

Homotopy methods are an emerging approach to particle filtering that avoid numerical deficiencies of standard particle methods using a particle flow. This correspondence develops a new filter with non-thresholded measurements (i.e., a track-before-detect log-homotopy particle filter). We show the performance by simulating a rotating pulsed radar forming nonthresholded Range/Doppler maps.

I. INTRODUCTION

A particle filter (PF) is a target tracking approach that approximates a probability distribution by a set of samples, or particles, and corresponding weights. The standard approach is known as the sequential importance resampling (SIR) particle filter [1]. In the motion update step of SIR, particle locations are propagated according to a motion model, while the weights remain fixed. In the information update step, particle locations remain fixed, while the weights are updated by Bayes' rule. This involves multiplying the weights by the likelihood function evaluated at the particle locations and then normalizing the sum to one.

Particle degeneracy is a well-known SIR pathology that occurs when the likelihood function is concentrated on only a few particles and the information-updated weights are mostly zeros [1]–[3]. The standard remedy is resampling, in which the particles with the largest weights are replicated in proportion to their weights, and the weights of the resampled particles are set to be equal. The main deficiency of this method is that the resampling step reduces particle diversity, leading to performance degradation, which is often catastrophic in practical problems. In particular, it often happens that exactly one particle remains to represent the entire density. This problem increases with the dimension of the state space.

In a series of papers [3]–[11], Daum has introduced a new homotopy-based method to implement the information update. In this approach, instead of using measurements to simply update the weights of particles from the prior,

Manuscript received February 11, 2018; revised May 11, 2018; released for publication May 13, 2018. Date of publication June 7, 2018; date of current version December 5, 2018.

DOI. No. 10.1109/TAES.2018.2845201

Refereeing of this contribution was handled by Y. Boers.

This work was supported by the U.S. Air Force Research Laboratory under contract FA8650-15-C-1865. Opinions, interpretations, conclusions, and recommendations are those of the authors and not necessarily endorsed by the U. S. Government.

Authors' addresses: C. Kreucher is with the Integrity Applications Incorporated, Ann Arbor, MI 48108 USA, E-mail (ckreuche@umich.edu); K. Bell is with the Metron, Inc., Reston, VA 20190 USA, E-mail: (kristine.bell@ieee.org). (*Corresponding author: Chris Kreucher.*)

measurements are used to flow prior particles to a posterior location. This avoids particle depletion entirely, which is particularly valuable when the observation likelihood is narrow compared to the prior.

Several practical implementations have been recently reported [12]–[21]. These implementations all assume a measurement model in which the observed data are a linear or nonlinear transformation of the state-space parameters plus measurement noise. This model is suitable for a conventional detect-then-track approach in which radar data are processed to form a surface in measurement space and then thresholded to obtain a set of detections. Detections are associated to the filter to produce state estimates over time.

In contrast, here we discuss the design and implementation of a track-before-detect (TBD) particle flow filter, which incorporates raw(-er) nonthresholded observations directly into the filter. In TBD, the tracker ingests the entire measurement surface and exploits raw intensities rather than just threshold exceedances. For that reason, the concept of missed detections and false alarms is not relevant. Related TBD particle filter work includes [22]–[25]. Previous work has shown the utility of PF methods over Kalman-based (extended Kalman filter (EKF), unscented kalman filter) approaches in similar radar TBD applications [26], [27].

Our main contribution is the derivation and implementation of the particle flow method for a nonthresholded Radar sensor model. The computational complexity is higher than for the detect-then-track model; however, we still show a meaningful improvement in tracking error for a fixed number of computations. There is a more significant gain in the performance per particle. The conference paper [28] contains a preliminary, abbreviated version of this material.

This paper proceeds as follows. In Section II, we review the homotopic approach to particle flow and the details of a special case referred to as Geodesic Flow. Next, Section III gives the details of our pixelated nonthresholded Radar measurement model. Third, Section IV specializes the flow equations to our model. Section V provides a simulation comparing the performance of the new filter to an SIR filter on a model Radar problem. Finally, Section VI concludes.

II. SUMMARY OF THE GEODESIC FLOW APPROACH

The homotopy-based approach to particle filtering [5], [6] starts by defining a flow of the conditional PDF on state vector \mathbf{x} with respect to a parameter λ

$$\log p(\mathbf{x}, \lambda) = \log g(\mathbf{x}) + \lambda \log h(\mathbf{x}) - \log K(\lambda) \quad (1)$$

where $g(\mathbf{x})$ is the prior density, $h(\mathbf{x})$ the likelihood, and $K(\lambda)$ a normalization. The conditional probability $p(\mathbf{x}, \lambda)$ moves between the prior and posterior as λ moves from 0 to 1.

