Limitations of Satellite AIS: Time Machine Wanted!

Eric Meger
CEO, Maerospace Corporation
eric.meger@!maerospace.com
3 October 2013

Maritime Domain Awareness is one of the most significant national and international interests worldwide. Tens of thousands of ships carry trillions of dollars of cargo between thousands of ports every year. Terrorism and piracy are in the news every week. Ports and logistics are in turmoil, expanding to accommodate new giant ships. The expansion of the Panama Canal and the opening of shorter routes through the Arctic are changing traditional trade routes. Trouble with cruise ships at sea affects the general public as well as the tourism industry. Wind farms and deep sea drilling rigs are creating new hazards for vessels. According to government of Canada estimates, the worldwide value of illegal, unreported and unregulated (IUU) catches is somewhere between $4 billion and $9 billion per year.

Everyone can see that monitoring and managing maritime traffic on a national and global view has never been more important; yet the tools available to those with responsibility are still in their infancy. Even the latest development, satellite AIS, has limitations that are often unseen.

This paper will survey the current state of MDA systems[1] and shed some light on some of the important but often overlooked challenges that have emerged with the advent of Satellite AIS.

1 State of Maritime Domain Awareness Systems

The job of those responsible for maintaining awareness of the maritime domain, whether for a global mandate or a small port, starts with simply knowing what ships are where in their Area of Interest (AOI). With this information, a user can detect anomalous behavior and take appropriate action (i.e. dispatch support resources, interdict vessels, etc.). Getting the basic information has historically been, and still remains a major challenge.

Today MDA relies mainly on three systems: Radar, VHF radio (voice) and AIS (Automatic Identification System). In addition, LRIT (Long Range Identification and Tracking Systems) have been deployed in many countries. Finally, Satellite AIS has burst onto the scene in the past 3 years. These last two are the only global systems. LRIT, for a variety of reasons has not been as useful as hoped and is today only used for MDA on a limited basis. Satellite AIS, however, is rapidly becoming the standard for wide area MDA.

[1] Due to space limitations, we are focused here on shore-based MDA and will not examine ship-to-ship MDA issues.

When looking beyond the range of coverage of shore AIS, VTS (an integrated system that includes Radar, AIS and VHF) or RADAR (typical ~25 nm), the detection of vessels drops dramatically. Although occasional targets are seen beyond this range, indeed to 70+ miles when atmospheric conditions are perfect, reliable tracking is typically limited. Interviews with maritime authorities around the world indicate that coastal AIS systems typically detect about 76% of all of the ships in their local area of interest. While this figure is surprisingly low, it represents what end users actually perceive rather than detailed engineering analysis. Presumably, this detection rate is much higher close to the AIS tower and decreases with distance to create the perceived average.

Today, VTS and AIS systems are being supplemented with Satellite AIS systems available from several vendors. These systems provide true global coverage but with the trade-off that there are delays in both detecting the ships and in delivering the messages to the end user display. As a result, maritime domain monitoring personnel can see the full EEZ for their country or even globally by just licensing data from a provider. Often, the end user's system will integrate coastal AIS data and sometimes multiple sources of satellite AIS.

2 Inherent Challenges with Satellite AIS

3 Revisit Issues

The first challenge of S-AIS is that a satellite must pass over the ship in order to have even a possibility of detecting the ship, a*nd the receiver must be turned on!* Another factor that impacts the usability of the data is that some AIS satellites don't have enough power to operate 100% of the time. In order to manage the cost of downlinks, some operators will deliberately operate their satellites at below 100% duty cycle. What this means for the end user is that there is no simple way to predict how often a satellite with an active receiver will pass over a given ship.

In addition, the orbits of the satellites make a huge difference. Today almost all AIS satellites are in Polar orbits – most in near a "10:30" orbit. Every current system follows a similar coverage pattern; for end users means there tends to be a series of satellite passes in short order, followed by a large gap –which can reach up to 9 hours today.

4 Latency

Latency is simply the time it takes from when the satellite or coastal base station detects an AIS message until the dot on the user's screen appears representing that AIS message. With the speed of light being the driver in coastal AIS systems, this talk latency is completely negligible in coastal systems. The opposite is true with satellite systems. When a ship transmits it's AIS messages, which it does continuously more or less, a satellite must be passing about it in the sky with its receiver turned on in order for such a message to be received. Once this message is received by the satellite, then the satellite must carry that data along its orbit until it reaches a compatible earth station and is able to downlink those messages to the ground. The downlink location in a remote location and then transmitted via the Internet back to a central processing center.

The first factor affecting latency is the processing time. With today's technology there is often a tradeoff: The highest detection rate is achieved through costly ground-based processing of signals from a Canadian company, exactEarth. Some of their satellites use very extensive ground-based processing to detect far more ships than is possible with on-board AIS processing, however all this processing takes time, leading to increased latency.

