MH370 data review, June 8, 2014

Since the Inmarsat ping data was released almost two weeks ago, I like many others have spent a good deal of time trying to discern what the data tells us. Particular thanks are due to Duncan Steel, Victor Iannello, Mike Exner, Don Thompson, Bill Holland and Brian Anderson, who've spent days and weeks performing numerous complex calculations and analysis of satellite and other data, much of which I've relied on in the write-up below. They have made considerable efforts to verify Inmarsat's calculations, and it looks like this is feasible for the BTO (burst time offset) information, which give the distance from the satellite, but much harder for the BFO (burst frequency offset) data, because the terminal on the plane pre-corrects the signal it transmits to remove most (but not all) of this frequency offset. In order to analyze the BFO data from MH370 (as shown in Figure 1 below), and understand issues like frequency drift in the transmissions, Inmarsat had to cross-check it against the equivalent data from other 777s flying at the same time, and of course none of the information from other planes has been made available. It remains unclear how accurate Inmarsat's model is, and the frequency pre-correction makes it difficult if not impossible to derive much useful information about the exact heading or speed of the plane when the pings were sent.

Figure 1: MH370 measured data against predicted tracks [Source: Inmarsat, March 25, 2014]

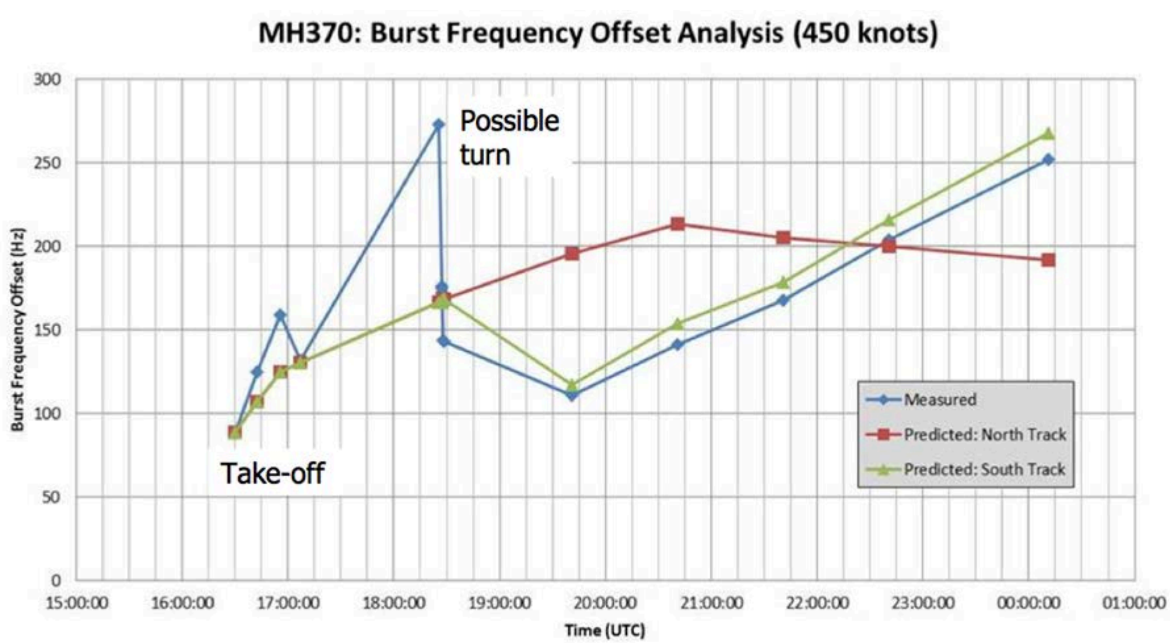

TMF

However, the raw ping data explains why (assuming that Inmarsat's BFO model accurately reflects the real world BFO data) Inmarsat could show the plane went south rather than north: the report notes that "the terminal type used on MH370 assumes a stationary satellite at a fixed orbital position" (i.e. above the equator), when as we know, the satellite was moving north and south in its orbit. This means that the Doppler shift due to the satellite's motion is not removed by the plane's terminal when it pre-corrects the frequency of transmission, and thus there is an asymmetry between the BFO which would be produced by a plane flying north of the equator (and thus situated to the north of the satellite) compared to one flying south on a similar track. In addition, the satellite position used in the plane's pre-correction algorithm is different from the actual position, and thus the frequency pre-correction that the terminal makes in respect of the plane's own speed relative to the satellite is not perfect, indeed it is worst when the satellite is furthest north in its orbit. This explains the confusion in Ari Schulman's The Atlantic article, which suggested incorrectly that there should be no difference in the residual frequency offset between the north and south tracks when the satellite was stationary at the extreme northern point of its orbit.

