

Chapter 5

Measurement Model, Satellite Communications

The Bayesian filter discussed in Chap. 3 relies on knowledge of three probability density functions: the state prior distribution, the state stochastic model, and the measurement conditional probability density. The prior used for the analysis in this book was discussed in Chap. 4. This chapter addresses the measurement probability density and Chaps. 6 and 7 discuss the state dynamics model.

The most general measurement model was defined in (3.5) and simply states that the measurement is some function of the aircraft state and measurement noise. In many systems it is not too restrictive to assume that the noise is additive, in which case (3.5) becomes

$$\mathbf{z}_k = \mathbf{h}_k(\mathbf{x}_k) + \mathbf{w}_k. \quad (5.1)$$

Prescribing the nonlinear function $\mathbf{h}_k(\mathbf{x}_k)$ and the statistical distribution of the measurement noise provides a complete description of the measurement probability density. The measurements available for the accident flight are timing and frequency logs of communication messages between the aircraft and a ground station. Details of the communication system software and hardware combine with the physics of the communication geometry to determine the nonlinear measurement function. The statistics of the noise were determined empirically by analysing a population of real measurements for known aircraft states.

This chapter gives a brief overview of the satellite communication system and then describes the nonlinear measurement functions and empirical noise models for the timing and frequency measurements. The majority of the communications messages available were automated signalling messages, but there were also two telephone calls made to the plane that remained unanswered. The first of these is particularly important because of when it occurred. The chapter concludes with a description of the measurement model for telephony. Further details may be found in [2].

5.1 Satellite Communications System

The accident aircraft was fitted with a satellite communications terminal that used the Inmarsat Classic Aero system [2]. This system uses a satellite to relay messages between the aircraft and a ground station. In the accident flight the messages were passed between the aircraft and a ground receiving unit located in Perth, Australia, via the Inmarsat-3F1 satellite. Figure 5.1 illustrates the satellite communication system in use during the accident flight. The aircraft is referred to as the Aircraft Earth Station (AES) and the ground receiving unit is referred to as the Ground Earth Station (GES). Inmarsat-3F1 is a satellite in geosynchronous orbit with longitude 64.5° East and was used exclusively for the duration of the accident flight.

An AES is equipped with a satellite data unit that comprises a satellite modem and auxiliary hardware and software. Transmission of data over the satellite is via bursts which are scheduled to arrive at the GES at a specified time and frequency. Communications from multiple users are coordinated by the allocation of different time and frequency slots to each user. Return channel (AES to GES) time slot boundaries are referenced to the forward channel (GES to AES) [2]. The duration of each time slot is sufficient to account for all possible positions of the aircraft with respect to the satellite. The width of each frequency slot is determined by the data rate and a guard width that accounts for possible variations in the satellite oscillator frequency and other possible frequency offsets. Frequency compensations applied onboard the aircraft (aircraft induced Doppler pre-compensation) and at the ground station serve to reduce the possible difference between the expected and actual frequency of the




Fig. 5.1 System model of the satellite communication system

messages received from the aircraft. The on-ground compensation makes use of a second ground station located in Burum, Netherlands that transmits a reference signal to the Inmarsat-3F1 satellite which is relayed to the Perth GES. Its purpose is to enable the receive modem in the GES to compensate for the Doppler frequency shift from the satellite to the Perth GES. This compensation process is referred to as Enhanced Automatic Frequency Correction.

After the Enhanced Automatic Frequency Correction process, the expected time of arrival of each communications burst is compared with the actual time of arrival and the difference between the two is referred to as the Burst Timing Offset (BTO). The BTO is minimised when the elevation angle to the satellite is 90° and increases as the aircraft moves away from the sub-satellite position. Hence, the BTO is a measure of how far the aircraft is from the sub-satellite position. Similarly, the difference between the expected frequency of each communications burst and the actual received frequency is referred to as the Burst Frequency Offset (BFO). The BFO is a function of the relative speed between the aircraft and the satellite. Given that the satellite position and speed are known, the BFO provides information about the aircraft velocity vector. The BTO and BFO are logged by the ground station for every communications burst. This logging was a relatively recent addition to the ground station following the Air France 447 accident [2, 45] and was intended to assist in locating an aircraft. Statistical models for these two measurement functions are now developed.

