SBAS Message Schemes to Support Inline Message Authentication

Todd Walter, Jason Anderson, and Sherman Lo
Stanford University

Abstract
The international community is considering the addition of authentication signatures to the Satellite Based Augmentation System (SBAS) Minimum Operational Performance Standards (MOPS). Authentication would protect the user against the possibility of undesired integrity data being mistaken for genuine State provided signals. Several authentication schemes have been proposed over the past years, but recently a preference has been expressed for inline signatures (i.e., signatures that are interleaved with the correction and integrity messages, all in the same data stream). An important concern is whether there is sufficient bandwidth to add these new messages and will their addition negatively affect performance. We have proposed a method that achieves this goal and evaluate its performance. Further, our method ensures rigorous protection for the user by ensuring that unauthenticated data is discarded and cannot cause harm to the user, while maintaining the required Time-To-Alert of the SBAS integrity messages.

Introduction
Our proposed method is based on the Timed Efficient Stream Loss-tolerant Authentication (TESLA) [1] protocol and requires sending a signature message every six seconds [2] [3] [4]. Thus, we need to ensure that ~17% of the possible message slots are available to be dedicated to this message. Further, such messages need to come out at fixed intervals and should not be delayed except in the rare event of an alert. Therefore, a key component of our proposed approach requires careful integration into the overall SBAS message schedule. We have recently developed a new system design that places messages on a rigid schedule to ensure that this regular update interval is achieved [5]. We have adapted the schedule to include these new authentication messages as occupying 1/6 of the total schedule. We have investigated how to configure the schedule to fit in these messages with the other required messages. This adaptation requires increasing the time interval between correction messages and this paper examines that impact.

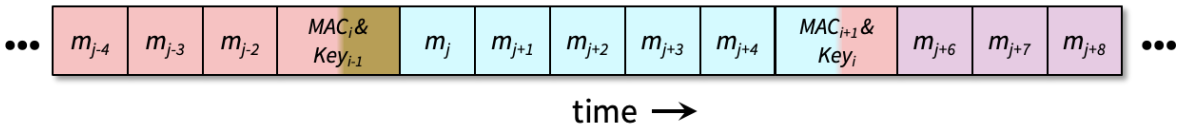

Figure 1. Interleaving of the signature messages in with the other SBAS messages

Figure 1 shows the basic concept of the TESLA approach [4]. A signature message is broadcast that contains Message Authentication Codes (MACs) and a cryptographic key [6]. The MACs use the information from the message and one of the keys to create an unpredictable sequence of bits. In order to verify a MACs and its corresponding message, the user must receive the

message, the MAC, and the key. Five MACs for the previous five messages are broadcast in the first part of the signature message. The key used to generate the MACS will be broadcast in the second part of the subsequent signature message. The color coding in Figure 1 indicates that the red Key_i is associated with the previous set of MACS and not the MACS contained in the same message. At the time that the red MACs are broadcast, only the true SBAS provider knows the key that was used to create them. Six seconds later the key is broadcast to everyone and can be used to authenticate messages *j*-6 through *j*-2. The key itself is verified because it is part of a long hash chain that connects it to the earlier keys [3].

We further investigate the influence of message loss on system performance. Ideally every broadcast message would be received by the aircraft. However, there is always the possibility that some messages may be lost, particularly at low signal to noise ratios [7]. The SBAS MOPS [8] [9] requires that signal and tracking design be such that the message loss rate is no worse one in a thousand. We investigate the impact of a variety of message loss rates all the way up to one in ten. We demonstrate that the TESLA based approach combined with the schedule design is robust for all rates within the MOPS limit. We examine the case where the SBAS only corrects GPS satellites along with one where both GPS and Galileo satellites are corrected.

Message Types and Message Scheduler

The L5 SBAS messages each contain a portion of the information needed by the dual-frequency SBAS user to obtain the satellite corrections and confidence bounds [9]. Each message contains 250 bits and is one second in duration. Table 1 lists the relevant message types used for this study. They include An alert message to discard previously receive content (MT 0) Satellite confidence parameters to be updated every six seconds (MT 35), Individual satellite correction messages (MT 32), a mask message to indicate which satellites are being corrected (MT 31), a bookkeeping message for specifying some integrity parameters (MT 37), almanac messages describing the orbits of the SBAS satellites (MT 47), the authentication signature messages (MT 50), messages to update the authentication keys (MT 51) [4], and messages with no content (MT 63).

