Day 23

Kalman Filter Examples

Static State Estimation

- recall the static state estimation problem we have been studying
  - the process or plant model

$$A_t = 1, \quad B_t = 0, \quad R_t = 0 \qquad \begin{aligned} x_t &= A_t x_{t-1} + B_t u_t + \varepsilon_t \\ &= x_{t-1} \end{aligned}$$

  - the observation model

$$C_t = 1, \quad Q_t = \sigma_t^2 \qquad z_t = x_t + \delta_t$$

Static State Estimation

- how well does the Kalman filter work

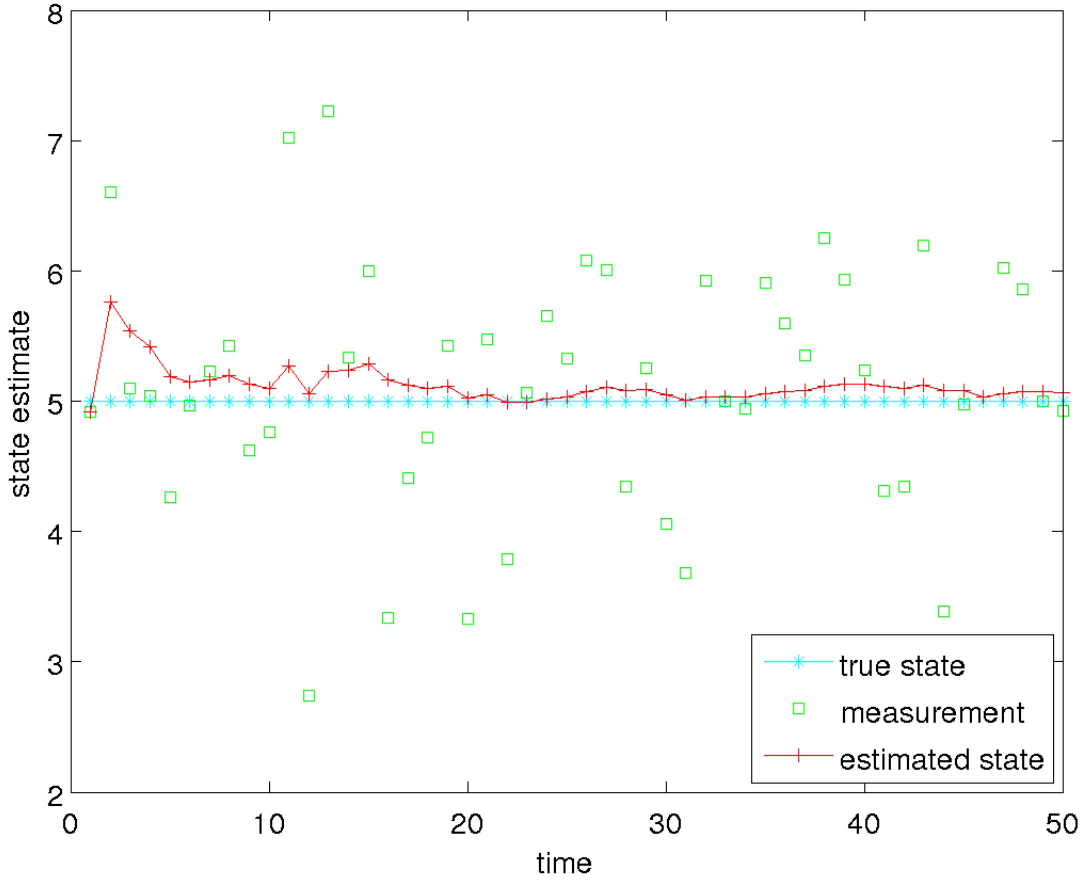

Static State Estimation

- notice that we need to specify the measurement noise covariance Q_t
- how sensitive is the Kalman filter to Q_t ?
  - e.g., what if we use a Q_t that is much smaller than the actual measurement noise?
  - e.g., what if we use a Q_t that is much larger than the actual measurement noise?