5.2 Burst Timing Offset

The Inmarsat Classic Aero system allocates a time slot for communications based on a nominal propagation delay that assumes a nominal satellite position and a nominal aircraft position. The nominal aircraft position is at zero altitude directly below the satellite's nominal orbital position of 64.5°E longitude, zero latitude and an altitude of 35788.122 km. The round trip delay is proportional to the distance from the ground station to the actual aircraft location via the actual satellite location. The actual satellite position is different from the nominal satellite position because Inmarsat-3F1 is not exactly motionless, but rather moves in a known way in a region about its nominal location. The Burst Timing Offset is the additional delay after the start of the allocated time slot at which the message is received [2]. The BTO is thus the difference between the round trip message delay and the nominal delay used for scheduling. In addition to the propagation delay the message delay includes the latency of the satellite data processing unit. Denoting a BTO measurement at time t_k as z_k^{BTO} , the BTO measurement function is

$$z_k^{\text{BTO}} = h_k^{\text{BTO}}(\mathbf{x}_k, \mathbf{s}_k) + w_k^{\text{BTO}}, \quad (5.2)$$

$$h_k^{\text{BTO}}(\mathbf{x}_k, \mathbf{s}_k) = T(\mathbf{x}_k, \mathbf{s}_k) - T^{\text{nom}} + T^{\text{channel}} - T_k^{\text{anomaly}}, \quad (5.3)$$

where

- $T(\mathbf{x}_k, \mathbf{s}_k)$ is the round trip propagation delay from the ground station to the aircraft via the satellite;
- \mathbf{s}_k is the state of the satellite at time t_k , that is its position and velocity in three dimensions, along with the satellite oscillator's state;
- T^{nom} is the nominal round trip delay;
- T^{channel} is a channel dependent bias term due to processing in the satellite data unit;
- T_k^{anomaly} is an anomaly correction term discussed below;
- w_k^{BTO} is a zero mean scalar noise process with statistics to be determined from measurement data logs.

The function $h_k^{\text{BTO}}(\mathbf{x}_k, \mathbf{s}_k)$ is assumed to be deterministic, so the measurement variance is the variance of the noise term w_k^{BTO} . The round trip delay can be expressed in terms of the distances from the satellite to the ground station and the aircraft as

$$T(\mathbf{x}_k, \mathbf{s}_k) = \frac{2}{c} \left(|\mathbf{H}^s \mathbf{s}_k - \mathbf{g}| + |\mathbf{H}^s \mathbf{s}_k - \mathbf{H}^x \mathbf{x}_k| \right), \quad (5.4)$$

where c is the speed of light; \mathbf{g} is the location of the ground station; \mathbf{H}^s selects the position elements of the satellite state; \mathbf{H}^x selects the position elements of the aircraft state; and the distance $|\cdot|$ is the three dimensional Cartesian distance, i.e. the l^2 norm.

Combining the nominal locations of the satellite and aircraft with the known location of the Perth ground station gives a value of $T^{\text{nom}} = 499,962 \mu\text{s}$.

There are a number of different channel types used that carry different traffic types and have different baud rates. Communications from the aircraft to the ground are typically over the R- and T-data channels with C-channel used for voice telephone calls. Communications from the ground to the aircraft are over the P-channel. The channel dependent calibration term T^{channel} is assumed to be constant over a single flight but can vary between flights. A fixed value for each flight assessed was empirically derived by comparing the communications logs with known aircraft positions: the calculated value of T^{channel} was the mean difference between the measured BTO and the expected BTO calculated using the known aircraft location. For the accident flight this calibration is only available for the time when the plane was at the tarmac and for the first half hour of flight. As such, values from the previous flight were also used in the calculation of T^{channel} . The majority of the messages available from the accident flight are R1200 messages for which $T^{\text{channel}} = -4,283 \mu\text{s}$.