Table 1. L5 SBAS Message Types

Message Type	Content	Update interval (seconds)	Timeout interval (seconds)
MT 0	Alert Message / Test Mode	-	-
MT 35	Satellite confidence bound (DFRE)	6	12
MT 32	Differential corrections for one satellite	≥ 15	≥ 30
MT 31	Specifies which satellites are in MT 35	120	600
MT 37	Specifies details about the confidence values	120	600
MT 47	SBAS satellite almanac (confirms PRN)	120	-
MT 50	TESLA signature message	6	-
MT 51	Authentication Over-The-Air-Rekeying (OTAR)	-	-
MT 63	Empty Message	-	-

Table 2. Example L5 SBAS Message Schedule

$t_m = 1$ MT 32 SV 01	$t_m = 2$ MT 32 SV 06	$t_m = 3$ MT 32 SV 11	$t_m = 4$ MT 31	$t_m = 5$ MT 35	$t_m = 6$ MT 32 SV 18
$t_m = 7$ MT 32 SV 02	$t_m = 8$ MT 32 SV 07	$t_m = 9$ MT 32 SV 12	$t_m = 10$ MT 32 SV 16	$t_m = 11$ MT 35	$t_m = 12$ MT 32 SV 20
$t_m = 13$ MT 32 SV 03	$t_m = 14$ MT 32 SV 08	$t_m = 15$ MT 32 SV 13	$t_m = 16$ MT 32 SV 17	$t_m = 17$ MT 35	$t_m = 18$ MT 32 SV 21
$t_m = 19$ MT 32 SV 04	$t_m = 20$ MT 32 SV 09	$t_m = 21$ MT 32 SV 14	$t_m = 22$ MT 32 SV 18	$t_m = 23$ MT 35	$t_m = 24$ MT 32 SV 22
$t_m = 25$ MT 32 SV 05	$t_m = 26$ MT 32 SV 10	$t_m = 27$ MT 32 SV 15	$t_m = 28$ MT 47 v1	$t_m = 29$ MT 35	$t_m = 30$ MT 32 SV 23

$t_m = 31$ MT 32 SV 01	$t_m = 32$ MT 32 SV 06	$t_m = 33$ MT 32 SV 11	$t_m = 34$ MT 37	$t_m = 35$ MT 35	$t_m = 36$ MT 32 SV 19
$t_m = 37$ MT 32 SV 02	$t_m = 38$ MT 32 SV 07	$t_m = 39$ MT 32 SV 12	$t_m = 40$ MT 32 SV 16	$t_m = 41$ MT 35	$t_m = 42$ MT 32 SV 20
$t_m = 43$ MT 32 SV 03	$t_m = 44$ MT 32 SV 08	$t_m = 45$ MT 32 SV 13	$t_m = 46$ MT 32 SV 17	$t_m = 47$ MT 35	$t_m = 48$ MT 32 SV 21
$t_m = 49$ MT 32 SV 04	$t_m = 50$ MT 32 SV 09	$t_m = 51$ MT 32 SV 14	$t_m = 52$ MT 32 SV 18	$t_m = 53$ MT 35	$t_m = 54$ MT 32 SV 22
$t_m = 55$ MT 32 SV 05	$t_m = 56$ MT 32 SV 10	$t_m = 57$ MT 32 SV 15	$t_m = 58$ MT 47 v2	$t_m = 59$ MT 35	$t_m = 60$ MT 32 SV 23

Our L5 SBAS rigid message scheduler has been designed use a predetermined pattern for when messages should be broadcast and ensures that the integrity message (MT 35) and the authentication message (MT 50) are each sent every six seconds [5]. Moreover, when authentication messages are used, the MT 50's immediately follows the MT 35's, minimizing the time it takes to authenticate the integrity information. Table 2 shows an example schedule using the rigid scheduler approach. In this example, the satellite corrections are updated every 30 seconds. The table shows two nearly identical 30 second blocks, whose only differences are highlighted in red. This schedule broadcasts MTs 31, 37, and two instances of 47 every 60 seconds. When authentication messages are sent for this schedule, the MT 32's in the last column of each table are replaced by MT 50's. Note that, this results in five fewer satellites that can be corrected within this 30 second update interval.