Static State Estimation

- specified Q_t = 0.01 * actual Q_t

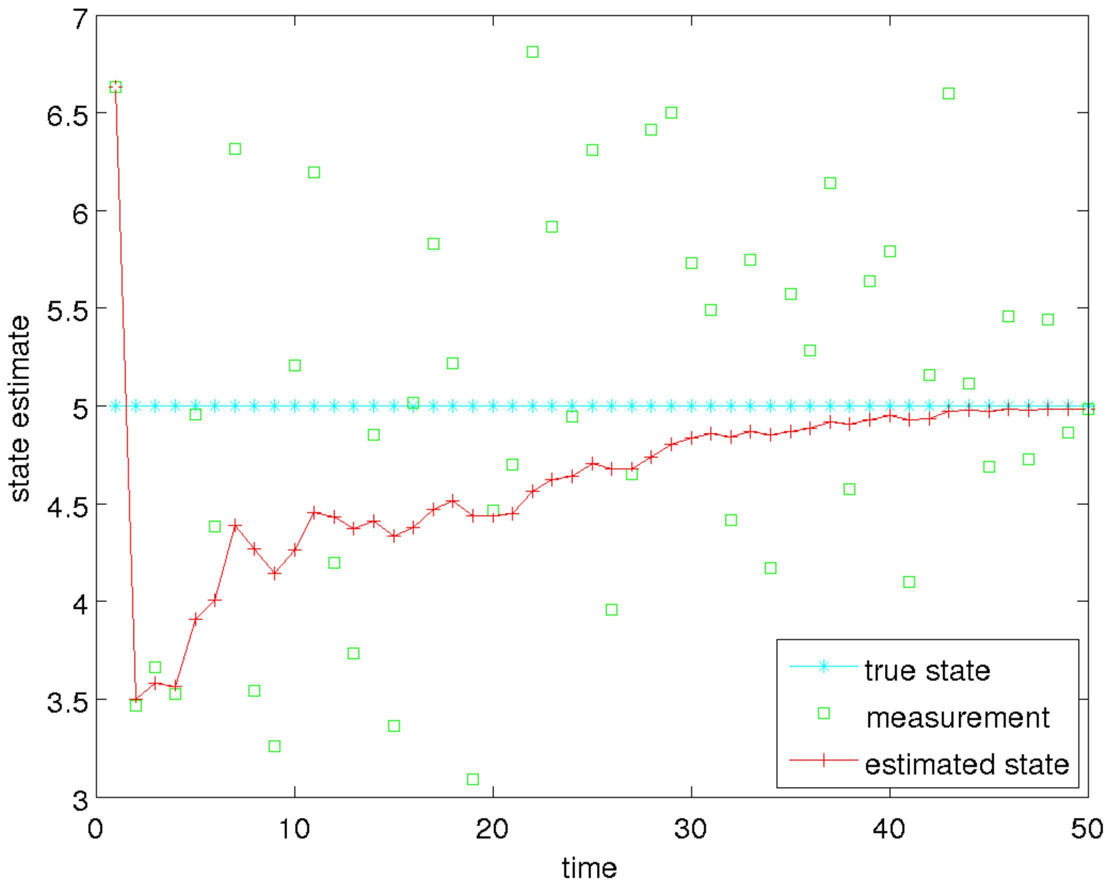

Static State Estimation

- specified Q_t = 100 * actual Q_t

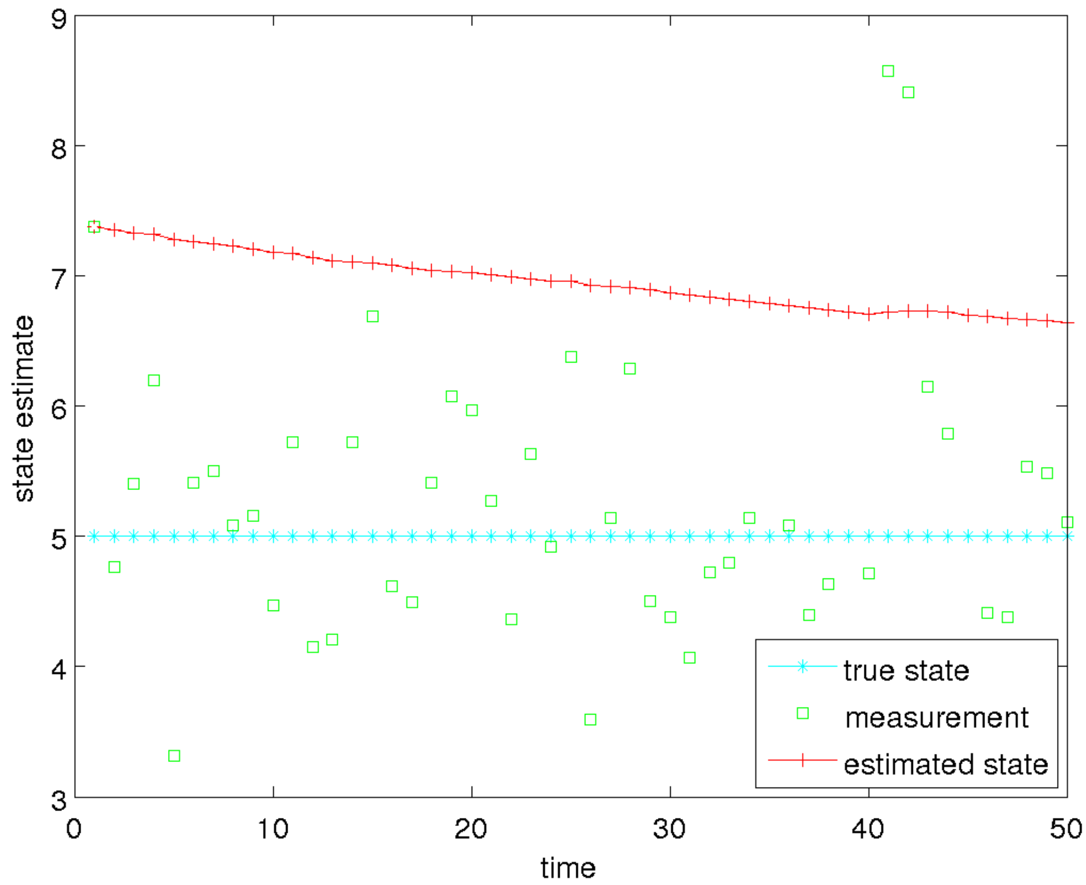

Static State Estimation

- suppose our measurements get progressively noisier over time

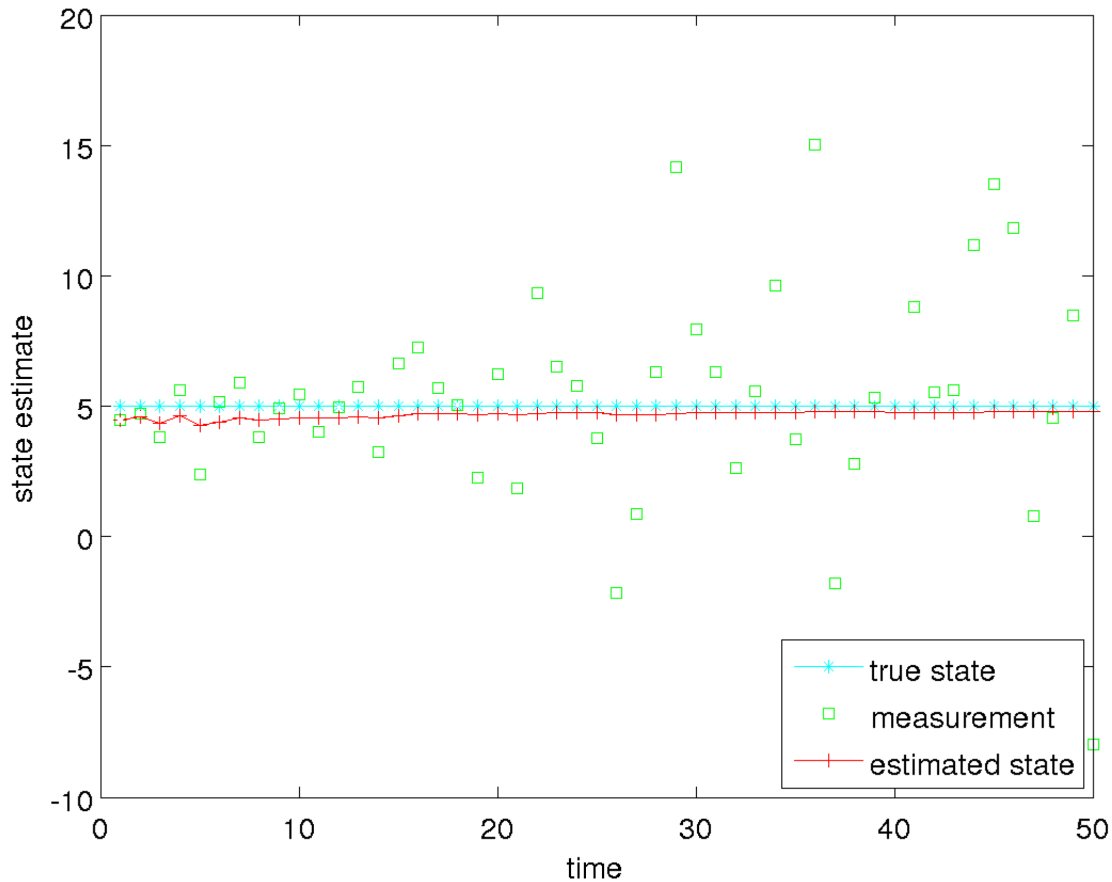

noise variance increases 10% for each successive measurement

Tank of Water

- estimate the level of water in the tank; the water could be
  - static, filling, or emptying
  - not sloshing or sloshing