Because elements of the BFO model remain unclear, for the moment we are therefore left with the BTO (and the derived distance from the satellite, aka the ping rings) as the principal source of useful data about the potential position of the plane at the various ping times. Although we know the distance of the plane from the satellite, the key question is what is the speed of the plane, since the faster the speed, the further south the plane will have traveled (the distance to travel between crossing each ping arc is greater if the plane was traveling south, compared to traveling southeast). Indeed, as shown in Figure 2 below, the primary search area was progressively shifted further and further to the northeast over time. This was based on a new assumption that the plane was flying slower than previously assumed (at between 323 and 350 knots, as shown in Figure 3), and the final "high priority" red search area was associated with the most northeasterly (slowest) of these tracks, apparently due to a re-evaluation of how much fuel was available and the length of time that the plane could therefore remain in the air.

Figure 2: Areas searched in the southern ocean [Source: AMSA, April 11, 2014]

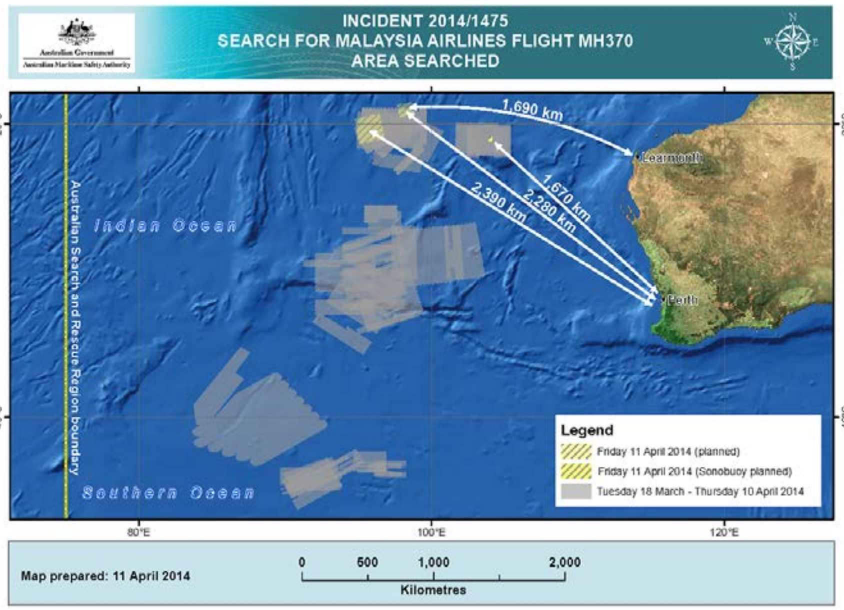

Figure 3: Most recent search areas [Source: Malaysian government statement, May 1, 2014]

