Chapter 4
Aircraft Prior Based on Primary Radar Data

The Bayesian approach described in the previous chapter is a recursive method that calculates the posterior state distribution at each measurement time from a distribution at the previous measurement time. It fundamentally requires knowledge of three probability density functions: the prior distribution of the state at initialisation, $p(\mathbf{x}(0))$; the state evolution $p(\mathbf{x}_k|\mathbf{x}_{k-1})$; and the measurement likelihood $p(\mathbf{z}_k|\mathbf{x}_k)$. Chapter 5 addresses the measurement probability density and Chaps. 6 and 7 discuss the state transition model. This chapter discusses the prior state distribution and the method used to define it. Intuitively, one might expect this prior to have a significant influence on the probability distribution at later times: a larger uncertainty in the prior might be expected to lead to a greater spread of uncertainty in the final pdf compared to a prior pdf with smaller uncertainty.

In the MH370 search, there are two data sources that are available to construct the prior. The aircraft reports its own location and other information to the ground via a satellite link using the Aircraft Communications Addressing and Reporting System (ACARS). Data from ACARS is available for the accident flight only up to the point where communications were lost: the final ACARS report was at 17:07:29. In Chap. 9, other flights with known aircraft locations are used to validate the models used for the accident flight. For these, ACARS data is available and this data is used to construct the prior.

The second source of prior information is radar. For the validation flights this radar data is not available and nor is it required given the presence of ACARS logs. For the accident flight, primary radar data provided by Malaysia is available from after the loss of communications up until 18:22:12. The radar data contains regular estimates of latitude, longitude and altitude at 10 s intervals from 16:42:27 to 18:01:49. A single additional latitude and longitude position was reported at 18:22:12. Figure 4.1 shows the radar data overlaid on a map. Under radar coverage, the aircraft turned sharply at approximately 17:24, crossed over Malaysia, and then turned to the North-West at 17:53.

The Bayes filter requires a prior over the full state space, not merely latitude and longitude. The development to follow will lead to a state vector containing several other parameters. Where possible it is preferable to specify a prior on these parameters

S. Davey et al., *Bayesian Methods in the Search for MH370*, SpringerBriefs
in Electrical and Computer Engineering, DOI 10.1007/978-981-10-0379-0_4