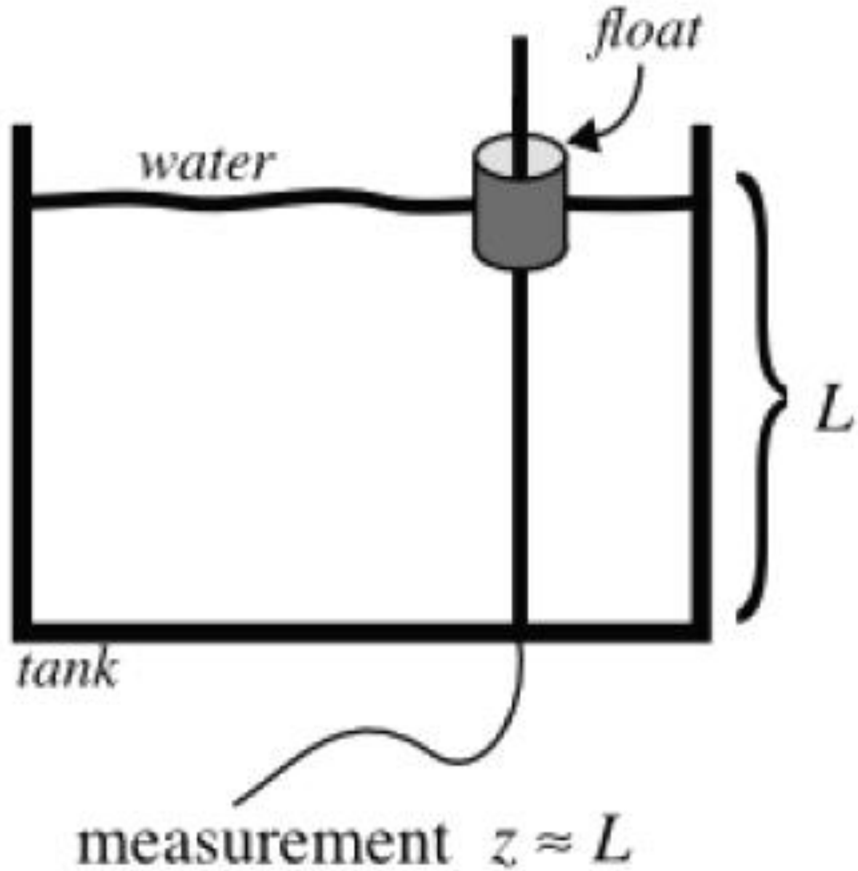

Tank of Water: Static and Not Sloshing

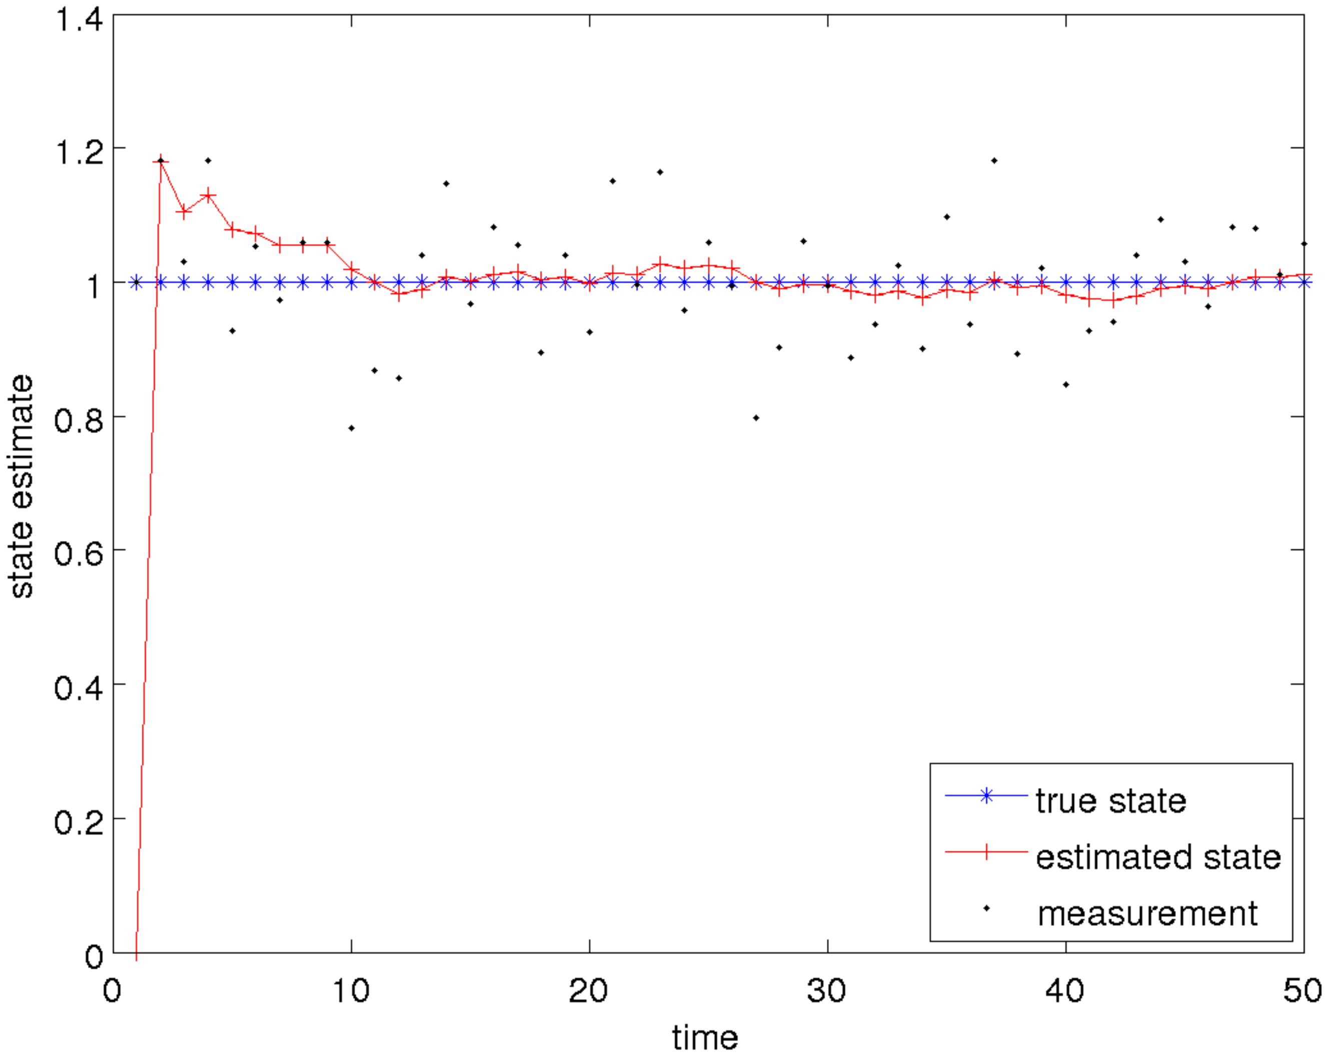

Tank of Water: Static and Not Sloshing

- notice that in this case the Kalman filter tends towards estimating a constant level because the plant noise covariance is small compared to the measurement noise covariance
  - the estimated state is much smoother than the measurements
