Northern Routes and Burst Frequency Offset for MH370, by Victor Iannello, May 16, 2015

Notice: The views expressed here are solely mine and do not represent the views of the Independent Group (IG), Jeff Wise, or any other group or individual.

Summary

In previous work, paths were reconstructed for MH370 using the available radar and satellite data. Paths to the north of Malaysia were studied by matching the measured Burst Timing Offset (BTO) data, but relaxing the constraint of matching the Burst Frequency Offset (BFO), which is appropriate if the BFO data was either corrupted or misinterpreted. It was found that there are paths to the north that end at airports that could be reached with the fuel that was loaded onto MH370.

In this work, the conventional interpretation of the BFO is challenged. In particular, the possibility that the operation of the SATCOM was deliberately modified so that a northern path would have the BFO signature of a southern path is studied. Some of the findings are:

- The Honeywell Thales MCS-6000 SATCOM used by MH370 has a frequency correction algorithm with the capability to correct for the Doppler shift caused by inclination of the satellite. This is known to the official investigation team but is not generally known by independent researchers.
- The value of inclination for the Inmarsat I3F1 satellite that was broadcast by the Ground Earth Station (GES) at Perth, Australia, to be used by SATCOMs logged into the satellite, was zero. The true inclination of the satellite was around 1.65°. The two parameters that describe the satellite inclination, the inclination angle and the time of the ascending node, are stored in the System Table of the SATCOM in non-volatile memory, and are used by the frequency compensation algorithm.
- If an individual obtained unauthorized access to the non-volatile memory of the SATCOM, the value of the inclination used by the frequency correction algorithm could be changed from 0 to 3.3°, or about twice the true inclination of the satellite. With this change, the BFO signature of a northern path that satisfied the BTO data would resemble the BFO signature of a southern path that satisfied the BTO data.
- The apparent turn to the south between 18:28 and 18:40 UTC that is suggested by the measured BFO data might have been caused by a change to the inclination parameters stored in the SATCOM's System Table during that time interval.
- The calculated values of BFO for northern paths with the inclination parameter changed to 3.3° match the measured BFO values with an RMS error less than 3.8 Hz. This is true for Mach numbers between 0.65 and 0.85 at FL350, with little variation in error seen in this speed range.
- At each log-on, the inclination parameters would be reset to zero. Therefore, the BFO data associated with the log-ons at 18:25 and 00:19 UTC should be evaluated with inclination parameters set to zero. The BFO data at times between these log-ons should be evaluated with the possibility that a change was made.
- The BFO value at 00:19 matches an aircraft along the northern part of the 7th arc on the ground and stationary once the BFO is adjusted for the log-on offset seen at 16:00 UTC. This suggests that if MH370 flew north, it might have successfully landed.

- Researchers have identified security vulnerabilities in other SATCOMs, including backdoors and access to memory, although the MCS-6000 has not been specifically studied. The possibility of "spoofing" the BFO to disguise location has been considered before.

Introduction

A number of analysts have studied the Burst Timing Offset (BTO) and Burst Frequency Offset (BFO) data from flight MH370, as relayed by Inmarsat I3F1 satellite and received by the Ground Earth Station (GES) in Perth, Australia. The aircraft was a Boeing B777-200ER registered as 9M-MRO. The satellite data suggests that the aircraft flew to the South Indian Ocean (SIO) and exhausted its fuel. The constraint of matching the BFO data within any reasonable limit on error eliminates the possibility of a northern path. Performance constraints such as fuel consumption and unattended (autopilot) navigation also limit the range of possible end points for southern paths.

More recently, I have reconstructed paths to the north that are allowed by the BTO data by ignoring the BFO data. The study was performed to assess the possibility of a successful landing in the event that the BFO data from MH370 is either corrupted or has been misinterpreted. I used the BTO data to identify the paths to the north that end at airports along the 7th arc at 00:19 UTC with the requirement that the BTO data is matched at all other handshake times. Three airports were identified that are located within the error bounds of the 7th arc. Of the three, there appears have been insufficient fuel to reach Kyzylorda Airport. On the other hand, there appears have been sufficient fuel to reach Almaty and Kuqa Qiuci Airports, although there would have been significantly less fuel margin to reach Kuqa Qiuci. Near to Almaty is a smaller airport named Boraldai that is also viable for landing. It is very unlikely that MH370 reached the runway at Yubileyniy, which was suggested in a scenario by Jeff Wise.

