Particle Filter for Seismic-Cycle Estimation

19

The particle filter is an effective data-assimilation method for low-dimensional, nonlinear systems. It is easy to implement, and it is straightforward to include model error, parameters, and controls in the state vector. This chapter demonstrates the use of a particle filter in the case of a parameter bias. We apply the standard particle filter with importance-resampling to estimate seismic cycles. Seismic-cycle estimation helps us assess and forecast fault slips in slow-slip events and earthquakes. In the simulation of seismic cycles, the parameter choice determines the periodicity of the seismic events. Consequently, having a parameter bias will result in errors in the simulated state trajectory. This chapter illustrates how we can compensate for errors in the state trajectory caused by a parameter bias using the particle filter. We include the model errors, and the joint estimation of state and parameters leads to more satisfactory results than we can obtain by estimating the state alone. We compare state-estimation results between the standard particle filter and EnKF. Furthermore, we illustrate the advantage of using a fully nonlinear data-assimilation method for systems with sudden transitions in the model trajectory.

19.1 Particle Filter for State and Parameter Estimation

We apply the particle filter with sequential importance resampling as described in Sect. 9.2.1 on an application for seismic-cycle estimation. This application provides an interesting data assimilation case as the state trajectory is sensitive to both parameters and state, which are generally uncertain.

The highly nonlinear behavior of seismic cycle models may limit the effectiveness of the EnKF or related methods. Within a set of realizations for seismic cycles, seismic events can occur at different times for each of the realizations. In this case,

G. Evensen et al., *Data Assimilation Fundamentals*,
Springer Textbooks in Earth Sciences, Geography and Environment,
https://doi.org/10.1007/978-3-030-96709-3_19

the particle filter has the advantage of propagating the full pdf of the state. The standard particle filter applied in this chapter only affects the weight of the particles. As we will illustrate for a quasi-geostropic ocean model in Chap. 20, it will not move the particles as the particle flow filter does.

19.2 Seismic Cycle Model

To simulate the stick-slip motion of a fault related to the movement of a tectonic plate, Burridge and Knopoff (1967) introduced a simplified representation of an assembly of springs and blocks. This Burridge–Knopoff (BK) model connects a set of blocks through springs. It also uses springs to connect the blocks to a loader plate. The loader plate moves with a uniform velocity and pulls the blocks. Here, we consider a zero-dimensional version of the BK model that represents a single spring-block system, as shown in Fig. 19.1. The contact surface, represented by the block and the rough surface over which the block can slide, simulates the fault.

The equations for this rate- and state-dependent friction law can be found in Ruina (1983). Erickson et al. (2011) shows how to solve these equations by rewriting them as three partial differential equations

$$\frac{\partial \theta}{\partial t} = -v(\theta + (1+\zeta)\ln v), \tag{19.1}$$

$$\frac{\partial u}{\partial t} = v - 1, \tag{19.2}$$

$$\frac{\partial v}{\partial t} = -F^2\left(u + \frac{1}{\xi}(\theta + \ln v)\right), \tag{19.3}$$

where F is the non-dimensional frequency of the resulting cyclic motion, and ξ is a non-dimensional spring constant. Furthermore, u is the block's non-dimensional

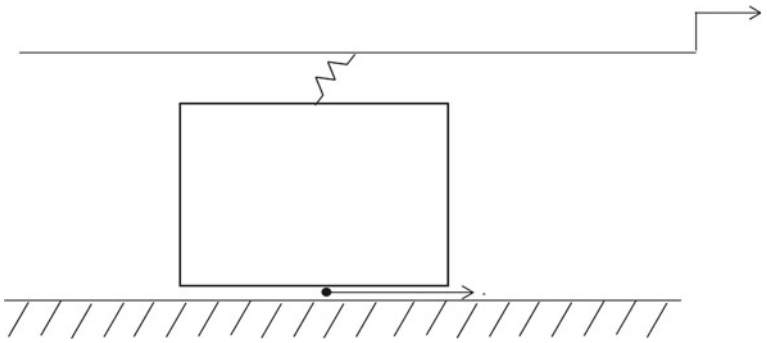

Fig. 19.1 The figure illustrates a zero-dimensional version of the spring-block system by Burridge and Knopoff (1967), consisting of a block connected with a spring to a driving plate moving with a constant velocity

slip, i.e., displacement, relative to the driver plate, and v is the corresponding slip rate. Finally, θ characterizes the state of the frictional surface. We interpret it as a measure of the average lifetime of the asperity contact population in a fault, a variable that influences the strength of the contacts. We obtain the fault shear stress τ can by multiplying the slip values with $-\xi$. Banerjee et al. (2022) give further details on the parameters used in the simulations.

