Chapter 7
Aircraft Manoeuvre Dynamics

As discussed in the previous chapter, the dynamics model used for this analysis consists of a sequence of deliberate manoeuvres interspersed with periods of cruise, in which the speed and control angle are almost constant. This chapter provides details of the statistical models used for manoeuvres. These models describe both the frequency of the aircraft manoeuvres and how the aircraft state changes as a result of each manoeuvre.

Three types of manoeuvre were incorporated into the model: turns, that result in changes to the aircraft control angle; accelerations, that result in changes to the aircraft Mach number; and vertical manoeuvres, that result in changes to the aircraft altitude. Each type of manoeuvre was assumed to occur independently. In practice it is common for multiple changes to occur together. The independent model does not preclude this but it does not favour it either. Without data to build a correlated model, independence is a pragmatic assumption.

7.1 Manoeuvre Frequency

Each of the aircraft manoeuvres can be described by three parameters: the start time of the manoeuvre, the rate of manoeuvre (which will be assumed to be constant), and the extent or duration of the manoeuvre. For example a turn can be defined using its start time, the rate of change of angle, and the angle turned through. We first describe the model for the start time of manoeuvres, which is the same for each manoeuvre type.

The time between manoeuvres was modelled as an exponential distribution. The exponential distribution has a single parameter, which can be interpreted as the average time between events. The different types of manoeuvre are modelled with

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_7

a single average manoeuvre period denoted by τ. The use of a single time constant implies an intuitive model where the aircraft can be thought to be either manoeuvring or not, based on the level of pilot interaction. If each manoeuvre was instantaneous, for example the speed changed from Mach 0.83 to Mach 0.81 in zero seconds, then the exponential distribution gives rise to a Poisson arrival process with an average number of T/τ accelerations in T seconds. In practice, the model takes a finite time to change from the pre-manoeuvre state to the post-manoeuvre state and the exponentially distributed delay to the next manoeuvre is applied after the end of the previous manoeuvre. The impact of this is to increase the effective average time between manoeuvres by the average manoeuvre duration. It also imposes a maximum number of manoeuvres in time T whereas the Poisson distribution has infinite support. We assume that each manoeuvre duration is short enough to omit it from the model description below.

The probability of making a turn at time t under the exponential model is given by

$$p(t;\tau) = \frac{1}{\tau}\exp\left\{-\frac{t}{\tau}\right\}, \tag{7.1}$$

and the probability of having no turns in the interval $[0, T]$ is

$$p(t > T;\tau) = \exp\left\{-\frac{T}{\tau}\right\}, \tag{7.2}$$

Using this model, and neglecting the turn duration, the prior likelihood of a sequence of N turns at times $0 \le t_1 \le \ldots t_N \le T$ is

$$\begin{aligned} p(t_{1:N};\tau) &= \left[\prod_{n=1}^{N} p\left(t_n|t_{1:n-1};\tau\right)\right] p\left(t_{N+1} > T|t_{1:N}\right) \\ &= \left[\prod_{n=1}^{N}\frac{1}{\tau}\exp\left\{-\frac{t_n - t_{n-1}}{\tau}\right\}\right]\exp\left\{-\frac{T - t_N}{\tau}\right\} \\ &= \tau^{-N}\exp\left\{-\frac{T}{\tau}\right\}, \end{aligned} \tag{7.3}$$

where $t_0 = 0$ is not a turn time but makes the notation more convenient.

Similar expressions hold for the accelerations and vertical manoeuvres. Since all three manoeuvres are assumed to occur independently, the probability of a sequence of turns, accelerations and vertical manoeuvres is the product of the turn sequence probability, the acceleration sequence probability and the vertical manoeuvre sequence probability.

