4.3 Estimation

4.3.1 Model and Data Consistency

In many scientific fields, the central issue is to develop understanding from the combination of a skillful model with measurements. The model is intended to encompass all one's theoretical knowledge about how the system behaves, and the data are the complete observational knowledge of the same system. Maxwell's Equations are one example of a model appropriate to electromagnetic phenomena. If done properly, and in the absence of contradictions between theory and experiment, inferences from the model/data combination should be no worse, and may well be very much better, than those made from either alone. It is the latter possibility that motivates the development of state estimation procedures. "Best-estimates" made by combining models with observations are often used to forecast a system (e.g., to land an airplane), but this is by no means the major application.

Such model/data problems are ones of statistical inference with a host of specific subtypes. Some powerful techniques are available, but like any powerful tools (a chain saw, for example), they can be dangerous to the user! In general, one is confronted one with a two-stage problem. Stage 1 involves developing a suitable model that is likely to be consistent with the data. "Consistent" means that within the estimated data errors, the model is likely to be able to describe the features of interest. Obtaining the data errors is itself a serious modeling problem. Stage 2 produces the actual estimate with its error estimates.

One can go very badly wrong at stage 1, before any computation takes place. If elastic wave propagation is modelled using the equations of fluid dynamics, estimation methods will commonly produce some kind of "answer", but one which would be nonsensical. Model failure can of course, be much more subtle, in which some omitted, supposed secondary element (e.g., a time-dependence), proves to be critical to a description of the data. Good technique alerts users to the presence of such failures, along with clues as to what should be changed in the model. But these issues do not however, apply only to the model. The assertion that a particular data set carries a signal of a particular kind, can prove to be false in a large number of ways. A temperature signal thought to represent the seasonal cycle might prove, on careful examination, to be dominated by higher or lower frequency structures, and thus its use with an excellent model of annual variation might prove disastrous. Whether this situation is to be regarded as a data or as a model issue is evidently somewhat arbitrary.

Thus stage 1 of any estimation problem has to involve understanding of whether the data and the model are physically and statistically consistent. If they are not, one should stop and reconsider. Often where they are believed to be generally consistent up to certain quantitative adjustments, one can combine the two stages. A model may have adjustable parameters (turbulent mixing coefficients; boundary condition errors; etc.) which could bring the model and the data into consistency, and then the estimation procedure becomes, in part, an attempt to find those parameters in addition to the state. Alternatively, the data error covariance, $\mathbf{R}(t)$, may be regarded as incompletely known, and one might seek as part of the state estimation procedure to improve one's estimate of it. (Problems like this one are taken up in the next Chapter under the subject of "adaptive filtering.")

Assuming for now that the model and data are likely to prove consistent, one can address what might be thought of as a form of interpolation: Given a set of observations in space and time as described by Equation (4.27), use the dynamics as described by the model (4.12) or (4.13) to estimate various state-vector elements at various times of interest. Yet another, less familiar, problem recognizes that some of the forcing terms $\mathbf{B}(t-1)\mathbf{q}(t-1)$ are partially or wholly unknown (e.g., no one believes that the windstress boundary conditions over the ocean are perfectly known), and one might seek to estimate them from whatever ocean observations are available and from the known model dynamics. Many real problems contain both of these issues.

The forcing terms—representing boundary conditions as well as interior sources/sinks and forces—almost always need to be divided into two elements: the known and the unknown parts. The latter will often be perturbations about the known values. Thus, rewrite (4.13) in the modified form

$$\mathbf{x}(t) = \mathbf{A}(t-1)\,\mathbf{x}(t-1) + \mathbf{B}(t-1)\,\mathbf{q}(t-1) + \mathbf{\Gamma}(t-1)\,\mathbf{u}(t-1), \tag{4.33}$$ {62001}

where now $\mathbf{B}(t)\,\mathbf{q}(t)$ represent the known forcing terms and $\mathbf{\Gamma}(t)\,\mathbf{u}(t)$ the unknown ones, which we will generally refer to as the "controls," or "control terms." $\mathbf{\Gamma}(t)$ is known and plays the same role for $\mathbf{u}(t)$ as does $\mathbf{B}(t)$ for $\mathbf{q}(t)$. Usually $\mathbf{B}(t)$, $\mathbf{\Gamma}(t)$ will be treated as time independent, but this simplification is not necessary. Almost always, we can make some estimate of the size of the control terms, as for example,

$$\langle \mathbf{u}(t) \rangle = \mathbf{0}, \qquad \langle \mathbf{u}(t)\,\mathbf{u}(t)^T \rangle = \mathbf{Q}(t)\,. \tag{4.34}$$ {62002}

The controls have a second, somewhat different, role: They can also represent the model error. All models are inaccurate to a degree—approximations are always made to the equations describing any particular physical situation. One can expect that the person who constructed

the model has some idea of the size and structure of the physics or chemistry, or biology, ..., that have been omitted or distorted in the model construction. In this context, $\mathbf{Q}(t)$ represents the covariance of the model error, and the control terms represent the missing physics. The assumption $\langle \mathbf{u}(t) \rangle = \mathbf{0}$ must be critically examined in this case, and in the event of failure, some modification of the model must be made or the variance artificially modified to attempt to accommodate what is a model bias error. But the most serious problem is that models are rarely produced with *any* quantitative description of their accuracy beyond one or two examples of comparison with known solutions. One is left to determine $\mathbf{Q}(t)$ by guesswork. Getting beyond such guesses is again a problem of adaptive estimation.

Collecting the standard equations of model and data:

$$\mathbf{x}(t) = \mathbf{A}(t-1)\mathbf{x}(t-1) + \mathbf{B}\mathbf{q}(t-1) + \mathbf{\Gamma}\mathbf{u}(t-1), \ 1 \leq t \leq t_f \tag{4.35}$$

$$\mathbf{E}(t)\mathbf{x}(t) + \mathbf{n}(t) = \mathbf{y}(t), \quad 1 \leq t \leq t_f, \quad \Delta t = 1 \tag{4.36}$$

$$\mathbf{n}(t) = \mathbf{0}, \quad \left\langle \mathbf{n}(t)\mathbf{n}(t)^T \right\rangle = \mathbf{R}(t), \quad \left\langle \mathbf{n}(t)\mathbf{n}(t') \right\rangle = \mathbf{0}, \ t \neq t' \tag{4.37}$$

$$\langle \mathbf{u} \rangle = \mathbf{0}, \quad \langle \mathbf{u}(t)\mathbf{u}(t) \rangle = \mathbf{Q}(t) \tag{4.38}$$

$$\tilde{\mathbf{x}}(0) = \mathbf{x}_0, \quad \left\langle (\tilde{\mathbf{x}}(0) - \mathbf{x}(0))(\tilde{\mathbf{x}}(0) - \mathbf{x}(0))^T \right\rangle = \mathbf{P}(0) \tag{4.39}$$

---

where t_f defines the endpoint of the interval of interest. The last equation, (4.39), treats the initial conditions of the model as a special case—the uncertain initialization problem, where $\mathbf{x}(0)$ is the true initial condition and $\tilde{\mathbf{x}}(0) = \mathbf{x}_0$ is the value actually used but with uncertainty $\mathbf{P}(0)$. Alternatively, one could write,

$$\mathbf{E}(0)\mathbf{x}(0) + \mathbf{n}(0) = \mathbf{x}_0, \ \mathbf{E}(0) = \mathbf{I}, \ \left\langle \mathbf{n}(0)\mathbf{n}(0)^T \right\rangle = \mathbf{P}(0), \tag{4.40}$$

and include the initial conditions as a special case of the observations—recognizing explicitly that one often obtains initial conditions from observations.

This general form permits one to grapple with reality. In the spirit of ordinary least squares and its intimate cousin, minimum-error variance estimation, consider the general problem of

finding state vectors and controls, $\mathbf{u}(t)$, that minimize an objective function,

{J1}

$$\begin{aligned} J = & [\mathbf{x}(0) - \mathbf{x}_0]^T \mathbf{P}(0)^{-1} [\mathbf{x}(0) - \mathbf{x}_0] \\ & + \sum_{t=1}^{t_f} [\mathbf{E}(t)\,\mathbf{x}(t) - \mathbf{y}(t)]^T \mathbf{R}(t)^{-1} [\mathbf{E}(t)\,\mathbf{x}(t) - \mathbf{y}(t)] \\ & + \sum_{t=0}^{t_f - 1} \mathbf{u}(t)^T \mathbf{Q}(t)^{-1} \mathbf{u}(t)\,, \end{aligned} \tag{4.41}$$

subject to the model, (4.35, 4.38, 4.39) As written here, this choice of an objective function is somewhat arbitrary but perhaps reasonable as the direct analogue to those used in Chapter 2. It seeks a state vector $\mathbf{x}(t)$, $0 \le t \le t_f$, and a set of controls, $\mathbf{u}(t)$, $0 \le t \le t_f - 1$, that satisfy the model and that agree with the observations to an extent determined by the weight matrices $\mathbf{R}(t)$ and $\mathbf{Q}(t)$, respectively. From the previous discussions of least squares and minimum-error variance estimation, the minimum-square requirement Eq. (4.41) will produce a solution identical to that derived from minimum variance estimation by the specific choice of the weight matrices as the corresponding prior uncertainties, $\mathbf{R}(t)$, $\mathbf{Q}(t)$, $\mathbf{P}(0)$. In a Gaussian system, it also proves to be the maximum likelihood estimate. The introduction of the controls, $\mathbf{u}(t)$, into the objective function represents an acknowledgment that arbitrarily large controls (forces) would not usually be an acceptable solution; they should be consistent with $\mathbf{Q}(t)$.