It is distinguished from standard particle filtering in that particles are flowed from a prior location to a posterior location using the measurements. In particular, the approach further supposes [6] that the flow of particles from prior

to posterior obeys the Ito stochastic differential equation (SDE)

$$d\mathbf{x} = \mathbf{f}(\mathbf{x}, \lambda)d\lambda + \mathbf{Q}(\mathbf{x}, \lambda)d\mathbf{w} \quad (2)$$

where $\mathbf{f}(\mathbf{x}, \lambda)$ is the flow function, \mathbf{w} is a Wiener process, and $\mathbf{Q}(\mathbf{x}, \lambda)$ is a diffusion matrix. The probability density function (PDF) of \mathbf{x} , $p(\mathbf{x}, \lambda)$, satisfies the Fokker–Plank equation [29]

$$\begin{aligned} \frac{\partial p(\mathbf{x}, \lambda)}{\partial \lambda} = & - \sum_{i=1}^N \frac{\partial}{\partial \mathbf{x}_i} [\mathbf{f}_i(\mathbf{x}, \lambda)p(\mathbf{x}, \lambda)] \\ & + \sum_{i=1}^N \sum_{j=i}^N \frac{\partial^2}{\partial \mathbf{x}_i \partial \mathbf{x}_j} [D_{ij}(\mathbf{x}, \lambda)p(\mathbf{x}, \lambda)] \end{aligned} \quad (3)$$

where

$$D_{ij}(\mathbf{x}, \lambda) = \frac{1}{2} \sum_{k=1}^N \mathbf{Q}_{ik}(\mathbf{x}, \lambda)\mathbf{Q}_{jk}(\mathbf{x}, \lambda). \quad (4)$$

Together, $\mathbf{f}(\mathbf{x}, \lambda)$ and $\mathbf{Q}(\mathbf{x}, \lambda)$ define an SDE in variables \mathbf{x} and λ that describes how particles flow from their prior location to their posterior location. Given the specification of the evolution of the PDF in (1), we can now in principle solve for the flow $\mathbf{f}(\mathbf{x}, \lambda)$. The solution is not unique, and Daum and Huang have developed many ways to find suitable solutions. Their recent work focuses on a special case [6], which results in the geodesic flow function

$$\mathbf{f}(\mathbf{x}, \lambda) = - \left[\frac{\partial^2 \log p(\mathbf{x}, \lambda)}{\partial \mathbf{x}^2} \right]^{-1} \left(\frac{\partial \log h(\mathbf{x})}{\partial \mathbf{x}} \right)^T. \quad (5)$$

The geodesic flow formulation requires the $\mathbf{Q}(\mathbf{x}, \lambda)$ function to have certain properties [6], and it is also not unique. It has been shown [4] that under a Gaussian approximation to the prior and measurement models, $\mathbf{Q}(\mathbf{x}, \lambda)$ has the form

$$\mathbf{Q}(\mathbf{x}, \lambda) = [\mathbf{P}^{-1} + \lambda \mathbf{H}^T \mathbf{R}^{-1} \mathbf{H}]^{-1} \mathbf{H}^T \mathbf{R}^{-1} \mathbf{H} [\mathbf{P}^{-1} + \lambda \mathbf{H}^T \mathbf{R}^{-1} \mathbf{H}]^{-1} \quad (6)$$

where \mathbf{P} is the prior covariance, \mathbf{H} is the standard EKF Jacobian matrix evaluated at \mathbf{x} , and \mathbf{R} is the measurement covariance. These quantities can be obtained easily when using a detect-then-track measurement model. For the nonthresholded model, the methodology is straightforward in principle but can be analytically intractable. The contribution of this paper is to develop a measurement model and derive the associated flow and diffusion functions for implementation in a Geodesic Flow particle filter.

A second choice is the diffusion-free approach [5] where

$$\mathbf{Q}(\mathbf{x}, \lambda) = \mathbf{0}. \quad (7)$$

Together (5) and (6) or (7) completely specify the flow of particles from prior to posterior.

III. NONTHRESHOLDED MEASUREMENT MODEL

This section defines our sensor measurements and the corresponding sensor likelihood $h(\mathbf{x})$ for nonthresholded sensor measurements. We assume a single target with state

$\mathbf{x} = [x \dot{x} y \dot{y}]$. A rotating monostatic pulsed radar collects coherent pulse intervals (CPIs) of raw data. Each CPI is a batch of pulse returns over an angular extent corresponding to the Radar pointing angle. The collected pulse data are aggregated and processed using Fourier methods [30] to generate a Range/Doppler map (RDM) at an angle θ , i.e., a map indexed by range and range-rate giving the intensity of the correlation at each. The measurement, thus, has information about range, range-rate, and azimuth angle. We do *not* propose to threshold the surface to deliver a detection, rather we want to fully exploit the entire data surface.

