Operator and User Perspective of Fractionated AIS Satellite Systems

Andreas Nordmo Skauen, Øystein Helleren, Øystein Olsen, Richard Olsen Norwegian Defence Research Establishment (FFI) PO Box 25, N-2027 Kjeller; +47 6380 7328 Andreas-Nordmo.Skauen@ffi.no

ABSTRACT

In 2013 AISSat-2 will join AISSat-1 in providing Norwegian authorities, and their partners, with the extended maritime situational awareness that satellite based AIS systems provides. Since neither satellite has propulsion and both will end up in very similar orbits, there will be times the satellites are covering the same area at the same time, and times where the satellites are separated up to half an orbit. Such a configuration gives rise to some interesting opportunities for the users and challenges for the operators. This paper investigates if an operator can use the variability to the users' advantage in fulfilling the mission objectives – extending and improving the maritime situational awareness.

Intuitively it stands to reason that the user would prefer the satellites to be spaced as far apart as possible, to minimize the mean time between updated information about vessel traffic in an area of interest. However, it is well known that in many areas of interest it is unlikely that a satellite AIS system detects every ship in the area in a single pass. With multiple systems covering the same area at the same time, the probability of vessel detection increases, and this is shown to improve value added products such as fused data products and verification of the AIS reported vessel position.

1

INTRODUCTION

Since the launch of AISSat-1¹ in 2010 Automatic Identification System (AIS) satellite systems have gone from being an experimental service to being a fully operational capability that authorities responsible for maritime safety and security worldwide have become more or less reliant on. The introduction of satellite AIS systems have provided an unprecedented ability to monitor vessel traffic on a global scale² and the user feedback indicates that now that the light has been switched on, no one wants to go back into the dark.

In order to maintain the satellite AIS service to Norwegian authorities a copy of AISSat-1, namely AISSat-2, has been built to serve as an in-orbit spare for AISSat-1, as well as to provide more vessel traffic data,. Both satellites will operate simultaneously, but since neither satellite has on-board propulsion, the orbits will drift independently from each other. The issues and challenges associated with this drift will be explored in this paper.

For completeness, a high level introduction to the AIS system is appropriate: AIS is a ship-to-ship and ship-to-shore reporting system intended to increase the safety of life at sea and to improve control and monitoring of maritime traffic. AIS equipped ships broadcast their identity, position, speed, heading, cargo, destination, etc. to vessels and shore stations within range of the

VHF transmission. On ground level AIS stations on shore can typically receive AIS messages at distances about 40-50 nautical miles off-shore. With a low noise, highly sensitive receiver capable of handling Doppler shifts up to 4000 Hz an AIS receiver on a satellite extends the range to global AIS message reception.

In addition to increasing the coverage area, a satellite viewpoint enables verification of the reported position in the AIS message, given that the payload provides a frequency shift estimate with each received AIS message. The uncertainty of the verification can be greatly reduced by having multiple satellites. The verification technique, and results from on-orbit testing, is discussed briefly later in the paper.

AISSAT MISSION ARCHITECTURE

The overall AISSat (comprising AISSat-1 and AISSat-2) mission architecture and main ground site locations are shown in Figure 1. The AISSat satellites receive AIS messages from vessels at sea globally and forward the messages to the Svalbard ground station (78°N). In principle, the Svalbard location permits contact with an AISSat satellite in all of the satellites' 15 daily orbits. In practice, using only one antenna for downlink and uplink respectively, only one satellite can be contacted at a time in passes where both AISSat-1 and AISSat-2 are within contact range simultaneously. The AIS messages are forwarded to the Norwegian Coastal

Administration (NCA) in Haugesund, and finally to the Mission Control Centre (MCC) at FFI, Kjeller. Commands for tasking and operation of AISSat satellites are sent in the opposite direction. The NCA is responsible for distributing the AIS data to the users.




Figure 1: AISSat-1 mission architecture.