Notation Note. As in Chapter 2, any values of $\mathbf{x}(t)$, $\mathbf{u}(t)$ minimizing J will be written $\tilde{\mathbf{x}}(t)$, $\tilde{\mathbf{u}}(t)$ and these symbols could be substituted into Eq. (4.41) if that is regarded as clearer.

Much of the rest of this chapter will be directed at solving the problem of finding the minimum of J subject to the solution satisfying the model. Notice that J involves the state vector, the controls, and the observations over the entire time period under consideration, $0 \le t \le t_f$. This type of objective function is the one usually of most interest to scientists attempting to understand their system—in which data are stored and employed over a finite time. In some other applications, most notably forecasting, and which is taken up immediately below, one has only the past measurements available; this situation proves to be a special case of the more general one.

Although we will not keep repeating the warning each time an objective function such as Eq. (4.41) is encountered, the reader is reminded of a general message from Chapter 2: *The assumption that the model and observations are consistent and that the minimum of the objective function produces a meaningful and useful estimate must always be tested after the fact.* That is, at the minimum of J, $\tilde{\mathbf{u}}(t)$ must prove consistent with $\mathbf{Q}(t)$, and $\tilde{\mathbf{x}}(t)$ must produce residuals consistent with $\mathbf{R}(t)$. Failure of these and other posterior tests should lead to rejection of

the model. As always, one can thus reject a model (which includes $\mathbf{Q}(t)$, $\mathbf{R}(t)$) on the basis of a failed consistency with observations. But a model is never "correct" or "valid", merely "consistent**."** (See Note 8, Chapter 1.)

4.3.2 The Kalman Filter

We begin with a special case. Suppose that by some means, at time $t = 0, \Delta t = 1$, we have an unbiased estimate, $\tilde{\mathbf{x}}(0)$, of the state vector with uncertainty $\mathbf{P}(0)$. At time $t = 1$, observations from Eq. (4.36) are available. How would the information available best be used to estimate $\mathbf{x}(1)$?

The model permits a forecast of what $\mathbf{x}(1)$ should be, were $\tilde{\mathbf{x}}(0)$ known perfectly,

{62009}
$$\tilde{\mathbf{x}}(1,-) = \mathbf{A}\tilde{\mathbf{x}}(0) + \mathbf{B}\mathbf{q}(0), \tag{4.42}$$

where the unknown control terms have been replaced by the best estimate we can make of them—their mean, which is zero, and $\mathbf{A}$ has been assumed to be time independent. A minus sign has been introduced into the argument of $\tilde{\mathbf{x}}(1,-)$ to show that *no data at* $t = 1$ *have yet been used to make the estimate* at $t = 1$, in a notation we will generally use. How good is this forecast?

Suppose the erroneous components of $\tilde{\mathbf{x}}(0)$ are,

{62010}
$$\boldsymbol{\gamma}(0) = \tilde{\mathbf{x}}(0) - \mathbf{x}(0)\,, \tag{4.43}$$

then the erroneous components of the forecast are,

$$\begin{aligned} \boldsymbol{\gamma}(1) \equiv \tilde{\mathbf{x}}(1,-) - \mathbf{x}(1) &= \mathbf{A}\tilde{\mathbf{x}}(0) + \mathbf{B}\mathbf{q}(0) - \big(\mathbf{A}\mathbf{x}(0) + \mathbf{B}\mathbf{q}(0) + \boldsymbol{\Gamma}\mathbf{u}(0)\big) \\ &= \mathbf{A}\boldsymbol{\gamma}(0) - \boldsymbol{\Gamma}\mathbf{u}(0)\,, \end{aligned} \tag{4.44}$$

that is, composed of two distinct elements: the propagated erroneous portion of $\tilde{\mathbf{x}}(0)$, and the unknown control term. Their second moments are

{error1}
$$\begin{aligned} \langle \boldsymbol{\gamma}(1)\,\boldsymbol{\gamma}(1)^T \rangle &= \langle \big(\mathbf{A}\boldsymbol{\gamma}(0) - \boldsymbol{\Gamma}\mathbf{u}(0)\big)\big(\mathbf{A}\boldsymbol{\gamma}(0) - \boldsymbol{\Gamma}\mathbf{u}(0)\big)^T \rangle \\ &= \mathbf{A}\langle \boldsymbol{\gamma}(0)\mathbf{P}\,(0)\;\boldsymbol{\gamma}(0)^T \rangle \mathbf{A}^T + \boldsymbol{\Gamma}\langle \mathbf{u}(0)\,\mathbf{u}(0)^T \rangle \boldsymbol{\Gamma}^T \\ &= \mathbf{A}\mathbf{P}\,(0)\,\mathbf{A}^T + \boldsymbol{\Gamma}\mathbf{Q}(0)\boldsymbol{\Gamma}^T \\ &\equiv \mathbf{P}(1,-) \end{aligned} \tag{4.45}$$

by the definitions of $\mathbf{P}(0)$, $\mathbf{Q}(0)$ and the assumption that the unknown controls are not correlated with the error in the state estimate at $t = 0$. We now have an estimate of $\mathbf{x}(1)$ with uncertainty $\mathbf{P}(1,-)$ and a set of observations,

{62013}
$$\mathbf{E}(1)\,\mathbf{x}(1) + \mathbf{n}(1) = \mathbf{y}(1)\,. \tag{4.46}$$

To combine the two sets of information, we use the recursive least-squares solution Eqs. (2.432-2.434). By assumption, the uncertainty in $\mathbf{y}(1)$ is uncorrelated with that in $\tilde{\mathbf{x}}(1,-)$. Making the appropriate substitutions into those equations,

$$\begin{aligned}\tilde{\mathbf{x}}(1) &= \tilde{\mathbf{x}}(1,-) + \mathbf{K}(1)\left[\mathbf{y}(1) - \mathbf{E}(1)\,\tilde{\mathbf{x}}(1,-)\right], \\ \mathbf{K}(1) &= \mathbf{P}(1,-)\,\mathbf{E}(1)^T\left[\mathbf{E}(1)\,\mathbf{P}(1,-)\,\mathbf{E}(1)^T + \mathbf{R}(1)\right]^{-1}\end{aligned} \tag{4.47}$$ {average1}

with new uncertainty,

$$\mathbf{P}(1) = \mathbf{P}(1,-) - \mathbf{K}(1)\,\mathbf{E}(1)\,\mathbf{P}(1,-)\,. \tag{4.48}$$ {62015}

(Compare to the discussion on P. 143.) Eq. (4.47) is best interpreted as being the average of the model estimate with that obtained from the data alone, slightly disguised by rearrangement.

Thus there are four steps:

1. Make a forecast using the model (4.35) with the unknown control terms $\mathbf{\Gamma u}$ set to zero.

2. Calculate the uncertainty of this forecast, Eq. (6.39), which is made up of two separate terms.

3. Do a weighted average (4.47) of the forecast with the observations, the weighting being chosen to reflect the relative uncertainties.

4. Compute the uncertainty of the final weighted average, Eq. (4.48).

Such a computation is called a "Kalman filter**.**"[113], and it is conventionally given a more formal derivation. $\mathbf{K}$ is called the "Kalman gain." At the stage where the forecast (4.42) has already been made, the problem was reduced to finding the minimum of the objective function,

$$\begin{aligned}J = &\left[\tilde{\mathbf{x}}(1,-) - \tilde{\mathbf{x}}(1)\right]^T\mathbf{P}(1,-)^{-1}\left[\tilde{\mathbf{x}}(1,-) - \tilde{\mathbf{x}}(1)\right] \\ &+ \left[\mathbf{y}(1) - \mathbf{E}(1)\,\tilde{\mathbf{x}}(1)\right]^T\mathbf{R}(1)^{-1}\left[\mathbf{y}(1) - \mathbf{E}(1)\,\tilde{\mathbf{x}}(1)\right],\end{aligned} \tag{4.49}$$ {J3}

which is a variation of the objective function used to define the recursive least-squares algorithm (Eq. 2.423). In this final stage, the explicit model has disappeared, being present only implicitly through the uncertainty $\mathbf{P}(1,-)$. After the averaging step, all of the information about the observations has been used too, and is included in $\tilde{\mathbf{x}}\,(t)$, $\mathbf{P}\,(t)$ and the data can be discarded. For clarity, tildes have been placed over all appearances of $\mathbf{x}\,(t)$.