Messages are only valid for a limited time. Table 1 lists the timeout interval for each. MT 31 and 37 are messages whose contents do not change often and are required in order to have valid service. By sending them every 60 seconds a user has ten opportunities to receive an update before older versions time out. Thus, timely versions of these messages should always available to users even for very high message loss rates. MTs 32 and 35 are typically only received twice before their timeout periods; thus, they are more sensitive to higher message loss rates. Losing both copies of an MT 32 means that an individual satellite will time out and temporarily not have a valid correction. Losing both copies of an MT 35 message means that all satellites' integrity data is temporarily unavailable which will lead to a brief loss of service. Therefore, losing MT 35 messages has greater consequences that losing an MT 32 message.

Application of the Authentication Messages

The MT 50 message is used to sign the previous five messages [4]. When it is received, the user stores the MACs corresponding to those messages. It also verifies the key against the previously verified key chain. If it is part of that chain, then it may be used with MACs from the previous MT 50 to verify the messages from more than six seconds earlier. In our scheme, most message content that has not yet been authenticated in this manner may not be applied by the

user. The only exceptions are the integrity data, namely MT 0 and the Dual Frequency Range Error Indicators (DFREIs) contained in MT's 32, 35, 36, and 40. All of these may all be used immediately. By doing so, we maintain the required Time-To-Alert (TTA) of the SBAS system [10]. Otherwise, such information could take an additional seven to eleven seconds to verify which would increase the TTA to unacceptable values. As a further detail any Dual Frequency Range Error Change Indicators (DFRECIs) contained in an MT 34 should also be treated as though they are DFREIs and be applied immediately [9]. All other data may not be used until authenticated. If it is not authenticated before timing out, then such data is never used.

Our proposal also requires that any resulting unauthenticated DFREI must be compared against the most recently received authenticated DFREI and the maximum of the two is to be applied by the user. In order to ensure that the authenticated DFREI is also timely, but that it also has two opportunities to be received, its timeout is set to 23 seconds. This value equals the 12 second timeout for the unauthenticated DFREI plus eleven seconds for the time it may take to become authenticated. In our scheme, a user must have received a DFREI within the last 12 seconds and an authenticated DFREI within the last 23 seconds. The maximum value of these two quantities is then used for the satellite. Our rationale for this approach is that an attacker cannot lower the DFREI value and therefore cannot create lower than intended protection levels.

Our European colleagues have suggested a small variation on this method [11] [12]. They propose that the user still must have received a DFREI within the last 12 seconds and an authenticated DFREI within the last 23 seconds. However, the most recent value is used instead of the maximum. Their rationale is that an attacker cannot meaningfully change the position estimate without altering the corrections, and this approach provides too brief of window to lower the protection level to be harmful. We will refer to this approach as the "Latest" method and our approach above as the "Max" method. Formally they are called Use-then-Authenticate (UtA) and Authenticate-then-Use (AtU) respectively.

Next, we will evaluate the performance impact of these two schemes relative to one that does not apply any authentication. Authentication creates two effects: it delays the application of the correction and integrity messages, and it makes them less likely to be successfully received under message loss. This can be seen more explicitly in Table 3. Each column represents a group of six messages culminating in the MT 50 used to sign the previous messages on that row. After the first six messages are received, none of initial four may yet be used. This is because while we do have the messages and MACs for the first five, we do not yet have the key that corresponds to these MACs. The DFREIs in the fifth message may be used as soon as they are received. Six seconds later we will receive the key with the last message on the second row. Only at that point can we verify and then use the contents of the first four messages and authenticate the DFREIs from the first row. Note that if we assume that the MT 35 on the second row is lost, we still can use the DFREIs from the prior row, but they will remain unauthenticated for longer. We also will be able to use the first four messages from this row once we receive the key because each message has its own MAC and the loss of one does not affect the others. But now note that if we lose the MT 50 on the third row, we lose all five

MACS for that row, and we will never be able to authenticate those five messages. Further, we will have to wait until we receive the MT 50 on the fourth row in order to determine the key that applies to the MACS from row two. Once we receive the last message in the table, all of the green shaded messages will have been authenticated, the red shaded ones will never be authenticated, and the yellow shaded ones will be authenticated in the future once the associated key is received. Therefore, we can see that the loss of any message in the first five columns (not an MT 50) will result in the loss of just one message. But the loss of an MT 50 results in the loss of four other messages and that the DFREIs in the MT 35's just before will never be authenticated.