The other Satellite AIS systems, notably Orbcomm, use on-board processing (OBP) where the AIS messages are extracted by the receivers on -board the satellite, thus requiring little time for ground-based processing. This decreases the latency but is not as robust to signal collisions (a.k.a. co-channel interference) as some other system.

The other key factor in latency is the time after the satellite receives the ship's signal and the time that the satellite later passes over a compatible earth station where the data can be sent back down to the satellite operator. This wait time depends on the number and location of the earth stations. Orbcomm benefits from a large network of ground stations used for its M2M business. exactEarth has also deployed and is expanding a significant network of ground stations but they are unlikely to match the global coverage available to Orbcomm due to the latter's M2M network support

Ships that are missed entirely by a satellite pass are not included in latency calculations. Often missed in latency considerations is the fact that latency is only measured for the messages actually received—in other words, average latency is calculated by averaging the time difference between from the transmission time of each *received message* with the time it is received by the end user and appears on the user's display. Some vessels are missed entirely! This critical factor will be discussed under Refresh Rate below.

5 The Impact of Probability of Detection

Probability of Detection (PoD) is one of the least understood aspects of AIS and Satellite AIS in particular.

Coastal AIS systems have a PoD that matches our common sense. The earth is round and beyond the horizon as seen from the AIS base station, we don't expect to detect ships. We also know that radio signals get weaker with distance, so we all expect that detections will drop off with distance. Regrettably, there is no hard line that indicates that a user will see everything inside of 25 miles and little else beyond that point. Many other factors affect whether a given ship will be detected by an AIS base station, including signal "shadowing" from terrain or buildings, radio interference from non-AIS sources, degraded antenna performance on the ship, and more.

This probability factor means that a coastal system will never reliably detect 100% of the vessels in the area "covered" by AIS. Indeed, in interviews with dozens of AIS operators I found two surprising facts. When asked "What is the range of your AIS?", they often brag about the maximum distance at which they have ever seen a ship, as much as 70 miles or more. Such ranges only occur when all atmospheric and radio conditions are just right. These optimistic numbers are fun to discuss but hardly useful for traffic management.

When asked "What percentage of the ships in your area of operation does AIS reliably detect?, the answer was a surprisingly low 76. And this figure was just for *coastal* AIS. We believe that this lower than expected perception is based primarily on the fact that the Area of Interest for most operators extends beyond the useful range of their AIS base stations. The operator often sees ships appear intermittently until they get closer to the base station.

With Satellite AIS, the situation is more complex. AIS signals were not designed to reach 1,000 miles out to a satellite, so the satellite systems need extremely sensitive receivers to even get the signals at all.

A far greater problem is the "collision" of signals at the satellite. A satellite's field of view (FoV) is about 3,000 miles in diameter. This is an area of more than 7 *million square nautical miles.* The satellite receives signals from ALL SHIPS in this huge area all at once. A study from John's Hopkins University in 2003 showed that if there are more than about 1,000 ships *anywhere within this field of view*, the probability of detecting even 1 message from a ship during a typical satellite pass drops dramatically. At 2000 ships in the FoV, the probability drops to near zero.

COMDEV and exactEarth have published several reports demonstrating that this effect is very real. In fact, exactEarth was formed to implement a ground-based technology that overcomes much of this problem. In order to get to market sooner than launch schedules allowed, exactEarth chose to deploy a hybrid constellation. Some of their satellites use their high-detection capabilities but a significant portion of their constellation uses traditional on-board processing. What the user sees is the combined performance of both classes of satellites.

This effect is often misunderstood or misstated, even by some AIS experts. There have been several articles from industry experts touting that signal collision is a major problem when there is a "high density of ships in the vicinity." Most readers will think about high density areas such as the Straits of Malacca or the Gulf of Mexico. The reality is that these dense areas (there are approximately 20 dense areas of marine traffic in the world) affect the PoD of satellites even for ships up to 3,000 miles away because the satellite may be seeing those ships at the same time. Looking at a globe, you will find that there are few places on Earth where a 3000-mile diameter circle does not cover at least one of the dense areas.

Satellite AIS operators have taken radically different approaches to addressing this issue. exactEarth uses patented ground-based software and huge computing power to extract messages from signals received by their satellites. Orbcomm is deploying a much larger constellation of satellites to increase the number of detection opportunities.

Both of these approaches work – exactEarth, when you consider only their spectrum-based satellites, will get a high percentage of ships on every pass of a satellite, *if* the satellite is one that uses their decollision technology – not all of them possess this capability. However, the cost and latency goes up compared with the on-board processing of Orbcomm and others. Orbcomm, because most of its satellites are funded through their other lines of business, can offer a lower end-user cost and lower latency but a user may have to wait for more satellite passes before a reasonably complete picture of the maritime domain is available.