A complete recursion can now be defined through the Equations (4.42)–(4.48), replacing all the $t = 0$ variables with $t = 1$ variables, the $t = 1$ variables becoming $t = 2$ variables, etc. In terms of arbitrary t, the recursion is,

;

$$\tilde{\mathbf{x}}(t,-) = \mathbf{A}(t-1)\,\tilde{\mathbf{x}}(t-1) + \mathbf{B}(t-1)\,\mathbf{q}(t-1), \tag{4.50}$$

$$\mathbf{P}(t,-) = \mathbf{A}(t-1)\,\mathbf{P}(t-1)\,\mathbf{A}(t-1)^T + \mathbf{\Gamma Q}(t-1)\,\mathbf{\Gamma}^T, \tag{4.51}$$

$$\tilde{\mathbf{x}}(t) = \tilde{\mathbf{x}}(t,-) + \mathbf{K}(t)\left[\mathbf{y}(t) - \mathbf{E}(t)\,\tilde{\mathbf{x}}(t,-)\right], \tag{4.52}$$

$$\mathbf{K}(t) = \mathbf{P}(t,-)\,\mathbf{E}(t)^T\left[\mathbf{E}(t)\,\mathbf{P}(t,-)\,\mathbf{E}(t)^T + \mathbf{R}(t)\right]^{-1}, \tag{4.53}$$

$$\mathbf{P}(t) = \mathbf{P}(t,-) - \mathbf{K}(t)\,\mathbf{E}(t)\,\mathbf{P}(t,-), \quad 1 \le t \le t_f. \tag{4.54}$$

These equations are those for the complete Kalman filter. Note that some authors prefer to write equations for $\tilde{\mathbf{x}}(t+1,-)$ in terms of $\tilde{\mathbf{x}}(t)$, etc. Eq. (2.37) permits one to rewrite Eq. (4.54) as,

{kalman6}

$$\mathbf{P}(t)^{-1} = \left[\mathbf{P}(t,-)^{-1} + \mathbf{E}(t)^T\,\mathbf{R}(t)^{-1}\,\mathbf{E}(t)\right], \tag{4.55}$$

and an alternate form for the gain is[114]

$$\mathbf{K}(t) = \mathbf{P}(t)\,\mathbf{E}(t)^T\,\mathbf{R}(t)^{-1}, \tag{4.56}$$

Other rearrangements are possible too. These re-written forms are often important for computational efficiency and accuracy. Note that in the special case where the observations are employed one-at-a-time, $\mathbf{E}(t)$ is a simple row vector, $\mathbf{E}(t)\,\mathbf{P}(t,-)\,\mathbf{E}(t)^T + \mathbf{R}(t)$ is a scalar, and no matrix inversion is required in Eqs. (4.50-4.54). The computation would then be dominated by matrix multiplications. Such a strategy demands that the noise be uncorrelated from one observation to another, or removed by "pre-whitening" which, however, itself often involves a matrix inversion. Various re-arrangements are worth examining in large problems.[115]

Notice that the model is being satisfied exactly; in the terminology introduced in Chapter 2, it is a hard constraint. But again, as was true with the static models, the hard constraint description is somewhat misleading, as the presence of the terms in $\mathbf{u}(t)$ means that model errors are permitted. Notice too, that $\mathbf{u}(0)$ has not been estimated.

Example

Consider again the mass-spring oscillator described earlier, with time history in Figure 4.2. It was supposed that the initial conditions were erroneously provided as $\tilde{\mathbf{x}}(0) = [10, 10]^T$, $\mathbf{P}(0) = \mathrm{diag}([100, 100])$, *but that the forcing was completely unknown. Observations of* $x_1(t)$ *were provided at every time step with a noise variance* $R = 50$. *The Kalman filter was computed by (4.50-4.54) and used to estimate the position at each time step. The result for part of the time history is in Figure 4.2a, showing the true value and the estimated value of component* $x_1(t)$. *The time history of the uncertainty of* $x_1(t)$, $\sqrt{P_{11}(t)}$, *is also depicted and rapidly reaches an symptote.*

Overall, the filter manages to track the position of the oscillator everywhere within two standard deviations.

If observations are not available at some time step, t, the best estimate reduces to that from the model forecast alone, $\mathbf{K}(t) = \mathbf{0}$, $\mathbf{P}(t) = \mathbf{P}(t, -)$ and one simply proceeds. Typically in such situations, the error variances will grow from the accumulation of the unknown $\mathbf{u}(t)$, at least, until such times as an observation does become available. If $\mathbf{u}(t)$ is purely random, the system will undergo a form of random walk [116]

Example.

Consider again the problem of fitting a straight line to data, as discussed in Chapter 2, but now in the context of a Kalman filter, using the canonical form derived from (4.50-4.54). "Data" were generated from the state transition matrix of Eq. (4.16) and an unforced model, as depicted in Figure 4.4. The observation equation is

$$y(t) = x_1(t) + n(t)\,,$$

that is, $\mathbf{E}(t) = \{1 \quad 0\}$, $R(t) = 50$, but observations were available only every 25th time step. *There were no unknown control disturbances—that is* $Q(t) = 0$, *but the initial state estimate was set erroneously as* $\tilde{\mathbf{x}}(0) = [30, 10]^T$ fs *with an uncertainty* $\mathbf{P}(0) = \text{diag}([900, 900)]$.*The result of the computation for the fit is shown in Figure 4.4 for 100 time steps. Note that in the Kalman filter estimate, the estimate diverges rapidly from the true value (although well within the estimated error) and is brought discontinuously toward the true value when the first observations become available.*

If the state vector is redefined to consist of the two model parameters a, b, *then* $\mathbf{x} = [a \quad b]^T$ and $\mathbf{A} = \mathbf{I}$. *Now the observation matrix is* $\mathbf{E} = [1 \quad t]$*–that is, time-dependent. The state vector has changed from a time-varying one to a constant. The same grossly incorrect estimates* $\tilde{\mathbf{x}}(0) = [10 \quad 10]^T$ *were used, with the same* $\mathbf{P}(0)$ *(the correct values are* $a = 1$, $b = 2$*) and with the time histories of the estimates depicted in Figure 4.2. At the end of 100 time steps, we have* $\tilde{a} = 1.85 \pm 2.0$, $\tilde{b} = 2.0 \pm 0.03$, *both of which are consistent with the correct values. For reasons the reader might wish to think about, the uncertainty of the intercept is much greater than for the slope.*

Example

For the mass spring oscillator in Fig. 4.2, it was supposed that the same noisy observations were available, but only at every 25th time step. In general, the presence of the model error, or control uncertainty, accumulates over the 25 time steps as the model is run forward without observations. The expected error of such a system is shown for 150 time steps in Figure 4.6e. Notice (1) the growing envelope as uncertainty accumulates faster than the observations can

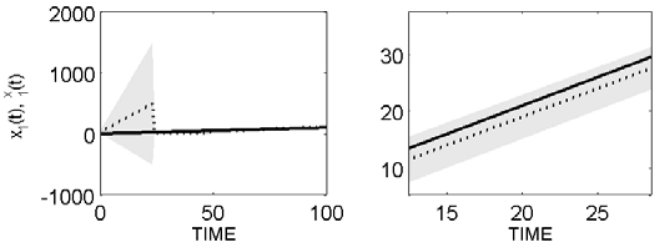

Figure 4.4: A straight line computed from the statespace model with **A** from Eq. (4.16) with no forcing. (a) Solid line shows the true values. Noisy observations were provided every 25th point, and the initial condition was set erroneously to $\tilde{\mathbf{x}}(0) = [30, 10]^T$ with $\mathbf{P}(0) = \text{diag}([900, 900])$. The estimation error grows rapidly away from the incorrect initial conditions until the first observations are obtained. Estimate is shown as the dashed line. (b) Result of applying the RTS smoother to the data and model in (a). Gray band is the one standard deviation error bar. {kfrts3.eps}

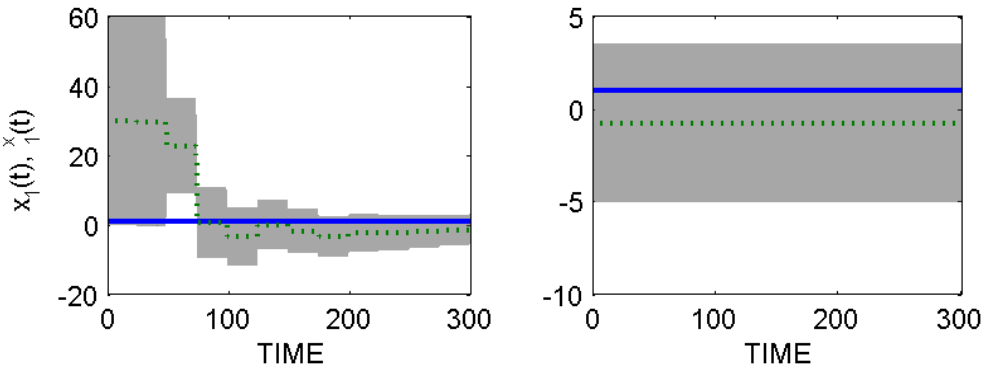

Figure 4.5: (a) $x_1(t)$ (solid) and $\tilde{x}_1(t) = \tilde{a}$ from the straight-line model and the Kalman filter estimate when the state vector was defined to be the intercept and slope, and $\mathbf{E}(t) = [1, t]$. (b) Smoothed estimate, $\tilde{x}_1(t, +)$, and its uncertainty corresponding to (a). {kfrts4.eps}

reduce it; (2) the periodic nature of the error within the growing envelope; and (3) that the envelope appears to be asymptoting to a fixed upper bound for large t. The true and estimated time histories for a portion of the time history are shown in Figure 4.6d. As expected, with fewer available observations, , the misfit of the estimated and true values is larger than with data at every time step. At every 25th point, the error norm drops as observations become available, but with the estimated value of $x_1(t)$ undergoing a jump when the observation is available.