The satellites were built, tested and prepared for flight by the Space Flight Laboratory at the University of Toronto Institute for Aerospace Studies (UTIAS/SFL), and are based on the 20 cm cube Generic Nano-satellite Bus (GNB). UTIAS/SFL is also responsible for launch arrangements. AISSat-1 was launched by the Indian Polar Satellite Launch Vehicle (PSLV) from southern India July 2010, while AISSat-2 is slated for a launch with the Soyuz launch vehicle from Baikonur in Kazakhstan in the second half of 2013. Both launches were as secondary "piggyback" satellites. The payload, identical for the satellites, is an AIS receiver developed and manufactured by Kongsberg Seatex AS, Trondheim Norway. The AIS payload is a software defined radio that supports, and has been subject to, in-orbit upgrade of the payload algorithms, enabling higher performance. The Norwegian Defence Research Establishment (FFI), Kjeller Norway developed the AISSat mission concept and has been responsible for managing the project, testing and preparing the AIS payload for flight, and analyzing the data.

With the Soyuz launch, AISSat-2 is likely to end up in an orbit with an altitude of 640 km, 98.4° inclination and a 09:00 hour local time on descending node (LTDN). AISSat-1 was launched into a 635 km altitude, 98.1° inclination sun-synchronous orbit with LTDN equal 09:30 hours, but has had a drift of roughly 8.5° per year. In Q4 2013 AISSat-1 is expected to be at an altitude of 628 km and inclination of 98.0° with LTDN equal 11:20 hours.

These different orbits have been simulated to investigate the impact they will have on operations and the satellite AIS service provided to the user.

ORBIT SIMULATIONS

The orbit parameters of AISSat-1 and AISSat-2 were put into the software package Systems Tool Kit 10 and simulated 6 months forward in time. Because no drift was simulated for AISSat-1 and the actual orbit of AISSat-2 is unknown, the results herein must be considered guidelines only. The absolute times, duration and frequency of events will differ for real operations. The results are however useful for operators and users in preparing for and modifying the operations concept.

Since the AIS payload onboard both AISSat-1 and AISSat-2 employs a monopole antenna with an omnidirectional field of view to the horizon, the high level effect on the average number of accesses to an area by adding AISSat-2 is illustrated in Figure 2: The average number of accesses per day doubles, which for the primary area of interest means an increase from between 10-15 by AISSat-1 alone to 20-30 with the addition of AISSat-2. Since the satellites are not in the same orbit and have no orbit keeping capability the relative access times to a specific area will vary with time.

Figure 3 shows the frequency of time between accesses to the Svalbard ground station for the AISSat constellation. The plot shows that the operators, given the ground station setup with only one antenna for downlink and uplink respectively, must be prepared to choose which satellite to contact for more than 11% of all contacts (22% of a single satellite contacts). This will be a very different scenario from the highly automated AISSat-1 operations, which in nominal operations only require operator intervention once a week. Even though anomalies are handled during normal office hours only by scientists with other responsibilities as well, an overall system uptime greater than 95% was achieved in 2012 (satellite uptime was 97%). The current operations scheme requires the antenna access times to be calculated for one week in advance. With the addition of AISSat-2, using the same operations scheme, in the event of simultaneous access one must then choose which satellite to contact a week in advance. If an anomaly occurs on the satellite one does not contact, the anomaly will naturally go undetected until contact is made, and may not be corrected until the next scheduled contact time after the anomaly detection. The periodicity and duration of phases with simultaneous AISSat-1 and AISSat-2 access to the ground station is shown in Figure 4. Typically the conflicts last nearly 7 days, with a 20 day gap between the end of a conflict phase and beginning of the next.




Figure 2: Average number of accesses per day globally (left) and in the area of interest (right) for AISSat-1 (top) and both AISSats combined (bottom). On the left hand side plots the colours run from 2 – 20+ in steps of 2, but on the right hand side the colours run from 7 to 28+ in steps of 3.




Figure 3: Distribution of the normalised frequency of revisit time to the Svalbard ground station for any of the AISSat satellites. The first bin indicates that both satellites are within ground station range within 0 – 15 minutes of each other. The bins increase in steps of 15 minutes up to 1 hour 30 min – 1 hour 45 min as a maximum.




Figure 4: Access times and duration, in seconds, for both AISSat satellites contact with Svalbard.