7.2 Manoeuvre Extent

The model requires a statistical description of how the aircraft manoeuvres as well as when it manoeuvres. As stated above, we assume that all manoeuvres happen at a constant rate; for example, the heading angle could change steadily at $0.5^\circ\,s^{-1}$. For real aircraft manoeuvres, the precise value of this rate can be different but in the context of the flight prediction required for MH370, it is possible to assume fixed values without losing diversity of the sampled paths. Different rates of turn and longitudinal accelerations were explored in earlier versions of the model but these changes simply resulted in a negligible increase in position uncertainty after the manoeuvre.

The aircraft turns were assumed to follow a simple bank model where the relationship between the bank angle and the angular velocity of a turn follows

$$\omega = \frac{\tan\left(\theta_{\text{bank}}\right) g}{v}, \tag{7.4}$$

where θ_{bank} is the aircraft bank angle and g is the acceleration due to gravity. Tight turns with a high angular velocity require low speed or high bank angle. Varying the bank angle was found to produce only minor variations in the overall aircraft trajectory since turns were only a small portion of the trajectories, so a fixed bank angle of 15° was chosen. This is nominally a steep bank for a commercial aircraft, although well within its performance limits. At an air speed of 500 kn, for example, the turn rate would be approximately $0.6^\circ\,s^{-1}$ and a 90° turn would take 2.5 min to perform.

The final parameter required to characterise a turn is the duration of the turn, or equivalently the total change in angle. The aircraft was assumed to not make turns of more than 180° and the change in angle was uniformly sampled from −180° to 180°. It is possible for the model to sample two or more turns in very quick succession and as such looping turns of more than 180° are possible, though unlikely. The likelihood of a quick succession of turns like this depends on the mean time τ.

The accelerations were also assumed to occur with a fixed rate of change of Mach. The assumed rate corresponds to a change of Mach of 0.1 in one minute. This nominal number was not matched to aircraft performance standards but instead simply acts to define a finite duration for each speed change. The Mach number after an acceleration was uniformly sampled from 0.73 to 0.84; this range was chosen on the basis of expert advice to ensure that the required flight endurance is achievable.

The vertical manoeuvres were assumed to occur at a nominal rate of 4,000 ft per minute. Again this rate is not meant to model the actual typical behaviour of commercial aircraft. Rather it provides a reasonable finite time extent to the manoeuvre. The altitude after manoeuvre was uniformly selected in steps of 1,000 ft between 25,000 and 43,000 ft.

Using finite time to execute manoeuvres rather than making instantaneous changes allows for the possibility that a measurement could have been collected during a manoeuvre. The model retains the ground velocity vector but not the altitude rate, so in principle it could describe measurements collected part way through a turn or acceleration but not during climbs or descents. The vertical rate will affect a BFO measurement but since the model only uses a nominal vertical rate it is unlikely to match any actual vertical manoeuvre in detail. In practice, if a measurement was collected during a turn or acceleration it would be very difficult for the filter to infer the trajectory since it would need to model the instantaneous ground speed at the measurement time to match the measured BFO. For the C-channel data at 18:39 and 23:15 there are clusters of BFO measurements and these appear to be statistically stationary, i.e., they do not support the premise that the aircraft is turning during collection.

The three types of manoeuvre are independent and the distribution of manoeuvre extents is uniform for each type, the overall prior probability of a sequence of turns $[(t_{\theta,1}, \theta_1), \ldots, (t_{\theta,N_\theta}, \theta_{N_\theta})]$, accelerations $[(t_{a,1}, v_1), \ldots, (t_{a,N_a}, v_{N_a})]$ and altitude changes $[(t_{h,1}, h_1), \ldots, (t_{h,N_h}, h_{N_h})]$ is given by

$$p\left(\{t_{\theta,1:N_\theta}, \theta_{1:N_\theta}\}, \{t_{a,1:N_a}, v_{1:N_a}\}, \{t_{h,1:N_h}, h_{1:N_h}\}; \tau\right) = \exp\left\{-\frac{3T}{\tau}\right\}\left(\tau\delta\theta\right)^{-N_\theta}\left(\tau\delta a\right)^{-N_a}\left(\tau\delta h\right)^{-N_h}, \tag{7.5}$$

where $\delta\theta = 2\pi$ rad is the span of the uniform turn angle distribution, $\delta a = 0.11$ (Mach) is the span of the uniform speed distribution, and $\delta h = 19$ is the number of possible discrete altitude values.