If the observation is that of the velocity $x_1(t)-x_2(t)=\xi(t)-\xi(t-1)$, then $E=\{1 \quad -1\}$. A portion of the time history of the Kalman filtered estimate with a velocity observation available only at every 25th point may be seen in Figure 4.7. Velocity observations are evidently useful for estimating position, owing to the connection between velocity and position provided by the model and is a simple example of how observations of almost anything can be used to improve a state estimate.

A number of more general reformulations of the equations into algebraically equivalent forms are particularly important. In one form, one works not with the covariances, $\mathbf{P}(t,-),\ldots,$ but with their inverses, the so-called information matrices, $\mathbf{P}(t,-)^{-1}$, etc. (See Eq. (4.55).) This "information filter" form may be more efficient if, for example, the information matrices are banded and sparse while the covariance matrices are not. Or, if the initial conditions are infinitely uncertain, $\mathbf{P}(0)^{-1}$ can be set to the zero matrix. In another formulation, one uses the square roots (Cholesky decomposition) of the covariance matrices rather than the matrices themselves. This "square root filter" can be of great importance, as there is a tendency for the computation of the updated values of $\mathbf{P}$ in Eq. (4.54) to become nonpositive-definite owing to round-off errors, and the square root formulation guarantees a positive definite result. [117]

The Kalman filter does *not* produce the minimum of the objective function Eq. (4.41) because the data from times later than t are not being used to make estimates of the earlier values of the state vector or of $\mathbf{u}(t)$. At each step, the Kalman filter is instead minimizing an objective function of the form Eq. (4.49). To obtain the needed minimum, we have to consider what is called the "smoothing problem," to which we will turn in a moment. Note too, that the time history of $\mathbf{x}(t)$ does not satisfy a known equation at the time observations are introduced. When no observation is available, the time evolution obeys the known model equation with zero control term; the averaging step of the filter, however, leads to a change between t and $t-1$ that compensates for the accumulated error. The evolution equation is no longer satisfied in this interval.

The Kalman filter is, nonetheless, extremely important in practice for many problems. In particular, if one must literally make a forecast (e.g., such filters are used to help land airplanes

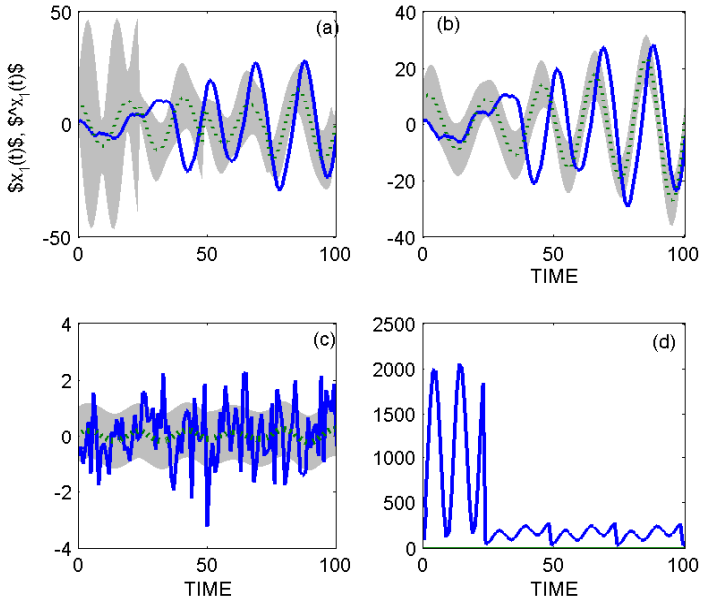

Figure 4.6: For the same model as in Fig. 4.2, except that noisy observations were available only every 25th point ($R = 50$). (a) Shows the correct trajectory of $x_1(t)$ for 100 time steps, the dotted line shows the filter estimate and the shaded band is the standard error of the estimate. (b) Displays the correct value of $x_1(t)$ compared (dotted) to the RTS smoother value with its standard error. (c) Is the estimated control (dotted) with its standard error, and the true value applied to mass spring oscillator. (d) Shows the behavior of $P_{11}(t)$ for the Kalman filter with very large values (oscillating with twice the frequency of the oscillator) and which become markedly reduced as soon as the first observations become available at the 25th point. {kfrts2.eps}

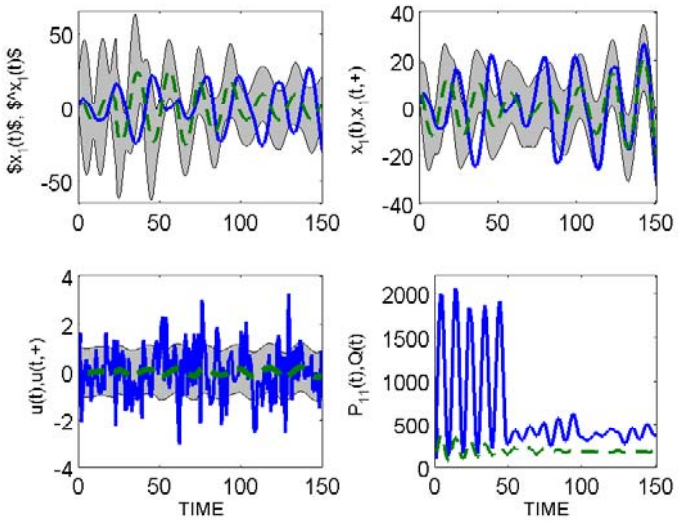

Figure 4.7: (a) $x_1(t)$ and Kalman filter estimate (dotted) when the noisy observations ($R = 50$) are of the velocity ($\mathbf{E} = [1, -1]$) every 25th point for the mass-spring oscillator. (b) RTS smoother estimate (dotted) and its uncertainty corresponding to (a). (c) shows the estimated control (dotted) and the correct value (solid). As seen previously, the high frequency variability in $u(t)$ is not detected by the moving mass, but only an integrated version. (d) $P_{11}(t)$ corresponding to the standard error in (a), and $P_{11}(t, +)$ corresponding to that in (b). {kfrts5.eps}

or, in a primitive way, to forecast the weather), then the future data are simply unavailable, and the state estimate made at time t, using data up to and including time t, is the best one can do.[118]

For estimation, the Kalman filter is only a first step—owing to its failure to use data from the formal future. It also raises questions about computational feasibility. As with all recursive estimators, the uncertainties $\mathbf{P}(t,-),\mathbf{P}(t)$ must be available so as to form the weighted averages. If the state vector contains N elements, then the model, (4.50), requires multiplying an N-dimensional vector by an $N \times N$ matrix at each time step. The covariance update (4.51) requires updating each of N columns of $\mathbf{P}(t)$ in the same way, and then doing it again (i.e., in practice, one forms $\mathbf{A}(t)\mathbf{P}\,(t)$, transposes it, and forms $\mathbf{A}\,(t)\left(\mathbf{A}(t)\mathbf{P}\,(t)\right)^T$, equivalent to running the model $2N$ times at each time step). In many applications, particularly in geophysical fluids, this covariance update step dominates the calculation, renders it impractical, and leads to some of the approximate methods taken up presently.

