the Gaussian random variables for the extended Kalman filter were chosen by trying several different variances experimentally and keeping the one that gave the best results.

Our algorithm has been thoroughly tested in Mathematica and C on this example under a variety of combinations of different σ , p_{v_k} , ρ , and number of grid points. The following two general conclusion have been made:

- 1. As the tails of p_{v_k} increase or σ decreases (i.e. the signal to noise ratio decreases) one has to represent the tails of $p_{j+1|j}$ better to maintain the same level of accuracy in the estimates. This is due to the increased probability that the observation will be in the tails of $p_{j+1|j}$ and means that a larger number of grid points (64, 128 or higher per dimension) should be used when the signal to noise ratio is very low.
- 2. When a reasonable number of grid points are used the non-linear filter always outperforms the extended Kalman filter in terms of root mean square (RMS) error from h_t over long enough time periods. As the number of grid points increases the performance of both filters increases but the ratio by which the non-linear filter outperforms the extended Kalman filter also increases.

For this note, we illustrate the operation of our algorithm with our example under the two cases of non-existent and heavy tails for v_k :

Table 1. Cases considered						
p_{v_k}	σ	No. Grid pts.	$p_ ho$	δt_j		
Uniform[-20,20]	4.1	$50\! imes\!50$	$\delta((x,y)-(2000,0))$	$4.8\mathrm{s}$		
$\operatorname{Cauchy}(0,15)$	5.3	$50\! imes\!50$	$\delta((x,y) - (2000, -5))$	$4.8\mathrm{s}$		

Table 1: Cases Considered

i.e. we started the system at 2000 meters and 0 or -5 meters/s and obtained a new observation every 4.8 s. In order to produce nice graphics we decided to use Mathematica which limited us to about 50×50 grid points on a SUN Sparcstation 10 and means that there will still be a noticeable difference between our discretized densities and the unrepresentable exact continuous densities. This will of course adversely affect our estimates and somewhat better performance occurs when our algorithm is run under C with a larger number of grid points. (Grid sizes up to 256×256 were easily used on 486 PC with C.)

Figures 1 and 2 at the end of this note illustrate use of our non-linear filter and an extended Kalman filter in mid-operation under the two cases documented above. Because it is not economical to represent $p_{j+1|j}$ and p_{j+1} as explicit functions of (x, y)on a computer (see Paragraph 4 of Section 5), the axes of graphs (a) and (b) of both Figures 1 and 2 are not position and velocity but rather an affine function of both. Similarly, the axes of (c) and (d) are not quite position and velocity but, in this case, the lines of constant position are readily discernible by the edges of the truncated densities.

In the uniform density prediction case, one can easily see that the top of the extended Kalman filter density Figure 1 (b) is slightly broader than that of the nonlinear filter (a) reflecting the fact that the rounding had to be modeled as additional noise in the Kalman filter case. Moreover, one can see that the base of the nonlinear filter density is a lot narrower in some directions reflecting (proper treatment of rounding and) the fact that the tails of this density decrease faster than Gaussian tails. This rapid decrease is a result of the fact that this density is derived from the convolution of a Gaussian with the previous tracking densities based upon the same observation (see graph (c)). The next step tracking densities based upon the same observation Y_{j+1} are given in graphs (c) and (d). Based upon these densities it is reasonable to expect the extended Kalman filter to give poorer estimates especially for position and to show "less confidence" in these estimates. When statistics other than the conditional mean are required it is obvious that the non-linear filter will do a far better job. Data for a typical run through ten time steps is given below:

Step	h_{t_k}	${y}_k$	$h_{t_k}^N$	$e_{t_k}^N \doteq h_{t_k} - h_{t_k}^N $	$h_{t_k}^K$	$e_{t_k}^K$
1	2043.55	2029.58	2036.57	6.98	2024.37	19.18
2	2098.23	2066.14	2081.2	17.03	2067.35	30.88
3	2224.34	2244.76	2238.59	14.25	2241.8	17.46
4	2330.82	2301.46	2319.8	11.02	2302.07	28.75
5	2500.43	2504.56	2518.42	17.99	2501.75	1.32
6	2667.98	2666.35	2680.99	13.01	2667.56	0.42
7	2805.93	2819.01	2800.38	5.55	2818.94	13.01
8	2926.39	2939.68	2919.96	6.43	2940.74	14.35
9	3108.75	3118.41	3119.35	10.60	3118.09	9.34
10	3275.35	3299.89	3279.36	4.01	3298.39	23.04

Table 2: Uniform Observation Noise Tracking Data

The non-linear filter performed better seven times, the extended Kalman filter performed better twice, and they performed similarly (within two meters) once. Moreover, the RMS error $\sqrt{\frac{1}{10}\sum_{k=1}^{10}(e_{t_k}^K)^2}$ for the extended Kalman filter was 18.57 m and only 11.65 m for non-linear filter. This experiment was repeated six times and the seven-run average RMS errors for the Kalman and non-linear filters were respectively 18.64 and 11.04 m or 68% more error for the extended Kalman filter.

In the Cauchy density prediction case, one can see from Figure 2 (a) and (b) that the extended Kalman filter has a broader top due roundoff handling but a narrower base reflecting the fact that the normal distribution has faster decreasing tails than the Cauchy distribution. Data for a typical run through ten time steps is given below:

Step	h_{t_k}	${y}_k$	$\widehat{h}_{t_k}^N$	$e_{t_k}^N \doteq h_{t_k} - \hat{h}_{t_k}^N $	$\widehat{h}_{t_k}^K$	$e_{t_k}^K$
1	1974.22	1937.56	1946.83	27.39	1947.15	27.07
2	1891.75	1867.08	1873.56	18.19	1870.23	21.52
3	1773.28	1755.97	1761.51	11.77	1758.83	14.45
4	1640.3	1640.1	1638.87	1.43	1639.43	0.87
5	1463.29	1520.89	1519.77	56.48	1520.56	57.27
6	1250.35	1253.34	1268.85	18.50	1260.4	10.05
7	1011.39	995.411	999.583	11.807	997.058	14.332
8	747.432	773.497	769.707	22.275	772.752	25.32
9	488.845	515.714	516.496	27.651	517.021	28.176
10	352.556	184.266	194.423	158.133	190.414	162.142

Table 3: Cauchy Observation Noise Tracking Data

One can see that the non-linear filter performed better five times, the extended

Kalman filter performed better once, and the two filters performed similarly three times. The extended Kalman filter has trouble correcting from bad observations which occur frequently with the Cauchy distribution and provides poorer estimates when the observations are good. Moreover, the RMS error $\sqrt{\frac{1}{10}\sum_{k=1}^{10}(e_{t_k}^K)^2}$ for the extended Kalman filter was 57.2 meters whereas it was only 55.82 meters for the nonlinear filter. Interestingly enough, if the two highly corrupted observations (steps 5 and 10) are excluded from the RMS calculation, the non-linear filter still outperforms the extended Kalman filter by a slightly larger amount. This same experiment was repeated six times with our Mathematica program and the average RMS errors including the highly corrupted observations (absolute errors of more than 50 meters between h_{t_k} and y_k) for the Kalman and non-linear filters were respectively 109.36 and 83.59 meters which corresponds to 31% more error for the extended Kalman filter. When the highly corrupted observations were excluded the average RMS errors for the Kalman and non-linear filters were respectively 41.32 and 16.76 which corresponds to 147% more error for the extended Kalman filter. The sizeable discrepancy between the data in Table 3 and the average run statistics can be explained by the fact that for two runs the extended Kalman filter was completely "fooled" by the highly corrupted observations and the non-linear filter continued to supply reasonable estimates. Data for one of these runs is given below:

Step	h_{t_k}	${y}_k$	$\widehat{h}_{t_k}^N$	$e_{t_k}^N \doteq h_{t_k} - \hat{h}_{t_k}^N $	$\widehat{h}_{t_k}^K$	$e_{t_k}^K$
1	2052.85	2041.33	2029.48	23.37	2026.42	26.43
2	2094.69	2096.04	2096.	1.31	2095.96	1.27
3	2239.67	2287.09	2269.88	30.21	2277.66	37.99
4	2408.07	1905.82	2364.38	43.69	2101.31	306.76
5	2519.31	2491.19	2484.9	34.41	2215.	303.69
6	2637.13	2655.8	2649.53	12.4	2629.45	7.68
7	2757.19	2763.21	2763.66	6.47	2774.82	17.63
8	2841.28	2760.29	2770.82	70.46	2771.1	70.18
9	2926.64	2918.42	2910.6	16.06	2914.84	11.8
10	2979.8	2822.5	2840.06	139.74	2828.2	151.6

Table 4: Extended Kalman filter failure

The extended Kalman filter was completely fooled by observation 4 and could not correct at the next observation. In the other five runs, when the observations were highly corrupted they were so far into the tails that they were not represented well by either filter and both filters basically failed due to numerical reasons. It is believed that with a larger number of grid points making better representation of the tails the performance of the non-linear filter relative to the extended Kalman filter would be even better.

5. PROGRAMMING NOTES

In this section, we briefly discuss some technical issues associated with implementing our algorithm on a computer. The Mathematica program used to implement the previous example according to our algorithm may be requested from the author. The C program and all information about its current use are the property of Lockheed Martin corporation and any inquires should be addressed to Lockheed Martin Tactical Defense Systems, P.O. Box 64525, St. Paul, MN 55164-0525, Mail Stop U1P28.

As was previously mentioned, the routine simplifies when (i) δt_j = constant or (ii) *m* is affine. There are other obvious simplifications that can be made if the "predictive" and/or "tracking" statistics are not required at each time step. The FFT routines are usually most efficient if the number of grid points in each dimension is a power of 2. For the previous example it was found that 32, 64, or 128 points worked well depending on the accuracy desired as well as the chosen p_{v_k} and σ . All integration was done using a multi-dimensional version Simpson's 1/3 rule. Unfortunately, this requires an odd number of grid points so the integration routine had to be "patched" by adding a single column of Trapezoidal rule for each dimension.

A little care is required when mixing Fourier transforms known in closed form with (multi-dimensional) FFTs for r_j . It was demonstrated at the end of Section 3 that the convolution could be evaluated using a known Fourier transform for \mathcal{N} , a FFT for \tilde{r}_j , and then an inverse FFT. When this done, one will usually have to premultiply \tilde{r}_j by a phase shift, postmultiply its FFT by a phase shift, and then perform a similar transformation around the inverse FFT. These phase shifts will depend on the exact implementation of the FFT routine. An alternate strategy which was not employed in our example might be to take FFTs of \mathcal{N} and the time lag instead of using known formulae. Of course, there would be a small increase in error with the alternate strategy. Regardless of which strategy is chosen the phase shift in (21) should be replaced by a different phase shift to provide nicely centered graphs.

The routine is very linear in the sense that once q_{j+1} is known there is no longer any need for $p_{j+1|j}$ etc. and the same storage can be reused many times over. In storing these objects one must often use two arrays: one for values of $\beta_j, p_{j+1|j}$, or like objects and one for the values of (x, y) for which each function can be evaluated at. Thus, in (e) of the algorithm the operation $\beta_j(U(\delta t_k) {x+\delta \choose y} - {q \choose (\alpha+\mu q)\delta t_j})$ is implemented as a manipulation of the array for the arguments. This method also helps us to maintain better usage of our uniform grids.

The Fourier transform for \tilde{r}_j is trivial to obtain from that of r_j and vice versa. However, the points (x, y) at which r_j is evaluated in (b) of the algorithm are not the exact points required for the convolution or FFT in (d) and, consequently, interpolation is required. This interpolation can be done on either \tilde{r}_j or its FFT. In the previous example it was done on the FFT.