Table 3. Example of Impact of Message Loss

MT 32	MT 32	MT 32	MT 31	MT 35	MT 50
MT 32	MT 32	MT 32	MT 32	X	MT 50
MT 32	MT 32	MT 32	MT 32	MT 35	X
MT 32	MT 32	MT 32	MT 32	MT 35	MT 50
MT 32	MT 32	MT 32	MT 47	MT 35	MT 50

Simulation Setup

To evaluate the impact of authentication on performance we enhanced our MATLAB Algorithm Availability Simulation Toolset (MAAST) [13] to implement the L5 SBAS message set. Figure 2 shows a block diagram of MAAST which can simulate the output of an SBAS. It is capable of running at 1 Hz producing an L5 message at each epoch and then having this message applied to a grid of users covering a region. These messages are accumulated over time to determine the location specific protection levels and corresponding availability and continuity. We also added the ability to inject message loss, causing some messages not to be received by the simulated users.

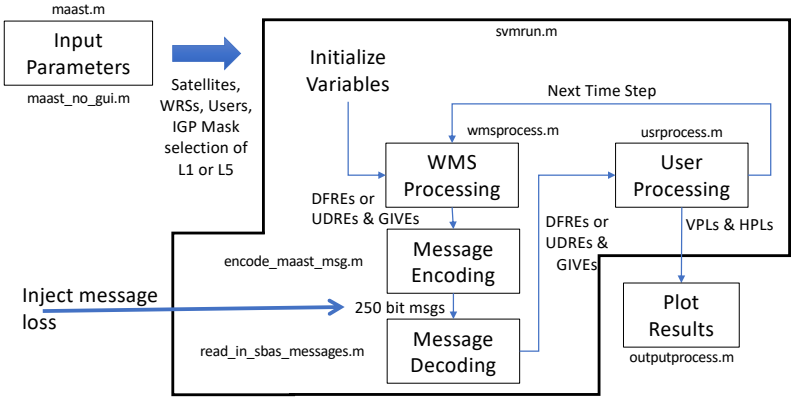

Figure 2. A block diagram of the MAAST Processing Chain

We created four different 24-hour message sequences:

- A GPS 24 satellite constellation without authentication messages
- A GPS 24 satellite constellation with authentication messages
- A combined GPS 24 satellite and 24 Galileo satellite constellation without authentication messages
- A combined GPS 24 satellite and 24 Galileo satellite constellation with authentication messages

Table 4 shows the update intervals and timeout intervals for each case. Adding satellites and/or adding authentication messages requires the satellite corrections to be updated less frequently. We need to leave enough room in the schedule to update all satellites in view of the system at any given time, as well as to have margin to handle alerts and broadcast the OTAR messages [4]. For sequences without authentication, unused message slots (i.e., scheduled MT 32's that don't have corrected satellite currently assigned to them) will be sent as MT 63's (empty messages). While for sequences with authentication, these unused slots will be used to broadcast MT 51. Notice that the timeout intervals in all cases have an extra 12 seconds (compared to two times the update interval). This margin accommodates alerts and the delays in becoming authenticated.

Table 4. Update and Timeout Intervals for the Four Message Sequences

Constellations Used	Update Interval (seconds)	Timeout Interval (seconds)
Without Authentication		
GPS only	30	72
GPS & Galileo	42	96
With Authentication Messages		
GPS only	30	72
GPS & Galileo	54	120

All sequences were tested with varying Word Error Rates (WERs) from zero (no messages lost) to 0.1 (one in ten messages on average). The MOPS requirement is that the operational WER shall be no worse than 10^{-3} [8] [9]. However, we chose to test against larger values to evaluate the robustness of each approach. For each WER, each message timeslot was given the same probability of being lost independently of all other time slots. Once it was determined which timeslots had lost messages, the same missing timeslots were used for all message sequences for a particular WER. The schedules with authentication messages were evaluated using the two different methods previously described: "Max" (AtU) and "Latest" (UtA). The sequences without authentication had the users apply all received messages immediately.

Message Capacity

The message schedules in Table 4 were each implemented in MAAST, and we evaluated how many available slots there were in each case. Figure 3 shows the percentage of the total available slots that each message type occupied. The top left chart corresponds to the GPS-only, No Authentication case. As expected, MT 35 occupied $1/6^{th}$ of all slots while MT 31, 37 and each of the two MT 47 versions occupied $1/60^{th}$. The MT 32's took up nearly half of all

messages and there was quite a bit of spare capacity remaining, nearly 1/3rd of all slots. This indicates that we could send the corrections more often than once every 30 seconds for this scenario. The bottom left chart corresponds to GPS and Galileo, No Authentication case. Now the MT32's occupy nearly 2/3 of the capacity and the empty messages are a much more reasonable 14%.