The parameter ζ is the parameter of interest here. It depends on the friction parameters in the rate- and state-dependent friction law and determines the stick-slip behavior of the slip events. In our zero-dimensional model, we choose the parameter ζ such, that the system has a stable limit cycle solution with periodic oscillations. We investigate the effectiveness of the particle filter in a case of a biased ζ. A bias in ζ affects the simulated amplitude and frequency of fault stress, slip and the state variable θ in the earthquake cycle.

19.3 Data-Assimilation Experiments

We apply a standard particle filter with sequential importance resampling in three cases:

Case A: state estimation,
Case B: state estimation with increased model error,
Case C: state- and parameter estimation.

In all three cases, we represent the model noise by an additive model error $d\mathbf{q}$ in Eq. 2.24). We add this model error to the Eq. (19.2) for $\frac{\partial u}{\partial t}$. For Case A and Case C, we draw the model error from a Gaussian distribution with zero mean and a standard deviation of 0.01. For Case B, the standard deviation of the model error is 0.1.

We run the forward model for a period of 500 time units. We define the system's state by the fault shear stress τ, the slip rate v, and the state variable θ at the fault interface. In Case A and Case B, the state vector contains the model state plus the model error $\mathbf{z} = [\theta, \tau, v, dq]^{\mathrm{T}}$, while Case C extends the state vector with the parameter ζ, i.e., $\mathbf{z} = [\theta, \tau, v, dq, \zeta]^{\mathrm{T}}$. The measurement vector has two elements and contains the (synthetic) measurements of fault shear stress and slip rate.

We sample synthetic measurements from the truth run with parameter $\zeta_t = 0.7$. We assume the state variables of shear stress (τ) and slip-velocity (v) to be measured every 4 time units. Measurement errors are sampled from a normal distribution and added to the truth-run samples. For shear stress, the standard deviation of this normal distribution is 0.6. The standard deviation of the measurement errors added to the slip velocity measurements is 1.15.

We generate the prior particles with a biased $\zeta = 0.6$. In the following, we denote this biased parameter with ζ' and the true parameter as $\hat{\zeta}$. We evaluate the data-assimilation performance by comparing the posterior of $\mathbf{z}$ to the values of $\mathbf{z}$ in the true trajectory. Figure 19.2 illustrates the state trajectories for $\hat{\zeta} = 0.7$ and $\zeta' = 0.6$.

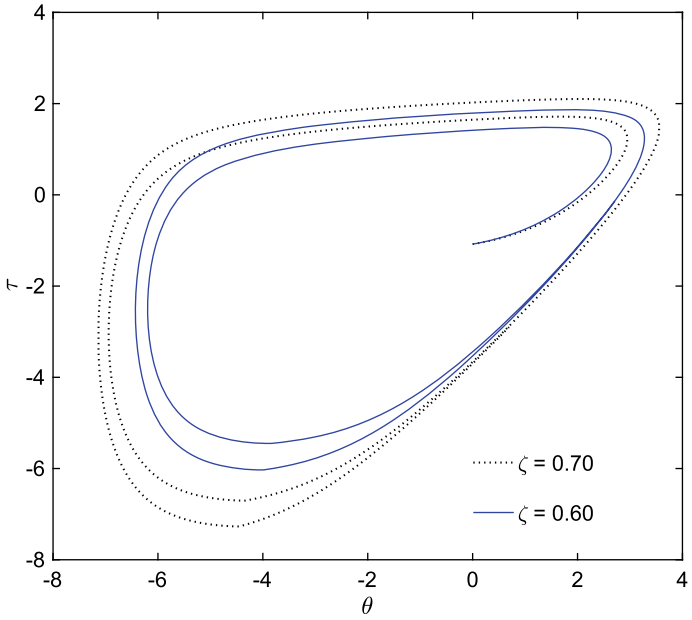

Fig. 19.2 Phase diagram for a Burridge–Knopoff (BK) model with rate-and-state-friction law when $\zeta = 0.6$ (bias case) and $\zeta = 0.7$, i.e., a case of no bias