We now define our model of the measurement statistics. We assume the azimuth angle measurement error is characterized by a Gaussian with variance σ_θ^2 . This statistical model is used, for example, with angle centroiding and monopulse methods [30, Ch. 18]. Let z_θ denote this azimuth measurement. Then, letting $\theta(\mathbf{x}) = \arctan(\frac{\dot{y}}{\dot{x}})$, we have

$$z_\theta \sim N\left(\theta(\mathbf{x}), \sigma_\theta^2\right) \quad (8)$$

where for simplicity the radar is at $(0, 0)$.

At this angle, received pulses are processed to generate a RDM. Let z_{ij} denote the envelope-detected value in the (i, j) th range/range-rate resolution cell. The number of cells and the cell resolutions are determined by the number of aggregated pulses, the pulse-repetition frequency and bandwidth of the radar [31]. Denote the sequence of MN intensities by $\mathbf{z}^{\text{RDM}} = [z_{11}, z_{12}, \dots, z_{MN}]$. The statistics of the measurement in cell (i, j) depend on its proximity to the true range and range-rate of the target in measurement space.

The target pixel in measurement space will be captured by the mapping

$$\mathbf{m}(\mathbf{x}) = [m_1(\mathbf{x}) \quad m_2(\mathbf{x})]^T \quad (9)$$

which projects the target state \mathbf{x} to the sensor space (pixels).

In our application, this is a nonlinear mapping from Cartesian position and velocity to range/range-rate pixel. The range pixel of target \mathbf{x} is

$$m_1(\mathbf{x}) = \frac{\sqrt{x^2 + y^2}}{\Delta_R} + R_0 \quad (10)$$

and the range-rate pixel is

$$m_2(\mathbf{x}) = \frac{x\dot{x} + y\dot{y}}{\sqrt{x^2 + y^2}\Delta_{\dot{R}}} + \dot{R}_0 \quad (11)$$

where Δ_R and $\Delta_{\dot{R}}$ denote the pixel spacing while R_0 and \dot{R}_0 are the zero-reference points for the pixels.

The distance between the projection of \mathbf{x} and an arbitrary sensor pixel (i, j) will be denoted

$$\delta_{ij}(\mathbf{x}) = \mathbf{m}(\mathbf{x}) - [i \quad j]^T. \quad (12)$$

The physical model we employ is that the statistics of the cells are Rayleigh with mode dictated by the distance between the cell and the projection of the target [30]. In particular, the expected intensity in pixel (i, j) will be given by its distance from the projection of the target state \mathbf{x} ,




Fig. 1. Example intensity map (left) and RDM measurement realization (right) with $s_t^2 = 100$, $s_b^2 = 1$, and $\mathbf{C} = \mathbf{I}$. The target maps to $(r = 225 \text{ m}, \dot{r} = 19 \text{ m/s})$

weighted by the target impulse response (IPR). We have elected to use an exponential for the target IPR and define its value in pixel (i, j) as

$$I_{ij}(\mathbf{x}) = e^{-\frac{1}{2}\delta_{ij}^T(\mathbf{x})\mathbf{C}^{-1}\delta_{ij}(\mathbf{x})}. \quad (13)$$

For this model, contours of constant intensity are elliptical and \mathbf{C} is a 2×2 matrix that defines the shape of the ellipse.

The Rayleigh intensity parameter in pixel (i, j) is then

$$s_{ij}^2(\mathbf{x}) = s_b^2 + (s_t^2 - s_b^2)I_{ij}(\mathbf{x}). \quad (14)$$

This model captures the fact that the intensity is s_b^2 (the background intensity) for pixels very far from the target projection, s_t^2 (the target intensity) for a pixel centered at the target, and falls off as dictated by the IPR matrix \mathbf{C} .

In summary, our statistical model is that a target with state $\mathbf{x} = [x \dot{x} y \dot{y}]$ maps to sensor pixel $\mathbf{m}(\mathbf{x}) = [m_1(\mathbf{x}) \quad m_2(\mathbf{x})]^T$ through range and range-rate expressions combined with the sensor resolutions. A measurement pixel (i, j) is modeled as having intensity drawn from a Rayleigh random variable whose mode is given by the pixel distance from the target projection, weighted by the target IPR. The target is only measured when illuminated, i.e., when it is within the Radar beamwidth characterized by σ_θ . Fig. 1 illustrates the mode and an example scan.

With this as background, we can finally write the full nonthresholded measurement model explicitly as

$$\begin{aligned} h(\mathbf{x}) &\doteq p(\mathbf{z}^{\text{RDM}}, z_\theta | \mathbf{x}) = p(z_\theta | \mathbf{x}) \prod_{ij} p(z_{ij} | \mathbf{x}) \\ &= \left(\frac{1}{\sqrt{2\pi\sigma_\theta^2}} e^{-(z_\theta - \theta(\mathbf{x}))^2 / 2\sigma_\theta^2} \right) \prod_{ij} \frac{z_{ij}}{s_{ij}^2(\mathbf{x})} e^{-\frac{z_{ij}^2}{2s_{ij}^2(\mathbf{x})}}. \end{aligned} \quad (15)$$

IV. GEODESIC FLOW FOR NONTHRESHOLDED MEASUREMENTS

This section derives the flow equations for our nonthresholded sensor model. Particle flow from prior to posterior is effected by solving the SDE given in (5) and (6) or (7) using the measurement model in (15). We discuss each of these terms in turn.