7.2.1 Parameter Selection

The exponential delay model is parameterised by the average time between events, τ. It is not clear what this value should be or even if it should be constant. In particular one could argue that the behaviour of the aircraft during the accident flight did not match typical commercial aircraft. The average delay is therefore treated as a constant hyperparameter that is potentially different for every hypothesised flight trajectory, but fixed over time.[1] A Jeffreys prior [24] was applied to the time constant; this prior is a non-informative distribution in the sense that it is invariant to the parameterisation used. For example, it yields the same result for the two parameterisations of the exponential distribution, i.e., the rate parameterisation and

[1] Additional experiments were conducted in which the parameter was permitted to change at 19:41 (well after the initial manoeuvre), but little change to the final distribution was evident.

the time-constant parameterisation. The Jeffreys prior is proportional to the square root of the determinant of the Fisher information; in the case of an exponential distribution, it is given by

$$p(\tau) = \begin{cases} K(\tau_1, \tau_2)\, \tau^{-1} & \tau_1 \leq \tau \leq \tau_2 \\ 0 & \text{otherwise} \end{cases}, \tag{7.6}$$

where $K(\tau_1, \tau_2) = \Big(\log(\tau_2) - \log(\tau_1)\Big)^{-1}$. A support range of $0.1 < \tau < 10\,\text{h}$ between events was chosen for each manoeuvre type, spanning the range of cases when manoeuvres occur every few minutes, to when manoeuvres are rare in the entire flight. A single, common parameter was drawn for all three types of manoeuvre.

The joint pdf of the full manoeuvre description is given by

$$\begin{aligned} &p\Big(\{t_{\theta,1:N_\theta}, \theta_{1:N_\theta}\}, \{t_{a,1:N_a}, v_{1:N_a}\}, \{t_{h,1:N_h}, h_{1:N_h}\}, \tau\Big) \\ &\quad = p(\tau) \exp\left\{-\frac{3T}{\tau}\right\} \Big(\tau\delta\theta\Big)^{-N_\theta} \Big(\tau\delta v\Big)^{-N_a} \Big(\tau\delta h\Big)^{-N_h}. \end{aligned} \tag{7.7}$$

7.2.2 Manoeuvre Model Summary

Figure 7.1 shows a graphical representation of the aircraft dynamics model. Ignoring measurements (which are not illustrated), the sequence of manoeuvre times and extents are conditionally independent from other kinds of manoeuvres given the

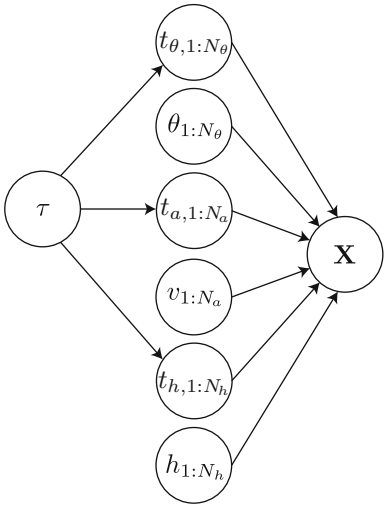

Fig. 7.1 Graphical model showing the conditional dependencies in the probabilistic model, where $\mathbf{X} = (\mathbf{x}_0, \ldots, \mathbf{x}_K)$ represents the entire trajectory

common manoeuvre time constant. The aircraft state forms a Markov chain conditioned on the manoeuvre parameters. The measurements (not illustrated) are conditionally independent from each other given the state sequence.