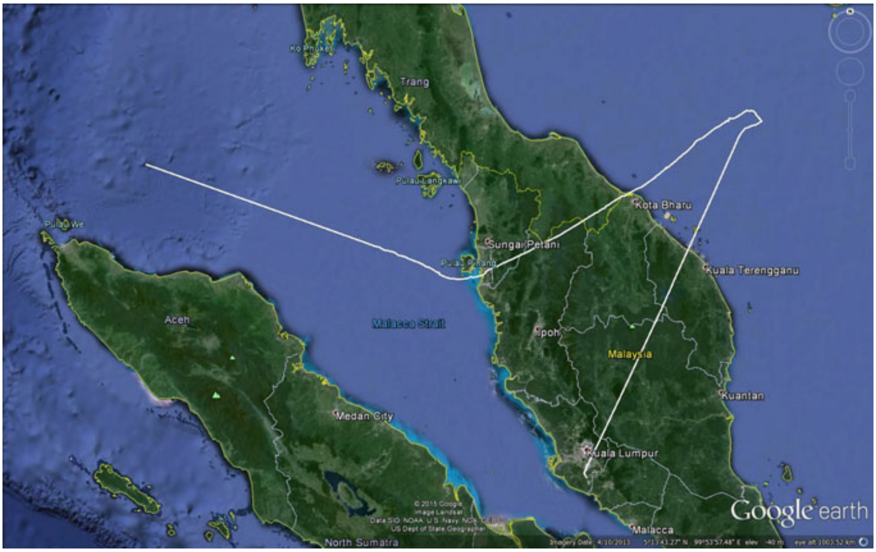

Fig. 4.1 Primary and secondary radar data available for MH370

using radar data or ACARS rather than subjective belief. Where this is not possible, the philosophy has been to use priors that are diffuse to avoid prejudicing the filter output. It is possible to derive the angle and speed of the ground velocity from the radar reports by assuming a simplified almost constant velocity model and applying a Kalman filter. This assumption is acceptable for the primary radar data because the reports are closely spaced in time. Figure 4.2 shows the derived speed and direction obtained from this filter.[1] The speed estimates vary dramatically during the first turn, which is not an accurate representation of the aircraft speed at this time. It is likely due to the mismatch between the assumed linear Kalman filter model and the high acceleration manoeuvre performed by the aircraft. Since these artefacts are localised to the time of the turn the influence on the state at the end of the sequence is negligible.

The final reported position from radar was at very long range from the sensor and there was a long time delay between it and the penultimate radar report. This report is at long range and it is likely to have rather poor accuracy because the angular errors translate to large location errors at long range. The radar report at 18:22 is closer to the penultimate report at 18:02 than the filter speed predicts. Also, it was observed that the range ring derived from the timing measurements at 18:25 and 18:28 are closer to the 18:02 report than predictions based on either the 18:02 filtered speed or the 18:22 filtered speed. Figure 4.3 shows the relative positions of the 18:25 arc and the filter predictions based on data up to 18:02. Collectively these data points suggest that the aircraft may have slowed down at some point between 18:02 and 18:22. In

[1]The measurement error was assumed to have a standard deviation of 0.5 nm and the process noise variance was $3.5 \times 10^{-4}\,\mathrm{nm}^2\,\mathrm{s}^{-3}$. The process noise was adjusted to minimise the mean squared prediction error.

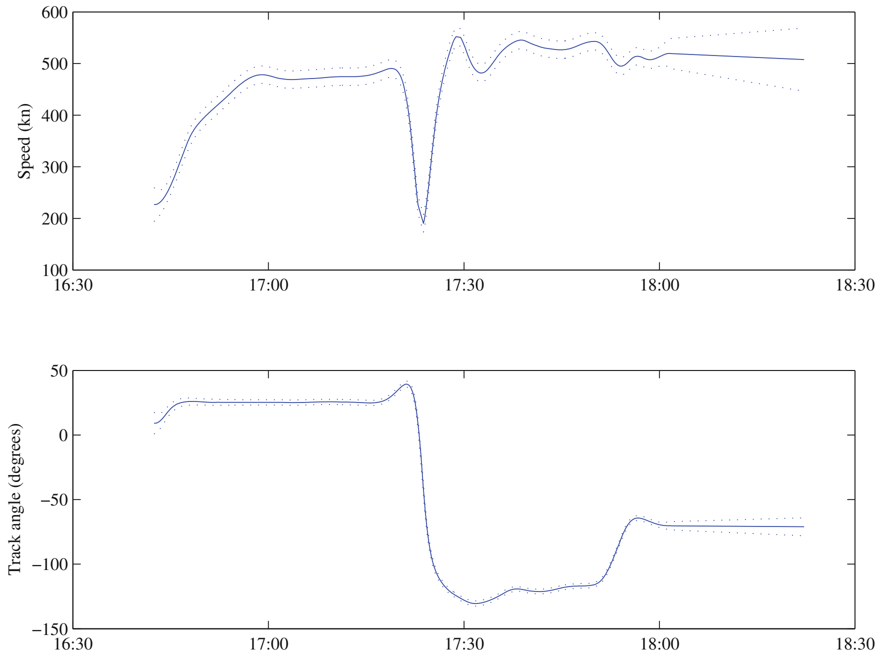

Fig. 4.2 Smoothed estimates of speed and direction derived from radar data. *Dotted lines* show covariance of estimates, illustrated as mean plus and minus three-sigma value

addition, the ground speed observed by the radar prior to 18:02 is relatively high and implies that the aircraft would have been at low altitude. This is likely to result in poor fuel efficiency, and in order to maintain flight for the duration indicated by the satellite data, the aircraft would have had to slow and increase altitude at some stage to conserve fuel. This is also consistent with a potential speed change between 18:02 and 18:22.