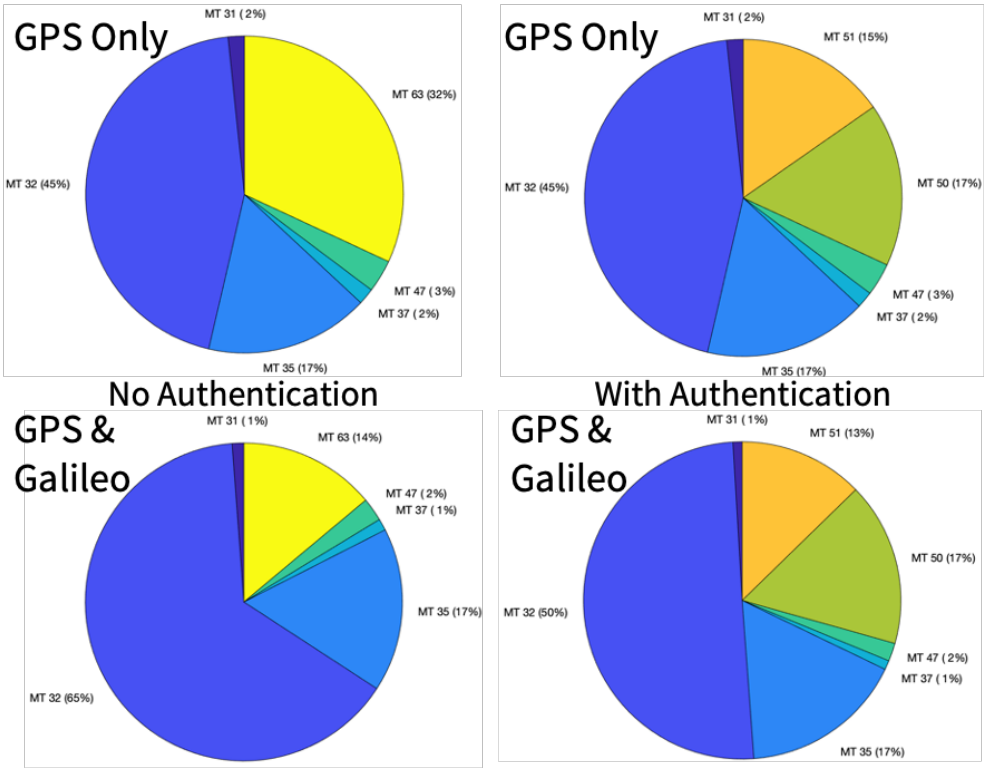

Figure 3. Percentage of total bandwidth used by message type for GPS-only (top) and GPS and Galileo together (bottom) for cases without (left) and with authentication (right)

The charts on the right side of Figure 3 show the cases with authentication. Now MT 50 also occupies 1/6th of all messages. MT 51 has 15% and 13% available for the two different cases. This is a good percentage to support OTAR and therefore these sequences seem to be well adapted for each scenario. Both support updating all OTAR data in well under five minutes.

Performance Results

Figure 4 shows LPV-200 availability for the GPS only case with a WER = 10^{-2} (ten times higher than allowed by the requirement). Both have excellent coverage, but the "Latest" performance is slightly better. This "Latest" performance also matches the No Authentication performance for WER = 10^{-2}. For all WERs below 10^{-2}, all three methods also have identical coverage to the "Latest" case presented in Figure 4. Figure 5 shows similar results for GPS and Galileo together, but now for WER = $3x10^{-2}$ (thirty times higher than allowed by the requirement). Here performance is even better. With more satellites in view, loss of an individual satellite has less of an impact on the user's availability.

Figure 6 shows the summary of the coverage of availability (percentage of region in the plots that achieve 99.9% availability or better). On the left is the GPS-only case and on the right is the GPS and Galileo case. As mentioned earlier all methods work well for WERs up to 10^{-3}. No Authentication and "Latest" have better performance for higher WER values. Since the MOPS requirement is for WER ≤ 10^{-3}, all methods work well for the required outage rates.