A. Deterministic Component of the Flow, $\mathbf{f}(\mathbf{x}, \lambda)$

The deterministic component of the Geodesic Flow formulation given in (5) for nonthresholded measurements is specified as follows. The first term is found starting with

the identity

$$\frac{\partial^2 \log p(\mathbf{x}, \lambda)}{\partial \mathbf{x}^2} = \lambda \frac{\partial^2 \log h(\mathbf{x})}{\partial \mathbf{x}^2} + \frac{\partial^2 \log g(\mathbf{x})}{\partial \mathbf{x}^2} \quad (16)$$

which comes immediately from (1). The second term of the flow is simply the gradient of the log measurement likelihood and is defined by (15).

With these observations, the deterministic component of the flow is completely specified through the partials of $\log h(\mathbf{x})$ (the measurement likelihood) and $\log g(\mathbf{x})$ (the prior).

The required partials of $\log h(\mathbf{x})$ can be found in a straightforward manner using the definitions of the likelihood in (15), the Rayleigh intensity in (14) and the IPR in (13).

First, we have directly

$$\begin{aligned} \log h(\mathbf{x}) = & -\log \sqrt{2\pi\sigma_\theta^2} - \frac{(z_\theta - \theta(\mathbf{x}))^2}{2\sigma_\theta^2} \\ & + \sum_{ij} \log z_{ij} - \log s_{ij}^2(\mathbf{x}) + \left(\frac{-z_{ij}^2}{2s_{ij}^2(\mathbf{x})} \right). \end{aligned} \quad (17)$$

Let a and b represent components of x . The (scalar) partials with respect to these variables are given as

$$\begin{aligned} \frac{\partial \log h(\mathbf{x})}{\partial a} = & \frac{(z_\theta - \theta(\mathbf{x}))}{\sigma_\theta^2} \frac{\partial \theta(\mathbf{x})}{\partial a} \\ & + \sum_{ij} \left(\frac{z_{ij}^2 - 2s_{ij}^2(\mathbf{x})}{2s_{ij}^4(\mathbf{x})} \right) \frac{\partial s_{ij}^2(\mathbf{x})}{\partial a} \end{aligned} \quad (18)$$

and

$$\begin{aligned} \frac{\partial^2 \log h(\mathbf{x})}{\partial a \partial b} = & \frac{z_\theta - \theta(\mathbf{x})}{\sigma_\theta^2} \frac{\partial^2 \theta(\mathbf{x})}{\partial a \partial b} - \frac{\partial \theta(\mathbf{x})}{\partial a} \frac{\partial \theta(\mathbf{x})}{\partial b} \frac{1}{\sigma_\theta^2} \\ & + \sum_{ij} \frac{z_{ij}^2 - 2s_{ij}^2(\mathbf{x})}{2s_{ij}^4(\mathbf{x})} \frac{\partial^2 s_{ij}^2(\mathbf{x})}{\partial a \partial b} \\ & + \frac{s_{ij}^2(\mathbf{x}) - z_{ij}^2}{s_{ij}^6(\mathbf{x})} \frac{\partial s_{ij}^2(\mathbf{x})}{\partial a} \frac{\partial s_{ij}^2(\mathbf{x})}{\partial b}. \end{aligned} \quad (19)$$

Strictly speaking, computation of (15) through (19) require $O(N_{\text{cells}}^2)$ calculations. In practice, only cells close to the predicted location meaningfully influence the sum and so far fewer calculations are actually required.

From (14), the partials of $s_{ij}^2(\mathbf{x})$ are scaled versions of the partials of the IPR $I_{ij}(\mathbf{x})$, i.e.,

$$\frac{\partial s_{ij}^2(\mathbf{x})}{\partial a} = (s_r^2 - s_b^2) \frac{\partial I_{ij}(\mathbf{x})}{\partial a} \quad (20)$$

and so on. The partials of $I_{ij}(\mathbf{x})$ are tedious but come directly from the definition in (9), (12), and (13), i.e.,

$$\frac{\partial I_{ij}(\mathbf{x})}{\partial \mathbf{x}} = -\mathbf{J}\mathbf{C}^{-1}\boldsymbol{\delta}_{ij}(\mathbf{x})I_{ij}(\mathbf{x}) \quad (21)$$

where \mathbf{J} is the $N \times N$ Jacobian of $\boldsymbol{\delta}_{ij}(\mathbf{x})$, which is computed using the definitions of $m_1(\mathbf{x})$ and $m_2(\mathbf{x})$ above.