In the current study, I investigate the possible corruption of the BFO data. In particular, I have researched a technical scenario in which the operation of the SATCOM was deliberately modified so that a path to the north would have the BFO signature of a path to the south. The technical feasibility of malicious intrusion into the SATCOM is also discussed.

BTO and BFO Calculations

The methodology to reconstruct northern paths is described in my previous work, and is similar to what has been presented by others, including the published work of Inmarsat's Chris Ashton and the IG's Richard Godfrey . A BTO value defines an arc on the surface of the earth, and paths can be reconstructed that cross these arcs at the appropriate time by matching the satellite-aircraft range. (The exact position of the arc depends on the altitude of the aircraft. At higher altitudes, the arc is located further from the subsatellite position.) The paths were reconstructed by forward integrating in time and exactly matching the satellite-aircraft range at handshake times as derived from the BTO values and the satellite position. Paths for which the Mach number was held constant were studied. The model includes an accurate parameterization of the satellite position and velocity, meteorological data, and the WGS84 model of the earth's ellipsoid geometry.

The explanation of the BFO data is more complicated. Each measured value represents the sum of six terms, as described by the following equation:

$$BFO = \Delta f_{up} + \Delta f_{AES} + \Delta f_{down} + \Delta f_{AFC} + \Delta f_{sat} + B \qquad (1)$$

where the SATCOM on the plane is referred to as an Aircraft Earth Station (AES), the corresponding Ground Earth Station (GES) is located in Perth, Australia, and

- Δf_{up} is the uplink Doppler shift due to relative motion between the aircraft and the satellite
- Δf_{AES} is the frequency correction applied by the AES to compensate for Δf_{up}
- Δf_{down} is the downlink Doppler shift due to the relative motion between the satellite and the GES
- Δf_{AFC} is the correction applied by the Enhanced Automatic Frequency Compensation at the GES to compensate for Δf_{down}
- Δf_{sat} is the frequency offset of the satellite due to thermal effects
- B is the fixed frequency bias

In accordance with Chris Ashton's work, a bias value of B = 150 Hz is used, which is a calibrated value based on the measured BFO when 9M-MRO was parked at Kuala Lumpur International Airport (KLIA). Geoff Hyman and Barry Martin have shown that the bias B shows a statistical correlation with the particular GES channel used for the AES-GES communications, with a value of 150 Hz for channels R4 and R11 and 154 Hz for channels T10, T12, R8, and T8. In fact, the handshakes at 19:41, 20:41, 21:41, 22:41, and 00:11 UTC all were on channel R4, so the use of a constant value B = 150 Hz is justified for the work here. (The satellite calls at 18:40 and 23:14 UTC produced BFO values using channel C6; unfortunately, there is no calibration data available for this channel to determine if a different bias is warranted.)

Of the six terms in Eq (1), four (Δf_{down}, Δf_{AFC}, Δf_{sat}, and B) are independent of the specific path of the aircraft and can be calculated based on the satellite position and velocity vectors, the location of the GES in Perth, and the measured performance of the EAFC at Perth. That leaves two other parameters, Δf_{up} and Δf_{AES}, that are dependent on the particular path that MH370 followed.