Starting from an initial condition of $\theta = 0$ and $\tau = -1$, both θ and τ increase, after which the shear stress τ stays constant and θ drops (the interseismic phase). Following this phase, the shear stress drops while θ remains stable in the coseismic phase. After that, we observe a very short post-seismic phase where τ remains low and θ rapidly becomes less negative. The system then moves to the interseismic phase in which both θ and τ increase again and the cycle repeats. For both $\zeta = 0.6$ and $\zeta = 0.7$, the trajectories are stable cycles. Figure 19.2 shows that the trajectory for $\zeta = 0.6$ has smaller amplitudes of τ and θ than the trajectory for $\zeta = 0.7$. A smaller value of ζ results in a shorter cycle.

The experiments use $N = 1000$ particles in a standard particle filter with importance resampling similar as described in Sect. 9.2.1. We assume the likelihood $f(\mathbf{d}|\mathbf{z}_j)$ in Eq. 9.4 to be known (see Algorithm 9). We use a Lorentz function rather than a Gaussian (see Algorithm 9), i.e., for a measurement d and the model state $\mathbf{z}$ in assimilation window l, the likelihood becomes

$$f\left(\mathrm{d}_l|\mathbf{z}_l\right) = \frac{1}{1 + \frac{\left(\mathrm{d}_l - H(\mathbf{z}_l)\right)^2}{\sigma_l^2}}, \tag{19.4}$$

with the variance σ_l^2 of the measurement noise. This function has the advantage of broader tails than the Gaussian function, resulting in fewer particles having zero weight (Vossepoel & Van Leeuwen, 2007). While this choice helps the particle filter

algorithm avoid degeneracy, it is slightly inconsistent with the actual measurement errors drawn from a Gaussian with the same width.

19.4 Case A: State Estimation

Figure 19.3 illustrates the evolution of the posterior mean τ, θ and v, in the presence of a bias ($\zeta' = 0.6$) compared to the truth case of no bias ($\hat{\zeta} = 0.7$). The figure illustrates that when performing only state estimation while not accounting for the parameter bias, the data assimilation cannot fully recover the true trajectory of these state variables. While the data assimilation appears to decrease the frequency of the events and brings the shear stress and θ closer to the truth, the resulting analysis does not capture the amplitude of the τ and θ cycles very well, and the reconstruction of the slip rate v is poor.

During the interseismic phase (e.g., at $t = 16$), the difference between prior and truth is not very large. But at time $t = 20$, i.e., in the coseismic phase, the distribution of the prior shear stress has a peak around $\tau = -5$, while the weights