The anomaly correction term T_k^{anomaly} was empirically derived through analysis of a collection of communications logs. For some communication messages, typically during initial log-on, there was a very large difference between the measured BTO and the nominal delay. Analysis showed that rather than simple outliers, these anomalous BTO measurements could be *corrected* by a factor of $N \times 7,820 \mu\text{s}$ where N is a positive integer. The origin of these anomalous BTO measurements has not been fully determined, but the empirical correction time is quite close to the

transmission interrupt clock period of $7,812.5 \mu\text{s}$ and the BTO collection process contains quantisation.

The channel dependent calibration and anomaly correction terms result in a residual measurement error that is approximately Gaussian. For the R1200 messages, the empirically derived standard deviation of the measurement noise w_k^{BTO} is $29 \mu\text{s}$, and for R600 messages, $62 \mu\text{s}$. For anomalous R1200 messages a standard deviation of $43 \mu\text{s}$ was used. Figure 5.2 shows a histogram of the residual BTO measurement errors for R1200 messages referenced to the 7 March 2014 T^{channel} value, and the moment-matched Gaussian approximation. The data used to construct the histogram and the empirical parameters were obtained from logs of the 20 flights of 9M-MRO prior to the accident flight.

The histogram in Fig. 5.2 has an underlying mean of $10 \mu\text{s}$. This is due to the channel dependent calibration term T^{channel} not being stationary. Over the span of a

Fig. 5.2 Histogram of BTO residual measurement errors

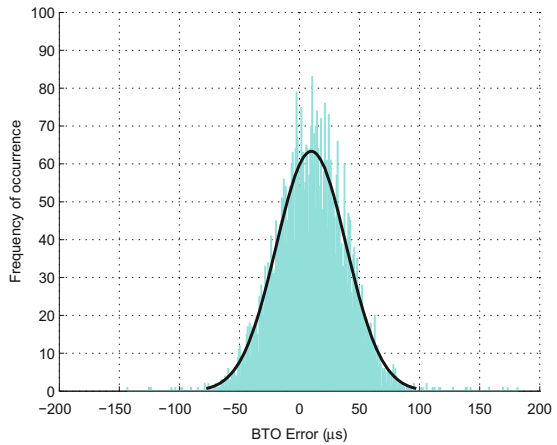



Fig. 5.3 Scatter plot of BTO errors for 20 flights prior to MH370



day it appears constant but in the context of the 20 flights represented in Fig. 5.2 there is a slow variation. As discussed above, different values were fitted for each flight. Figure 5.3 shows a scatter plot of the BTO errors against time. The channel dependent calibration term T^{channel} was matched to the final flight before the accident flight, MH371 (and the beginning of the accident flight) and the BTO errors from MH371 on 7 March 2014 are marked as red crosses. The variation in bias is sufficiently slow that assuming it is the same for the accident flight as the previous flight is satisfactory.

5.3 Burst Frequency Offset

The Burst Frequency Offset is a function of the Doppler shifts imparted on the communication signal due to the motion of the satellite and the aircraft. The relationship is more complicated than a direct Doppler calculation because the aircraft software contains Doppler compensation that offsets the Doppler shift due to the aircraft motion. Although the aircraft attempts to compensate for its own motion, it does this under the assumption that the communications satellite is in motionless geostationary orbit and it does not include the vertical component of the aircraft velocity (which is non-zero when it is ascending or descending) [2]. Since Inmarsat-3F1 is not exactly geostationary, the compensation is unable to completely remove Doppler effects. Empirical analysis of the BFO was conducted for the 20 flights of 9M-MRO prior to the accident flight. This analysis used the same Doppler correction software as the 9M-MRO satellite data unit to determine the expected BFO given a known reported aircraft position and velocity and compared this with the observed measurements.