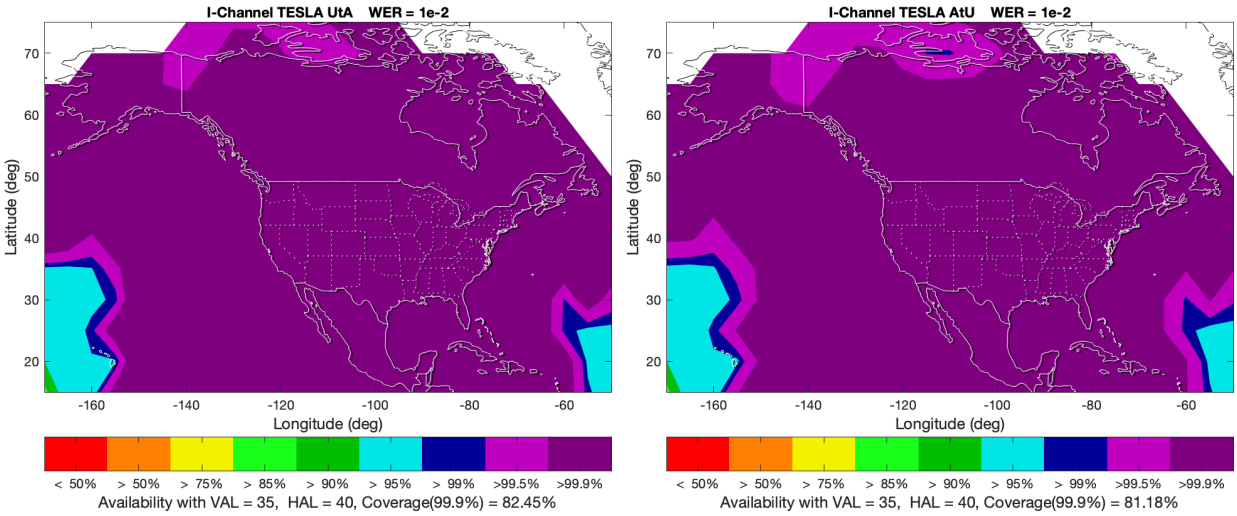

Figure 4. Availability of LPV-200 for "Latest" (left) and "Max" (right) when WER = 10^{-2} when using GPS-only

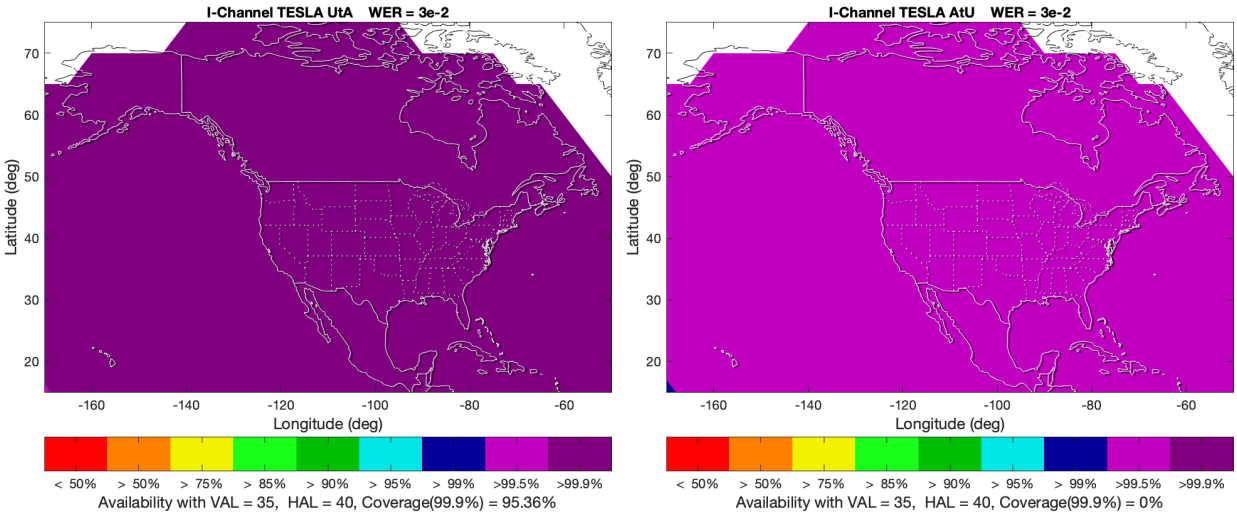

Figure 5. Availability of LPV-200 for "Latest" (left) and "Max" (right) when WER = 3x10^{-2} when using GPS and Galileo