Of importance to the users is that they can expect more frequent updates, with nearly 50% of contacts with the ground station happening within 30 – 75 minutes after the previous contact, shown in Figure 3. This significant improvement in update rate from the single AISSat-1 satellite will greatly improve the recognized maritime picture for the users. Furthermore, the system reliability is expected to increase with the addition of AISSat-2. If a satellite experiences an anomaly, the users may still receive updated data within a short timeframe from the other satellite. The exception will be in cases where the anomaly happens on both satellites or during a ground station conflict on the primary satellite. Assuming that satellite anomalies are random, as all evidence suggests, and that AISSat-2

will achieve the same satellite uptime of 97% as AISSat-1, the probability of both satellites experiencing an anomaly at the same time is very low at about 0.1%. Assuming the primary satellite for ground station contact alternates during conflict phases, the probability that the primary satellite has experienced an anomaly during a conflict phase is estimated to be 0.33%. Given that ground segment issues caused the system uptime to be less than that of the satellite at 95% in 2012, these estimated probabilities can be considered a lower bound.

Selecting the primary satellite during a contact conflict phase is a new issue for the operators. Besides simply alternating which satellite is the primary for a pass, several other options exists: One may either always select one satellite as the primary, or one may select the satellite that first is within contact range. One could also select the satellite with the longest contact time, or select the satellite that maximizes the AIS detection probability from a viewing geometry standpoint. Finally, one may even divide the contact time into two parts, spending the first half communicating with one satellite and the second half with the other.

The effect of adding AISSat-2 on the frequency of time between accesses for other areas of the world is illustrated in Figure 5 and Figure 6 for a point at 40°N and at 0°N respectively. While a single AISSat satellite may access a point at any latitude in subsequent passes because of the large field of view and orbit geometry, the revisit time for non subsequent passes at lower latitudes increases significantly and with increasing latitude. At the equator for example the typical revisit time for non subsequent passes is 10.5 hours, maximum 12 hours. The addition of AISSat-2 reduces the maximum revisit time to just over 9.5 hours while the typical revisit time is more spread out and reduced.




Figure 5: Distribution of normalised frequency of revisit time to a point at $40^{\circ}N$ for any of the AISSat satellites. The first bin indicates that both satellites are within view of the same point at $40^{\circ}N$ within 0-30 minutes of each other, while the second bin is for access within 30 min -1 hour. The third bin indicates access within 1-2 hours, while the remaining bins then increase in steps of 2 hours up to 10-12 hours as a maximum.




Figure 6: Distribution of normalised frequency of revisit time to a point at the equator for any of the AISSat satellites. The first bin indicates that both satellites are within view of the same point at the equator within 0 – 30 minutes of each other, while the second bin is for access within 30 min – 1 hour. The third bin indicates access within 1 – 2 hours, while the remaining bins then increase in steps of 2 hours up to 12–14 hours as a maximum.

The gap between the end of a phase with simultaneous AISSat-1 and AISSat-2 access and the beginning of the next for point access at latitudes south of the Svalbard ground station roughly increases by 0.5 days per 10° southward change in latitude. For the examples at 40°N and 0°N the typical gap is 22 and 24 days, though this is not an exact measure. The duration of phases not only decreases with decreasing latitude, but also in number of simultaneous accesses during a phase as shown in Figure 7 and Figure 8. For a user, this means not only

an increase in update rate globally in times when the satellites are spaced apart, but also improved instantaneous vessel detection probability during simultaneous coverage times. Simulation work has shown, and in orbit data has verified, that a satellite AIS system is seldom capable of detecting AIS message from all transponders within its field of view during a particular pass. Having multiple satellites viewing the same area at the same time increases the vessel detection probability, providing the user with an improved product.




Figure 7: Access times and duration, in seconds, for simultaneous AISSat-1 and AISSat-2 contact with a point at 40°N




Figure 8: Access times and duration, in seconds, for simultaneous AISSat-1 and AISSat-2 contact with a point at 0°N

Note that for all figures shown in this section, save Figure 2, a 2° minimum elevation was required between a satellite and a point target. This means that some simultaneous coverage periods only briefly overlap in a small geographic area, evident from some of the short conflicts shown in Figure 4, Figure 7 and Figure 8.

While this section has focused on the intuitive issues and opportunities an operator and user can expect from the fractionated AISSat-1 and AISSat-2 constellation

the next section investigates more advanced techniques and strategies for leveraging the new operations scheme to provide value added products to the users.