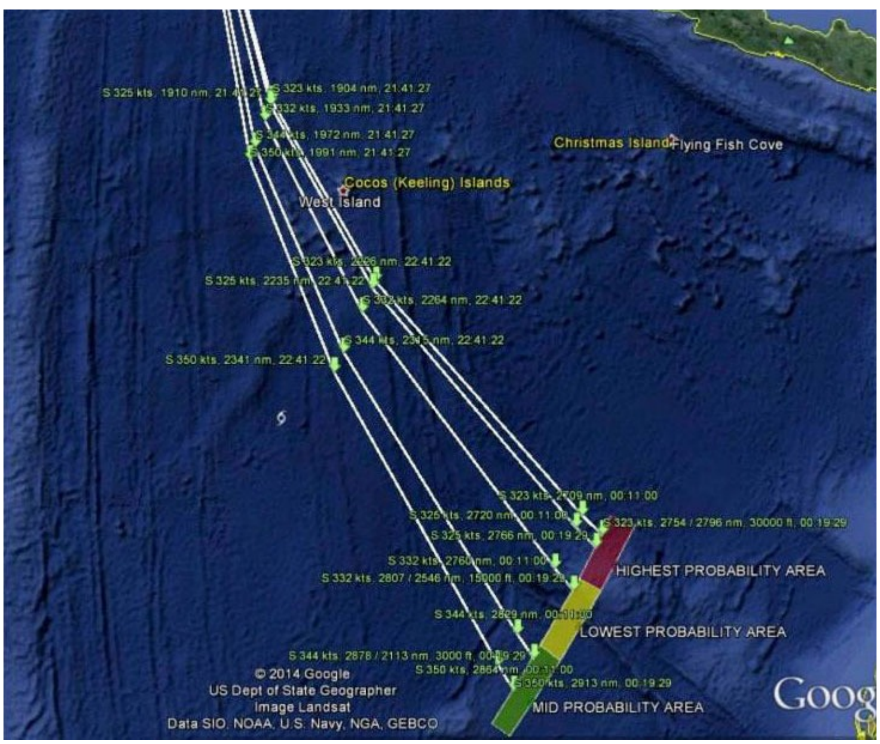

It is clear that the precise tracks depicted in Figure 3 above cannot possibly be correct, because the track shows a turn at each waypoint and none of the systems (or people) onboard the plane could have had knowledge of the exact time at which these pings were sent. The plotted track may be an approximation, but a curved path (as opposed to the great circle or magnetic heading on which a plane would normally fly) is equally implausible, and the speed is of course subject to possible effects from headwinds or tailwinds. Moreover, the high priority search area has now been searched fairly comprehensively, without any sign of the plane being found.

If we take a step back, then there are three unknown variables to consider:

1) where did the plane start on its route to the Southern Ocean (i.e. when it made a turn to the south near Sumatra)?
2) what speed did it fly at (and was this even constant)?
3) what heading did the plane fly on and did it make any turns?

The process apparently followed by the search team, and documented in Figure 3 above and Figure 4 below, appears to be:

1) assume a start point based on a turn at 18:27, just after the time when the ping data shows the satellite terminal (and perhaps other navigation equipment) was rebooted,
2) consider what speeds would give fuel exhaustion at 00:19 UTC when the last “partial ping” was received,
3) plot a straight line path between each ping arc with a distance equal to the assumed speed.

Figure 4: Assumed turn point and tracks [Source: Malaysian government statement, May 1, 2014]

Fundamentally, this process appears to be the result of certain independent pieces of engineering data being combined in an apparently logical way, but without any consideration of what flight paths could actually be flown in practice. As noted above, assumption 3 is of course wrong, because the plane simply could **not** have turned at each ping arc time. If the plane was going to make turns, then these would be expected to happen at waypoints (or other locations that could be programmed into the flight computer if the plane was being flown on autopilot) and it appears that this has simply not been taken into account. Some have suggested that following the M641 airway through various waypoints towards Perth could be a possible track, but this would not match the ping ring arcs unless the speed of the plane changed significantly over time. Thus it remains a mystery where any "real world" turning points might have been.

Secondly, with respect to assumption 1, although the sharp change in BFO at around 18:25 UTC was originally identified as being associated with a "possible turn" (as shown in Figure 1), it now appears from the released ping data that this fluctuation was a result of the terminal being rebooted (the power cycling off and on again) and thus it is unclear exactly when the turn to the south actually took place.

Broadly speaking, the further the plane flew west beyond 18:25 UTC, the more its path would ultimately have headed back to the southeast, to maintain the same distance of closest approach to the satellite (which is defined by the ping arcs).

Finally, no data has been made available on the fuel exhaustion calculations which are critical to assumption 2. If the plane had flown in a straight line, then the speed would certainly be much faster than assumed in determining the search areas. In fact, as shown in Figure 5, we can actually estimate the ground speed of the plane from the ping arc data assuming it flew in a straight line across the 19:41 and 20:41 ping arcs, by estimating the point of closest approach to the satellite and using high school geometry with the known adjacent ping times and ranges as shown in the diagram below (thanks to Brian Anderson for suggesting this approach). Note however, that the actual *airspeed* would be different due to headwinds and tailwinds which vary by location and altitude and the plane would be expected to maintain a fixed airspeed not a fixed ground speed. It is unclear if the impact of winds was taken into account in defining the search areas shown in Figure 3.

