

8 Filling In and Forecasting DSM2 Tidal Boundary Stage

8.1 Introduction

An adequate characterization of stage at Martinez is critical in order for DSM2 to simulate delta dynamics accurately. This section discusses a method of modeling Martinez stage, applicable to short-term prediction and the filling of historical records. The model combines a traditional astronomical tide with a model of the residual tide (the part of the observed tide that is not explained by an astronomical model).

8.2 Data and Preliminary Observations

A comparison of the bay tides with astronomical tides provides an interesting introduction to the boundary stage estimation problem. Figure 8-1 shows the tide at four stations: San Francisco Presidio, San Pablo Bay (RSAC045), Martinez (RSAC054) and Mallard (RSAC075). A harmonic mean tide has been fit to each of the stations:

$$z(t) = \sum_{i=1}^N f_i a_i \cos(\omega_i t + \phi_i + u_i) \quad (0.1)$$

where $z(t)$ is the harmonic mean tide, ω_i are known frequencies associated with astronomical motion, ϕ_i is a local phase, and f_i and u_i are slowly varying amplitude and phase adjustments attributable to the 19-year-cycle of the lunar node, as tabulated by the National Ocean Service (see Schureman, 1941).

The choice and the number N of constituents came from the standard NOS menu of 37 common constituents. In accordance with traditional NOS practice, constituents that had estimated amplitudes smaller than 0.03 feet were dropped – this may err toward over-fit, which is acceptable because the predictors are orthogonal and any “extra” coefficient estimates do not tend to hurt the “significant” ones. The final selection was slightly different for each station – between 21 (RSAC075) and 28 constituents (RSAC054) were retained per station. Many common constituents were used in all of the fits (the most important were M_2 , K_1 , S_2 , N_2 , O_1 , Mf , $(L2)$, Ssa , $2MK_3$, and $2SM_2$). The differences were in minor species and shallow water constituents, all of which had small amplitudes a few hundredths above the “cutoff” of 0.03 ft. Fewer species could be fit to RSAC075 than to RSAC054, despite RSAC075 being higher in the estuary where more shallow

water species have had a chance to develop. This suggests that distortion and noise play a limiting role on harmonic fits higher up in the estuary.

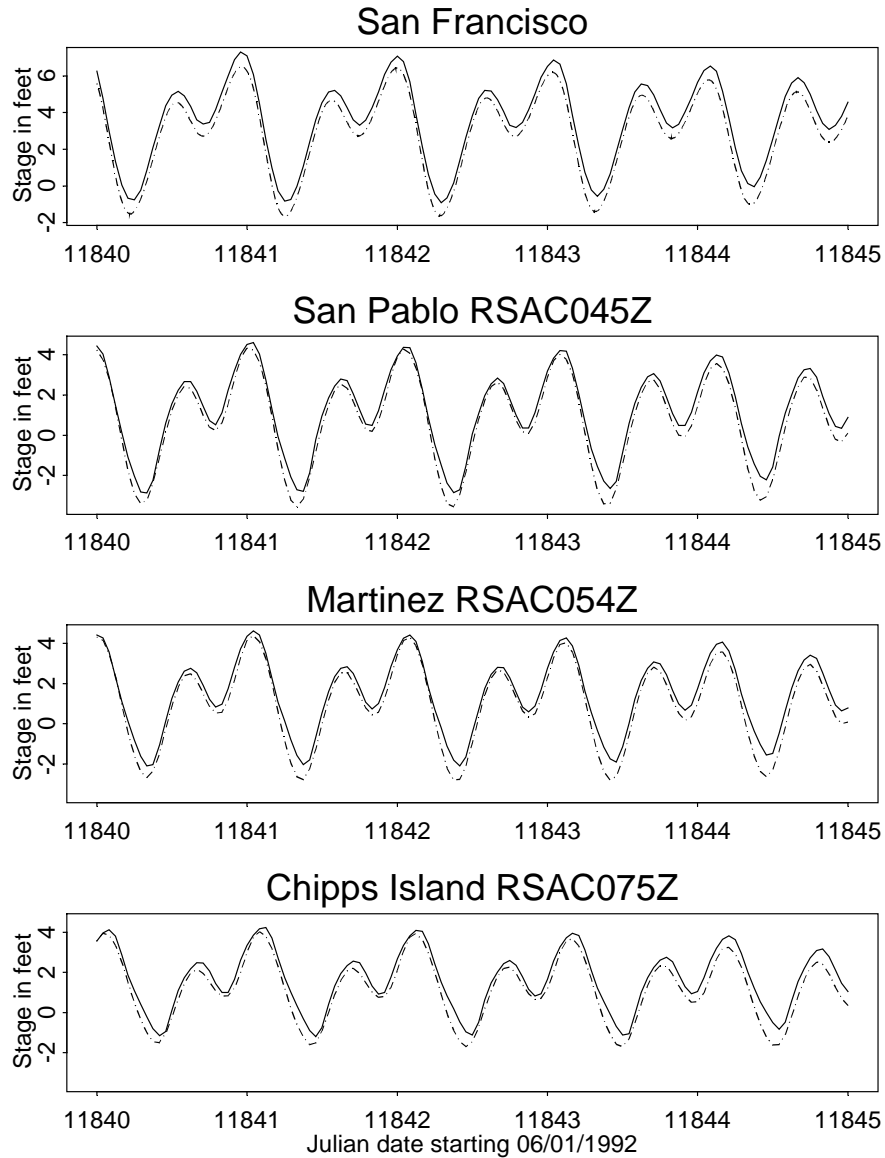


Figure 8-1: Tides (solid) and astronomical fits (dashed) at four stations.