7.3 Example Realisations

In order to illustrate the behaviour of the complete dynamics model, Fig. 7.2 shows example realisations of trajectories randomly sampled from the model. We emphasise that these do not incorporate information from the SATCOM measurements, but rather represent samples of the prior distribution over trajectories (e.g., $p(\mathbf{x})$ in (3.3)). Each trajectory starts with independent samples of the manoeuvre time constant and initial location and velocity sampled from the accident flight prior, and then draws an independent set of sampled turns, and speed and altitude changes. The prior starts the particles at 18:02 and the figure shows each trajectory through to 00:19, the time of the final SATCOM message. As the figure shows, some of the trajectories exhibit many manoeuvres, while others do not turn or change speed at all.

The histogram of the number of turns and speed changes hypothesised by the prior model is shown in Fig. 7.3. The model is in effect a mixture of Poisson distributions, consisting of a continuum of components ranging from very low times between manoeuvre (six minutes) to very high times between manoeuvre (ten hours). The distribution is the same for each type of manoeuvre since the prior distributions of time constants are the same.

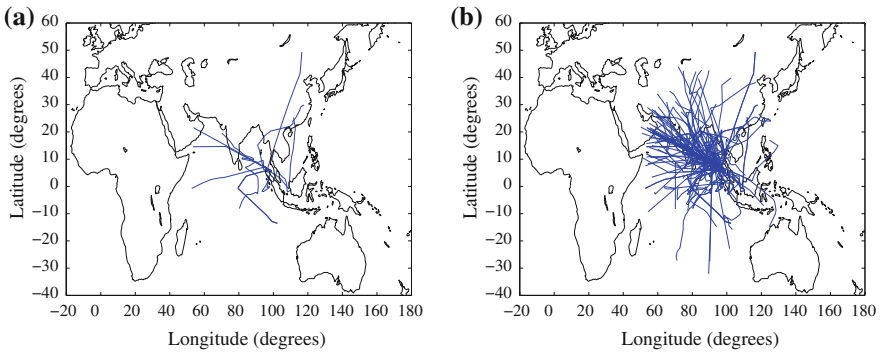

Fig. 7.2 Example trajectories sampled from prior distribution of flight dynamics: **a** 10 random realisations; **b** 100 random realisations

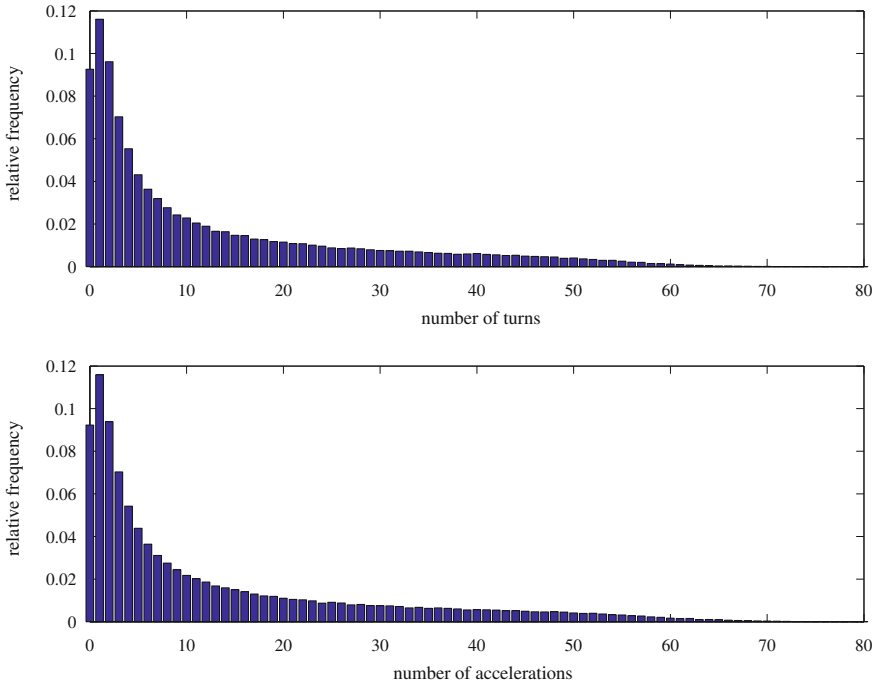

Fig. 7.3 Histograms of the number of manoeuvres selected by 200,000 independent realisations of the model (i.e., not using any BTO/BFO measurements)