- what happens if we increase the plant noise covariance?

Tank of Water: Filling and Not Sloshing

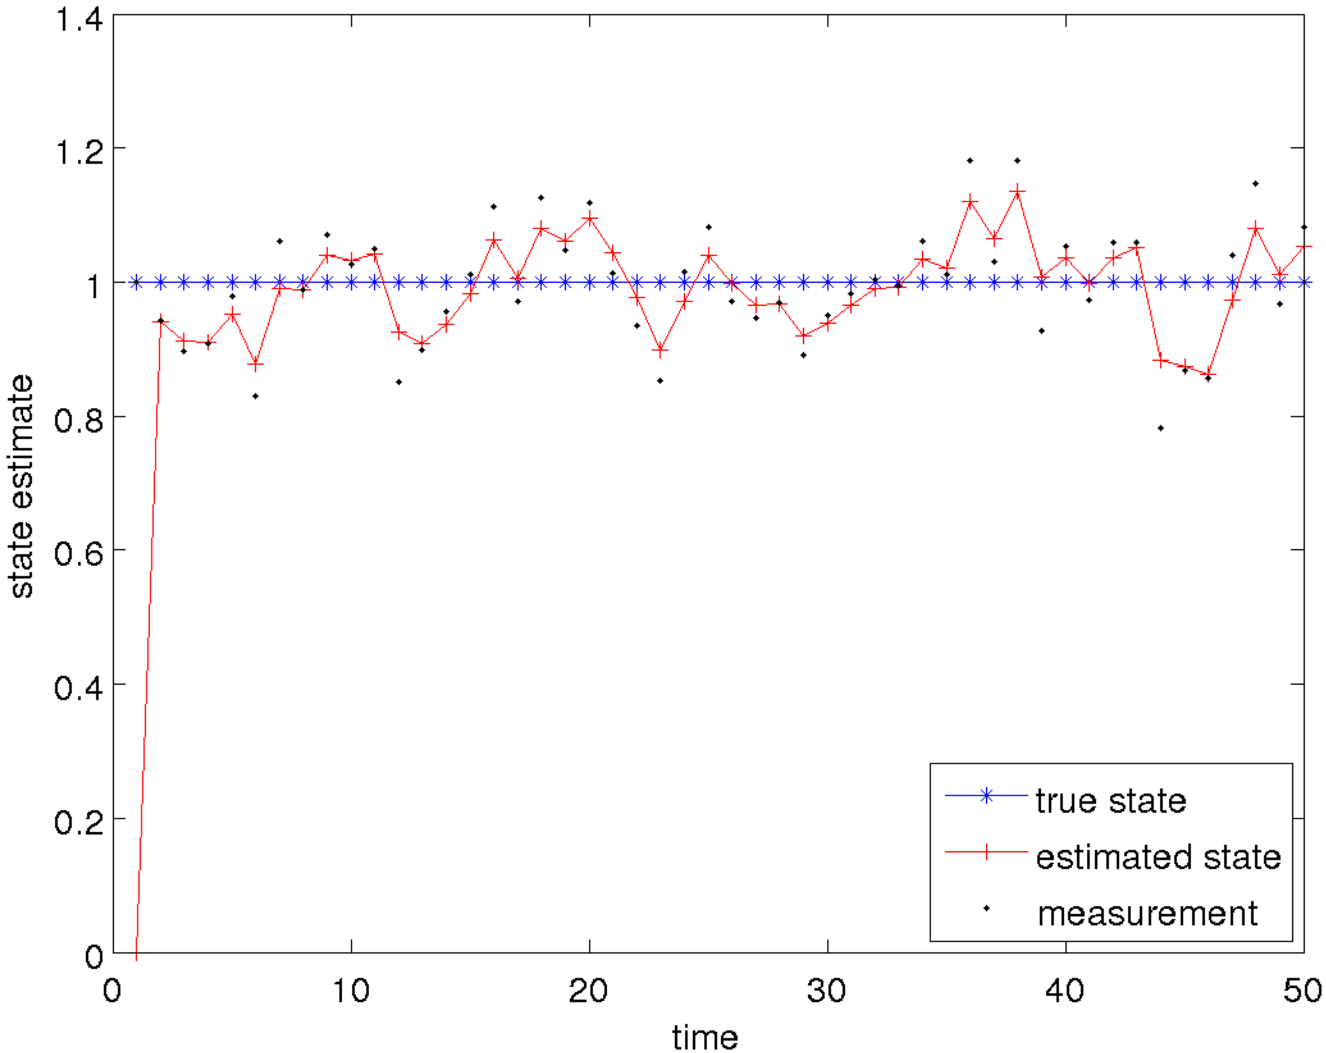

Tank of Water: Static and Not Sloshing

- notice that in this case the Kalman filter tends towards estimating values that are closer to the measurements
- increasing the plant noise covariance causes the filter to place more weight on the measurements