Figure 5: Geometry for estimating straight line speed [Source: Brian Anderson]

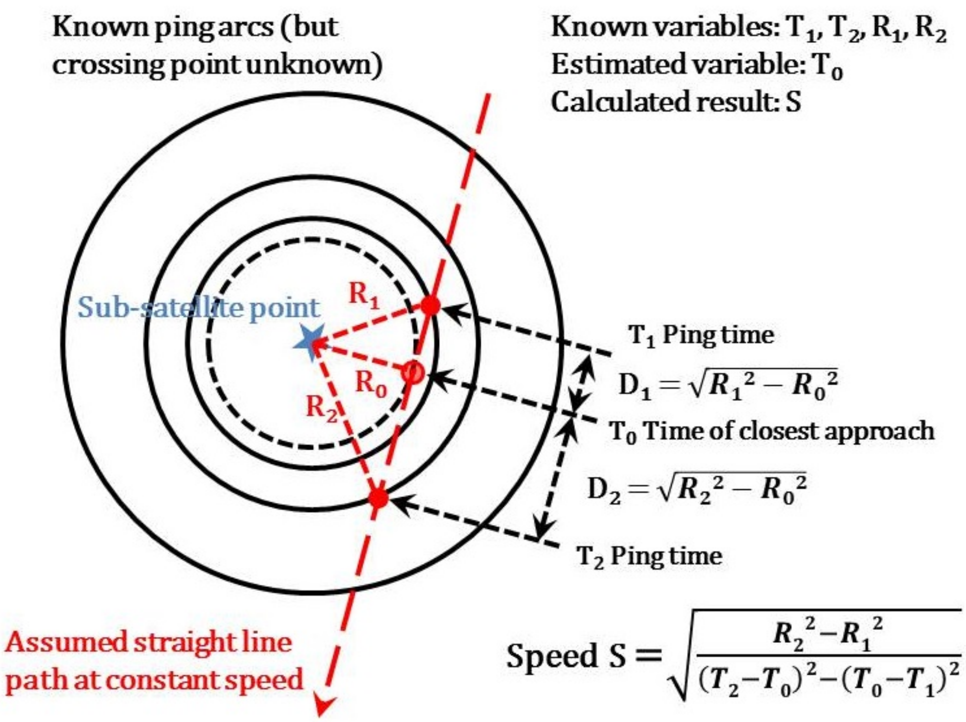

Based on the estimated ping range of 1755 nmi at 19:41 and 1794 nmi at 20:41 and a closest approach at roughly 19:54 (by curve fitting the ping ranges), this calculation gives a much higher speed of about 493 knots (though of course this is using planar geometry and the answer is likely to be around 10 knots lower using spherical geometry (plus a true earth model) and when allowances are made for the motion of the satellite during these 60 minutes). Indeed, as shown in Figure 6 below, there is also a fair amount of error in the speed, depending on exactly when the point of closest approach was reached, and somewhat lower speeds may therefore be possible. However, for a much slower speed (below 350 knots), then the plane would have had to reach its closest point to the satellite before 19:41 and already be heading away from the satellite at that time.

Figure 6: Variation in straight line speed with time of closest approach [Source: Brian Anderson]