Finally, let $\boldsymbol{\mu}$ and \mathbf{P} be the empirical mean and covariance of the particles used to represent the prior $g(\mathbf{x})$.

Under a Gaussian approximation [4], we find $\frac{\partial \log g(\mathbf{x})}{\partial \mathbf{x}} = -\mathbf{P}^{-1}(\mathbf{x} - \boldsymbol{\mu})^T$ and $\frac{\partial^2 \log g(\mathbf{x})}{\partial \mathbf{x}^2} = -\mathbf{P}^{-1}$. This now completely specifies the particle flow $\mathbf{f}(\mathbf{x}, \lambda)$ for the nonthresholded model.

B. Diffusion, $\mathbf{Q}(\mathbf{x}, \lambda)$

One choice of the diffusion term is a Gaussian approximation [4], given in (6). \mathbf{P} is the empirical covariance of the particles used to represent the prior as before. \mathbf{H} is the EKF Jacobian matrix, i.e.,

$$\mathbf{H} = \begin{pmatrix} \frac{\partial m_1(\mathbf{x})}{\partial x} & \frac{\partial m_1(\mathbf{x})}{\partial \dot{x}} & \frac{\partial m_1(\mathbf{x})}{\partial y} & \frac{\partial m_1(\mathbf{x})}{\partial \dot{y}} \\ \frac{\partial m_2(\mathbf{x})}{\partial x} & \frac{\partial m_2(\mathbf{x})}{\partial \dot{x}} & \frac{\partial m_2(\mathbf{x})}{\partial y} & \frac{\partial m_2(\mathbf{x})}{\partial \dot{y}} \\ \frac{\partial \theta(\mathbf{x})}{\partial x} & \frac{\partial \theta(\mathbf{x})}{\partial \dot{x}} & \frac{\partial \theta(\mathbf{x})}{\partial y} & \frac{\partial \theta(\mathbf{x})}{\partial \dot{y}} \end{pmatrix} \quad (22)$$

and \mathbf{R} is matched to the covariance of the measurements. Here, we use the IPR and bearing variance so $\mathbf{R} = \begin{pmatrix} \mathbf{C} & 0 \\ 0 & \sigma_\theta^2 \end{pmatrix}$.

C. On the Numerical Solution of the SDE

Our fundamental approach to solving the Geodesic Flow SDE is the Euler method [32], i.e.,

$$\mathbf{x}_{n+1}^i = \mathbf{x}_n^i + f(\mathbf{x}_n, \lambda)\Delta\lambda + \mathbf{Q}(\mathbf{x}_n, \lambda)\Delta\mathbf{W}_n \quad (23)$$

where i is the particle. There are N steps in the discretization with variable sized $\Delta\lambda$, which moves λ between 0 and 1. Since most of the important flow happens near $\lambda = 0$, we space the discretization points logarithmically between 10^{-5} and 1. Specifically, the solution is effected by discretizing λ log-spaced with N points between $\lambda = 10^{-5}$ to $\lambda = 1$ (i.e., $\lambda = 10^{-5+5d/N}$, $d = 0 \dots N-1$).

Geodesic Flow SDEs are stiff [11] and, thus, require special attention for this numerical integration. We find, empirically, that even specially crafted stiff solvers such those packaged with MATLAB often require step sizes so small as to be make the solution intractable. Adaptive SDE solution strategies for homotopic flow is an active area of investigation [11].

We have looked at two approaches to address the stiffness without completely intractable N : 1) trap for events by simply looking at the numerical gradient of the particle position during the solution. When a large movement happens, the particle is reflowed from the prior state to its posterior state; 2) Perform a resample-like procedure to eliminate those particles that have been flowed in a numerically unstable way. In practice, we find option 2) gives the best trade between compute time and performance.

V. SIMULATION

This section shows a simulation comparing the new nonthresholded Geodesic Flow particle filter with an SIR particle filter on a model Radar tracking problem. Both trackers use the exact same temporal prediction, measurement model, and measurements. The only distinction is that the Geodesic Flow filter flows particles to a posterior




Fig. 2. New method flows particles from prior location to posterior locations as λ goes from 0 to 1. The true target range is 230 m and range-rate is 18 m/s. Posterior locations are consistent with these values given the resolutions and IPR.