By ICAO specification, an AES is required to compensate for the uplink Doppler shift, Δf_{up}, by applying a frequency correction Δf_{AES} that modifies the transmitted frequency so that the residual offset L is less than 100 Hz, where the residual is given by

$$L = \Delta f_{up} + \Delta f_{AES} \qquad (2)$$

The uplink Doppler shift Δf_{up} is dependent on the following

- The position and velocity of the satellite
- The position, groundspeed, track, and vertical speed of the aircraft
- The frequency of the uplink signal, which is in the L-band, and equal to 1.6466525 GHz

One way to compensate for the uplink Doppler shift is for the AES to calculate what it expects the uplink Doppler shift to be, and then applying a correction to the transmitted frequency which "pre-compensates" for this expected shift. If this correction algorithm were to use the exact position and velocity of the satellite, as well as the exact position, groundspeed, track, and vertical speed of the aircraft, then the residual offset L would be close to zero, and the BFO could not be used to discriminate between different paths followed by the aircraft.

The SATCOM in 9M-MRO was a Honeywell Thales MCS-6000. In this SATCOM, the AES correction algorithm uses a simplification of the full model that would be required to exactly compensate for the

uplink Doppler shift, and the result is a non-zero residual offset. As described by Chris Ashton, the AES correction algorithm makes the following simplifications:

- The satellite is assumed to be geostationary, i.e., relative to the earth, the velocity vector is zero, the altitude above the Earth is constant at 35,786 km, and the subsatellite position remains at a constant longitude at the equator.
- The vertical speed of the aircraft is assumed to be zero.

In fact, we know that Inmarsat's I3F1 satellite is not geostationary. The satellite's orbit is indeed geosynchronous, i.e., one orbit about the earth lasts exactly one sidereal day (23.9345 hr), but the satellite orbit is "inclined" relative to the equator and the orbit is also slightly "eccentric". Because of the inclination, the satellite is not stationary relative to the earth, and because of its eccentricity, the altitude above the earth is not constant. The consequence of the inclination is that the satellite follows a track across the surface of the earth, and its declination (i.e., the latitude) varies sinusoidally with a period of one sidereal day. As a consequence of the inexact cancelation of the uplink Doppler due to the orbit simplification in the AES correction algorithm, the residual L varies according to the path of the aircraft. It is this variation that allowed investigators in the days after the disappearance to conclude that MH370 followed a southern rather than a northern path.

To show how the BFO can be used to discriminate between paths to the north and south, I reconstructed two representative paths in this way:

- For both northern and southern paths, there was an assumed turn at 18:34 UTC.
- After 18:34, the speed was held constant at $M = 0.8$ at an altitude of 35,000 ft (FL350).
- Meteorological data was incorporated so that the effects of wind and temperature aloft were included.
- The BTO was matched exactly at each handshake by allowing turns at each handshake and great circle routes were flown between handshakes.
- At 00:11 UTC, the speed was reduced so that the plane reached the 7th arc at 00:19 at sea level.
- A bias of $B = 167$ Hz was used to calculate the BFO at the log-on at 00:19 UTC on channel R10 since that this bias was observed for the log-on at 16:00 UTC, also on channel R10.
- For the southern path, a steep descent was introduced to match the measured BFO at 00:19. For the northern path, the plane was assumed to be on the ground and stationary.

Figure 1 shows how the BFO varies for the representative northern and southern paths. As can be seen in the figure, the calculated BFO curve for the southern path agrees with the measured BFO values, and the calculated BFO curve for the northern path does not agree with the measured data. Analyses similar to this are what led the official investigative team to conclude that the plane followed a course to the south. The conclusion is independent of the details of the modeled paths, i.e., whether or not MH370 followed a constant track after the turn, the exact airspeed, the time of the turn, etc. These details have been the subject of vigorous debate, and affect the final location of MH370 if it turned to the south. However, for the work presented here, these details are not important.