The Kalman filter was derived heuristically as a simple generalization of the ideas used in Chapter 2. Unsurprisingly, the static inverse results are readily recovered from the filter in various limits. As one example, consider the nearly noise-free case in which both process and observation noise are very small, i.e. $\|\mathbf{Q}\|, \|\mathbf{R}\| \to 0$. Then if $\mathbf{P}(t,-)$ is nearly diagonal, $\mathbf{P}(t,-) \sim \delta^2\mathbf{I}$, and

$$\mathbf{K}(t) \longrightarrow \mathbf{E}^T\left(\mathbf{E}\mathbf{E}^T\right)^{-1},$$

assuming existence of the inverse and,

{limit1}

$$\begin{aligned}\tilde{\mathbf{x}}(t) &\sim \mathbf{A}\tilde{\mathbf{x}}(t-1)+\mathbf{B}\mathbf{q}(t-1)\\ &\quad+\mathbf{E}^T(\mathbf{E}\mathbf{E}^T)^{-1}\{\mathbf{y}(t)-\mathbf{E}\left[\mathbf{A}\tilde{\mathbf{x}}(t-1)+\mathbf{B}\mathbf{q}(t-1)\right]\}\\ &=\mathbf{E}^T(\mathbf{E}\mathbf{E}^T)^{-1}\mathbf{y}(t)\\ &\quad+\left[\mathbf{I}-\mathbf{E}^T\left(\mathbf{E}\mathbf{E}^T\right)^{-1}\mathbf{E}\right]\,\left[\mathbf{A}\tilde{\mathbf{x}}(t-1)+\mathbf{B}\mathbf{q}(t-1)\right]\,.\end{aligned} \tag{4.57}$$

$\mathbf{E}^T\left(\mathbf{E}\mathbf{E}^T\right)^{-1}\mathbf{y}(t)$ is just the expression in Eq. (2.96) for the direct estimate of $\mathbf{x}(t)$ from a set of underdetermined full-rank, noise-free observations. It is the static estimate we would use at time t if no dynamics were available. The columns of $\mathbf{I}-\mathbf{E}^T\left(\mathbf{E}\mathbf{E}^T\right)^{-1}\mathbf{E}$ are the nullspace of $\mathbf{E}$ (recall the definition of $\mathbf{H}$ in Eq. (2.98)) and (4.57) thus employs only those elements of the forecast lying in the nullspace of the observations—a sensible result given that the observations here produce perfect estimates of components of $\mathbf{x}(t+1)$ in the range of $\mathbf{E}$. Thus, in this particular limit, the Kalman filter computes from the noise-free observations those elements of $\mathbf{x}(t+1)$ that it can, and for those which cannot it cannot, it forecasts them from the dynamics. The reader ought to examine other limiting cases–retaining process and/or observational noise–including the behavior of the error covariance propagation.

Example.

It is interesting to apply some of these expressions to the simple problem of finding the mean of a set of observations, considered before on P. 137. The model is of an unchanging scalar mean,

$$x(t) = x(t-1)$$

observed in the presence of noise,

$$y(t) = x(t) + n(t)$$

where $\langle n(t)^2 \rangle = R$, *so* $E = 1$, $A = 1$, $Q = 0$. *In contrast to the situation on P. 137, the machinery we have developed here requires that the noise be uncorrelated:* $\langle n(t)n(t') \rangle = 0$, $t \neq t'$, *although as already mentioned, methods exist to overcome this restriction. Suppose that the initial estimate of the mean is* 0—*that is,* $\tilde{x}(0) = 0$, *with uncertainty* $P(0)$. *Eq. (4.51) is* $P(t,-) = P(t-1)$, *and the Kalman filter uncertainty, in the form (4.55), is*

$$\frac{1}{P(t)} = \frac{1}{P(t-1)} + \frac{1}{R},$$

a difference equation, with known initial condition, whose solution by inspection is,

$$\frac{1}{P(t)} = \frac{t}{R} + \frac{1}{P(0)}.$$

Using (4.52) with $E = 1$, *and successively stepping forward produces,*[119]

$$\tilde{x}(t) = \frac{R}{R + tP(0)} \left\{ \frac{P(0)}{R} \sum_{j=1}^{t} y(j) \right\}, \tag{4.58}$$ {P1}

whose limit as $t \to \infty$ *is*

$$\tilde{x}(t) \longrightarrow \frac{1}{t} \sum_{j=1}^{t} y(j),$$

the simple average, with uncertainty $P(t) \to 0$, *as* $t \to \infty$. *If there is no useful estimate available of* $P(0)$, *rewrite Eq.* (4.58) *as,*

$$\tilde{x}(t) = \frac{R}{R/P(0) + t} \left\{ \frac{1}{R} \sum_{j=1}^{t} y(j) \right\}, \tag{4.59}$$

and take the agnostic limit, $1/P(0) \to 0$, *or*

$$\tilde{x}(t) = \frac{1}{t} \left\{ \sum_{j=1}^{t} y(j) \right\}, \tag{4.60}$$

which is wholly conventional. (Compare these results to those on P. 138. The problem and result here are necessarily identical to that on P. 137, except that now we explicitly identified $x(t)$ as *a statevector rather than as a constant. Kalman filters with static models are just recursive least-squares solutions.)*

4.3.3 The Smoothing Problem

The Kalman filter permits one to make an optimal forecast from a linear model, subject to the accuracy of the various assumptions being made. Between observation times, the state estimate evolves smoothly according to the model dynamics. But when observations become available, the averaging can draw the combined state estimate abruptly towards the observations, and in the interval between the last unobserved state and the new one, model evolution is not followed. To obtain a state trajectory that is both consistent with model evolution and the data at all times, we need to eliminate the state estimate jumps at the observation times, and solve the problem as originally stated. Minimization of J in Eq. (4.41) subject to the model is still the goal. Begin the discussion by again considering a one-step process[120]. Consider the problem Eqs. (4.35-4.39), but where there are only two times involved, $t = 0,\ 1$. There is an initial estimate $\tilde{\mathbf{x}}(0)$, $\tilde{\mathbf{u}}(0) \equiv 0$ with uncertainties $\mathbf{P}(0), \mathbf{Q}(0)$ for the initial state and control vectors respectively, a set of measurements at time-step 1, and the model. The objective function is,

{J4}
$$\begin{aligned} J = &\left[\tilde{\mathbf{x}}(0,+) - \tilde{\mathbf{x}}(0)\right]^T \mathbf{P}(0)^{-1} \left[\tilde{\mathbf{x}}(0,+) - \tilde{\mathbf{x}}(0)\right] \\ &+ \left[\tilde{\mathbf{u}}(0,+) - \tilde{\mathbf{u}}(0)\right]^T \mathbf{Q}(0)^{-1} \left[\tilde{\mathbf{u}}(0,+) - \tilde{\mathbf{u}}(0)\right] \\ &+ \left[\mathbf{y}(1) - \mathbf{E}(1)\,\tilde{\mathbf{x}}(1)\right]^T \mathbf{R}(1)^{-1} \left[\mathbf{y}(1) - \mathbf{E}(1)\,\tilde{\mathbf{x}}(1)\right] , \end{aligned} \tag{4.61}$$

subject to the model,

{62028}
$$\tilde{\mathbf{x}}(1) = \mathbf{A}(0)\,\tilde{\mathbf{x}}(0,+) + \mathbf{B}(0)\mathbf{q}(0) + \mathbf{\Gamma}\tilde{\mathbf{u}}(0,+)\,, \tag{4.62}$$

with the weight matrices again chosen as the inverses of the prior covariances. A minimizing solution to this objective function would produce a new estimate of $\mathbf{x}(0)$, denoted $\tilde{\mathbf{x}}(0,+)$, with error covariance $\mathbf{P}(0,+)$; the + denotes use of future observations, $\mathbf{y}(1)$, in the estimate. On the other hand, we would still denote the estimate at $t = 1$ as $\tilde{\mathbf{x}}(1)$, coinciding with the Kalman filter estimate, because only data prior to and at the same time would have been used. The estimate $\tilde{\mathbf{x}}(1)$ must be given by Eq. (4.47), but it remains to improve $\tilde{\mathbf{u}}(0)$, $\tilde{\mathbf{x}}(0)$, while simultaneously eliminating the problem of the estimated statevector jump at the filter averaging (observation) time.

The basic issue can be understood by observing that the initial estimates $\tilde{\mathbf{u}}(0) = \mathbf{0}$, $\tilde{\mathbf{x}}(0)$ lead to a model forecast that disagrees with the final best estimate $\tilde{\mathbf{x}}(1)$. If either of $\tilde{\mathbf{u}}(0)$, or $\tilde{\mathbf{x}}(0)$ were known perfectly, the forecast discrepancy could be ascribed to the other one, permitting ready computation of the new required value. In practice, both are somewhat uncertain, and the modification must be partitioned between them; one would not be surprised to find that the partitioning proves to be proportional to their initial uncertainties.

To find the stationary point (we will not trouble to prove it a minimum rather than a maximum), set the differential of J with respect to $\tilde{\mathbf{x}}(0,+)$, $\tilde{\mathbf{x}}(1)$, $\tilde{\mathbf{u}}(0,+)$ to zero,

$$\begin{aligned}\frac{dJ}{2} &= d\tilde{\mathbf{x}}(0,+)^T\mathbf{P}(0)^{-1}\left[\tilde{\mathbf{x}}(0,+)-\tilde{\mathbf{x}}(0)\right]\\ &\quad + d\tilde{\mathbf{u}}(0,+)^T\mathbf{Q}(0)^{-1}\left[\tilde{\mathbf{u}}(0,+)-\tilde{\mathbf{u}}(0)\right]\\ &\quad - d\tilde{\mathbf{x}}(1)^T\mathbf{E}(1)^T\mathbf{R}(1)^{-1}\left[\mathbf{y}(1)-\mathbf{E}(1)\tilde{\mathbf{x}}(1)\right]=0\,.\end{aligned} \tag{4.63}$$

The coefficients of the differentials cannot be set to zero separately because they are connected via the model, Eq. (4.62), which provides a relationship

$$d\tilde{\mathbf{x}}(1) = \mathbf{A}(0)\,d\tilde{\mathbf{x}}(0,+) + \mathbf{\Gamma}(0)\,d\tilde{\mathbf{u}}(0,+). \tag{4.64}$$ {62030}