Similar to the BTO, the BFO measurement function consists of a nonlinear function of the aircraft and satellite states and several bias terms [2]

$$z_k^{\text{BFO}} = h_k^{\text{BFO}}(\mathbf{x}_k, \mathbf{s}_k) + w_k^{\text{BFO}} \quad (5.5)$$

$$h_k^{\text{BFO}}(\mathbf{x}_k, \mathbf{s}_k) = \Delta F_k^{\text{up}}(\mathbf{x}_k, \mathbf{s}_k) + \Delta F_k^{\text{down}}(\mathbf{s}_k) + \delta f_k^{\text{comp}}(\mathbf{x}_k) + \delta f_k^{\text{sat}}(\mathbf{s}_k) + \delta f_k^{\text{AFC}}(\mathbf{s}_k) + \delta f_k^{\text{bias}}(\mathbf{x}_k, \mathbf{s}_k), \quad (5.6)$$

where

- $\Delta F_k^{\text{up}}(\mathbf{x}_k, \mathbf{s}_k)$ is the uplink (aircraft to satellite) Doppler shift;
- $\Delta F_k^{\text{down}}(\mathbf{s}_k)$ is the downlink (satellite to ground station) Doppler shift;
- $\delta f_k^{\text{comp}}(\mathbf{x}_k)$ is the frequency compensation applied by the aircraft;
- $\delta f_k^{\text{sat}}(\mathbf{s}_k)$ is the variation in satellite translation frequency: the satellite uses a local oscillator to translate the carrier frequency of the message;
- $\delta f_k^{\text{AFC}}(\mathbf{s}_k)$ is the frequency compensation applied by the ground station receive chain;
- $\delta f_k^{\text{bias}}(\mathbf{x}_k, \mathbf{s}_k)$ is a slowly varying bias due to errors in the aircraft and satellite oscillators and processing in the satellite data unit;
- w_k^{BFO} is a zero mean scalar noise process with statistics to be determined from measurement data logs.

This function was described in detail in [2], we review it briefly and elaborate where the analysis herein makes different modeling assumptions to those in [2]. Again, the function $h_k^{\text{BFO}}(\mathbf{x}_k, \mathbf{s}_k)$ is assumed to be deterministic, so the measurement variance is the variance of the noise term w_k^{BFO} . This is a less reliable assumption than for BTO because the bias term $\delta f_k^{\text{bias}}(\mathbf{x}_k, \mathbf{s}_k)$ changes. To compensate for this, the measurement variance was inflated from the empirically derived w_k^{BFO} variance.

The uplink Doppler shift can be expressed as a product of the uplink frequency with the inner product of the relative velocity (normalised by the speed of light, c) and a unit vector along the relative displacement between the aircraft and the satellite

$$\Delta F_k^{\text{up}}(\mathbf{x}_k, \mathbf{s}_k) = \frac{F^{\text{up}}}{c} \frac{(\mathbf{V}^s \mathbf{s}_k - \mathbf{V}^x \mathbf{x}_k)^\top (\mathbf{H}^s \mathbf{s}_k - \mathbf{H}^x \mathbf{x}_k)}{|\mathbf{H}^s \mathbf{s}_k - \mathbf{H}^x \mathbf{x}_k|}, \quad (5.7)$$

where F^{up} is the uplink carrier frequency; \mathbf{V}^s selects the velocity elements of the satellite state; and \mathbf{V}^x selects the velocity elements of the aircraft state. Similarly the downlink Doppler shift and the frequency compensation are given by

$$\Delta F_k^{\text{down}}(\mathbf{s}_k) = \frac{F^{\text{down}}}{c} \frac{(\mathbf{V}^s \mathbf{s}_k)^\top (\mathbf{H}^s \mathbf{s}_k - \mathbf{g})}{|\mathbf{H}^s \mathbf{s}_k - \mathbf{g}|}, \quad (5.8)$$

$$\delta f_k^{\text{comp}}(\mathbf{x}_k) = \frac{F^{\text{up}}}{c} \frac{(\tilde{\mathbf{V}}^x \mathbf{x}_k)^\top (\tilde{\mathbf{H}}^x \mathbf{x}_k - \tilde{\mathbf{s}})}{|\tilde{\mathbf{H}}^x \mathbf{x}_k - \tilde{\mathbf{s}}|}, \quad (5.9)$$

where $\tilde{\mathbf{s}}$ is the nominal satellite position assumed by the aircraft, and F^{down} is downlink carrier frequency. The aircraft frequency compensation term $\delta f_k^{\text{comp}}(\mathbf{x}_k)$ is determined using the aircraft's own knowledge of its position and velocity but with an assumed altitude of zero and an assumed vertical speed of zero. The modified position matrix $\tilde{\mathbf{H}}^x$ selects only the horizontal location variables and sets the altitude to zero, and similarly for $\tilde{\mathbf{V}}^x$. The compensation also assumes a motionless satellite at its nominal satellite location of 64.5°E. The satellite altitude used in the correction is 422 km higher than the nominal 35788.12 km value.