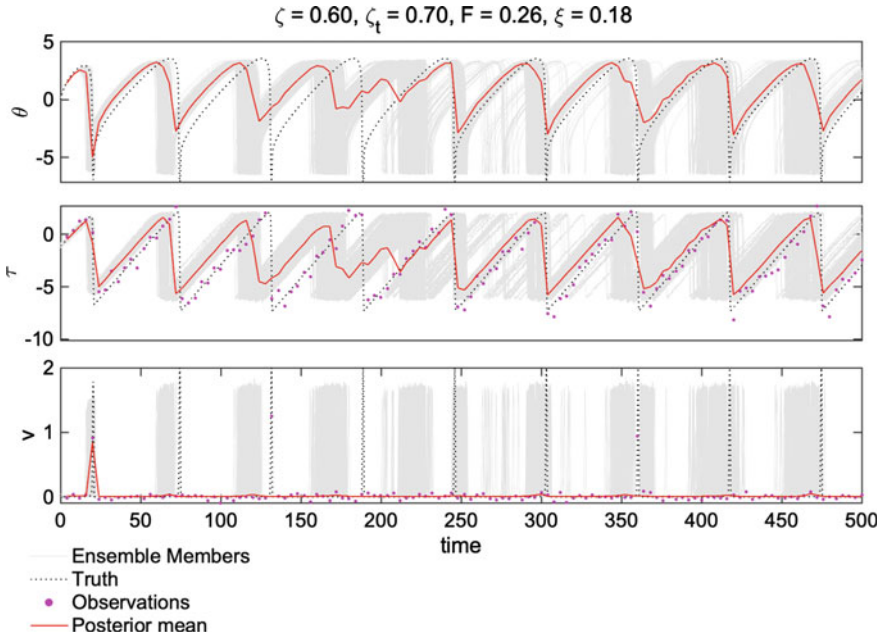

Fig. 19.3 Case A: Evolution of 1000 particles and the posterior mean and truth of state variables θ, shear stress τ, and slip velocity v in an experiment with state estimation for a synthetic experiment where the prior has a biased parameter relative to the truth. The dotted blue lines provide the actual θ, τ, and v, and cyan dots represent synthetic measurements. Grey lines represent the prior particles, and the red line is the analysis given by the weighted mean of the posterior distribution

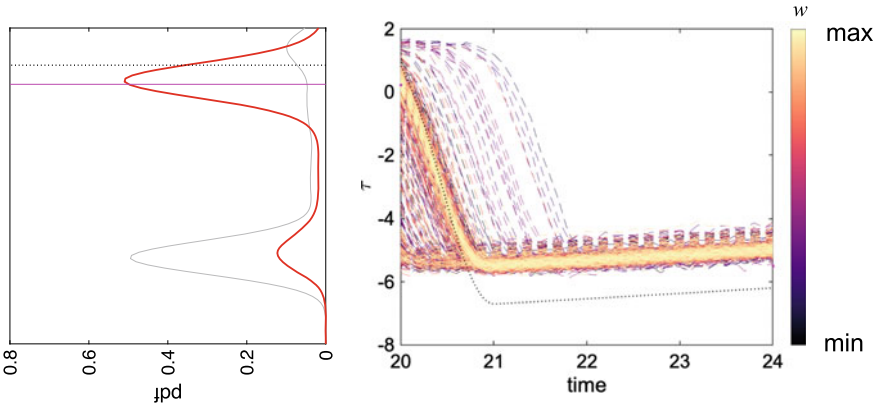

Fig. 19.4 Time evolution of 1000 particles (right panel), as shown in Fig. 19.3, zoomed in for $t = 20$. The colors of the lines represent the weights for $t = 20$. The lines have a lighter color when the particle they represent has a higher weight. The neighboring left panel graph gives the distribution of prior particles (in grey) and the distribution after resampling (in red)

at that moment in time are highest for the shear stress value around $\tau = 0$ (see Fig. 19.4). The weighting of the particles favors the particles whose shear stress has not yet dropped. The relatively late change in shear stress as measured only occurs in a few particles, and as a result, the filter reaches the criterion for resampling. Figure 19.4 illustrates the particles' distribution before and after resampling.

The different behavior of the particle filter compared to the ensemble Kalman filter (EnKF) is illustrated in Fig. 19.5. For the times $t = 20$, $t = 24$, $t = 72$, and $t = 76$, this figure presents the pdfs for Case A with the standard particle filter and the EnKF. It shows that the posterior distribution of τ in the case of a particle filter can feature small secondary peaks next to the dominant peak, and it can redistribute the particles in such a way that sudden transitions in the trajectory, such as in the coseismic phase at $t = 20$, are captured in the posterior. In the EnKF experiment, the posterior distribution of τ remains very close to its prior distribution. For the interseismic period, the results of the particle filter and the EnKF are virtually the same.

Figure 19.6 illustrates the particle results further with a phase diagram of the θ and τ values of the particles. These plots give the expected phase trajectories, as shown in Fig. 19.2. The lower value of ζ' in the prior particles implies a phase difference compared to the truth. We can only partly compensate for this difference by favoring particles that quickly enter the coseismic phase. We also observe that for $t = 20$, the particles closest to the truth in the τ-θ space obtain the highest weight. At $t = 24$, $t = 72$, and $t = 76$, this is less so because of the match of the slip rate v to the truth (not shown). By adjusting the state, the assimilation partly compensates for the bias in the parameter, but the resulting posterior fails to follow the actual trajectory.