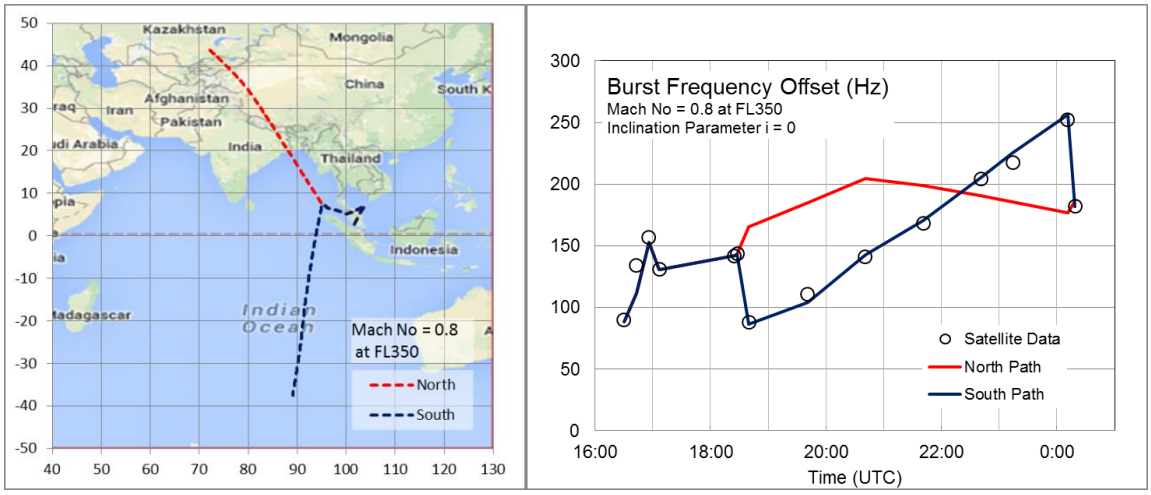

Figure 1. BFO for representative paths to the north and south.

Ways to Alter the Measured Value of BFO

I was interested in exploring ways that the input values to the AES correction algorithm might have been altered so that the modified frequency correction Δf_{AES} caused the BFO for a northern path to look like a southern path. A constraint is that by changing the input parameter, the steering of the high gain antenna (HGA) should not be affected since we know from the class of service that the HGA was in use. With that in mind, I evaluated two schemes for altering the input parameters to the AES correction algorithm:

- Intercepting and replacing navigational parameters sent to the SATCOM over the ARINC 429 bus, as proposed in Jeff Wise's scenario
- Altering values stored in the SATCOM's memory related to the satellite orbit

If one or more navigational parameters were intercepted and replaced, it would require a processor to be placed inline before the ARINC 429 bus input to the SATCOM. This processor would take the actual value of position, speed, and track and replace one or more values such that it would cause the SATCOM to apply a frequency correction Δf_{AES} in a way that the residual Doppler shift L is consistent with a southern path, even though the actual uplink Doppler shift Δf_{up} and the frequency correction Δf_{AES} are individually different for southern and northern paths. (Only the residual L can be derived from the measured BFO if there is no knowledge of the path or the frequency correction.)

Because of the constraint of continuous antenna steering, it would be difficult to change the position or track fed to the SATCOM without the satellite falling outside of the lobe of the HGA's antenna pattern. What parameter is left is the groundspeed. In fact, I have found that by changing the groundspeed used in the AES correction algorithm, the BFO can be varied, and so a northern path could be disguised as a southern path. However, this scheme suffers from two shortcomings:

- The sensitivity of the BFO to speed is low at 19:41 UTC, requiring that the actual speed of the aircraft be replaced with one unrealistically high. This is because the plane was flying tangentially to the ping arc around this time and the satellite was at its point of peak declination.

- The complexity in designing the hardware the software, getting it on the plane, and installing it in-flight is high.

I next considered the possibility that one or more of the values related to the satellite's orbit stored in the SATCOM's memory were altered. This scheme is inherently simpler than the previous one because values do not need to be continuously changed based on the aircraft's position and speed and so there would be no need for an inline processor. However, there would have to be a method for accessing the SATCOM's memory and altering one or more stored values.

If the model used to describe the satellite's orbit was the geostationary representation described by Chris Ashton and others, there is only one parameter needed to specify the position of the satellite: its longitude, which is 64.5° for satellite I3F1. However, the constraint that steering of the HGA remains operational severely limits the possible range for the modified value of the longitude. Independent of this constraint, I could not find a modified value of longitude that would produce a BFO signature for a northern path that is consistent with the measured values.