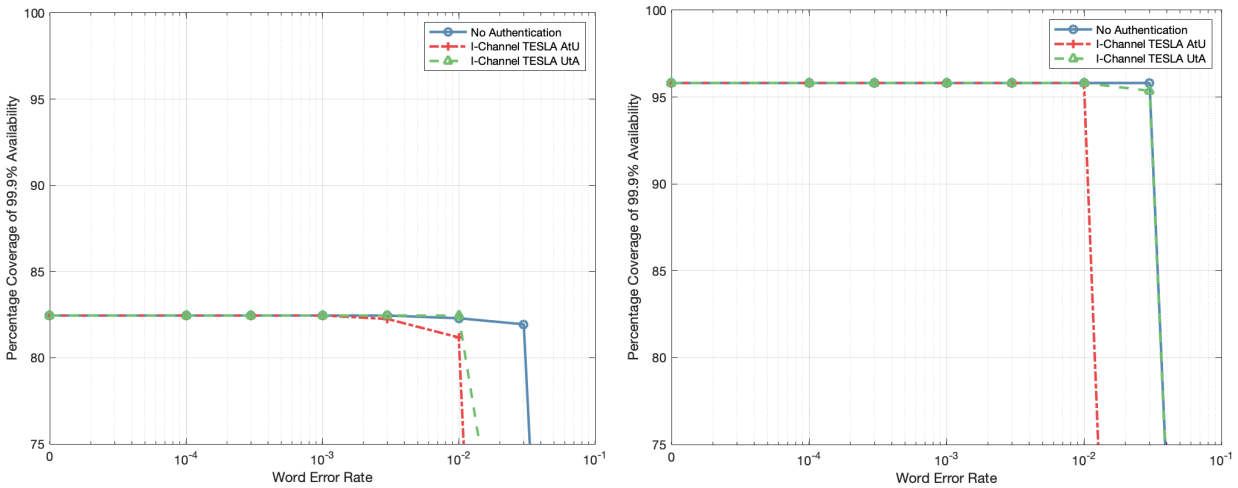

Figure 6. Coverage of LPV-200 for GPS Only (left) and GPS and Galileo (right) for different values of word error rate

Figure 7 shows the impact on the Vertical Protection Level (VPL) [9]. Delayed messages can result in either increased or decreased protection level values. However, on average the delays should increase these confidence values more often than decrease them. The Figure shows the 95% values (i.e. 95% of the time the VPL changed by this much or less). On the left is the GPS-only case and on the right is the GPS and Galileo case. All methods work well for WERs up to $3x10^{-2}$. No Authentication and "Latest" again show better performance. For GPS only 95% of VPL values increased by no more that 8 cm under "Latest" and no more than 16 cm under "Max". For GPS and Galileo together, 95% of VPL values increased by no more that 5 cm under "Latest" and no more than 8 cm under "Max". As seen in the earlier figures this has a negligible impact on availability. Although MAAST cannot evaluate accuracy, the expected impact would be even smaller than these values.

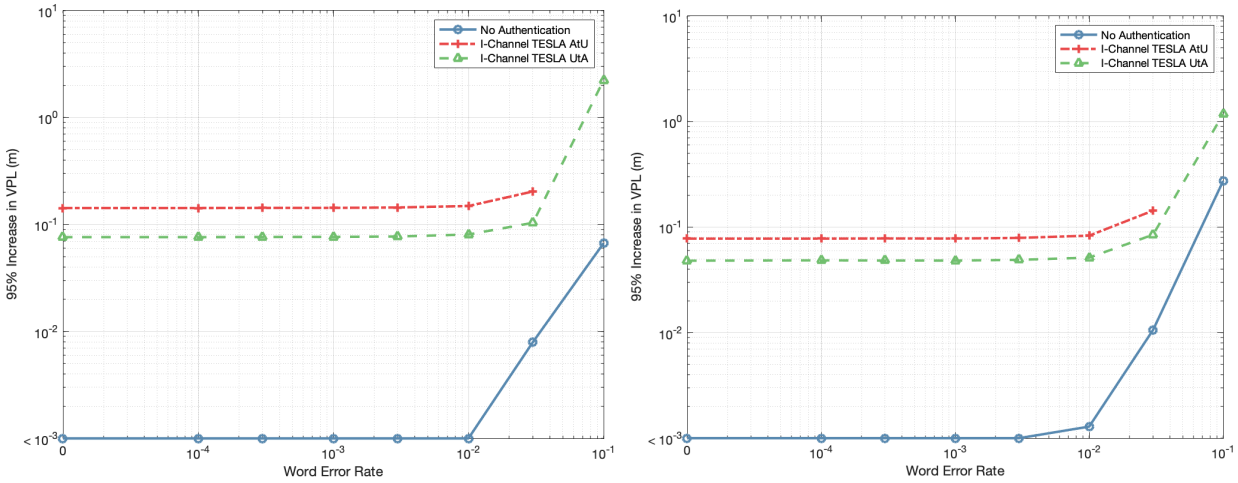

Figure 7. Increase in VPL for GPS Only (left) and GPS and Galileo (right) for different values of word error rate