VALUE ADDED PRODUCTS

Independent Position Verification

Since AIS is a cooperative system, some users are interested in verifying the broadcasted information. There are multiple examples for instance of vessels reporting to be at impossible locations such as on land, at the south pole, or simply positions outside the field of view of the satellite. Additionally, without position input, the standard reported position in the AIS system is 91°N, 181°E. Using the AIS messages and orbit geometry alone FFI has developed an algorithm to estimate the AIS transponder position in order to locate the transponder and verify the reported position.

The relative motion of the satellite with respect to AIS transponders introduces a Doppler shift in the measured transmission frequency. The magnitude of the frequency shift is determined by the relative location and velocity of the satellite with respect to the AIS transponder, illustrated in Figure 9.




Figure 9: Induced Doppler shift (Hz) over the field of view of AISSat-1

Since FFI specified that the AIS payload should provide a frequency shift estimate along with the AIS message and time of reception, it is possible to estimate or verify the AIS transponder position independent of the reported position.

An example of the algorithm results are shown in Figure 10. In January 2011 a ship was captured by pirates in the Indian Ocean. In the figure, the ship has come down through the Red Sea and out the Gulf of Aden before being captured and taken southwards, then westward towards Somalia. Shortly after the capture, the ship AIS transponder reported positions of 91°N, 181°E effectively disappearing from a map. The green squares in Figure 10 indicate the reported ship position, while the red squares indicates the position FFIs algorithm estimated for the ship. When the ship once

again reported valid positions, the estimate and reported position was in excellent agreement.




Figure 10: Example of position verification and estimation using AISSat-1.

Despite the good results exemplified in Figure 10, the procedure is not without challenges. Fundamentally, in order to achieve good accuracy, the frequency estimate must be precise and the transponder emitted frequency must be stable. Frequency residuals of 20 Hz are equivalent to an accuracy of 10 km in the estimated position. Many transponders do not transmit at the nominal frequencies such that a frequency bias must be estimated for each transponder. Furthermore, as Figure 9 shows, Doppler shifts are symmetrical about the ground track such that additional information to a single pass is required for find a unique solution. Finally, transponders that are only detected during a small part of the pass reduce the achievable accuracy. A typical distribution of the semi-major axis of the uncertainty ellipse for AISSat-1 position estimates is shown in Figure 11.




Figure 11: Distribution of position estimation uncertainty using AISSat-1.

The typical accuracy achieved using AISSat-1 data is in the 16-36 km range. For some vessels the uncertainties are larger, up to 100 km, since Doppler measurements are not well suited to constrain the estimated position perpendicular to the ground track and as previously stated, the Doppler shifts are symmetrical about the

ground track. Having a second AIS satellite platform in a slightly different orbit should drastically reduce the uncertainty estimate, and indeed this has been verified by combining AISSat-1 data with NORAIS data. NORAIS is an AISSat-1 type AIS payload installed on the Columbus module of the International Space Station (ISS). The NORAIS instrument is also operated by FFI on behalf of the European Space Agency. The typical accuracy is reduced to 8-20 km, and the number of larger uncertainties is significantly reduced.




Figure 12: Distribution of position estimation uncertainty using AISSat-1 and NORAIS combined.

Using the independent Doppler shift position verification one can create a database of "trusted" vessels that report a correct position, leaving the users more time to focus on vessels with anomalies. Of course, since the NORAIS instrument is installed on the ISS, the primary area of interest for Norwegian users is not accessible, and users interested in Norwegian areas will not benefit from the two platform accuracy improvement until AISSat-2 is in orbit.

Improved Image Fusion

Another means for position verification is through fusion with independent sources, such as satellite imagery, be it optical or radar. A significant challenge for fusion of independent sources is the difference in data age and the fact that in many areas of the world it is very unlikely that a satellite AIS system detects every ship in the area in a single pass^{3,4}. This makes fusion with imagery suboptimal. One example is the introduction of vessel identification ambiguities because of erroneous position extrapolation due to data time differences. Another example is false alarms that are triggered if vessels are not detected by the satellite AIS system, indicating that the vessel is not reporting on AIS. The end result for the users is suboptimal resource allocation and support further investigation.