Figure 8-2 (a) compares observed tide to astronomical over a single tidal cycle. Figure 8-2 (b) shows the astronomical residue over a fourteen-day cycle beginning on the same day, and Figure 8-3 compares the residuals at four stations. Two features are striking. First, there are local “trends” with period from several days to several weeks, often with magnitudes of a foot or more. On the right plot of Figure 8-2 long-wave variation is about one foot up over ten days (which is not an extreme example). The main physical contributors to these trends are barometric changes and 14-29 day oscillations induced by the interaction between lunar and solar constituents.

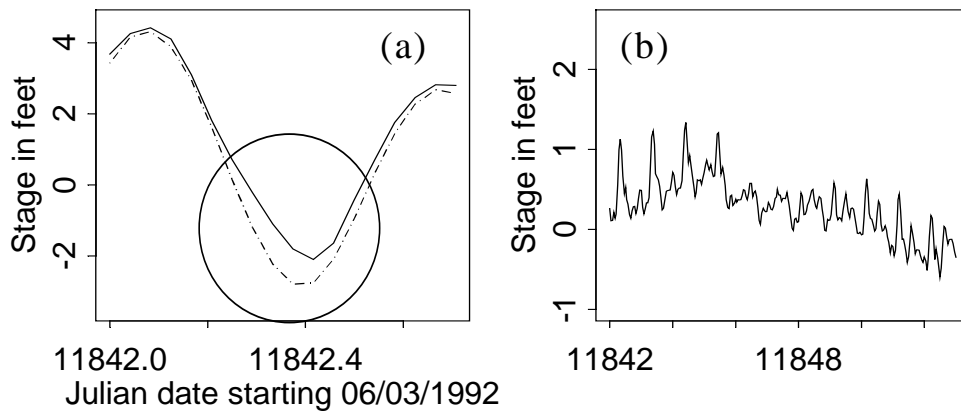


Figure 8-2: (a) Tidal distortion (solid observed vs. dashed astronomic) at Martinez over one tide cycle; (b) tidal residue during a 14-day spring-neap cycle.

The second source of error is that the tide is distorted (from sinusoidal) as it travels upstream, because of bottom friction, circulation in bays and effects induced by the shape of the basin. The episodes of greatest distortion cause instantaneous errors of about one foot at Martinez, and the complexity of the error develops as the tide moves upstream. Distortion is often manifest as patterns or “slivers” between the harmonic and observed tides (as in the circle inset in Figure 8-2), which last for several days. Common patterns are periodic, but recalcitrant to deterministic modeling – the patterns sometimes disappear and reappear for months at a time under circumstances which are sometimes not easy to tie to an external cause (season, delta outflow, pressure etc.).