Tank of Water: Filling and not Sloshing

- suppose the true situation is that the tank is filling at a constant rate but we use the static tank plant model
  - i.e., we have a plant model that does not accurately model the state transition

Tank of Water: Filling and not Sloshing

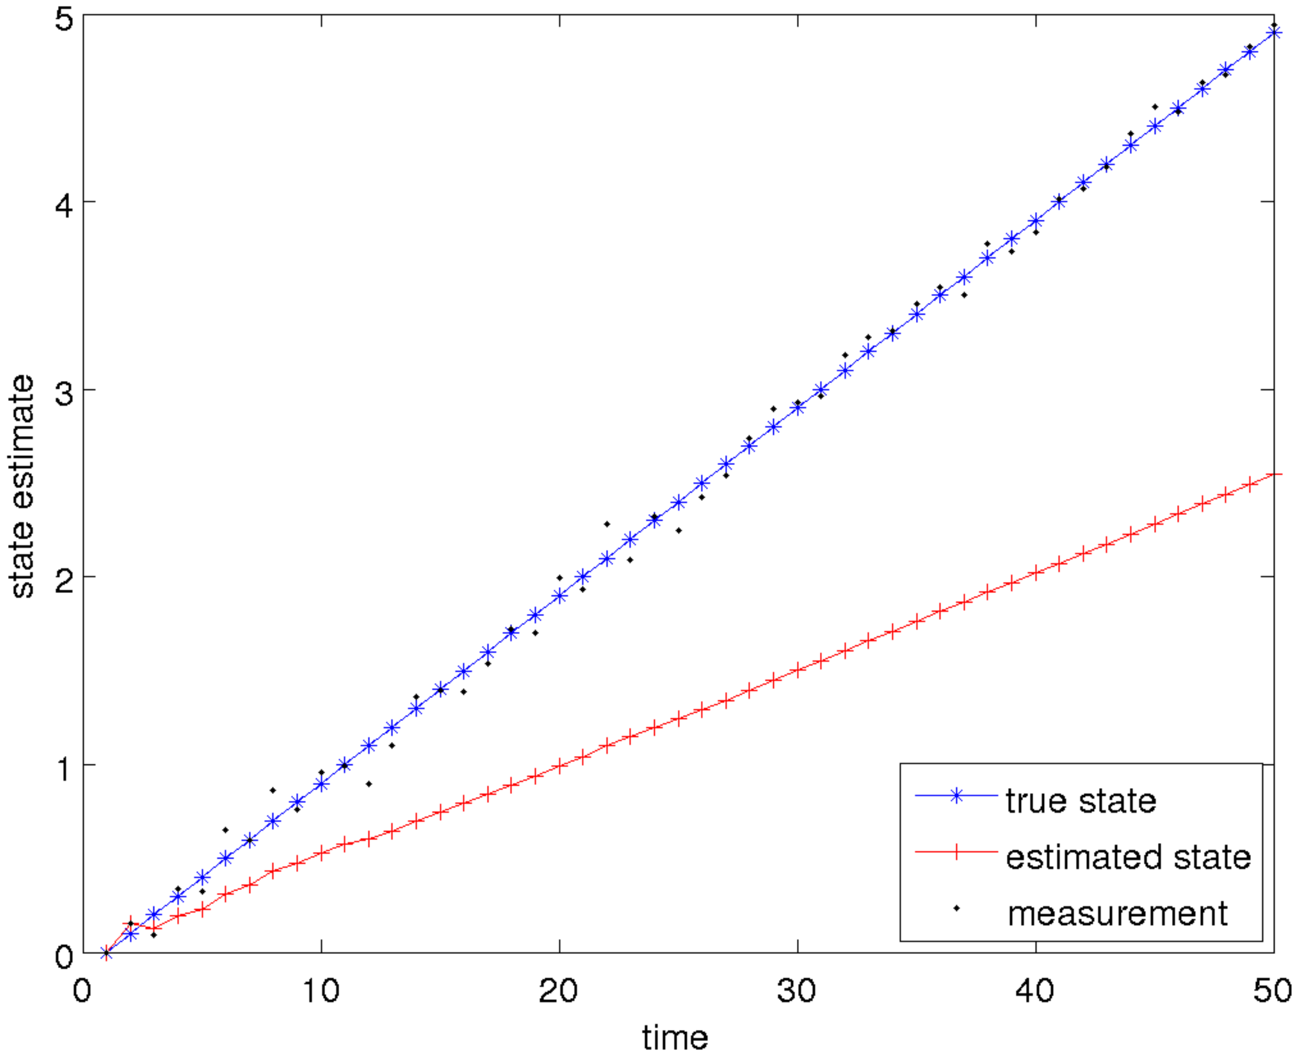

Tank of Water: Filling and not Sloshing

- notice that in this case the estimated state trails behind the true level
  - estimated state has a much greater error than the noisy measurements
- if the plant model does not accurately model reality than you can expect poor results
  - however, increasing the plant noise covariance will allow the filter to weight the measurements more heavily in the estimation…