The satellite translates the frequency of messages using a local oscillator that is maintained in a temperature-controlled enclosure to improve its stability. During eclipse periods, when the satellite passes through the Earth's shadow, the satellite temperature drops, resulting in a small variation in translation frequency [2]. An eclipse period occurred during the accident flight and some of the validation flights were also affected by eclipses. The oscillator temperature also varies with time of day as the satellite orientation to the sun changes and as the temperature control system applies its controls. All of these thermal effects are included in the term $\delta f_k^{\text{sat}}(\mathbf{s}_k)$. The specific details of the functions that define $\delta f_k^{\text{sat}}(\mathbf{s}_k)$ and $\delta f_k^{\text{AFC}}(\mathbf{s}_k)$ are proprietary of Inmarsat.

The bias term $\delta f_k^{\text{bias}}(\mathbf{x}_k, \mathbf{s}_k)$ is time varying. In the BTO case the variations in bias were slow enough to be ignored within a single flight and we were able to assume that T^{channel} was constant for each flight. This is not the case for the BFO bias term.




Fig. 5.4 Results for the 2-Mar-2014 flight from Mumbai to Kuala Lumpur

The mean bias is different between flights and even within a single flight there is evidence of structured variation. Figure 5.4 shows an example of the BFO measurement errors for a flight between Mumbai (BOM) and Kuala Lumpur (KUL) on 2 March 2014. The figure shows the difference between the measured BFO values and predicted values (using the actual SDU software for determining $\delta f_k^{\text{comp}}(\mathbf{x}_k)$ in the SATCOM system model) based on the known geometry and aircraft velocity vector at the time. The bias used for the plot was obtained by analysing BFO measurements while the aircraft was on the tarmac. The residual error is clearly not zero-mean, and the mean varies with time. Substantial effort was made to characterise this structured bias. It was found to have a geographic dependency but it has not been possible to determine a quantitative function to compensate for this change in bias.


The variations in bias shown in Fig. 5.4 happen over a timescale of minutes rather than hours. In the accident flight the available BFO values are generally at least an hour apart. This is a relatively long time compared with the correlation structure of the error, so the model does not use a coloured noise model for the BFO. However, the drift of the BFO bias means that it is not sufficient to assume that $\delta f_k^{\text{bias}}(\mathbf{x}_k, \mathbf{s}_k)$ will be the same in flight as on the tarmac before takeoff. The potential variations were incorporated by modeling the BFO bias as an unknown constant with a prior mean given by the tarmac value and a standard deviation of 25 Hz. Since the BFO measurement Eq. (5.6) is linear in the bias its distribution conditioned on the other states can be estimated with a Kalman filter. This is the Rao-Blackwellised particle filter described in Sect. 3.3.

Empirical statistics of the residual measurement noise w_k^{BFO} were determined using the previous 20 flights of 9M-MRO. Data points corresponding to when the aircraft was climbing or descending were excluded. Table 5.1 shows the empirical statistics of the BFO measurements for R1200 and R600 messages. The ‘in-flight only’ statistics show the combined effects of noise and bias variation without the influence of ‘on-tarmac’ outliers (potentially due to taxiing). The ‘including tarmac’ statistics on the other hand are also influenced by the BFO bias value applied to keep the BFO error at the source tarmac for R1200 messages close to zero. The mean BFO error was close to zero in all cases, indicating appropriate $\delta f_k^{\text{bias}}(\mathbf{x}_k, \mathbf{s}_k)$ values were chosen for each flight. The statistics show that even when outliers are discarded a standard deviation of about 4.3 Hz is applicable. As discussed above,

Table 5.1 Statistics of BFO errors for 20 flights of aircraft 9M-MRO prior to MH370