Near Geostationary Model for the Satellite in the AES Correction Algorithm

It occurred to me that independent of what was presented by Chris Ashton and others, the AES correction algorithm might use a more accurate model for the satellite orbit than the geostationary representation we have assumed. I considered a model that is incrementally higher in complexity, corresponding to a satellite with a circular, but "inclined orbit". We know that the inclination of I3F1 was about 1.65° at the time of the disappearance of MH370, and this inclination produces the discrimination in BFO between northern and southern paths. A satellite with a small value of inclination is sometimes referred to as "near geostationary". I was very interested to see what effect a change in the inclination parameters would have on the predicted BFO values.

A satellite orbit that is near geostationary can be described by three parameters: longitude, inclination angle, and time of the ascending node, which is the time of the day when the satellite passes over the equator from south to north. With the observation that antenna steering was not disturbed, I held the longitude of the orbit constant and studied what would occur if the AES correction algorithm had the capability to model satellite inclination and the two inclination parameters (inclination angle and ascending node) where changed. In particular, I studied whether there were inclination parameters that would cause the BFO values for northern paths to look like southern paths.

The excellent results are shown in Figure 2 for northern paths between $M = 0.65$ and $M = 0.85$. For all speeds, the modified value for the inclination angle parameter used in the AES correction algorithm was set to 3.3°, or about two times the true inclination of 1.65°, and the time of the ascension node was set to 14:12 UTC. These modified values were used for the BFO calculations between 18:40 and 00:11 UTC, inclusive. For the log-on at 00:19, the value of inclination angle was again set to zero. The reason for this reset to zero is described in the next section.

The RMS error in BFO for any of the northern paths studied is less than 3.8 Hz. The error in the calculated BFO at 00:19, which is based on the assumption that MH370 was on the ground and stationary, is about 5 Hz. This lends support to the theory that MH370 successfully landed.

In order to further investigate whether a modification of the AES frequency correction was the mechanism to disguise a path to the north, I had to discover whether a near geostationary model for