Tank of Water: Filling and not Sloshing

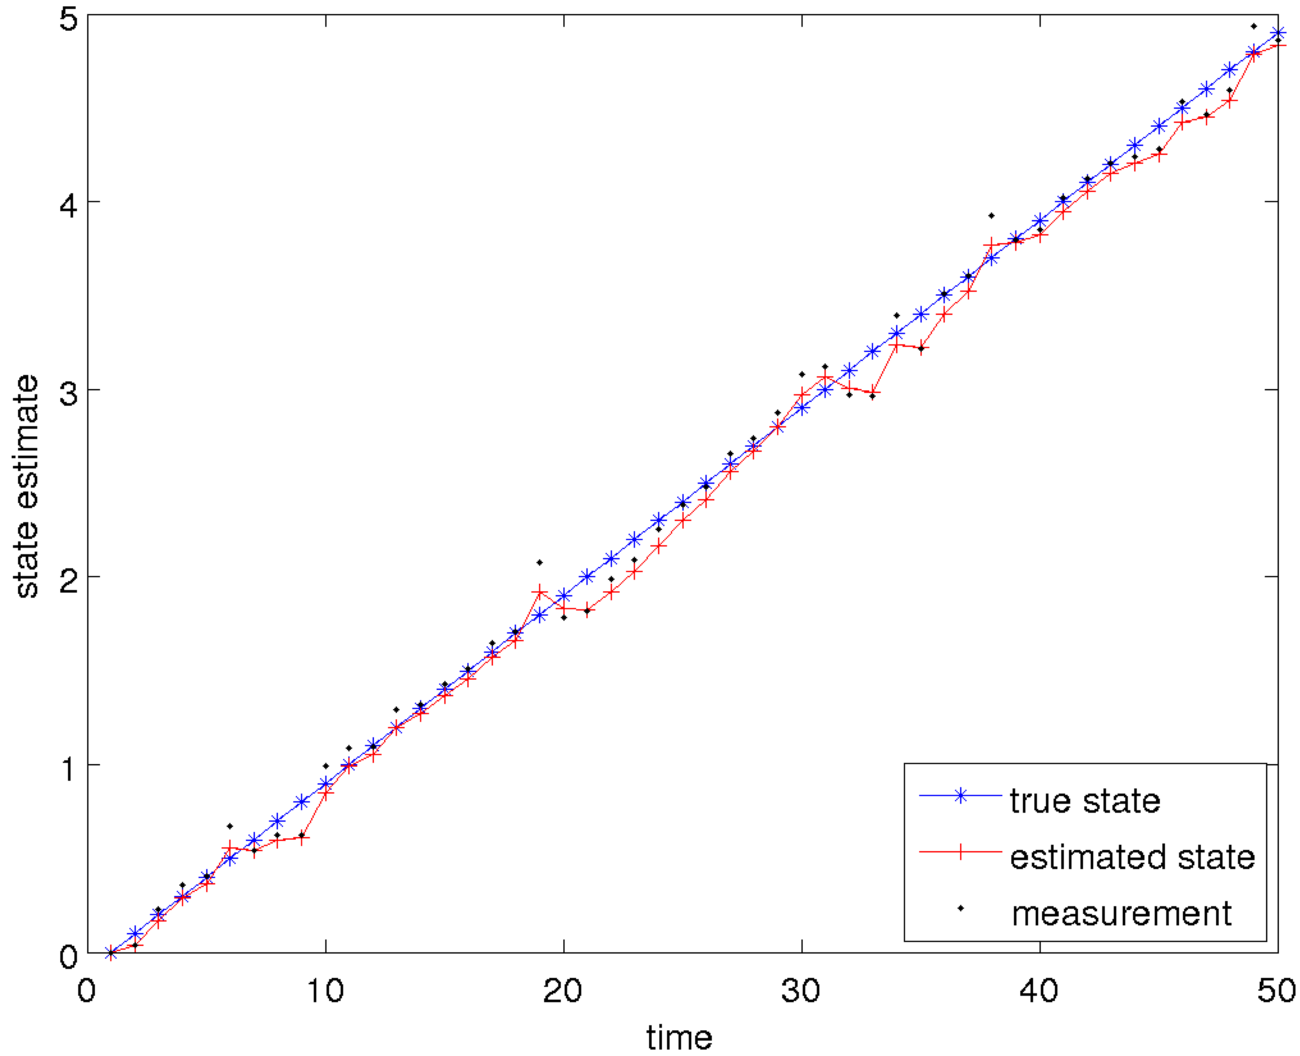

Tank of Water: Filling and not Sloshing

- it is not clear if we have gained anything in this case
  - the estimated state is reasonable but it is not clear if it is more accurate than the measurements
- what happens if we change the plant model to more accurately reflect the filling?