Having multiple satellite systems that drift in and out of synchronisation can improve the fusion process in several ways. When the satellite systems cover the same area simultaneously, the probability of vessel detection increases. This reduces or potentially eliminates the

number of undetected vessels and the associated false alarms. Conversely, when the satellite systems are distributed, the time difference between the image product and the AIS data can be minimised. With more frequent AIS data updates, it should also be possible to improve any necessary extrapolation. While having an AIS payload on the same platform as the imaging payload would be beneficial in the future, one is still left with the issues relating to the false alarms caused by undetected vessels at the time of the image in many areas of the world. Having a fractionated satellite AIS constellation instead of pooling all resources in a single co-located AIS and imaging payload satellite may also enable the users to leverage the range of available imaging systems better suited for a particular task. Sometimes the wide swath and all weather capability of a Synthetic Aperture Radar (SAR) product is preferred, while other times a high resolution optical image is desired.

FFI has developed a sensor fusion tool implemented as a multiple hypothesis tracker (MHT) to enable fusion of AIS and imagery data amongst other inputs. The value added from additional sensors is evident comparing Figure 13 and Figure 14 that are both showing the output from the MHT fusion between a satellite SAR image from the Barents Sea and several vessel identification data sensors. Both figures show the same image fusion example, but Figure 14 features additional vessel identification data from AISSat-1. In the figures the red targets identifies vessel detections in the SAR imagery that are unidentified, while the green targets identify vessel detections in the SAR image that have been associated with identifications from the vessel identification data. Only SAR detections are plotted in the figures. If there are additional detections from the vessel identification data they are not plotted in the examples. The vessel identification data in the examples are from the Vessel Monitoring System, Long Range Identification and Tracking and AISSat-1. Of the remaining unidentified targets after fusion, three are likely false alarms, as they are very close to small islands / rocks from comparison with nautical charts, but the final four are likely vessels.




Figure 13: Multi hypothesis tracker output from satellite SAR imagery (dotted rectangle) fusion with Vessel Monitoring System (VMS) and Long Range Identification and Tracking (LRIT) data. Red targets are unidentified SAR vessel detections, green targets are VMS or LRIT identified SAR detections.



Figure 14: Multi hypothesis tracker output from satellite SAR imagery (dotted rectangle) fusion with Vessel Monitoring System (VMS), Long Range Identification and Tracking (LRIT) and AISSat-1 data. Red targets are unidentified SAR vessel detections, green targets are VMS, LRIT or AISSat-1 identified SAR detections.

SUMMARY

This paper has discussed the future challenges and opportunities for the operators and users of the fractionated AISSat AIS satellite system constellation. While the operators will face ground station conflict issues, it is considered that the benefits of a fractionated system outweigh the challenges. Multiple satellites drifting in and out of synchronisation has been shown to enable times of reduced data latency, improved sensor fusion possibilities and position verification and estimation by virtue of the variability.

Given the impressive system uptime achieved with the AISSat-1 satellite and ground segment, the probability of anomalies during ground station conflict times appears low. In the primary area of interest nearly 50% of contacts with the ground station are expected to happen within 30-75 minutes of the previous contact, significantly reducing the previous 95 minute latency. It was also shown that the position estimation uncertainty is expected to reduce from 16-36 km to 8-20 km.

Overall the paper has shown how the operators can use the relative satellite orbit variability to the users' advantage in fulfilling the mission objectives – extending and improving the maritime situational awareness

Acknowledgments

The authors would like to acknowledge important contributions from, and great cooperation with, our partners in the project, in alphabetical order: Kongsberg Seatex, Kongsberg Satellite Services, Norwegian Coastal Administration, Norwegian Space Centre and the Space Flight Laboratory at the University of Toronto Institute for Aerospace Studies.

References

- Narheim, B.T et. al., "AISSat-1 Early Results," 25th Annual AIAA/USU Conference on Small Satellites, 2011.
- 2. Eriksen, T et. al., "Tracking Ship Traffic With Space-Based AIS: Experience Gained in First Months of Operations", Waterside Security Conference paper, 2010.
- 3. Hannevik, T.N et.al., "Ship detection using high resolution satellite imagery and space-based AIS", Waterside Security Conference paper, 2010.
- 4. "Technology Reference and Proof-of-Concept for a Space-based Automatic Identification System for Maritime Security", ESA ESTEC contract 20492/06/NL/JA, RES-270-21