There is another, more insidious, result of these probabilities, the effect on the refresh of the location data for individual ships. When there is a dense region anywhere within the satellite Field of View, any given satellite may only detect between 15% and 80% of the ships actually present. .

What does this all mean for an AIS user? It means that the dots on the screen will be updated (after the latency wait) for only that percentage of the ships that were detected on a particular pass. Any ship that was not detected will stay in place and not move--the actual location "dot" becomes outdated by the amount of the latency period summed with the amount of time required by the wait until the next satellite pass for which the location of the vessel *is* detected. Even with low latency systems, this means that the average age of the location dots on the display map may still be several hours old!

To summarize, on board processing takes many passes at lower probability of detection to build up a complete picture. Ground based processing also takes multiple passes but starts with a much fuller picture, however ground based processing imposes more serious latency issues. Regardless of the system approach, both types of systems cause long average ship refresh times and create a wide diversity of location ages on the ships displayed. This corresponds to an inaccuracy in positions.

Errors in the visual display are not at all obvious. The average user generally does not notice this time influenced effect because many ship location dots are updating and moving all the time. Only when a particular target is selected does the timestamp appear.

6 Complicated and confusing partial data gives a WRONG picture

A key challenge for an end-user is to understand what is being shown on his monitor when looking at a fused coastal and satellite AIS display. The dominant feature of any such display is the location "dots," arrows or icons on the map representing the locations of vessels. Typically, by clicking on one of these dots the track history of that vessel is displayed on the screen along with a sidebar providing to provide the details of all of the contents of the AIS message representing that ships position report (See Figure 1). Some systems also search a database and can display additional information about a particular chosen vessel from third party or internal databases. Some systems go further and provide a geofencing capability: With this feature the user can draw passage lines or polygons onto their map and establish alert criteria whereby the system will generate an alert when a given vessel crosses that geofence line. Alerts are often escalated via an audible warning, a change of color on the ship's display screen icon, by an email message, or an SMS text message. On some systems it is also possible to have these alerts sent to another software system to manage dispatch and interdiction operations.

All of these functions work fine with a normal AIS display system with a set of AIS base stations connected through a network directly to the AIS display center. However, introduce satellite data and the system gets much more complex.

In an AIS display system containing both coastal data and satellite data the dots on the map **do not represent the current location of all of the ships** in the area displayed. Rather, the dots indicate the *most recently reported* position of all of those ships *that were detected* by the satellite. This creates a problem that is not always obvious to the end user.

Each dot for a given ship stays on the display until one of two things happens:

- Either a new position report is received from that same ship by the system. When this happens, the old location is erased and the new position is displayed on the screen. The earlier dot can be recalled as part of the track history of a given vessel when accessed;

- Or, a location dot representing a ship *will disappear* when there has been no reported position from that vessel in a defined period of time. With most AIS a display using coastal AIS, this period is set to a fairly short time, typically 15 minutes. When possible, satellite system displays must be configured to the particular system delays, which can be much longer—a period of 24 hours is typically set into the system. This means that a ship detected via satellite and not detected again for the next 24 hours, will have its location dot erased from the operator's screen and remain invisible until a new message from the ship is detected. This happens far more often than most users can track.

If the satellites don't detect all the ships within the 24-hour display window, then the picture is incomplete. This is not a critical problem because most systems detect a high percentage of the ships when they have a whole day to work on it.

But there is a bigger, and easily misunderstood problem with the satellite data enhanced systems: When looking at wide-area, there are usually a great number of targets on the screen. Because of our past experience with coastal AIS systems, we tend to assume that we're looking at all of the ships of that area and *where they are presently located.* After all, that's what it would be displayed in a normal, ground based AIS system. However, several facts remain important:

1) Despite the large number of ships, many ships of interest may not be shown at all
2) Specific reported ship data may be more than 24 hours old and hence not visible
3) Most importantly, the reported positions of ships shown are actually from very different times and therefore the ships positions shown are WRONG.

This last point is critically important. The position of a single ship is interesting mainly in relation to the position of other ships or land—On a display with satellite data presented onscreen, the ships shown *have never been in the configuration shown.*

7 User Implications of Mixed Timestamps in AIS Displays

Let's look at a few examples.

On the AIS display screen are two ships that appear to be 1 mile apart and on a collision course moving at 15 knots. This is an emergency, radio contact will be initiated and perhaps other preparatory actions. Yet, if one ship's position was reported ten minutes ago and the other was reported 10 hours ago, then the ships are actually more than 300 miles past each other. Not only is the emergency a waste of time and resources, it's *embarrassing.*

Another example shows why geofences must be designed for this latency. Suppose a fence is defined at a range of 25 miles around a protected area such as a coral reef. If a ship is headed for this area at 10 knots, the user would expect to get an alert when it crosses the fence. Then the user would have 2 ½ hours to contact the ship and warn her before she grounds and causes damage or worse. However, if the ship is detected 5 miles before it passes the fence and the next received report is 3 hours later, then the ship may have already run aground before the geofence can generate the alert.