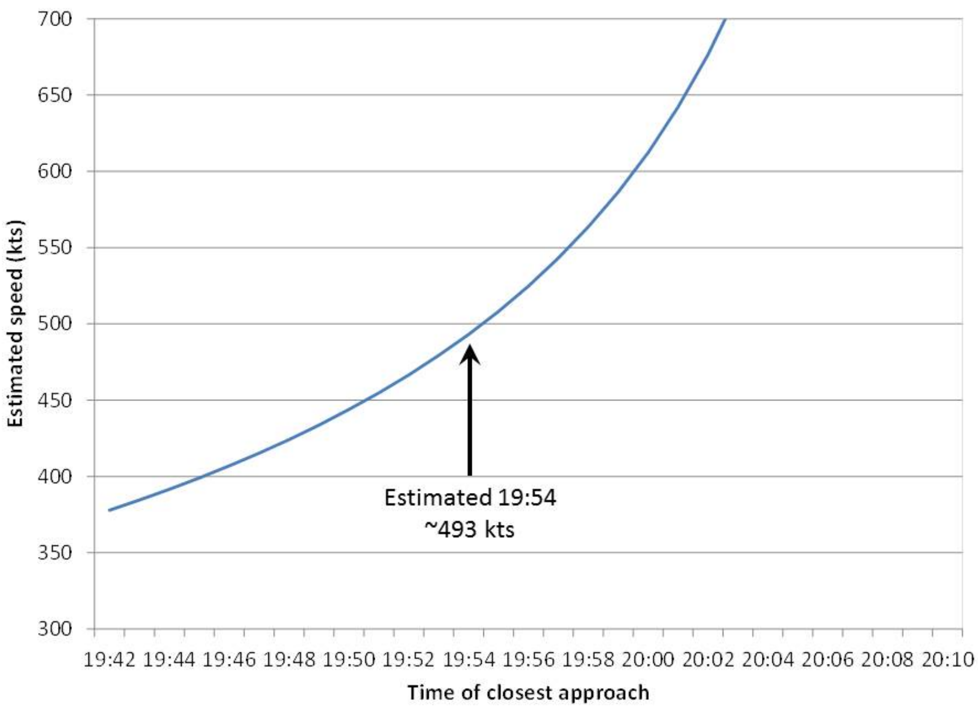

It also appears that the low speed solution has not been checked against Inmarsat's own BFO model, which is based on a 450 knot speed that is relatively close to a straight line track, as shown below. Although many uncertainties remain as to the exact formulas used by Inmarsat, initial attempts to reverse engineer Inmarsat's model by Victor Iannello suggest that the 325 knot track used to define the "high priority" search area is a less good fit to the actual BFO data than Inmarsat's own 450 knot track shown in Figure 7. Depending on the accuracy of Inmarsat's BFO model, it would certainly be expected to be useful in narrowing the search area along the southern arc to some degree. As a result, it is very surprising that the BFO model has been used by the investigation team to prove the impossibility of a northern track, but apparently has not been utilized as a cross-check for the 325 knot hypothesis.

Figure 7: Inmarsat example 450 knot track used in BFO model [Source: Inmarsat, March 25, 2014]