Tank of Water: Filling and not Sloshing

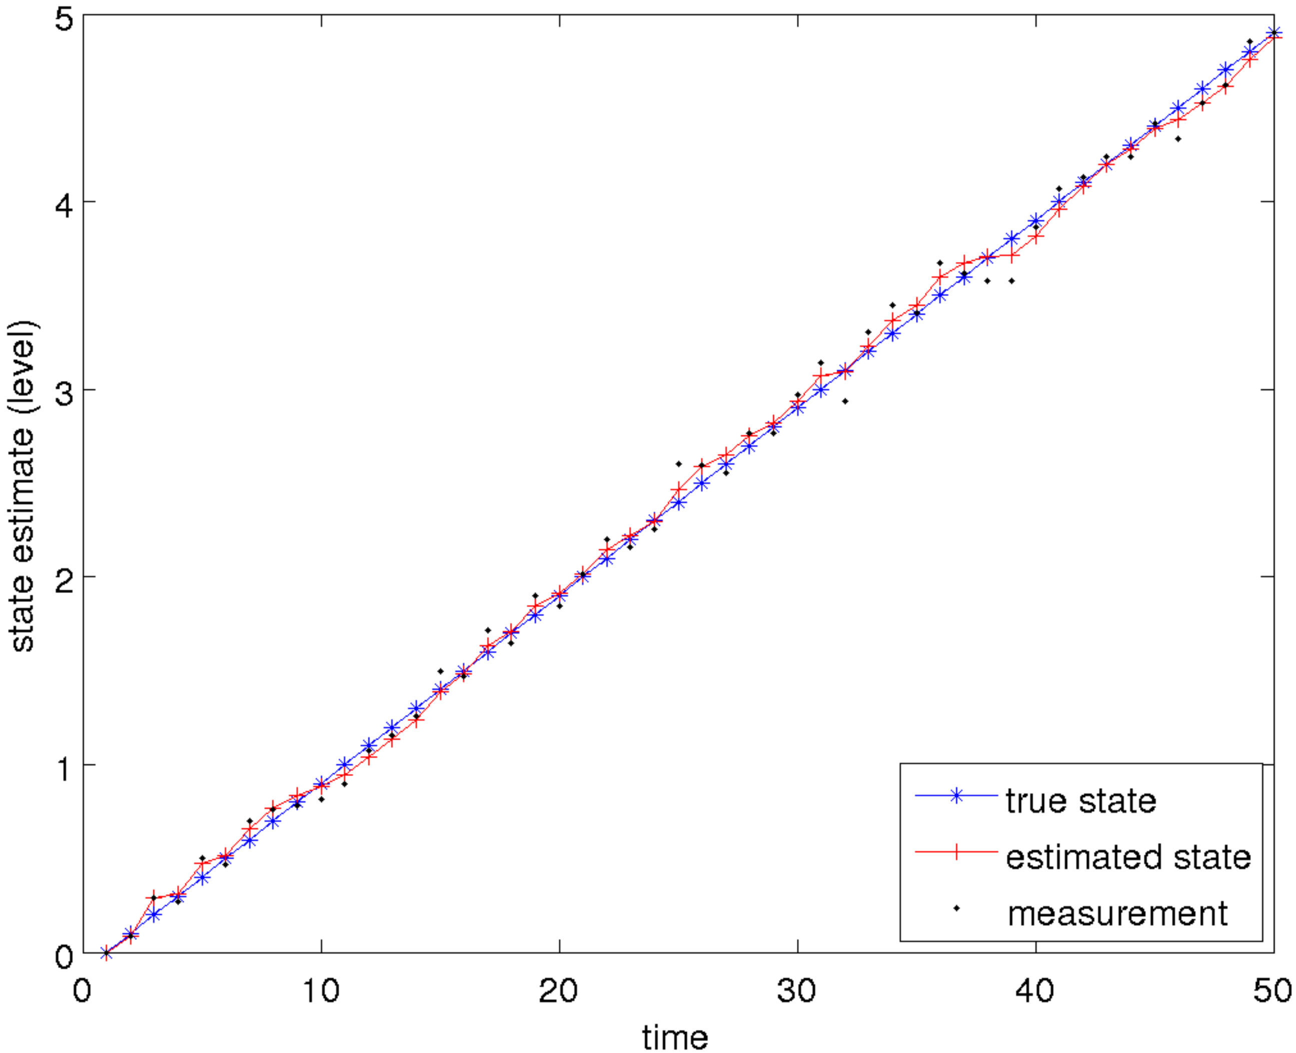

Tank of Water: Filling and not Sloshing

- notice that the estimated state is more accurate and smoother than the measurements
- what about the filling rate?