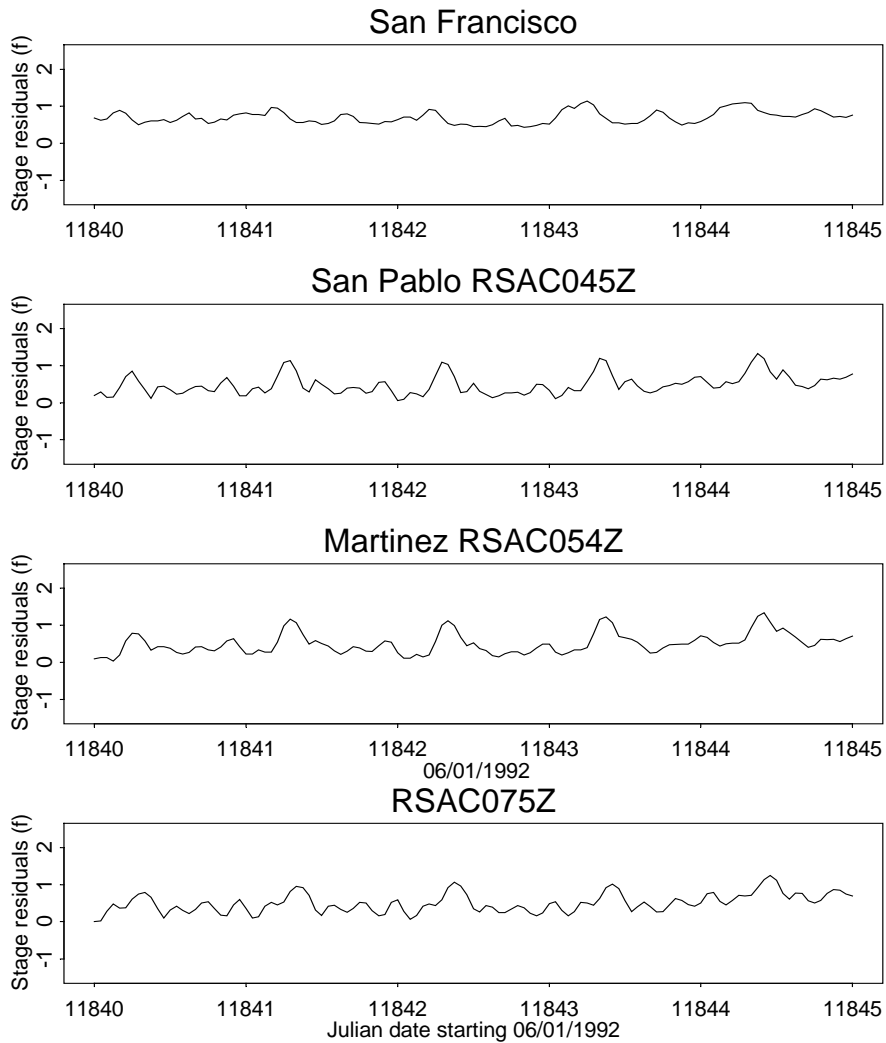


Figure 8-3: Residual tide at four stations (stage minus astronomical estimate).

8.3 VAR Tide Residue Model

8.3.1 Background in Linear Tide Models

The simplest method of linking tide observations at several stations is the linear model or convolution filter. Munk and Cartwright (1966) used linear (and quadratic Volterra-type) filtrations of an astronomical forcing function to produce tide predictions, and the same approach could be adapted to model one station on another in the estuary. The time domain representation of an appropriate filter model can be lengthy. For their application, Munk and Cartwright used a filter with lags spaced at $\Delta t = 2$ days in a two-sided filter spanning from -6 to +6 days, (thirteen terms). Their design was based on “equilibrium” tides filtered into narrow tidal bands. If their reasoning and spacing

criteria were applied to the full tidal spectrum (say 0 to 3 cycles/day), 4-hour spacing would be required over thirteen days.

Coefficient error and the short-term nature of shallow water distortion make a filter of such size unattractive for the present application. In a model of Martinez tide on San Francisco or San Pablo tide, 1-2 hour spacing over a maximum of 27 hours yielded much better fit (based on a Bayes Information Criterion) than did longer or more widely spaced filters. The only really important terms seem to be a few lags in the immediate past and another couple from the previous tidal day. Unidirectional filtration (modeling upstream stations on downstream) worked better than spatially two-sided filters, despite some bi-directionality in the underlying physical processes.

Removing the astronomic tide from the original tidal series improves the performance of the filtering approach and allows the use of a shorter filter. For instance, the best linear filter model of Martinez on San Francisco has a standard error of 0.27 feet if the full series is used; if instead residuals are used, the error is only 0.20 feet. It has been argued that with a true linear system, the decision as to whether to pass the mean tide through the filter or remove it and add it back later is arbitrary. This argument, neglects the effects of estimation error in determining the filter. When the mean tide is removed beforehand (eliminating line spectra), a much less complicated filter is required for the residuals and thus fewer filter coefficients have to be estimated.

For the remainder of this report, the astronomical tide will be taken to be an exogenous predictor – in other words it has already been computed before the rest of the analysis is carried out. An alternate approach would be to estimate the astronomical and residual components simultaneously. It is the author's experience from tinkering with the model that the residual model is not particularly sensitive to the astronomical model, and this at least is partial, informal confirmation of exogeneity. More detailed is deferred to later work.

8.3.2 Vector Autoregressive Model

Having made the above preliminary arguments, it is possible to specify a model that performs well. The model is a vector autoregressive (VAR) model based on astronomical residuals.

Because distortion appears to develop in an upstream direction, each station is modeled on its own past plus that of its downstream neighbor (except San Francisco, which is purely autoregressive), i.e.:

$$\begin{aligned}
z_1(t) &= c_{1,1,1} z_1(t-1) + c_{1,1,2} z_1(t-2) + \dots + c_{1,1,27} z_1(t-27) + \varepsilon_1(t) \\
z_2(t) &= c_{2,1,1} z_1(t-1) + \dots + c_{2,1,27} z_1(t-27) + c_{2,2,1} z_2(t-1) + \dots + c_{2,2,27} z_2(t-27) + \varepsilon_2(t) \\
\dots & \quad \dots \\
z_4(t) &= c_{4,3,1} z_3(t-1) + \dots + c_{4,3,27} z_3(t-27) + c_{4,4,1} z_4(t-1) + \dots + c_{4,4,27} z_4(t-27) + \varepsilon_4(t)
\end{aligned} \tag{0.2}$$