Although it is unlikely that any given trajectory will travel for more than six hours with no manoeuvre, those that do all end up in roughly the same place. Figure 7.4 shows a contour representation of the pdf of latitude and longitude at 00:19 that was generated using 200,000 particle draws from the prior model with no measurements and an isotropic kernel. The red diamond in the figure shows the mean of the prior. There is a clear peak corresponding to paths that make no manoeuvre. This is smeared around an arc corresponding to paths that turn once but maintain a steady speed and also radially from the start point corresponding to paths that do not turn but do change speed. The model samples a hyperparameter for the time between manoeuvres from a prior that covers 0.1 to 10 h. When a particle samples a very high value for the time constant it is unlikely to choose to manoeuvre. Where manoeuvres were selected it was again unlikely to choose to change speed and direction, leading to the contours shown. The altitude behaviour of the model is not visualised in the figure. There is an approximately circular region around the initialisation point where the pdf is elevated. This corresponds to particles that have sampled a very small time between manoeuvres: these turn so much that they circle back on themselves repeatedly and are effectively trapped close to their starting location. Beyond this circle, the pdf is

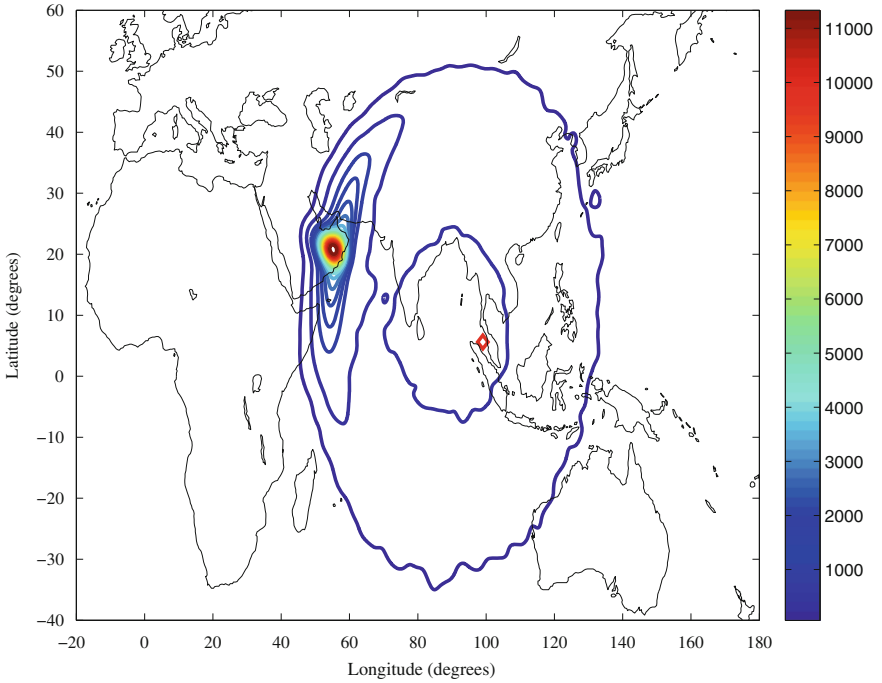

Fig. 7.4 Distribution of position at 00:19 using only model predictions (i.e., not using any of the BTO/BFO measurements). The *red* diamond marks the mean on the prior

lower behind the start point than in front of it. It is also clear that the search region is in a very low probability part of the pdf: this indicates that the prior is not biased towards selecting a particular part of the 00:19 arc. The search region is inside the lowest probability contour used in the map.