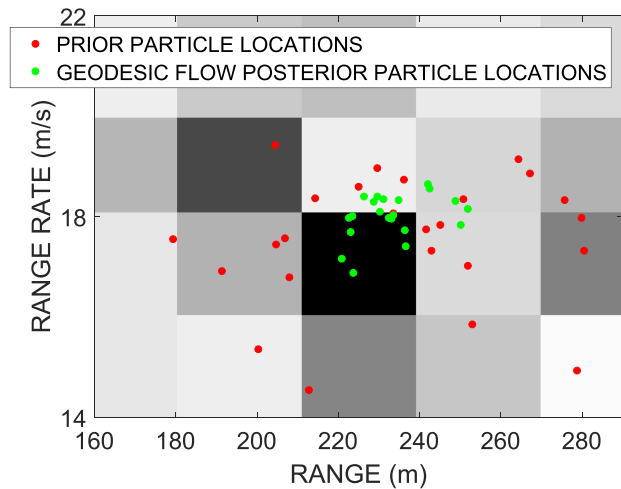


Fig. 3. Illustration of the prior and posterior locations of particles under the Geodesic Flow method. The homotopic approach moves particles from their prior location to their posterior location using the measurements, whereas the SIR method simply updates the weights of the prior particles.

position by solving an SDE. In contrast, in the SIR filter, particles stay at their prior position and have weights updated via Bayes' rule.

A. Illustrative Images

Fig. 2 shows that how the geodesic approach flows particles from prior to posterior location as the SDE is solved from $\lambda = 0$ to $\lambda = 1$. The figures illustrate how a relatively broad prior is concentrated to a narrower posterior.

Next, Figs. 3 and 4 illustrate why the geodesic flow approach improves performance over SIR. In both figures, particle locations are superimposed on the observation surface. The use case is the situation with a broad prior and narrow observation likelihood. The SIR approach suffers from the well-known particle degeneracy pathology.

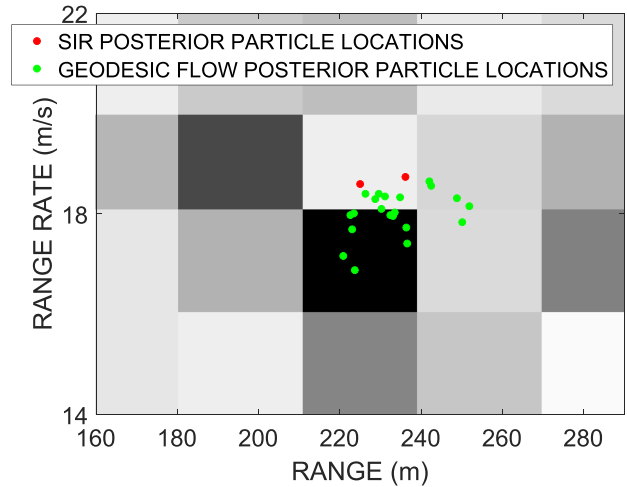


Fig. 4. SIR method can only update prior particle weights, and resamples to discard low-impact particles. We find in this case that only two particles survive the resampling and neither are in the main part of the measurement distribution.

First, in Fig. 4, we see that the prior particle locations do not correspond to the dominant measurement location. While the SIR approach can only reflect this mismatch via weighting the particles, the Geodesic approach can flow particles from these poor prior locations to the more likely posterior location.

Next, in Fig. 4, we see that only the few particles close to the measurement likelihood peak get all of the weight and survive SIR resampling. On the other hand, the Geodesic Flow method is able to move particles from the prior to the correct posterior location to more accurately represent the posterior.

B. Simulation Description

The Radar tracking simulation is described as follows. A target is characterized by its four-dimensional (x, \dot{x}, y, \dot{y}) state. The true state evolves with nearly constant velocity for 100 time steps.

At each time, nonthresholded measurements are made on a 51×51 range/range-rate grid with IPR $\mathbf{C} = \mathbf{I}$ at a pointing angle θ . The parameters of the grid are matched to a rotating S-band radar with $F_c = 3$ MHz and a 10-ms CPI. This yields a range bin size of $\Delta_R = 30$ m and range-rate bin size of $\Delta_{\dot{R}} = 5$ m/s. The antenna beam width is selected as 5° . The Rayleigh intensity in each grid cell is given by (14) with $s_t^2 = 100$, i.e., it is defined by the true target location and the IPR.

C. Results

The performance of the tracker is measured by the root mean-square error (RMSE) between true state and predicted state over the simulation window. We compare the SIR and geodesic flow filters as a function of number of particles in Fig. 5 and as a function of processing time in Fig. 6. The Geodesic flow filter uses $N = 11$ discretization steps.

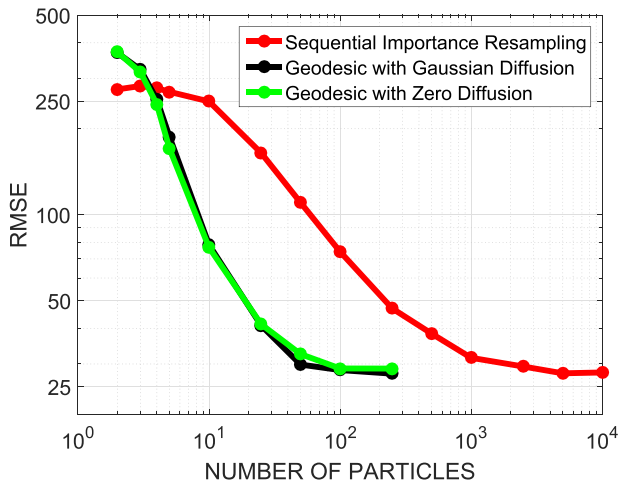


Fig. 5. Geodesic flow approach needs about 25× fewer particles than the SIR approach to achieve the same RMSE.