where the $z_i(t), i = 1 \dots 4$, respectively, represent water levels at the four stations San Francisco, San Pablo Bay (near the Carquinez Bridge), Martinez and Chipps Island/Pittsburg and the $c_{i,j,k}$ are coefficients in equation i of the j th station at lag k . The error terms ε_i are assumed to be Normal, to have a covariance matrix \mathbf{Q} , which includes correlation between stations but is independent and invariant over time (adequacy of the Normal model will be assessed shortly). The formulation of Equation (0.2) is general, but in practice only subsets of lags 1, 2, 3, 24, 25, 26, and 27 hours were found to be useful (using the Bayes Information Criterion, or BIC).

Coefficient estimates were obtained using Gaussian maximum likelihood, conditional on the first 27 hours of data. Given a complete tidal record, the likelihood calculations are equivalent to regressing each station on its own past lags and those of its downstream neighbor; i.e. fitting Equation (0.2) by least squares. The matrix \mathbf{Q} is estimated by:

$$\hat{\mathbf{Q}} = \frac{1}{T} \sum_{t=1}^T \hat{\varepsilon}(t) \hat{\varepsilon}(t)^T \tag{0.3}$$

where the $\hat{\varepsilon}(t)$ represent one-step prediction errors.

Coefficient estimates are given in Table 8-1. The errors estimated from the constituent regression calculations were three to four orders of magnitude smaller than the coefficients. These regression estimates are not valid error estimates for the problem, but suggest that the coefficient estimates are “tight” and that significance is not an issue for the lags employed. Formal estimates of coefficient error were not made because such estimates are not a natural product of the EM algorithm (see below). The models have changed slightly in real-time applications because RSAC045 in the San Pablo Bay is no longer monitored.

lag	S.Francisco	S. Pablo	Martinez		Chipps Is.		
	on self	SF terms	San Pablo	San Pablo	Martinez	Martinez	Chipps Is.
1	0.713	0.303	0.722	0.233	0.854	0.442	0.890
2	0.203		-0.047		-0.144	-0.118	-0.322
3							0.108
24	0.302		0.255		0.290		0.291
25	0.186		0.398		0.366	-0.185	0.346
26	-0.206	-0.208	-0.394	-0.225	-0.439	-0.124	-0.395
27	-0.219		-0.034		0.060		0.072

Table 8-1: VAR model coefficients.

8.3.3 Missing Data

Missing data complicates the likelihood calculation (the fairly simple least squared problem described above). A significant fraction of the historical records at the tidal stations are missing – about 15-25 percent of the observations at San Pablo Bay, Martinez and Pittsburg. Since each “squared error” expression in the likelihood involves several stations and lags, any one of which could be missing, the fraction of expressions affected by missing data is very high indeed. Fortunately, many of the data are “missing at random,” which means that the reason the data are missing is administrative or mechanical and is not directly related to the quantity being measured.

This section outlines how the missing data were treated using the EM algorithm. The EM (Expectation-Maximization) algorithm is an iterative numerical method for Maximum Likelihood Estimation (MLE), in which the original data are supplemented by “added” data or variables to form a “complete” data set. The added data are chosen in such a way that calculation of the likelihood using the complete data is easy. The added “data” are often in contrived forms.

The EM algorithm iterates between two steps:

- E:** calculate the expected value of the “complete” (log) likelihood expression using parameter estimates from the last iteration (or a starting value). Often, the likelihood expression is linear in the “added” data, in which case the E-step is equivalent to:
 - i. using the model to calculate the expected values of the “added” data necessary to “complete” the data set, and
 - ii. inserting the expected value of the “added” data into the expression for the likelihood.

M: maximize the “complete” likelihood expression from the E-step with respect to the parameters. In the VAR model this is just least-squares for the coefficients plus the estimate of \mathbf{Q} , as described above.

The naïve implementation of this idea would be to use the VAR to estimate missing stage residue observations, use the estimates to “complete” the historical record, and then re-estimate the VAR model. This turns out not to be an EM algorithm or even a good idea. Why? The likelihood is *not* linear in the raw stage values, but rather in sums of squares and cross-products – terms that look like $z_i z_j$ (i and j don’t have to be distinct). In standard regression terminology, these are the terms usually denoted as $\mathbf{X}^T \mathbf{y}$ (or $\mathbf{X}^T \mathbf{Y}$) and $\mathbf{X}^T \mathbf{X}$. See Mardia (1979) for the applicable equations or Hamilton (1994) for the application of the likelihood to VAR models. Because the likelihood is not linear in the tide residue, substitution of even a good, unbiased estimate of the tide residue into the likelihood does not give an unbiased estimate of the likelihood.

Instead, we should choose our “added” data to be the missing cross-products. Here, “missing cross-products” refers to products that are needed for the likelihood but cannot be calculated because they involve at least one piece of missing data. The problem, then, is to calculate the expected value of cross-products.