Eliminating $d\tilde{\mathbf{x}}(1)$,

$$\begin{aligned}\frac{dJ}{2} &= d\tilde{\mathbf{x}}(0,+)^T\{\mathbf{P}(0)^{-1}\left[\tilde{\mathbf{x}}(0,+)-\tilde{\mathbf{x}}(0)\right]\\ &\qquad - \mathbf{A}(0)^T\mathbf{E}(1)^T\mathbf{R}(1)^{-1}\left[\mathbf{y}(1)-\mathbf{E}(1)\tilde{\mathbf{x}}(1)\right]\}\\ &\quad + d\tilde{\mathbf{u}}(0,+)^T\{\mathbf{Q}(0)^{-1}\left[\tilde{\mathbf{u}}(0,+)-\tilde{\mathbf{u}}(0)\right]\\ &\qquad + \mathbf{\Gamma}^T(0)\mathbf{E}(1)^T\mathbf{R}(1)^{-1}\left[\mathbf{y}(1)-\mathbf{E}(1)\tilde{\mathbf{x}}(1)\right]\}\,.\end{aligned} \tag{4.65}$$

dJ vanishes, producing a stationary value of J, only if the coefficients of $d\tilde{\mathbf{x}}(0,+)$, $d\tilde{\mathbf{u}}(0,+)$ separately vanish yielding,

$$\tilde{\mathbf{x}}(0,+) = \tilde{\mathbf{x}}(0) + \mathbf{P}(0)\mathbf{A}(0)^T\mathbf{E}(1)^T\mathbf{R}(1)^{-1}\left[\mathbf{y}(1)-\mathbf{E}(1)\tilde{\mathbf{x}}(1)\right], \tag{4.66}$$ {62032}

$$\tilde{\mathbf{u}}(0,+) = \tilde{\mathbf{u}}(0) + \mathbf{Q}(0)\mathbf{\Gamma}(0)^T\mathbf{E}(1)^T\mathbf{R}(1)^{-1}\left[\mathbf{y}(1)-\mathbf{E}(1)\tilde{\mathbf{x}}(1)\right], \tag{4.67}$$ {62033}

and

$$\begin{aligned}\tilde{\mathbf{x}}(1) &= \tilde{\mathbf{x}}(1,-)\\ &\quad + \mathbf{P}(1,-)\mathbf{E}(1)^T\left[\mathbf{E}(1)\mathbf{P}(1,-)\mathbf{E}(1)^T+\mathbf{R}(1)\right]^{-1}\left[\mathbf{y}(1)-\mathbf{E}(1)\tilde{\mathbf{x}}(1,-)\right],\end{aligned} \tag{4.68}$$ {kalman7}

using the previous definitions of $\tilde{\mathbf{x}}(1,-)$, $\mathbf{P}(1,-)$. As anticipated, Eq. (4.68) is recognizable as the Kalman filter estimate. At this point we are essentially done: An estimate has been produced not only of $\mathbf{x}(1)$, but an improvement has been made in the prior estimate of $\mathbf{x}(0)$ using the future measurements, and we have estimated the control term. Notice the corrections to $\tilde{\mathbf{u}}(0)$, $\tilde{\mathbf{x}}(0)$ are proportional to $\mathbf{Q}(0)$, $\mathbf{P}(0)$, respectively, as anticipated. We still need to examine the uncertainties of these latter quantities.

First rewrite the estimates (4.66)–(4.67) as

$$\begin{aligned}\tilde{\mathbf{x}}(0,+) &= \tilde{\mathbf{x}}(0) + \mathbf{L}(1)\,\left[\tilde{\mathbf{x}}(1)-\tilde{\mathbf{x}}(1,-)\right], \quad \mathbf{L}(1)=\mathbf{P}(0)\mathbf{A}(0)^T\mathbf{P}(1,-)^{-1},\\ \tilde{\mathbf{u}}(0,+) &= \tilde{\mathbf{u}}(0) + \mathbf{M}(1)\,\left[\tilde{\mathbf{x}}(1)-\tilde{\mathbf{x}}(1,-)\right], \quad \mathbf{M}(1)=\mathbf{Q}(0)\mathbf{\Gamma}(0)^T\mathbf{P}(1,-)^{-1},\end{aligned} \tag{4.69}$$

which can be done by extended, but uninteresting, algebraic manipulation. The importance of these latter two expressions is that both $\tilde{\mathbf{x}}(0,+)$, $\tilde{\mathbf{u}}(0,+)$are expressed in terms of their prior estimates in a weighted average with the difference between the prediction of the state at $t = 1$, $\tilde{\mathbf{x}}(1,-)$ and what was actually estimated there following the data use, $\tilde{\mathbf{x}}(1)$. (But the data do not appear explicitly in (4.69).) It is also possible to show that,

{kalman7b}

$$\begin{aligned} \mathbf{P}(0,+) &= \mathbf{P}(0) + \mathbf{L}(1)\left[\mathbf{P}(1) - \mathbf{P}(1,-)\right]\mathbf{L}(1)^T, \\ \mathbf{Q}(0,+) &= \mathbf{Q}(0) + \mathbf{M}(1)\left[\mathbf{P}(1) - \mathbf{P}(1,-)\right]\mathbf{M}(1)^T. \end{aligned} \tag{4.70}$$

Based upon this one-step derivation, a complete recursion for any time interval can be inferred. Suppose the Kalman filter has been run all the way to a terminal time, t_f. The result is $\tilde{\mathbf{x}}(t_f)$ and its variance $\mathbf{P}(t_f)$. With no future data available, $\tilde{\mathbf{x}}(t_f)$ cannot be further improved. At time $t_f - 1$, we have an estimate $\tilde{\mathbf{x}}(t_f - 1)$ with uncertainty $\mathbf{P}(t_f - 1)$, which could be improved by knowledge of the future observations at t_f. But this situation is precisely the one addressed by the objective function Eq. (4.61) with $t = 1 \rightarrow t_f$, and $t = 0 \rightarrow t_f - 1$. Now having improved the estimate at $t_f - 1$ and calling it $\tilde{\mathbf{x}}(t_f - 1, +)$ with uncertainty $\mathbf{P}(t_f - 1, +)$, this new estimate is used to improve the prior estimate $\tilde{\mathbf{x}}(t_f - 2)$, and we step all the way back to $t = 0$. The complete recursion is,

---

$$\begin{aligned} \tilde{\mathbf{x}}(t,+) &= \tilde{\mathbf{x}}(t) + \mathbf{L}(t+1)\left[\tilde{\mathbf{x}}(t+1,+) - \tilde{\mathbf{x}}(t+1,-)\right], \\ \mathbf{L}(t+1) &= \mathbf{P}(t)\mathbf{A}(t)^T\mathbf{P}(t+1,-)^{-1} \end{aligned} \tag{4.71}$$

$$\begin{aligned} \tilde{\mathbf{u}}(t,+) &= \tilde{\mathbf{u}}(t) + \mathbf{M}(t+1)\left[\tilde{\mathbf{x}}(t+1,+) - \tilde{\mathbf{x}}(t+1,-)\right], \\ \mathbf{M}(t+1) &= \mathbf{Q}(t)\Gamma(t)^T\mathbf{P}(t+1,-)^{-1}, \end{aligned} \tag{4.72}$$

$$\mathbf{P}(t,+) = \mathbf{P}(t) + \mathbf{L}(t+1)\left[\mathbf{P}(t+1,+) - \mathbf{P}(t+1,-)\right]\mathbf{L}(t+1)^T, \tag{4.73}$$

$$\mathbf{Q}(t,+) = \mathbf{Q}(t) + \mathbf{M}(t+1)\left[\mathbf{P}(t+1,+) - \mathbf{P}(t+1,-)\right]\mathbf{M}(t+1)^T, \tag{4.74}$$

with $\tilde{\mathbf{x}}(t_f,+) \equiv \tilde{\mathbf{x}}(t_f)$, $\mathbf{P}(t_f,+) \equiv \mathbf{P}(t_f)$.

This recipe, which uses the Kalman filter on a first forward sweep to the end of the available data, and which then successively improves the prior estimates by sweeping backwards, is called the "RTS algorithm" or smoother.[121] The particular form has the advantage that the data are

not involved in the backward sweep, because all of the available information has been used in the filter calculation. It does have the potential disadvantage of requiring the storage at each time step of $\mathbf{P}(t)$. ($\mathbf{P}(t,-)$ is readily recomputed, without $\mathbf{y}(t)$, from (4.51) and need not be stored.) By direct analogy with the one-step objective function, the recursion (4.71)–(4.74) is seen to be the solution to the minimization of the objective function (4.61) subject to the model. Most important, assuming consistency of all assumptions, the resulting state vector trajectory $\tilde{\mathbf{x}}(t,+)$ now satisfies the model and no longer displays the jump discontinuities at observation times of the Kalman filter estimate.