Conclusion and Future Work

Two methods have been presented to implement an in-line TESLA based navigation message authentication scheme for SBAS. Both methods preserve performance provided that the word error rate does not significantly exceed the 10^{-3} MOPS requirement. Some degradation in performance has been shown when this WER is significantly exceeded. Both methods also support the required SBAS time-to-alert by having the user immediately react to any broadcast that increases any DFREI or tells the user to discard previous data. All other information, in particular the differential corrections, are only used after they have been authenticated. The "Max" method provides the highest degree of protection against potential spoofing attacks by using the maximum of the received authenticated and unauthenticated DFREIs. It is not clear that this level of protection is actually required. The "Latest" method prevents any unauthorized tampering with the position solution and would only allow a few second window for an attacker to reduce the computed protection levels. It is unclear that any actual harm could be caused by such an attack. This question is being examined by certification authorities. And pending the outcome, we recommend using the "Latest" method if allowed due to its increased resilience to message loss. Otherwise, the "Max" method would be applied which still provides excellent performance for all allowed values of WER.

References

[1] A. Perrig, R. Canetti, J. D. Tygar, and Dawn Song, "Efficient authentication and signing of multicast streams over lossy channels," in *Proceeding 2000 IEEE Symposium on Security and Privacy. S&P 2000*, pp. 56–73, doi: 10.1109/SECPRI.2000.848446.

[2] A. Neish, T. Walter, and J. David Powell, "SBAS data authentication: A concept of operations," 2019, doi: 10.33012/2019.17086.

[3] A. Neish, "Establishing Trust through Authentication in Satellite Based Augmentation Systems," Stanford University, 2020.

[4] J. Anderson, S. Lo, A. Neish, and T. Walter, "On SBAS Authentication with OTAR Schemes," 2021.

[5] T. Walter, A. Neish, and J. Blanch, "A Rigid Message Scheduler for SBAS," 2020, doi: 10.1109/PLANS46316.2020.9109849.

[6] I. Fernández-Hernández, V. Rijmen, G. Seco-Granados, J. Simon, I. Rodríguez, and J. D. Calle, "A Navigation Message Authentication Proposal for the Galileo Open Service," *Navig. J. Inst. Navig.*, vol. 63, no. 1, pp. 85–102, 2016, doi: 10.1002/navi.125.

[7] M. J. Hirschberger, S. Lo, and T. Walter, "Reevaluating the Message Loss Rate of the Wide Area Augmentation System (WAAS) in Flight," in *International Technical Meeting of The Institute of Navigation*, Feb. 2021, pp. 218–228, doi: 10.33012/2021.17819.

[8] "RTCA DO-229F Minimium Operational Performance Standards (MOPS) for Global Positioning System/Satellite-Based Augmentation System Ariborne Equipment," 2020.

[9] EUROCAE WG-62, "ED-259 Minimum Operational Performance Standard for Galileo / Global Positioning System / Satellite-Based Augmentation System Airborne Equipment," 2021.

[10] ICAO, "Annex 10 STandards and Recommended Practices (SARPS) Volume I Radio Navigation Aids," 2018.

[11] I. Fernández-Hernández *et al.*, "SBAS Message Authentication: A review of Protocols,

Figures of Merit and Standardization Plans," in *International Technical Meeting of The Institute of Navigation*, Feb. 2021, pp. 111–124, doi: 10.33012/2021.17829.

[12] C. Wullems, L. Tosato, A. D. Chiara, O. Pozzobon, G. F. Serrano, and M. Mabilleau, "Management of Active Data and Authentication in Future SBAS Receivers," in *International Technical Meeting of The Institute of Navigation*, Feb. 2021, pp. 84–97, doi: 10.33012/2021.17827.

[13] S.-S. Jan, W. Chan, and T. Walter, "MATLAB algorithm availability simulation tool," *GPS Solut.*, vol. 13, no. 4, 2009, doi: 10.1007/s10291-009-0117-4.