In the case where both contributors to the product are missing, recall this basic identity for covariance:

$$E(z_i z_j) = E(z_i)E(z_j) + \text{Cov}(z_i z_j) \quad (0.4)$$

where i and j represent times or stations and are not necessarily distinct, in which the covariance is a variance. The left side of Equation (0.4) is the “expected value of the product,” which is what we want. The first term on the right side is the “product of the expected values.” The second term on the right is the covariance between the two observations. This is the term that would be neglected under the “naïve” scheme. The covariance term tends to be particularly important for cross-products involving neighboring stations and time steps. This is usually the case. Cross-products arising from the least-squares fit of Equation (0.2) involve highly correlated lags and stations, and the missing data often occur in groups involving contiguous blocks of time or stations administered by the same agency.

In the simpler case, when cross-products involving one missing piece of data, the equivalent expression is:

$$E(z_i z_j) = E(z_i) z_j \quad (0.5)$$

Since z_j is known, it just acts as a scalar. In this case covariance does not come into play and the expression is linear in $E(z_i)$.

The expectations of cross-products can be calculated using Equation (0.4) or (0.5) using stage estimates and covariance estimates. To get both, the VAR model is put in state space form, and a Kalman filter and Kalman (Rauch-Tung-Striebel) smoother are used to produce the stage observations, cross-covariance and squared error terms conditional on all the data. Since the Kalman filter/smoothing is used implementing the VAR model for historical filling, the EM algorithm can be employed without production of tools that would be of no further use once the parameter estimation is finished. The Kalman filter is not discussed further here – it is a computational device discussed in detail in any book on linear systems or state space modeling. See, for instance, Harvey (1994) or Hamilton (1994). The details of using the Kalman filter to obtain cross-covariance terms for use in EM are discussed in Watson and Engle (1983).

The use of the EM algorithm makes the VAR fitting complicated, and this process was never automated. Should the VAR model be refit later due to a change in circumstances (e.g., a change in the list of available stations), it would probably be adequate to use a least squares fit to Equation (0.2) plus Equation (0.3) using whatever data are available. It is certainly better to just throw away time steps whose expressions involve missing data at any lag than to implement an *ad hoc* iterative scheme in which the covariance between observations is ignored. As long as the data are missing at random, throwing out missing data still leads to an asymptotically efficient estimator.

8.3.4 Diagnostics

Rudimentary diagnostics were carried out to assess the “whiteness” and Gaussianity of the innovations process of the VAR model. The empirical autocovariance at Martinez of one-step prediction errors is shown at the top of Figure 8-4, along with a 95 percent confidence interval around zero (assuming a stationary Gaussian process). Residual periodicity is evident, but is much smaller than that of alternative filter lengths and spacing considered.

The bottom of Figure 8-4 is a distributional (q-q normal) plot based on 4,600 one-step prediction errors; there is little evidence of severe skewness or thick tails which would indicate that a normal distribution (and least squares) is inappropriate. Nevertheless, Chi-squared (p-value 0.0038) and two-sided Kolmogorov-Smirnoff tests (p-value of 5.0E-4) reject the hypothesis that the residuals are Gaussian, even after trimming a few outliers. Such results are typical because of the high power of goodness-of-fit tests based on large data sets – with so much data it is easy to prove that the data are not perfectly normal.

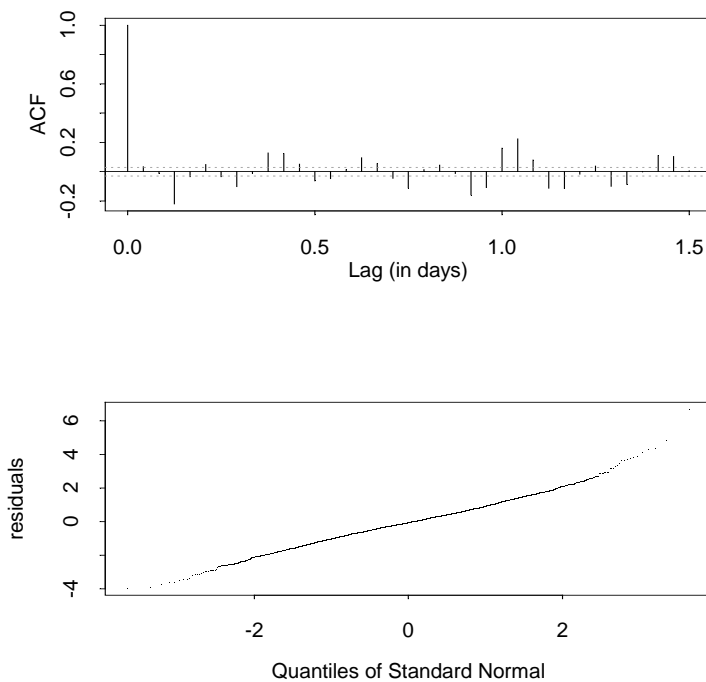


Figure 8-4: Diagnostics of the innovations process at Martinez. (Top) Empirical autocovariance. (Bottom) Quantile comparison with normal distribution.