The 18:22 radar observation was not used quantitatively because the latitude and longitude derived from it are likely to be less accurate at long range and the aircraft may have manoeuvred prior to 18:22. The radar observation was deemed to indicate that the aircraft did not turn between 18:02 and 18:22, but the numerical values were not used. Instead, a prior was defined at 18:01 at the penultimate radar point using the output of the Kalman filter described above. The position standard deviations were set to 0.5 nm and the direction standard deviation to 1°. Figure 4.3 shows predictions of the mean of this prior from 18:02 to 18:25, shown in yellow, and one-sigma lines at $\pm 1^\circ$. The 18:22 radar point, at the end of the radar track, is clearly within the azimuth fan. As described above, the filtered speed at the output of the Kalman filter is not consistent with the 18:25 measurement and predictions based purely on this will have a likelihood very close to zero. In addition, the model discussed in Chap. 6 specifies air speed in terms of Mach number. The manoeuvre model described in Chap. 7 allows for speed changes and these will be randomly sampled by the proposal distribution.

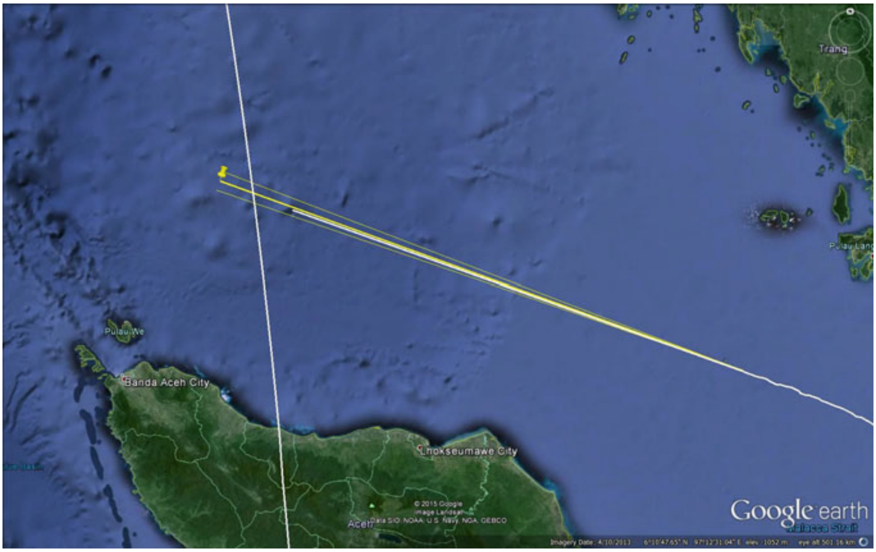

Fig. 4.3 18:02 prediction to 18:25, shown in *yellow*. The Malaysian military radar track is shown in *white*, on the *right*. The near-*vertical white line* on the *left* corresponds to the 18:25 BTO arc

Rather than trying to specify when the speed change occurred, the filter was expected to learn this information. This provides a richer description of the trajectory since the timing of a speed change and the new speed selected are coupled together to arrive at an appropriate position at 18:25. As will be seen, the filter had no difficulty finding paths that agreed with the measurement data. An initial Mach number was selected from a uniform prior between 0.73 and 0.84; this was chosen on the basis of expert advice to ensure that the required flight endurance is achievable.

It is possible to define a much earlier prior using only the ACARS data from early in the flight. In this case, the turns illustrated in the radar data become part of the unknown aircraft flight path to be estimated. Chapter 10 illustrates that this approach leads to a larger search zone, including the search zone resulting from the use of the radar data, and extending further North. This broadening occurs because the flight path from the final ACARS report to 18:25 is much less constrained.