Tank of Water: Filling and not Sloshing

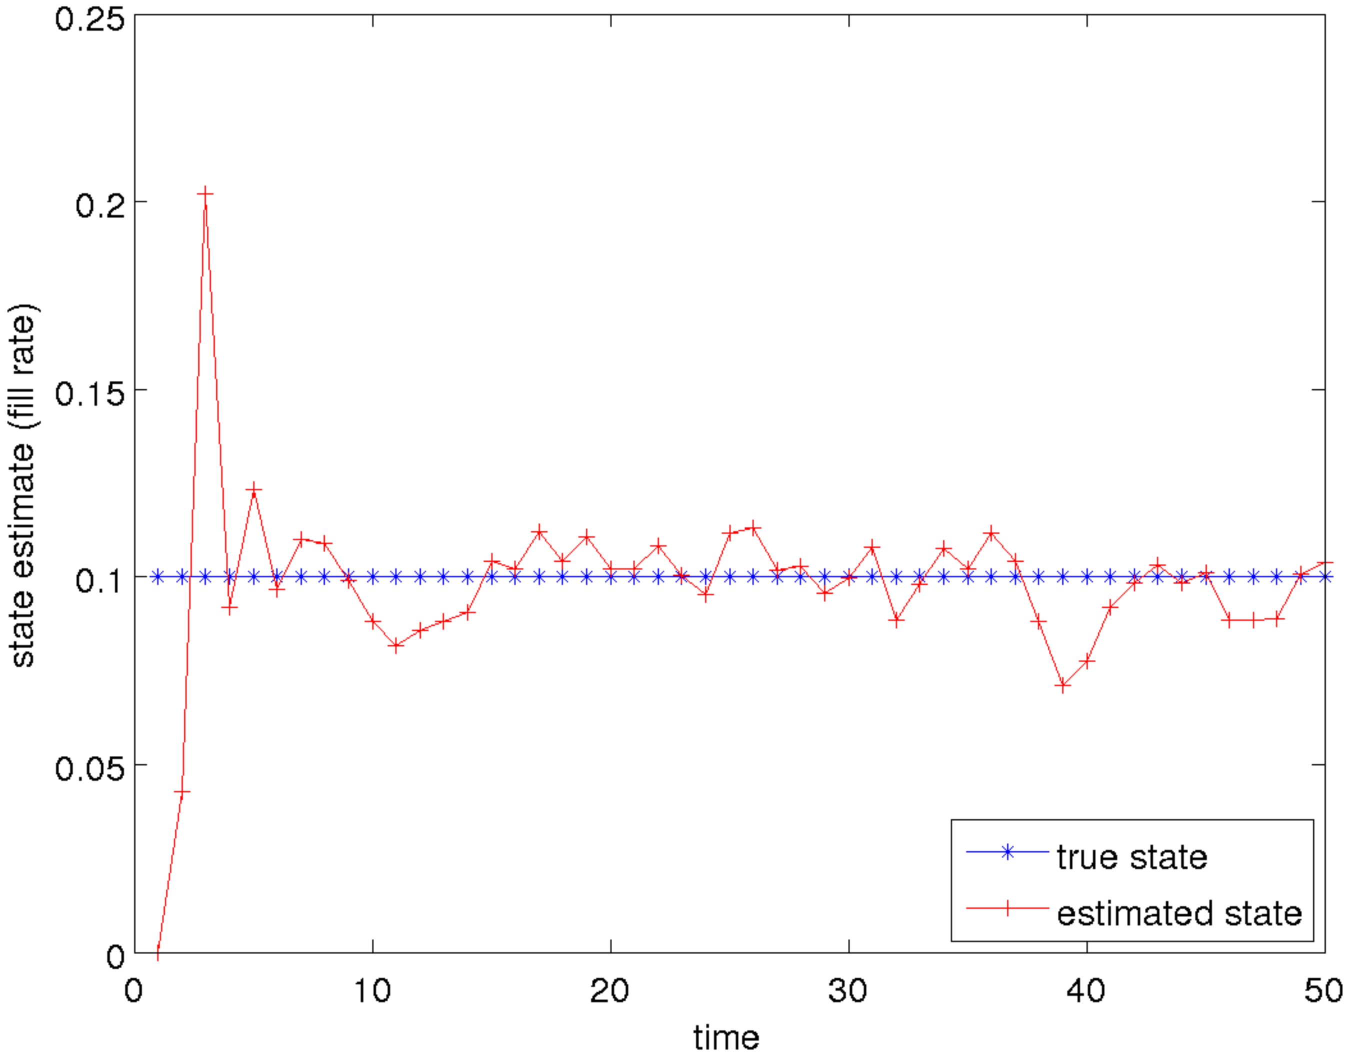

Tank of Water: Filling and not Sloshing

- notice that the estimated filling rate seems to jump more than the estimated level
  - this should not be surprising as we never actually measure the filling rate directly
    - it is being inferred from the measured level (which is quite noisy)

Tank of Water: Static and not Sloshing

- can we trick the filter by using the filling plant model when the level is static?
  - hopefully not, as the filter should converge to a fill rate of zero!

Tank of Water: Static and not Sloshing

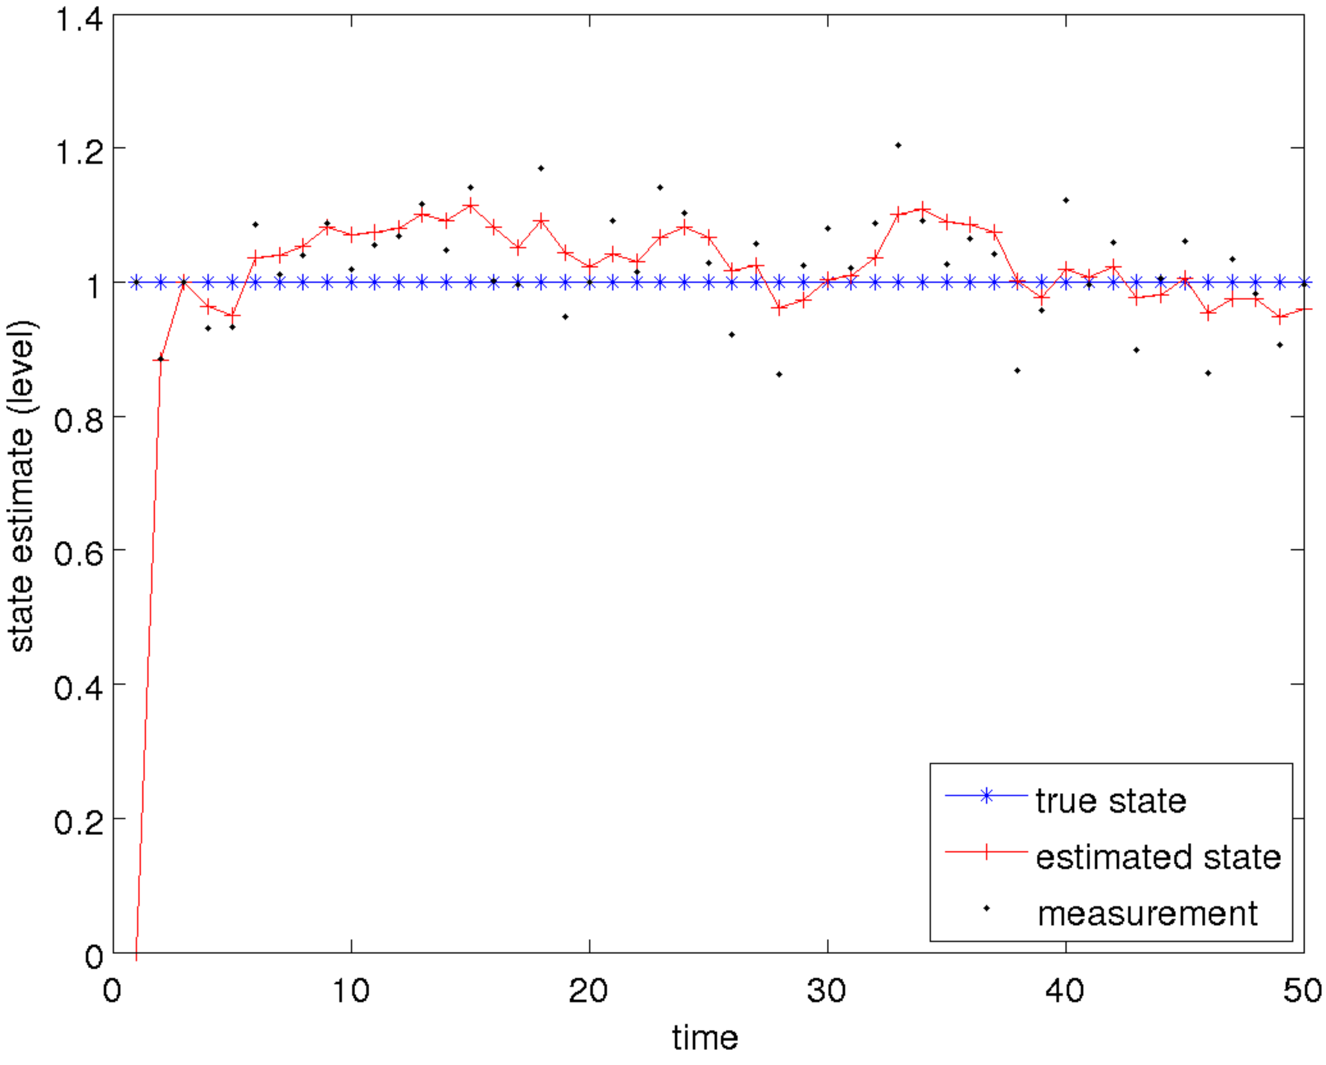

Tank of Water: Static and not Sloshing