8.3.5 Tide Series Reconstruction and Interpolation

After the residue has been predicted or filled, the final step is to add back the astronomical tide. The VAR component of the model (the part that works on the residue) operates on a 1-hour time step. To recover a 15-minute series for use in DSM2, the steps are as follows:

1. Interpolate the 1-hour estimate of the residue to 15-minutes, and
2. Add the result back to a 15-minute astronomical forcing.

Notice that only the “small” residual part is interpolated between the hours. The rest of the “shape” comes from the astronomical prediction, which can be resolved at any time step desired. Interpolating the hourly tide in this order is accurate, causing about ~0.01 foot error between the hour. Two alternatives would be to reconstruct the series first at a 1-hour time step and then interpolate the whole series, or to just reconstruct it at 1-hour and let DSM2 do the interpolating later. When linear interpolation is used, these alternatives are equivalent and result in errors at peaks of more than 0.1 foot, which is enough to spoil the accuracy of historical filling. Higher order interpolation is possible off-line, but not in DSM2.

In the production version of these tools, a fourth-order monotonicity-preserving spline due to Huynh (1993) was used for the interpolation, resulting in a more accurate and more visually pleasing series.

8.4 Applications

8.4.1 Prediction

The primary role of the VAR model in real-time modeling is to predict future stage. Predictions are produced by recursive application of the VAR Equations (0.2) to the tide residue with the error term set to zero. In practice, the synthesis of the VAR model is carried out with a Kalman filter, which is capable of producing variance estimates as well as a mean forecast. The Kalman filter is also required as the part of the historical filling algorithm. See Harvey (1994) or Hamilton (1994) for a discussion of how a vector autoregressive model is expressed in state space form for implementation in a Kalman filter.

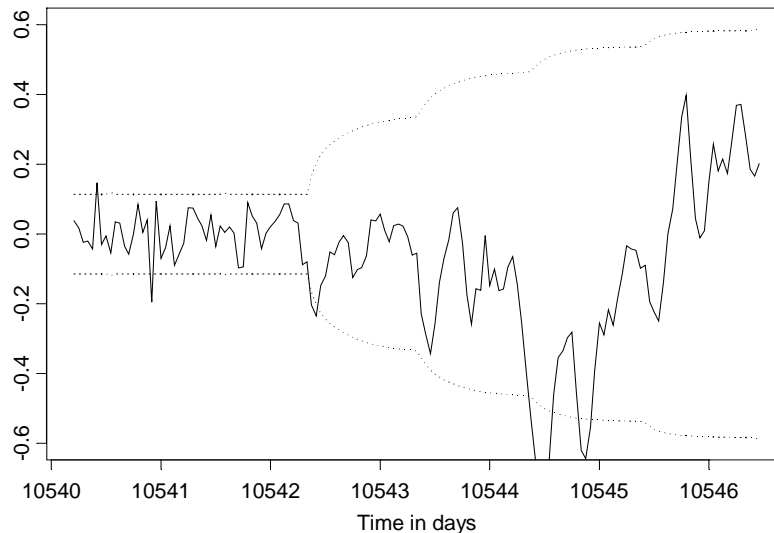


Figure 8-5: Errors and 95 percent confidence bound for simulated prediction using the VAR residue model (note sign: expected minus observed).

Figure 8-5 shows the results of simulated short-term forecasting at Martinez. All station records were artificially terminated at Julian date 10541 (November 10 1988-89). Until this point, the prediction errors and 95 percent confidence bounds represent one-step prediction error; after the cutoff date, the prediction error accumulates. Every tidal day, a discontinuous increase occurs in the confidence bounds, as real observations are replaced by estimates in the lagged components of the model. The confidence region is point-

wise: the probability of error continuously following the envelope is much smaller. The errors are reversed in sign from the usual for residuals (they are model minus observed).

The prediction error shown in Figure 8-5 takes on its worst values in the middle of Julian day 10,544. This can be explained in terms of long-period “trend” on that day. Figure 8-6 shows tidally averaged stage for the period. An unanticipated event occurred several days into the prediction period, causing water surface elevations to rise. This event is not captured in the VAR model, which can “remember” events that are in progress at the time of prediction, but cannot anticipate new barometric events. The attenuation portion of the event is captured well by the VAR model (as is a period of further attenuation not pictured). Further consideration of barometric events is given shortly; this example – a new event beginning soon within the prediction period – will be cited as a “tough case”.




Figure 8-6: Tidally filtered stage during the simulated prediction period.

8.4.2 Filling missing records

The ability to fill historical records is a strong point of the VAR model (it was originally designed for this purpose), but is a task of only secondary importance in real time modeling. What distinguishes the filling (smoothing) problem is that supporting information is available from other stations. Since low-frequency fluctuations are felt fairly uniformly over the whole delta, the algorithm does well even during extreme events.