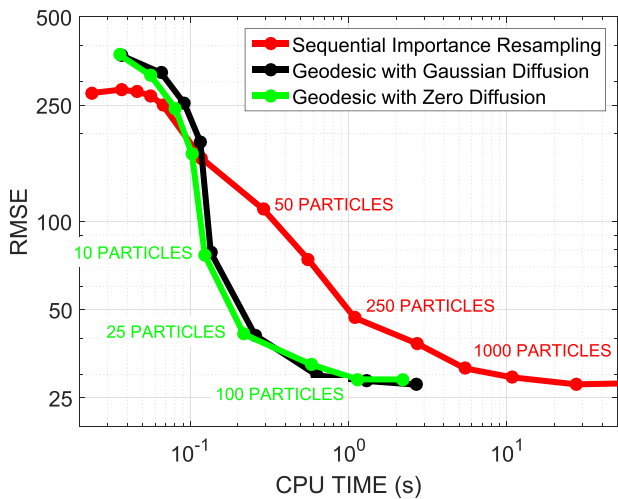


Fig. 6. Geodesic approach needs about 10× less CPU time than the SIR approach to achieve the same RMSE.

We find significant improvement in both per-particle and per-flop performance with the homotopy approach. We find no significant utility to using the Gaussian diffusion in this case, most likely due to the poor match between our highly nonlinear measurement and the EKF used in the Gaussian approximation. While empirically the zero-diffusion covariances are more clustered (smaller) than the Gaussian, we find empirically the performance to be similar.

VI. CONCLUSION

This paper has developed and illustrated by simulation a TBD log-homotopy particle filter, which is an implementation of the work of Daum and Huang. We have elected to use the Geodesic approach to particle flow coupled with Gaussian or zero diffusion. While the homotopy approach requires more computations per particle, on balance, it still provides a significant improvement per flop on our Radar tracking problem. Furthermore, in this problem, the zero-diffusion approximation performed similarly to the Gaussian approximation. Since it saves analytical

effort and some computations, the zero-diffusion approach provides the best performance of the techniques studied. An interesting next step in this paper is to extend it to multiple interacting targets.

ACKNOWLEDGMENT

The authors would like to thank F. Daum for his helpful comments during the initial stages of this investigation.

CHRIS KREUCHER ^{ID}
 Integrity Applications Incorporated, Ann
 Arbor, MI USA
KRISTINE L. BELL ^{ID}, Fellow, IEEE
 Metron, Inc., Reston, VA USA

REFERENCES

- [1] B. Ristic, S. Arulampalam, and G. Gordon *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Boston, MA, USA: Artech House, 2004.
- [2] L. Stone, T. Streit, T. Corwin, and K. Bell *Bayesian Multiple Target Tracking, 2nd ed.* Norwood, MA, USA: Artech House, 2014.
- [3] F. Daum and J. Huang Particle degeneracy: Root cause and solution *Proc. SPIE*, vol. 8050, 2011, Art. no. 80500W.
- [4] F. Daum, J. Huang, and A. Noushin Gromov's Method for Bayesian Stochastic Particle Flow: A Simple Exact Formula for Q in *Proc. IEEE Int. Conf. Multisensor Fusion Integration Intell. Syst.*, Baden-Baden Germany, Sep. 2016, pp. 540–545.
- [5] F. Daum, J. Huang, and A. Noushin Exact particle flow for nonlinear filters *Proc. SPIE*, vol. 7697, 2010, Art. no. 769704.
- [6] F. Daum and J. Huang Particle flow for nonlinear filters, Bayesian decisions and transport *Proc. IEEE Int. Conf. Inf. Fusion*, Istanbul, Turkey, 2013, pp. 1072–1079.
- [7] F. Daum and J. Huang How to avoid normalization of particle flow for nonlinear filters, Bayesian decisions, and transport *Proc. SPIE*, vol. 9092, 2014, Art. no. 90920B.
- [8] F. Daum and J. Huang Proof that particle flow corresponds to Bayes rule: necessary and sufficient conditions *Proc. SPIE*, vol. 9474, 2015, Art. no. 94740I.
- [9] F. Daum and J. Huang A plethora of open problems in particle flow research for nonlinear filters, Bayesian decisions, Bayesian learning, and transport *Proc. SPIE*, vol. 9842, 2016, Art. no. 98420I.
- [10] F. Daum Seven dubious methods to compute optimal Q for Bayesian stochastic particle flow in *Proc. IEEE Int. Conf. Inf. Fusion*, Heidelberg, Germany, 2016, pp. 2237–2244.
- [11] S. Mori, F. Daum, and J. Douglas Adaptive step size approach to homotopy-based particle filtering Bayesian update in *Proc. IEEE Conf. Inf. Fusion*, Heidelberg, Germany, 2016, pp. 2035–2042.
- [12] S. Choi, P. Willett, F. Daum, and J. Huang Discussion and application of the homotopy filter *Proc. SPIE*, vol. 8050, May 2011, Art. no. 805021.
- [13] T. Ding and M. Coates Implementation of the Daum-Huang exact-flow particle filter