	Mean BFO error (outliers included) (Hz)	Standard deviation of BFO (outliers included) (Hz)	Mean BFO error (outliers excluded) (Hz)	Standard deviation of BFO (outliers excluded) (Hz)
Including tarmac	0.2246	4.9592	0.2745	4.0192
In-flight only	0.1079	5.4840	0.1755	4.3177

Fig. 5.5 Histogram of BFO errors for 20 flights prior to MH370 (using only in-flight BFOs)



to be conservative and allow for potential variation in the $\delta f_k^{\text{bias}}(\mathbf{x}_k, \mathbf{s}_k)$ value on the accident flight, our model assumes a noise standard deviation of 7 Hz. Section 5.5 illustrates the sensitivity of the BTO and BFO measurements to variations in the aircraft state.

Figure 5.5 shows a histogram of the 3392 in-flight BFO errors. On-tarmac BFO errors were excluded due to the pre-biasing described above. A Gaussian fit to the distribution is shown as a black line. It can be seen that the distribution shows some non-Gaussian features and the tails of the distribution for negative errors are somewhat heavier than those for positive errors.

5.4 C-Channel Telephone Calls

There were two unanswered telephone calls from the ground to MH370 after the loss of radar data. These communications use the C-channel and result in measurements of BFO but not BTO. Initially the C-channel data was not included in the flight prediction but analysis from DST Group highlighted that the first of these calls provides critical information. The first call occurred from 18:39:53 to 18:40:56 and is important because the measured BFO is significantly different from the BFO on




Fig. 5.6 Implied track angles for MH370 with different assumed ground speeds at 18:40, 7 March 2014

the R1200 measurement preceding it at 18:28:15. The R1200 BFO value is consistent with the speed and direction of the aircraft while under radar coverage whereas the later C-channel BFO value is not. Assuming that the change in BFO implies a turn, the difference between the BFO predicted by using a MATLAB model of the SDU software¹ and the measured BFO on the C-channel was analysed as a function of post-turn direction and for a range of aircraft speeds and turn times between 18:28:15 and 18:39:53. Figure 5.6 shows the residual error and it clearly demonstrates that only Southerly track angles are consistent with the C-channel measurements. The model predicted BFO values of Northerly paths are more than 10 standard deviations away from the measured BFO.

The BTO measurements from the R1200 messages at 18:28:15 and 19:41:03 are not consistent with the velocity vector before 18:28:15. The only way to satisfy these measurements and maintain a feasible air speed is for the aircraft to have turned. However, the time window for this turn is more than an hour. The 18:39:53 C-channel measurement is critical because when combined with the 18:28:15 BFO measurement it significantly restricts this turning time window to a little over 11 minutes. Using the C-channel data restricts the aircraft trajectories much more tightly than using only R-channel data.

¹Note The difference between the MATLAB model output and the SDU software output was found to be inconsequential to this analysis for determining $\delta f_k^{\text{comp}}(\mathbf{x}_k)$ in the SATCOM system model.

5.5 Information Content of Measurements

The information content of the BTO and BFO measurements is illustrated in Fig. 5.7. The figure shows a small part of the likelihood function of the BTO and BFO measurements at 19:41:02. The plots were created assuming an altitude of 30,000 ft. The BFO diagram used an assumed aircraft position of 1°S and 93.6°E and a bias of 150Hz, which is the tarmac value for the accident flight. The BFO contour shape varies slowly with aircraft position. The figure used the measurement error model described earlier in this chapter, namely zero-mean Gaussian noises with standard deviation 29 μ s for BTO, and 7 Hz for BFO.




Fig. 5.7 Likelihood functions for BTO and BFO measurements at 19:41:02

The diagrams show that the BTO provides reasonable localisation along a circle of a given range from the satellite. The information provided by the BFO is less precise, providing information on speed, with standard deviation on the order of 50 kn, and direction on the order of 40°.

Open Access This chapter is distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, duplication, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, a link is provided to the Creative Commons license and any changes made are indicated.

The images or other third party material in this chapter are included in the work's Creative Commons license, unless indicated otherwise in the credit line; if such material is not included in the work's Creative Commons license and the respective action is not permitted by statutory regulation, users will need to obtain permission from the license holder to duplicate, adapt or reproduce the material.