6. Conclusions

- A new method for implementing continuous-state, discrete-observation nonlinear filters has been introduced.
- This method is particularly useful when the observations have gross (non-smooth) non-linearities, non-additive noise, or non-Gaussian noise.
- The method is known to handle certain diffusions with non-linear drifts (Proposition 1).
- The method outperforms an extended Kalman filter on the simple example given here.
- The method requires at least 32 grid point per dimension which would limit its use to problems of five or less dimensions on most personal computers built in 1996. However, this compares favorably to current on-line/off-line splitting methods and direct on-line partial differential equation solving methods.

References

- V. Benesh. Exact finite-dimensional filters for certain diffusions with nonlinear drift. Stochastics, 5:65 - 92, 1981.
- [2] A. Bensoussan, R. Glowinski, and R. Rascanu. Approximations of the Zakai equation by splitting up method. SIAM J. Control and Optimization, 28:1420 – 1431, 1990.
- [3] A. Budhiraha and G. Kallianpur. Approximations to the solution of the Zakai equation using multiple Wiener and Stratonovich integral expansions. Technical Report No. 447, Center for Stochastic Processes, University of North Carolina (Chapel Hill), 1995.
- [4] Z. Cai, F. Le Gland, and H. Zhang. An adaptive local grid refinement method for nonlinear filtering. IRISA Campus de Beaulieu Internal Report No. 954, 1995.
- [5] F. Daum. Exact finite dimensional nonlinear filters for continuous time processes with discrete time measurements. Proceedings of the 1984 IEEE Conference on Decision and Control, pp. 16-22, 1984.

- [6] F. Daum. Exact finite dimensional nonlinear filters. *IEEE Trans. Auto. Control*, AC-31:616-622, 1986.
- [7] F. Daum. New exact nonlinear filters. in Bayesian Analysis of Time Series and Dynamic Models, ed. J.C. Spall, New York: Marcel Dekker Inc., pages 199–226, 1988.
- [8] A. Friedman. Uniqueness for the Cauchy problem for degenerate parabolic equations. Pacific Journal of Mathematics, 46(1):131 - 147, 1973.
- [9] A. Friedman. Stochastic Differential Equations and Applications, Volume 1. Academic Press, New York, 1975.
- [10] A. Jazwinski. Stochastic Processes and Filtering Theory. Academic Press, New York and London, 1970.
- [11] I. Karatzas and S. E. Shreve. Brownian Motion and Stochastic Calculus. Springer, Berlin-Heidelberg-New York, 1988.
- [12] K. Kastella, M.A. Kouritzin, and A. Zatezalo. A nonlinear filter for altitude tracking. Under preparation for publication in the *Proceedings of the 41st Air Traffic Control Association Annual Meeting* (abstract accepted).
- [13] R. Lototsky, R. Mikulevicius, and B.L. Rozovskii. Nonlinear filtering revisited: A spectral approach. preprint, 1995.
- [14] R. Mikulevicius and B.L. Rozovskii. Separation of observations and parameters in nonlinear filtering, *Proceedings of the 32nd IEEE Conference on Decision and Control*, 2:1564–1569, 1989.
- [15] D. Mirković. N-Dimensional Finite Element Package. preprint, Department of Mathematics, Iowa State University, 1996.
- [16] D. Ocone, J. Baras, and S. Marcus. Explicit filters for diffusions with certain nonlinear drifts. *Stochastics*, 8:1–16, 1982.
- [17] I. Sonin. On a class of degenerate diffusion processes. Theory Prob. Appl., 12:490-496, 1967.





Figure 1: Uniform Observation Noise with 50x50 mesh for: (a) Non-Linear Filter Prediction Density; (b) Extended Kalman Filter Prediction Density; (c) Non-Linear Filter Tracking Density; (d) Extended Kalman Filter Tracking Density



Figure 2: Cauchy Observation Noise with 50x50 mesh for: (a) Non-Linear Filter Prediction Density; (b) Extended Kalman Filter Prediction Density; (c) Non-Linear Filter Tracking Density; (d) Extended Kalman Filter Tracking Density