- in *Proc. IEEE Statist. Signal Process. Workshop*, Ann Arbor, MI, USA, 2012, pp. 257–260.
- [14] V. Jilkov, J. Wu, and H. Chen
Performance comparison of GPU-accelerated particle flow and particle filters
in *Proc. IEEE Int. Conf. Inf. Fusion*, Istanbul, Turkey, 2013, pp. 1095–1102.
- [15] K. Bell and L. Stone
Implementation of the homotopy particle filter in the JPDA and MAP-PF multi-target tracking algorithms
in *Proc. IEEE Conf. Inf. Fusion*, Spain, Jul. 2014, pp. 1–8.
- [16] M. Khan and M. Ulmke
Improvements in the implementation of log-homotopy based particle flow filters
in *Proc. IEEE Conf. Inf. Fusion*, Washington, DC, Jul. 2015, pp. 74–81.
- [17] S. Gupta, J. Yu, M. Mallick, M. Coates, and M. Morelande
Comparison of angle-only filtering algorithms in 3D using EKF, UKF, PF, PFF, and Ensemble KF
in *Proc. IEEE Int. Conf. Inf. Fusion*, Washington, DC, Jul. 2015, pp. 1649–1656.
- [18] N. Moshtagh and M. Chan
Multisensor fusion using homotopy particle filter
in *Proc. IEEE Int. Conf. Inf. Fusion*, Washington, DC, Jul. 2015, pp. 1641–1648.
- [19] Y. Li, L. Zhao, and M. Coates
Particle flow auxiliary particle filter
in *Proc. Int. Workshop Comput. Adv. Multi-Sensor Adaptive Process.*, Cancun, Mexico, Dec. 2015, pp. 157–160.
- [20] M. Khan, M. Ulmke, and W. Koch
A log homotopy based particle flow solution for mixture of Gaussian prior densities
in *Proc. IEEE Int. Conf. Multisensor Fusion Integration Intell. Syst.*, Baden-Baden, Germany, Sep. 2016, pp. 546–551.
- [21] Y. Li and M. Coates
Particle filtering with invertible particle flow
IEEE Trans. Signal Process., vol. 65, no. 15, pp. 4102–4116, Aug. 2017.
- [22] D. Salmond and H. Birch
A Particle Filter for Track-before-detect
in *Proc. Amer. Control Conf.*, Arlington, VA, USA, Jun. 2001.
- [23] S. Musick, J. Greenwald, C. Kreucher, and K. Kastella
A Comparison of particle method and finite difference nonlinear filters for low snr target tracking
in *Proc. IEEE Conf. Inf. Fusion*, Montreal, CA, Aug. 2001, pp. 1–8.
- [24] C. Kreucher, K. Kastella, and A. Hero
Tracking multiple targets using a particle filter representation of the joint multitarget probability density
Proc. SPIE, vol. 5204, Aug. 2003, pp. 258–269.
- [25] Y. Boers and J. Driessen
Multitarget particle filter track before detect application
IEE Proc. Radar, Sonar Navig., vol. 151, no. 6, pp. 351–357, Dec. 2004.
- [26] G. Newstadt, E. Zelnio, L. Gorman, and A. Hero
Detection/tracking of moving targets with synthetic aperture radars
Proc. SPIE, vol. 7699, 2010, Art. no. 76990I.
- [27] C. Kreucher, K. Bell, and D. Sobota
A comparison of tracking algorithms for supermaneuverable targets
in *Proc. IEEE Conf. Inf. Fusion*, Washington DC, 2015, pp. 534–541.
- [28] C. Kreucher and K. Bell
A geodesic flow particle filter for nonthresholded measurements
in *Proc. IEEE Radar Conf.*, Seattle, WA, USA, May 2017, pp. 0891–0896.
- [29] A. Jazwinski
Stochastic Processes and Filtering Theory. 0New York, NY, USA: Dover, 1970.
- [30] M. Richards, J. Scheer, and W. Holm, eds.
Principles of Modern Radar: Basic Principles. Rayleigh, NC: SciTech Publishing, 2010.
- [31] J. Toomay and P. Hannen
Radar Principles for the Non-Specialist. 3rd. ed. Berlin, Germany: Springer, 2004.
- [32] P. Kloeden and E. Platen
Numerical Solution of Stochastic Differential Equations. Berlin, Germany: Springer-Verlag, 1992.