As with the Kalman filter, it is possible to examine limiting cases of the RTS smoother. For example, suppose again that $\mathbf{Q}$ vanishes, and $\mathbf{A}^{-1}$ exists. Then,

$$\mathbf{L}(t+1) \longrightarrow \mathbf{P}(t)\mathbf{A}^T\left(\mathbf{A}\mathbf{P}(t)\mathbf{A}^T\right)^{-1} = \mathbf{A}^{-1}, \tag{4.75}$$ {62041}

and Equation (4.71) becomes

$$\tilde{\mathbf{x}}(t,+) \longrightarrow \mathbf{A}^{-1}\left[\tilde{\mathbf{x}}(t+1,+) - \mathbf{B}\mathbf{q}(t)\right], \tag{4.76}$$ {62042}

a sensible backward estimate obtained by simply solving,

$$\tilde{\mathbf{x}}(t+1) = \mathbf{A}\tilde{\mathbf{x}}(t) + \mathbf{B}\mathbf{q}(t), \tag{4.77}$$ {62043}

for $\tilde{\mathbf{x}}(t)$. Other limits are also illuminating but are left to the reader.

Example.

The smoother result for the straight-line model (4.15) *is shown in Figures 4.4, 4.5 for both forms of state vector. The time-evolving estimate is now a nearly perfect straight line, whose uncertainty has a terminal value equal to that for the Kalman filter estimate, as it must, and reaches a minimum near the middle of the estimation period, before growing again toward* $t = 0$*, where the initial uncertainty was very large. In the case where the state vector consisted of the constant intercept and slope of the line, both smoothed estimates are seen, in contrast to the filter estimate, to conform very well to the known true behavior. It should be apparent that the best-fitting, straight-line solution of Chapter 2 is also the solution to the smoothing problem, but with the data and model handled all at once, a whole-domain method, rather than sequentially.*

Figures 4.2, 4.6, 4.7 shows the state estimate for the mass-spring oscillator made from a smoothing computation run backward from $t = 300$*, and its variance. On average, the smoothed estimate is closer to the correct value than is the filtered estimate, as expected. The standard error is also smaller for the smoothed estimate. The figures display the variance,* $Q_{11}(t,+)$*, of the estimate one can make of the scalar control variable* $u(t)$*.* $\tilde{u}(t)$ *does not show the high frequency variability present in practice, because the mass-spring oscillator integrates the one*

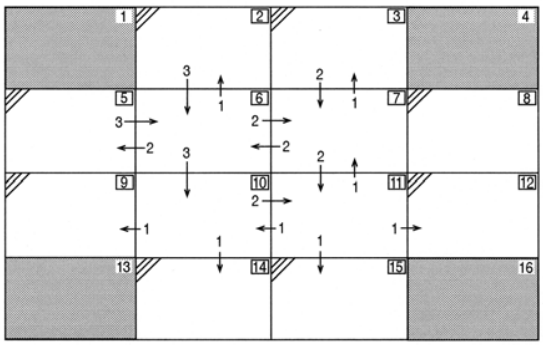

Figure 4.8: Tracer box model where J_{ij} represent fluxes between boxes and are chosen to be mass conserving. Boxes with shaded corners are boundary boxes with externally prescribed concentrations. Numbers in upper right corner are used to identify the boxes. Stippled boxes are unconnected and completely passive here.

{ocip_6a.tif}

time-step variability in such a way that only an integrated value affects the state vector. But the estimated value nonetheless always within two standard errors of the correct value.

Example.

Consider a problem stimulated by the need to extract information from transient tracers, C, in a fluid, which are assumed to satisfy an equation„

{62044}

$$\frac{\partial C}{\partial t} + \mathbf{v} \cdot \nabla C - \kappa \nabla^2 C = -\lambda C + q(\mathbf{r}, t), \tag{4.78}$$

where q represents sources/sinks and λ is a decay rate if the tracer is radioactive. To have a simple model that will capture the structure of this problem, the fluid is divided into a set of boxes as depicted in Figure 4.8. The flow field, as depicted there, is represented by exchanges between boxes given by the $J_{ij} \geq 0$. That is, the J_{ij} are a simplified representation of the effects of advection and mixing on a dye C. (A relationship can be obtained between such simple parameterizations and more formal and elaborate finite-difference schemes. Here, it will only be remarked that J_{ij} are chosen to be mass conserving so that the sum over all J_{ij} entering and

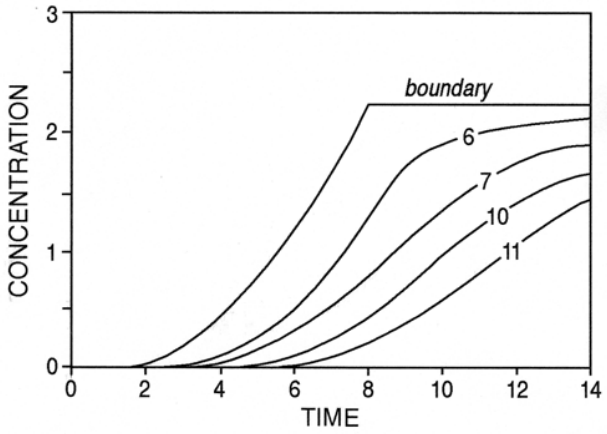

Figure 4.9: Time histories of the forward computation in which boundary concentrations shown were prescribed, and values computed from the forward model for boxes 6,7,10,11. These represent the "truth".

{cip6_6b.tif}

leaving a box vanishes.) The discrete analogue of (4.78) *is taken to be,*

$$C_i(t+1) = C_i(t) - \lambda \Delta t C_i(t) - \frac{\Delta t}{V} \sum_{j \in N(i)} C_i(t) J_{ij} + \frac{\Delta t}{V} \sum_{j \in N(i)} C_j(t) J_{ji} \tag{4.79}$$

where the notation $j \in N(i)$ *denotes an index sum over the neighboring boxes to box* i*,* V *is a volume for the boxes, and* Δt *is the time step. This model can easily be put into the canonical form,*

$$\mathbf{x}(t) = \mathbf{A}\mathbf{x}(t-1) + \mathbf{B}\mathbf{q}(t-1) + \mathbf{\Gamma}\mathbf{u}(t-1)\,, \qquad \mathbf{Q} = \mathbf{0}\,, \tag{4.80}$$

{62046}

with the state vector being the vector of box concentrations $C_i(t)$*,* $C_i(t-1)$*.*

A forward computation was run with initial concentrations everywhere of 0*, using the boundary conditions depicted in Figure 4.9, resulting in interior box values as shown. Based upon these correct values, noisy "observations" of the interior boxes only were constructed at times* $t = 5$*,* 9*,* 12*.*

An initial estimate of tracer concentrations at $t = 0$ *was taken (correctly) to be zero, but this estimate was given a large variance (diagonal* $\mathbf{P}(0)$ *with large norm). The boundary box concentrations were set erroneously to* $C = 2$ *for all* t *and held at that value. A Kalman filter computation was run as shown in Figure 4.10. Initially, the interior box concentration estimates rise erroneously (owing to the dye leaking in from the high nonzero concentrations in the boundary boxes). At* $t = 5$*, the first set of observations becomes available, and the combined*

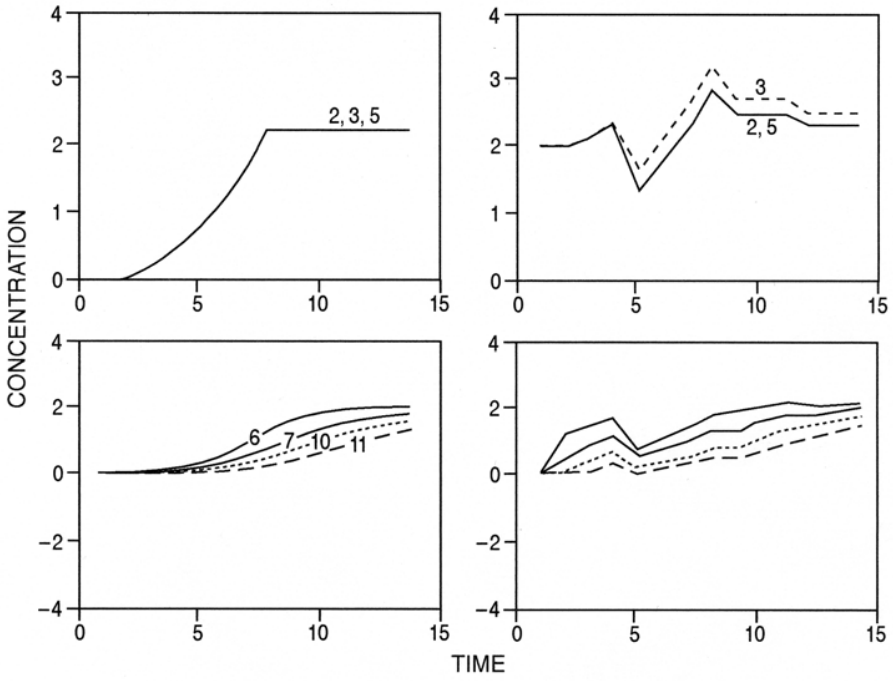

Figure 4.10: The left two panels show the correct values for boundary boxes (upper panel) and interior ones (lower panel). The right panels show the results of the Kalman filter for boundary (upper panel) and interior boxes (lower panel) when noisy observations were provided at $t = 5, 9, 12$. At the observation times, estimates are pulled abruptly toward the observations. By the time the last observations are used, estimated and correct values are quite close. Although not displayed here, there is an uncertainty estimate at all times.