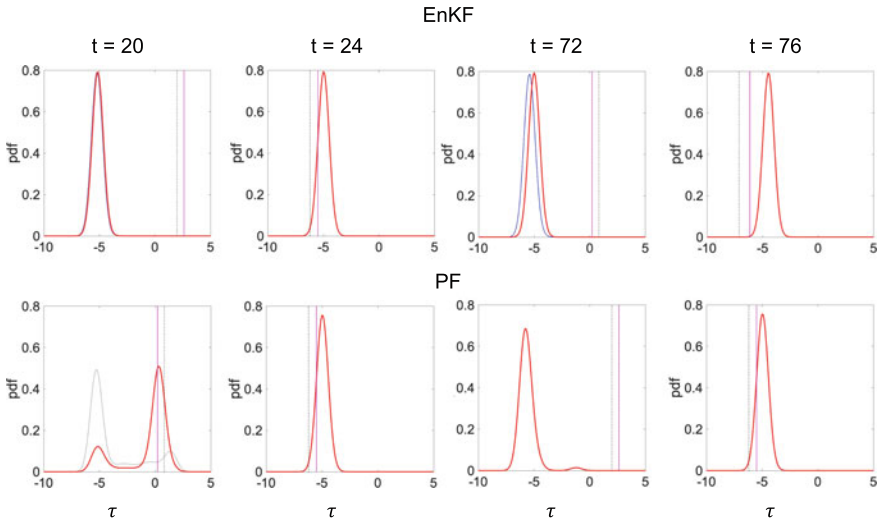

Fig. 19.5 Case A: Comparison of EnKF (top) and particle filter (bottom) for state estimation in the presence of a biased prior parameter for $t = 20$, $t = 24$, $t = 72$ and $t = 76$. The panels give the prior distribution (grey for the particle filter, blue for EnKF) and posterior distribution (red) of τ in the ensemble members/particles. The value of the measured shear stress is indicated with a pink line, the true shear stress with a dotted black line

19.5 Case B: State Estimation with Increased Model Error

In many real-life applications, it is not clear beforehand whether the difference between model and measurements comes from errors in the model state, in the model parameters, the model controls, or in the model itself. In these situations, one typically applies the model as a weak constraint. In Case B, we investigate whether we can compensate for a bias in the prior parameter by increasing the prior model error in Eq. (19.2) from 0.01 to 0.1.

Figure 19.7 shows the results. Not surprisingly, a more significant model error increases the spread in the prior compared to Case A (Fig. 19.6). The most noticeable difference occurs at $t = 72$, just before the coseismic phase. Some particles still have high shear-stress values, while the shear stress has dropped for others. By giving the high-shear-stress particles a relatively high weight, the particle filter with this weak-constraint formulation partly compensates for the bias in the parameter.

19.6 Case C: State- and Parameter Estimation

In Case C, we investigate state- and parameter estimation in the case of a biased prior parameter. The state vector contains the state variables τ, θ, and v and the parameter ζ. We generate the initial ensemble by sampling each particle from a log-