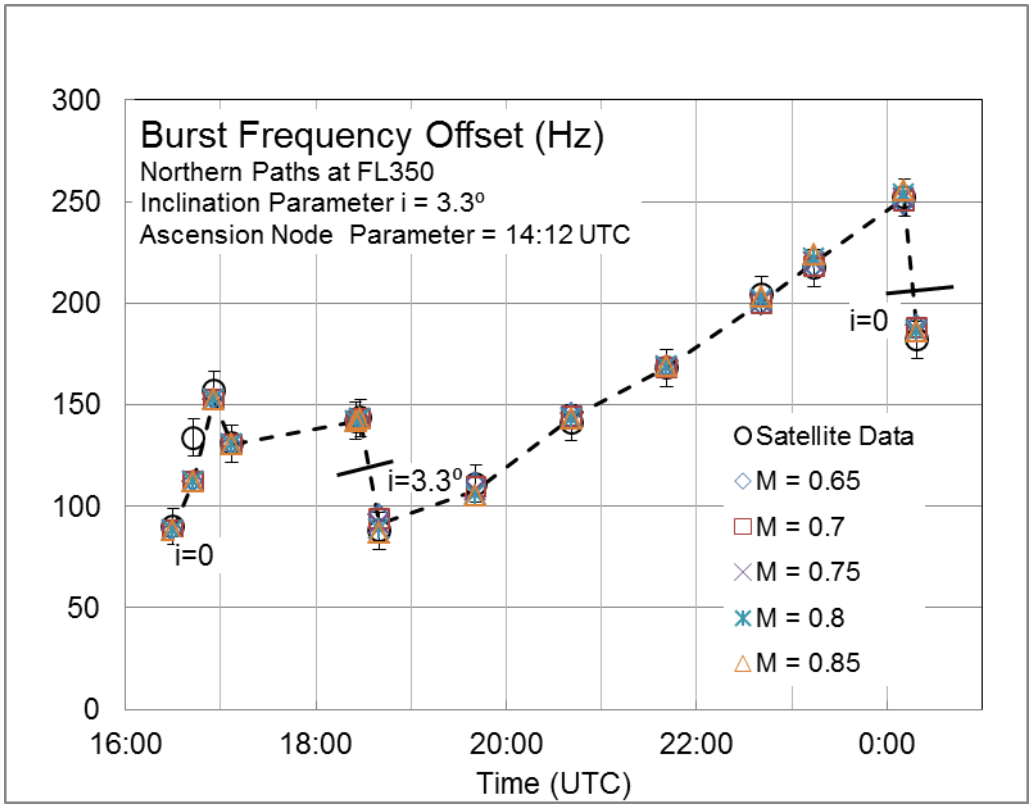

Figure 2. BFO for northern paths with altered inclination parameters.

the satellite was indeed implemented in the AES correction algorithm. If it was not implemented, a much more elaborate hack involving reloading new firmware into the SATCOM would be required to disguise the northern path. A modification of input parameters to the AES correction algorithm is much simpler than a modification of the algorithm itself.

Confirmation of the Near Geostationary Model in the AES Correction Algorithm

Readers might recall that ICAO requires that the residual frequency offset *L* (what is referred to as AES AFC error in the ICAO specification) is limited to +/- 100 Hz. For a satellite that is following a near geostationary orbit, its speed (relative to the earth) will be at a maximum value at the time of the ascending node, and the speed will be proportional to the inclination of the orbit. Therefore, if the AES correction does not account for satellite inclination, the offset *L* grows with inclination. Normally, a satellite's inclination is kept small by applying orbit maneuvers for "station-keeping". However, there is a finite amount of thrust propellant onboard a satellite that is available for station-keeping. As a result, as a satellite nears the end of its life and the supply of propellant is depleted, the satellite's inclination is allowed to grow. Such is the case for I3F1 which at the time of the disappearance of MH370 had an inclination of about 1.65°.

For satellite I3F3, which is a sister satellite to I3F1, I was able to find the Technical Description for the application for its FCC license. As part of that application, it specifies that the satellite will be operated with an inclination of less than 2.7°. The problem is that if the satellite's inclination does grow with time to the maximum allowed value of 2.7° AND the AES correction algorithm does not

compensate for satellite inclination, then the residual offset *L* just due to the inclination could be greater than the allowable limit of 100 Hz, depending on the position of the aircraft relative to the satellite. Since it is unlikely that Honeywell would produce a SATCOM that does not meet the ICAO specification, I reasoned that the actual model implemented in the AES correction algorithm probably did allow for a near geostationary satellite orbit, i.e., it compensates for the effect of inclination of the satellite's orbit.

Based on these suspicions, I queried a knowledgeable individual who is close to the investigation. My suspicions were confirmed—I received a statement that Honeywell's MCS-6000 SATCOM "*does indeed have a capability to correct its Doppler for satellite inclination: this is controlled by a table broadcast to all terminals, and this table was broadcast in no inclination for the 3F1 satellite at the time of the incident. Hence AES was not compensating for satellite Doppler.*"

So there was the confirmation that the capability existed for the SATCOM of MH370 to compensate for inclination, but there was also the belief by some close to the official investigation that this capability was not used, and certainly not used with modified inclination parameters as I proposed.

Each AES maintains a System Table stored in non-volatile memory in the Satellite Data Unit (SDU). The System Table contains the data required for an AES to interface with the Inmarsat network, including channel frequencies, channel data rates, satellite locations, and the GESs associated with each satellite. As part of the log-on process, an AES updates its System Table based on the values that are broadcast by the GES. The inclination and time of ascending node are broadcast as part of the System Table. But since the broadcast value was zero, there is the belief that no compensation in frequency was made for the satellite inclination.