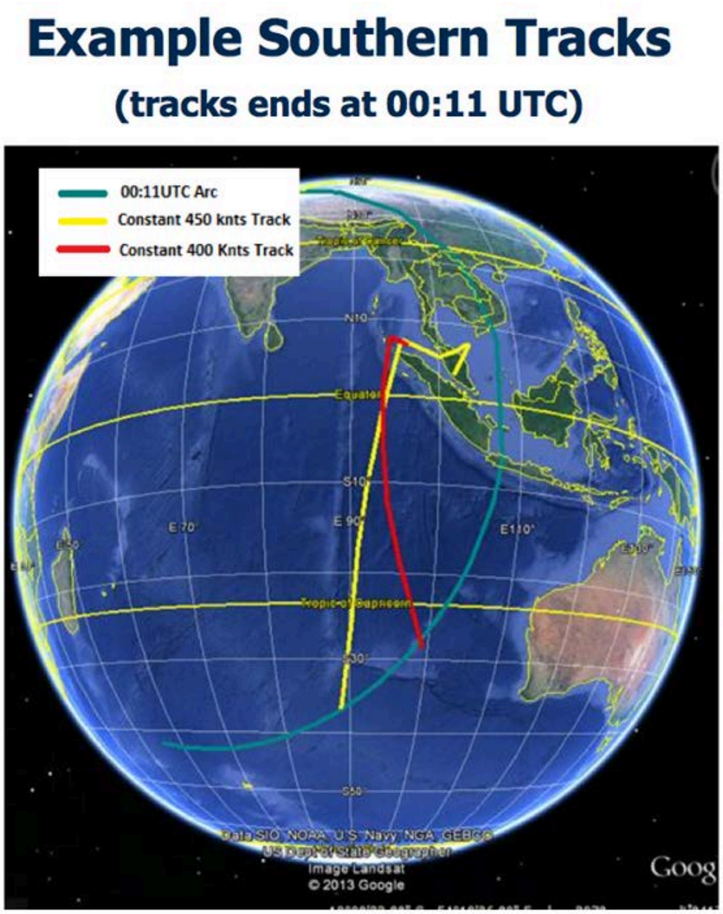

So it seems that the key next step is to better understand how certain it is that fuel exhaustion drives a lower speed solution, since it's unclear why a choice would be made to fly at this speed. If this hypothesis holds up, then it would make sense to look for potential turning points rather than simply assuming a turn at each ping time, and come up with a convincing rationale why either a pilot or autopilot might make a turn at that point. On the other hand, if it is possible that the plane might have been traveling at a rather higher cruise speed, then without any evidence why a turn should have been made, the default solution would be to find the straight line speed that fits with the ping arcs and extrapolate that speed (adjusting for headwinds and tailwinds) through the remainder of the flight. That would inevitably lead to a search area far to the southwest of the most recent search zone.

Figure 8: MH370 underwater search planning map [Source: ATSB, June 4, 2014]

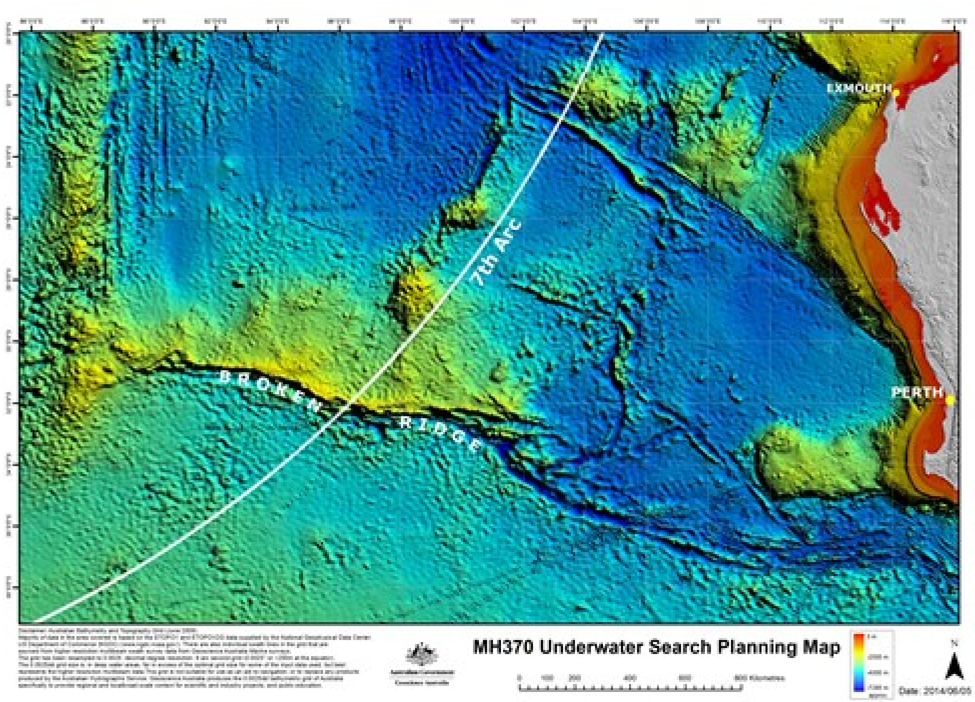

It appears that efforts are now being made to evaluate a much wider range of possible tracks, in order to give a statistical "best fit" to the known satellite data. A new (June 4) release from the ATSB seems to indicate that the full range of the 00:19 UTC "7th arc" from 20˚S to 39˚S is now under potential consideration, as shown in Figure 8 above, with the highest priority 650km length to be identified for searching (corresponding to 60,000 sq km of ocean floor) in the near future. This suggests that all three assumptions that were used to define the most recent search areas are now subject to question, including the hypothesis of a low speed track driven by fuel exhaustion, which may result in a dramatic widening of the potential search area. The total area of ocean floor that is expected to be searched is dictated by the available budget, and appears to correspond to the combined size of the most recent red, yellow and green search areas. Thus any attempt to identify new areas to be searched will require some part of the previously identified search areas to be dropped. The total length of the ping arc that may need to be considered as viable potential end points could now exceed 2000km in length, and thus less than one third of this area would be searched. Unfortunately, while some end points may be identified as more probable than others, it is clear that an increased level of uncertainty about the location at which the plane reached the last ping arc is likely to result in a lower probability of success in the search.