Although the same VAR formulation is used for prediction and filling missing records, the filling problem involves more complicated “smoothing” calculations. This is because we want to condition our estimate on all data available, which will include data before and after the historical gap. A single recursion marching forward in time is not sufficient, because later observations will not be used to improve earlier estimates. Instead, we need to make use of a bi-directional algorithm, and the one used here is a forward Kalman

filter coupled with a fixed-interval backward Kalman smoother (see Harvey, 1994). The result is an estimate for each missing value based on all available data.

Figure 8-7 shows errors resulting from two examples of simulated data filling, in which artificial gaps were created in the Martinez record. In series (a), which comes from the same period as the prediction problem of last section, Martinez goes off-line for two days. During the middle of this period San Pablo also goes off-line for several hours, causing a bump in the confidence interval. Note that the event that caused an episode of high error in the prediction problem is not even detectable in the filling problem.

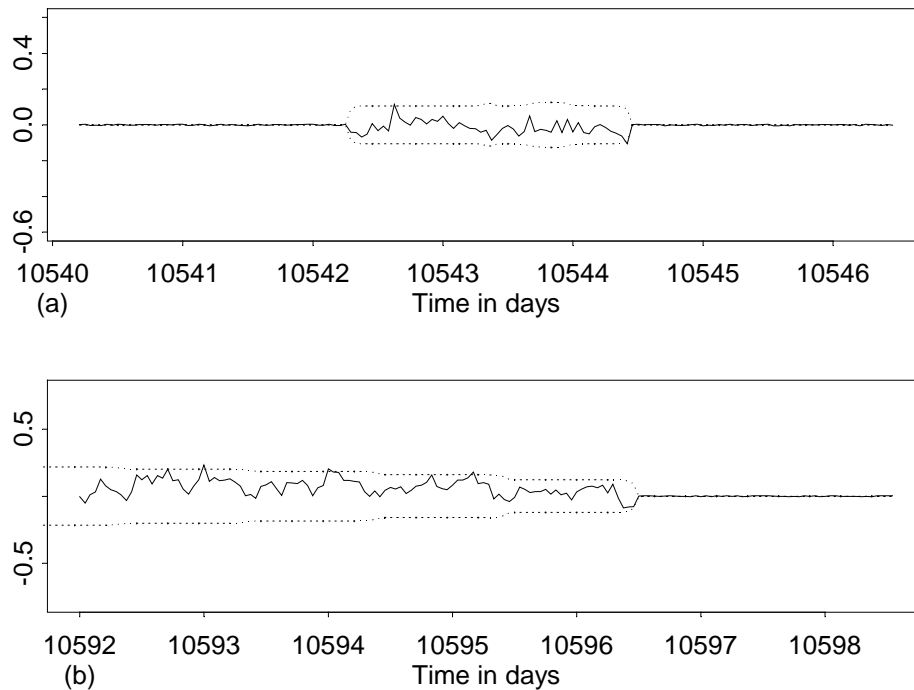


Figure 8-7: Errors for the smoothing of simulated gaps, plus 95 percent confidence bounds.

The main difficulty in filling data for real-time use seems to be the lack of supporting station data. Station RSAC045 was recently dismantled or moved. The lack of supporting station data will not matter much for small gaps, when the auto-regressive Martinez component does most of the work, but will make the model less accurate for long gaps. The San Francisco Presidio station is always available, and can be relied upon to pick up long wave events such as storms, but does not include information about shallow water distortion higher in the estuary.

Series (b) is a synthetic gap preceded by a month-long real gap in the Martinez record, and thus the model operates without autoregressive information from this station. San Pablo was also off-line at the beginning of the plot, coming on-line just before Julian day

10,595. Even during the worst long gaps, the smoothed Martinez record has a standard error of less than one tenth of a foot and a 95 percent confidence bound of about two-tenths of a foot. This level of accuracy has been verified with cross-validation over a large number of artificial gaps (every other week in 1996-1997), including episodes of moderate to high winds.

8.5 Discussion

8.5.1 Atmospheric Events

Correspondents who inquire about the stage boundary model are usually concerned with performance during atmospheric anomalies. The VAR model does not explicitly include barometric input, though the addition of such a term is a possibility for future development. Under most circumstances, the VAR residue model can perform just as well as a model that includes atmospheric terms. This is true in the following situations:

- Any historical filling scenario:
Stations in the Delta respond very similarly to low-frequency events (lower than 1 cycle per day). For this reason, stage data at supporting stations give us almost all the low-frequency information required for Martinez. To illustrate the similarity between stations, Figure 8-8 shows the Martinez (RSAC054) and San Pablo (RSAC045) tidally-averaged stage during December 1995, one of the biggest wind events in recent history. The two stations rise and fall neatly in unison. The distance between the two curves (shown as a separate line in the plot) varies only 0.04 feet in each direction from a flat line, at the height of the event. Describing the effect of such an event at Martinez to this precision using a causal model between atmosphere and tide would be out of the question.
- Prediction made near the peak of an atmospheric event (point (c) on Figure 8-8):
The VAR model captures and attenuates most events, as long as they are “in progress” at the time the prediction is made.
- Prediction made many days before the onset of an atmospheric event (point (a) on Figure 8-8):
In this case, the atmospheric effects cannot be predicted at all. Thus, there is no way that atmospheric information can be put into the model successfully even if there were a term to accommodate it.
- Seasonal atmospheric characteristics (e.g. summer winds):
The effects of the current season are manifest in the residual at the start of the prediction and are implicitly “in the model.” Of course, this assumes that the VAR model is used to predict over an interval that is short compared to the length of the season.




Figure 8-8: Tidally filtered stage during a major atmospherically-driven event.

Where improvement of the model is possible is when a prediction is made several days before an important, forecastable atmospheric event, as in point (b) on Figure 8-8 (recall that this was the case in the earlier example shown in Figure 8-5). The author is investigating the relationship between stage changes and barometric forecasts to see if the inclusion of a pressure term is warranted.

8.5.2 14-Day Cycling

A nonlinear interaction between solar and lunar constituents excites oscillations with a period of just over 14 days (this interaction is not the same as the “spring-neap” cycle, which is a linear interaction, but tends to be synchronized with this cycle). The oscillations are largely responsible for the “filling and draining” of the Delta which has been observed to be important in salinity transport. Oscillations of this periodicity are handled in equilibrium tide theory, standard NOS and DWR models, and in the model presented above inasmuch as they are captured in long-period constituent Mf.

Figure 8-9 shows low-passed stage over six months for RSAC054. Two-week cycling is apparent, but a sinusoidal representation seems unsatisfactory. In fact the Mf harmonic term fit to the full record (on stations where it was deemed significant) has an amplitude of about 0.1 foot, which means that it is not picking up variations of the magnitude that are obvious from inspection of Figure 8-9.

An alternate way of incorporating the 14-day term is to absorb it as a cyclical autoregressive part of the model. From Figure 8-9, we can see that characteristics tend to

repeat between adjacent two week cycles and this would be represented using standard VAR or ARMA terms – a common practice in seasonal hydrologic models (Bras, 1993). Figure 8-9 suggests that the improvement in medium-term prediction accuracy (forecasts of 14-20 days) might be as much as a half a foot. The downside to including 14-day periodicity is that the model will become much larger and more cumbersome.




Figure 8-9: Tidally averaged stage at Martinez over six months.

8.6 Conclusions

Adding a tide residue model to augment ordinary astronomical models increases the quality of prediction and smoothing. The VAR model suggested here produces fill-in (smoothed) values that are extremely accurate, regardless of events which take place during the fill-in period. The model also improves prediction of tides, with the largest improvements during the first week. Both types of estimates tend to have close to zero error when compared to the observed data, so that the transition between real and estimated data is smooth.

This section has also examined two potential shortcomings of the VAR approach. The first is that the VAR does not employ a causal representation of atmospheric forcing on tides. This affects accuracy if the prediction is made just before the onset of a foreseeable storm. In almost all other situations involving astronomical events, it has been shown why explicit consideration of such forcing cannot benefit the model.

The second shortcoming of the VAR approach is that it uses a naïve representation of 14-day period variations. These oscillations are not really sinusoidal and attempts to model them as such usually fail to capture their full amplitude. A possible remedy for

this shortcoming is to include a cyclical but slowly varying 14-day term in the VAR model. This will improve medium forecasts.

8.7 References

- Bras, R.L. and I. Rodriguez-Iturbe. (1993). *Random Functions and Hydrology*. Dover.
- Godin, Gabriel. (1972). *The Analysis of Tides*. University of Toronto Press.
- Hamilton, J. D. (1994). *Time Series Analysis*. University Press.
- Harvey, A.C. (1994). *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press.
- Huynh, H.T. (1993). "Accurate Monotone Cubic Interpolation." *SIAM J. Numer. Analysis*. **30**(1), pp 57-100.
- McLachlan, G.J. and Thriyambakam, K. (1997). *The EM Algorithm and Extensions*. Wiley.
- Mardia, K.V., J.T. Kent, and J.M. Bibby. (1979). *Multivariate analysis*. Academic Press.
- Munk, W.H. and D.E. Cartwright. (1966). "Tidal Spectroscopy and Prediction." *Phil. Trans. Roy. Soc. (London)*, Ser. A **259**,533-81.
- Schureman, P. (1941). *Manual of Harmonic Analysis and Prediction of Tides, Special Publication*. No. 98, U.S. Coast and Geodetic Survey.
- Zetler, B.D. (1982). *Computer Applications to Tides in the National Ocean Survey, Supplement to Manual of Harmonic Analysis and Prediction of Tides (Special Publication No. 98)*. National Ocean Survey (NOAA).