It occurred to me that since the values of the System Table are overwritten each time the SATCOM logs on, if a deliberate change was made to the inclination parameters, then that change would have occurred AFTER the log-on that initiated at 18:25 UTC. In fact, the apparent turn to the south is predicted to have occurred between 18:28 and 18:40 UTC, as suggested by the BFO. What has been interpreted as a turn to the south could have been an indication that the inclination parameters were deliberately altered. If so, the AES was applying a frequency correction, but the correction was based on value of inclination that was about twice the actual inclination in order to disguise the northern route. At the log-on at 00:19, the inclination parameters stored in the System Table of the SDU would again be overwritten to zero. Hence, the BFO value at 00:19 should be interpreted assuming the inclination parameters were zero, as was presented in the previous section.

How the Inclination Parameters Might Have Been Altered

If the inclination parameters were altered in the SATCOM, it is hard to imagine a scenario in which that modification was not deliberate. More specifically, it would indicate there was unauthorized access to the non-volatile memory of the SDU.

Ruben Santamarta, a security consultant associated with IOActive, studied the potential for the malicious attack of SATCOMs. In a landmark white paper entitled A Wakeup Call for SATCOM Security, he studied the vulnerabilities of SATCOMs offered across many sectors, including maritime, land communications, industrial control, civil aviation, and military, and presented his results at Blackhat 2014. His method of study was to obtain the firmware for the SDU and then

reverse engineer (disassemble) the code using an Interactive Disassembler (IDA). Santamarta found that every SATCOM he studied had vulnerabilities that could be exploited by hackers.

To be clear, Santamarta did not study the Honeywell Thales MCS series of SATCOMs, although the Cobham Aviator 700, which offers Classic Aero service like the MCS-6000 on MH370, was part of the study. For the Aviator 700, he found backdoors, a weak password reset, insecure protocols, and hardcoded credentials. He discovered a backdoor that could be entered through the Multi-Controller Display Unit (MCDU) in the cockpit that would allow parameters to be changed and the SDU to be rebooted. A similar method might have been used to alter the inclination parameters stored in the non-volatile memory in the SDU of MH370.

The prospect of BFO spoofing was considered by Gerry Soejatman, an aviation expert in Jakarta, Indonesia. I collaborated with Gerry in analyzing the data from the Indonesia Air QZ8501 incident, and I found him to be knowledgeable and honest. On his blog, Gerry writes:

> *In 2011, I led the Aerospace and Defence Solutions department at one of the local Inmarsat resellers here in Indonesia. I told him [Jeff Wise] that back then I have heard rumours of 2 Indonesian guys who have managed to remote spoof the Doppler while they worked for Inmarsat during integrity testings of the Inmarsat 3 system. And then in several defence related meetings in 2011, I was also told that the other guys who can spoof the Doppler (remote or through the satcom terminal) are Israelis (using Russian immigrant engineers), the Chinese (using the Israeli expertise) and the Russians too, but obviously my sources didn't want to go into details. The other interesting thing is that the Israelis do have their own set of satcom engineers dealing with "new innovations" for Inmarsat satcom, through one of the Inmarsat Distribution Partners, so, nothing surprising there if anyone can spoof the BFO… all you need to do is spoof the Doppler.*

In this blog post, Gerry described the technical details of a scenario similar to that of Jeff Wise, in which the inertial data fed to the SDU is replaced by an inline processor feeding altered inertial data, except in his scenario, the tampering was performed by access to the SATCOM in the cabin instead of in the Electronics/Equipment (E/E) bay as was suggested by Jeff. For reasons I outlined above, I don't think this is as likely as changing the value of the inclination parameters stored in the SDU. However, the fact that others had considered ways to spoof the BFO is very relevant.

Conclusion

This work follows the previous work in which possible northern paths for MH370 was studied by constraining the paths based on the BTO data and the available fuel. The present works extends the past work by offering an explanation of how the plane might have flown north yet exhibit the BFO signature of a southern path in order to disguise the true path of the plane. To do this, parameters related to the description of the satellite's orbit and stored in the memory of the SATCOM would need to be changed via a backdoor or some other means of unauthorized access to the SATCOM's memory. Other SATCOMs have been shown to possess vulnerabilities of this kind, and spoofing of the BFO may have been considered before theMH370 incident. This work challenges the conventional belief that MH370 flew south to the SIO.