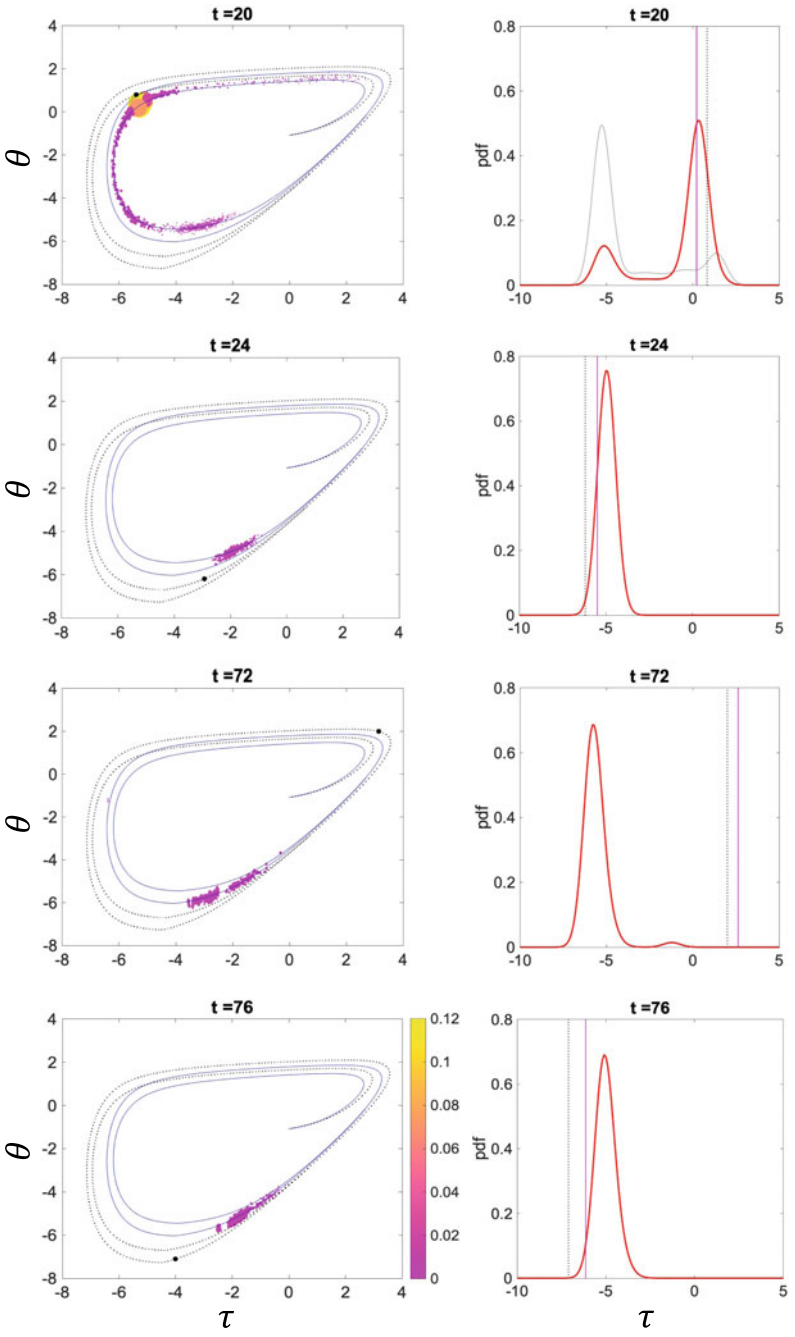

Fig. 19.6 Case A: Left column: phase diagrams showing the θ-τ trajectories of the true state (black, $\hat{\zeta} = 0.7$) and the biased prior (blue, $\zeta' = 0.6$) as in Fig. 19.2. Particles of the Case A experiment are indicated with dots, whose size and color indicate their weights for $t = 20$, 24, 72, and 76. We indicate the true values of θ and τ for these times with a black dot. Right column: the posterior (red) and prior (grey) shear stress distributions. For $t = 24$, 72, and 76, the difference between the prior and the posterior is so small that the two lines overlap