cip6_6c.tif}

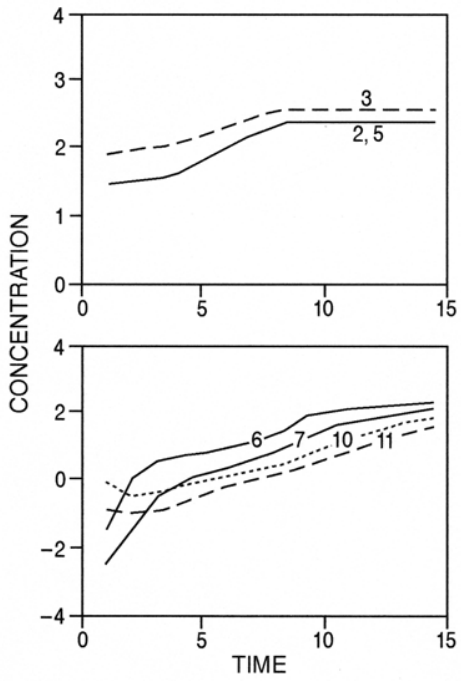

Figure 4.11: Smoothed estimates of boundary (upper panel) and interior (lower panel) values of tracers. The unphysical negative values for small t are indistinguishable from zero within the error estimates.

cip6_6d.tif}

estimate is driven much closer to the true values. By the time the last set of observations is used, the estimated and correct concentrations are quite close, although the time history of the interior is somewhat in error. The RTS algorithm was then applied to generate the smoothed histories shown in Figure 4.11 and to estimate the boundary concentrations (the controls). As expected, the smoothed estimates are closer to the true time history than are the filtered ones (the uncertainty estimates are not shown, but the results are consistent with the "truth" within statistical expectation). Unless further information is provided, no other estimation procedure could do better, given that the model is the correct one.

Other versions of smoother algorithms exist. Consider one other approach. Suppose the Kalman filter has been run forward to some time t_c, producing an estimate $\tilde{\mathbf{x}}(t_c)$ with uncertainty $\mathbf{P}(t_c)$. Now suppose, perhaps on the basis of some further observations, that at a *later* time t_f, an independent estimate $\tilde{\mathbf{x}}(t_f)$ has been made, with uncertainty $\mathbf{P}(t_f)$. The independence is crucial—we suppose this latter estimate is made without using any observations at time t_c or earlier so that any errors in $\tilde{\mathbf{x}}(t_c)$ and $\tilde{\mathbf{x}}(t_f)$ are uncorrelated.

Let us run the model *backward* in time from t_f to $t_f - 1$:

{62047}
$$\tilde{\mathbf{x}}_b(t_f - 1) = \mathbf{A}^{-1}\tilde{\mathbf{x}}(t_f) - \mathbf{A}^{-1}\mathbf{B}\mathbf{q}(t_f - 1) \tag{4.81}$$

where the subscript b denotes a backward-in-time estimate. The reader may object that running a model backward in time will often be an unstable operation; this objection needs to be addressed, but ignore it for the moment. The uncertainty of $\tilde{\mathbf{x}}(t_f - 1)$ is,

{62048}
$$\mathbf{P}_b(t_f - 1) = \mathbf{A}^{-1}\mathbf{P}(t_f)\mathbf{A}^{-T} + \mathbf{A}^{-1}\mathbf{\Gamma}\mathbf{Q}(t_f - 1)\mathbf{\Gamma}^T\mathbf{A}^{-T}, \tag{4.82}$$

as in the forward model computation. This backward computation can be continued to time t_c, at which point we will have an estimate, $\tilde{\mathbf{x}}_b(t_c)$, with uncertainty $\mathbf{P}_b(t_c)$.

The two independent estimates of $\mathbf{x}(t_c)$ can be combined to make an improved estimate using the relations Chapter 2, Eq. (2.442),

{62049}
$$\tilde{\mathbf{x}}(t_c, +) = \tilde{\mathbf{x}}(t_c) + \mathbf{P}(t_c)\big(\mathbf{P}(t_c) + \mathbf{P}_b(t_c)\big)^{-1}\big(\tilde{\mathbf{x}}_b(t_c) - \tilde{\mathbf{x}}(t_c)\big) \tag{4.83}$$

and (Eq. 2.444),

{P2}
$$\begin{aligned}\mathbf{P}(t_c) &= \Big\langle [\tilde{\mathbf{x}}(t_c, +) - \mathbf{x}(t_c)]\ \tilde{\mathbf{x}}\,[(t_c, +) - \mathbf{x}(t_c)]^T \Big\rangle \\ &= \big[\mathbf{P}(t_c)^{-1} + \mathbf{P}_b(t_c)^{-1}\big]^{-1}\ .\end{aligned} \tag{4.84}$$

This estimate is the same as would be obtained from the RTS algorithm run back to time t_c—because the same objective function, model, and data have been employed. The computation has

been organized differently in the two cases. This backwards running computation can be used at all points of the interval, as long as the data used in the forward and backward computations are kept disjoint so that the two estimates are uncorrelated.

Running a model backward in time may indeed be unstable if it contains any dissipative terms. A forward model may be unstable too, if there are unstable elements, either real ones or numerical artifacts. But the expressions in (4.83-4.84) are stable, because the computation of $\mathbf{P}_b(t)$ and its use in the updating expression (4.83) automatically downweights unstable elements whose errors will be very large, and which will not carry useful information from the later state concerning the earlier one. The same situation would occur if the forward model had unstable elements—these instabilities would amplify slight errors in the statement of their initial conditions, rendering them difficult to estimate from observations at later times. Examination of the covariance propagation equation and the filter gain matrix shows that these elements are suppressed in the Kalman filter estimate, with correspondingly large uncertainties. The filter/smoother formalism properly accounts for unstable, and hence difficult-to-calculate parameters, by estimating their uncertainty as very large, thus handling very general ill-conditioning. In practice, one needs to be careful, for numerical reasons, of the pitfalls in computing and using matrices that may have norms growing exponentially in time. But the conceptual problem is solved. As with the Kalman filter, it is possible to rewrite the RTS smoother expressions (4.71)–(4.74) in various ways for computational efficiency, storage reduction, and improved accuracy.[122]

The dominant computational load in the smoother is again the calculation of the updated covariance matrices, whose size is square of the state-vector dimension, at every time step, leading to efforts to construct simplified algorithms that retain most of the virtues of the filter/smoother combination but with reduced load. For example, it may have already occurred to the reader that in some of the examples displayed, the state vector uncertainties, $\mathbf{P}$, in both the filter and the smoother appear to rapidly approach a steady state. This asymptotic behavior in turn means that the gain matrices, $\mathbf{K}$, $\mathbf{L}$, $\mathbf{M}$ will also achieve a steady state, implying that one no longer needs to undertake the updating steps–fixed gains can be used. Such steady-state operators are known as "Wiener filters" and "smoothers" and they represent a potentially very large computational savings. One needs to understand the circumstances under which such steady states can be expected to appear, and we will examine the problem on P. 246. {pagesteadysta

4.3.4 Other Smoothers

The RTS algorithm is an example of what is usually called a "fixed-interval" smoother because it assumed that the results are required for a particular stored interval $0 \leq t \leq t_f$. Other forms for other purposes are described in the literature, including "fixed-lag" smoothers in which one

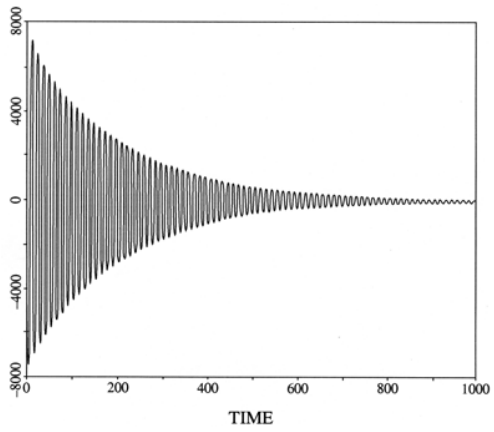

Figure 4.12: Mass-spring oscillator model with friction ($r = 0.01$) run backwards in time from conditions specified at $t = 1000$. The system is unstable, and small uncertainties in the starting conditions would amplify. but the Kalman filter run backwards remains stable because its error estimate also grows—systematically downweighting the model forecast relative to any data that become available at earlier times. A model with unstable elements in the forward direction would behave analogously when integrated in time with growing estimated model forecast error. {ocip6_8.tif}

is interested in an estimate at a fixed time $t_f - t_1$ as t_f advances, usually in realtime. A "fixed point" smoother addresses the problem of finding a best estimate $\tilde{\mathbf{x}}(t_1)$ with t_1 fixed and t_f continually increasing. When $t_1 = 0$, as data accumulates, the problem of estimating the initial conditions is a special case of the fixed point smoother problem.