As a final example, consider a security operation where a ship is suspected of smuggling. When the ship appears at a reasonable range, say 60 miles from the security base, the authorities dispatch 2 interdiction patrol vessels. At 20 knots, these vessels will take 3 hours to arrive at the targeted location. However, if the target ship's reported position was 2 hours old when it was used for the dispatch, then the ship actually has 5 hours in which to rendezvous with other ships and to move up to 75 miles in that time. The local authorities will not be able to provide even a rough idea of the maritime situation as they race into harm's way.

8 Coming Improvements in S-AIS won't solve *these* problems

There are several important developments underway in the satellite AIS industry. These efforts are primarily focused on enhancing detection rates for on-board processing systems. exactEarth continues to add both satellites and earth stations to address and improve their latency. Orbcomm is launching their next generation constellation which will improve their revisit times. The ITU is proposing improvements in radio channel allocations as well as new system level transmission protocols for AIS to improve detection rates.

Several other countries including Denmark, Canada, Japan , the UK and Russia are planning their own AIS satellites. These tend to be either experimental or single satellite systems that are unlikely to provide operational services of use to maritime authorities. The European Space Agency has an ambitious AIS satellite constellation plan but it will take quite a few years to reach deployment and the industry will evolve substantially before then.

Canada is launching a radar satellite constellation that will include AIS payloads. Again, these will be available only years from now.

All of these initiatives are important and will add or improve satellite AIS industry capabilities in important ways.

Yet none of these enhancements will address the inherently complex and confusing display and analysis challenges described in this paper, especially the negative implications of mixed timestamps.

9 Data Fusion Exacerbates the Problem

When an M.D.A. technician uses a display containing both coastal and satellite AIS, the challenge of time spreads is even worse. Ships at the edge of coastal coverage may be detected first by coastal and then the same ship may shift to satellite based detections with the corresponding latency.

A typical fuse display will have ships shown by indicators ranging in time from a few seconds to 24 hours or more. With no consistency in the time stamps of displayed ships, operators must check each ship for its timestamp and make mental estimates of the current position of the target. Each ship in a region must be checked manually, one by one, in order to gain a reasonable estimate to true surface picture.

10 Innovative Solutions Needed

Despite the concerns raised in this paper, nothing here takes away the tremendous increase in maritime domain awareness provided by Satellite AIS. Rather, it focuses our attention on the need for additional operator training. Some of the visualization tools in use today *appear* to behave in the same manner as earlier coastal AIS systems. It is essential that operators understand the differences as outlined in this report. With proper training and a clear understanding of the limitations that come with vastly expanded range of coverage, operators will be able to exercise good judgment and make decisions with a clear understanding of the information at hand.

A potential solution that some have considered is a very large satellite constellation to provide persistent coverage over the planet. Such a system is technically feasible but the costs involved are enormous. Even if such a system were developed, it will take many years to bring it into operation due to both the complexities and the difficulties in obtaining launch opportunities.

Yet another solution would be software that would display all ships at a common point in time. A system could accumulate at high percentage of ships and then display only those location dots that are close in time to each other.

A more advanced solution would be to project the positions of all ships of interest to the present time and to display the current positions of all ships in the area. There are numerous interesting challenges in doing this well that are beyond the scope of this paper, but two approaches are worth mentioning. Consider a ship that has not been detected in 6 hours: If it was travelling east at 12 knots, then a “straight line” projected position would put it 72 miles east of its most recently reported position. If this ship was most recently detected 70 miles offshore to the west, then a projected position would show a position 2 miles inland! Clearly a naïve approach to position forecasting will not suffice.

A second factor concerns data fusion. Projecting positions using only one data feeds will likely create far more problems than it solves. Most users of S-AIS data incorporate at least one terrestrial feed along with their satellite feed while many include multiple data feeds. Any system that intends to provide a common operating picture of the maritime domain will necessarily operate on all of these available data feeds.

11 Conclusion

While satellite AIS has become a tremendous step forward for Maritime Domain Awareness, it is still early in the development of this technology. Revisit times, latency and single-pass detection rates are all key metrics defining the performance of satellite AIS systems in a highly competitive market.

None of these metrics addresses one of the key challenges inherent in all S-AIS systems, the variability in delays that cause AIS displays to portray an out-of date picture of the maritime area of interest.

We titled this paper Time Machine Wanted because of the clear need to improve the Satellite AIS systems so that they can provide a more timely and current accurate view, a way to "look into the future" as well as the past. We suggested several possible avenues of improvement and we have no doubt that even better innovations will emerge in the near future.