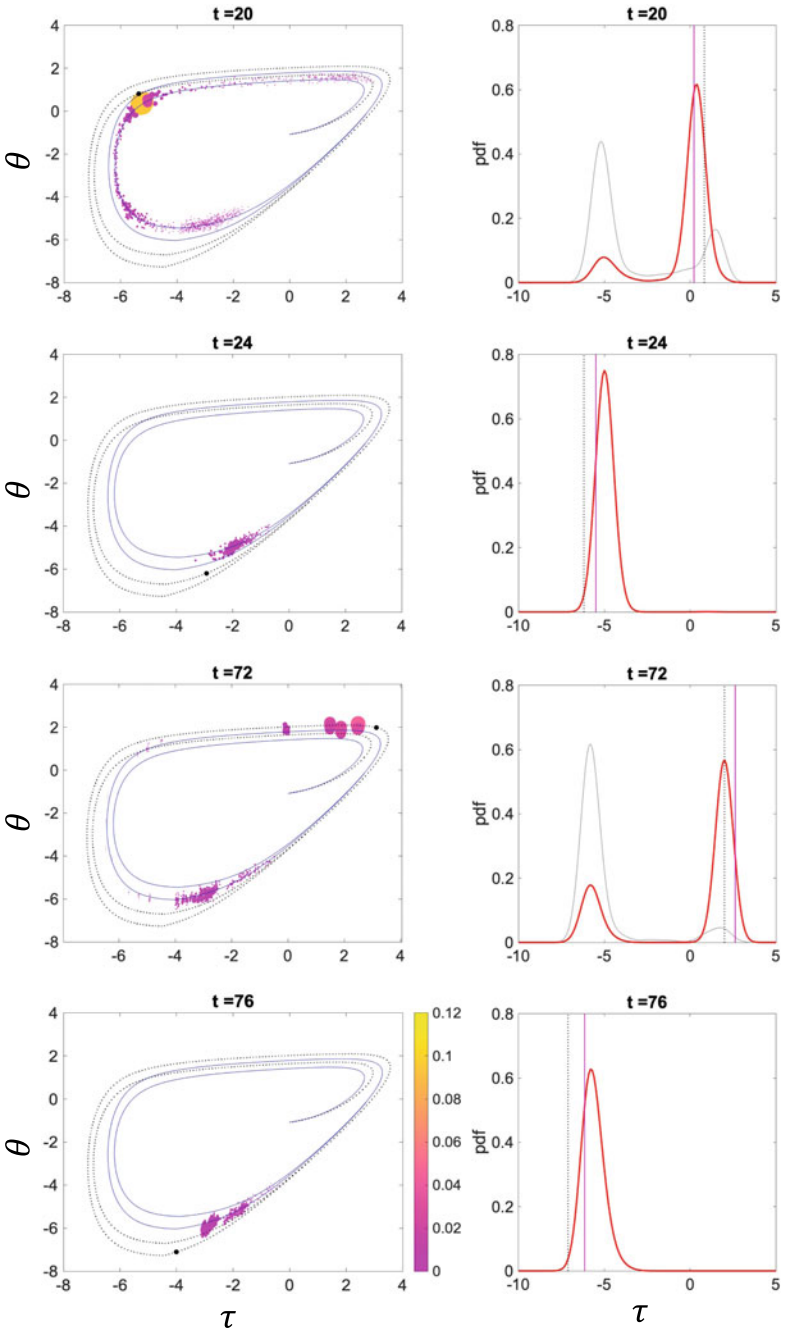

Fig. 19.7 As in Fig. 19.6, but now for Case B

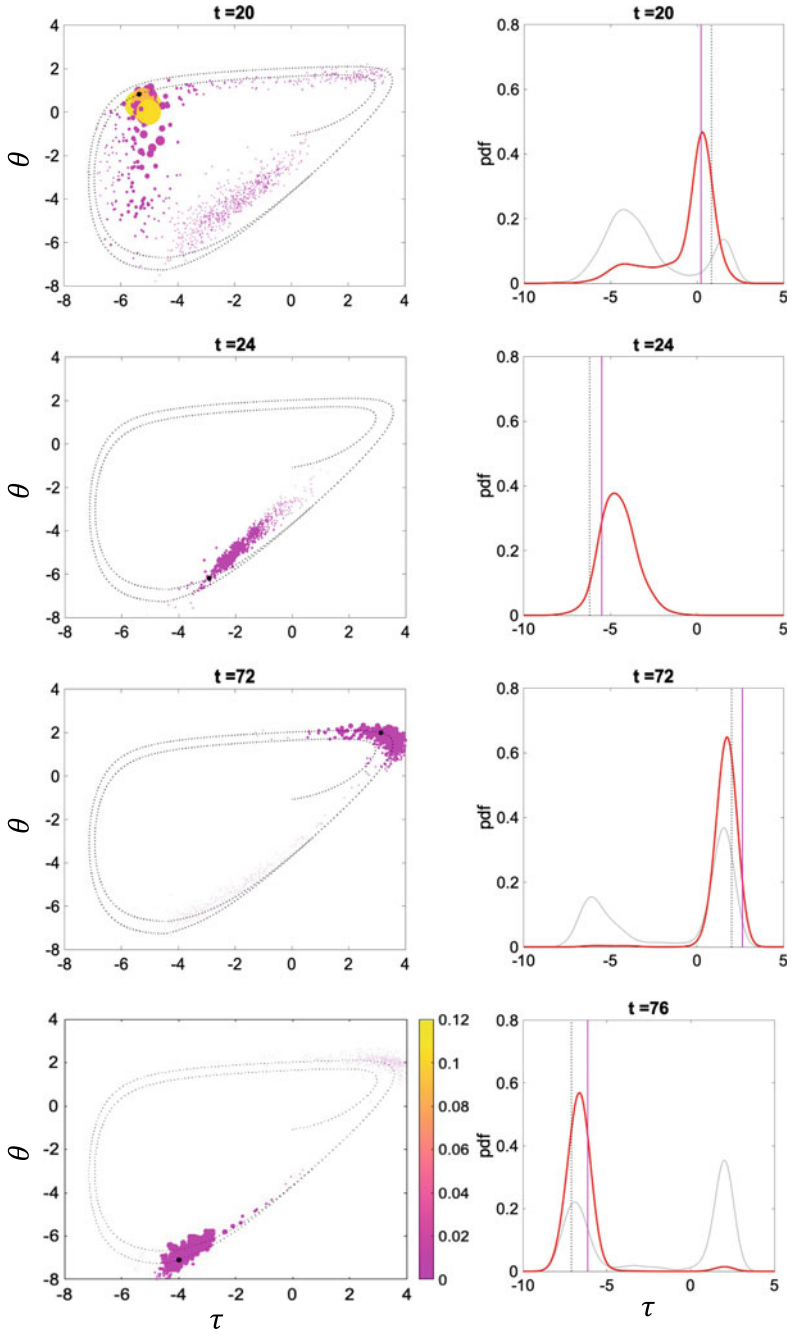

Fig. 19.8 As in Fig. 19.6, but now for Case C

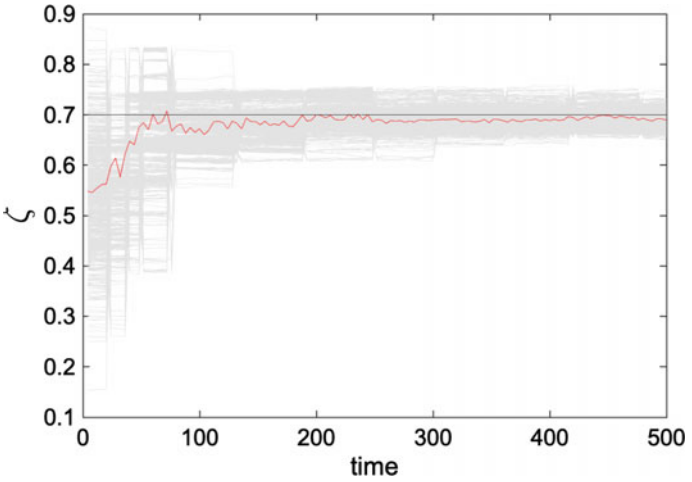

Fig. 19.9 Case C: evolution of the parameter estimation

normal distribution with mean ζ' and standard deviation 0.01. To ensure that the state of each particle is different but consistent with its parameter value, we sample its state from a single realization with this parameter at a random moment in time. This approach is similar to the lagged forecasting by Hoffman and Kalnay (1983). The resulting N particles have N parameter values centered around ζ' and different initial states that are consistent with the parameter of each particle.

Figure 19.8 provides the phase diagrams for Case C's state- and parameter estimation. The distributions in this figure demonstrate that the state- and parameter estimation results in a posterior closer to the truth than in Case A or B. By varying parameter values within the ensemble, the resulting states and phases within the seismic cycle cover a more significant part of the θ-τ domain. The results suggest that, for this case, state- and parameter estimation provides an effective reconstruction of the true τ and θ. However, for the time steps $t = 24$, 72, and 76, none of the particles represented v accurately, and as a result, the weights remained relatively low. Figure 19.9 shows the reconstruction of ζ. After a period with significant adjustment of the posterior parameter value, the estimation of the parameter fluctuates around the true value, and the true value is within the ensemble spread. For this particular case, there is no degeneracy of the ensemble. An artificial evolution equation for the parameters could have resolved the problem if we observed degeneracy.

19.7 Summary

These results indicate that for the particular case of seismic-cycle estimation presented in this chapter, the combined state- and parameter estimation provides a more favorable data-assimilation outcome in the presence of a parameter bias than state

estimation alone. Increasing model error in state estimation can only partly compensate for this bias. We also see that a particle filter handles bimodal pdfs and provides physically acceptable solutions while the EnKF fails. The low dimension of the current example makes the use of a standard particle filter straightforward.