A Rigid Message Scheduler for SBAS

Todd Walter, Andrew Neish, and Juan Blanch
Stanford University

Abstract

Satellite Based Augmentation Systems (SBASs) broadcast a series of 250-bit messages that together contain a full set of differential correction and integrity information. Each message requires one second to send and contains only a small subset of the full set of required information. Different types of data have different update periods, ranging from six seconds for the satellite integrity information, to five minutes for the ionospheric corrections. A key component of an SBAS is deciding which message needs to be sent at any given second. Each message has an associated time-out period that is approximately two times the update period. It is possible to get into a situation where several messages will time out within a very short interval and it may not be possible to update all of them within the desired time-frame. Further, if there is a change in integrity status that requires an alert to be broadcast, these alert messages will preempt and delay the broadcast of otherwise intended messages.

The Wide Area Augmentation System (WAAS) currently uses a dynamic message scheduler that actively determines which messages should go out over the next several seconds. As the corrections or confidence levels change, some messages may change in priority. This scheduler has been carefully tuned for the current number of GPS satellites and ionospheric grid points, but the algorithm is not adaptable should either value change appreciably. Further, with the advent of L5 messages that have their content organized very differently (and that have no ionospheric corrections) a new message scheduling algorithm is needed.

Introduction

The message scheduler is an unheralded algorithm in SBAS, however, it can have a large impact on the performance of the system. Each SBAS message is capable of broadcasting only a small portion of the overall correction and integrity data [1][2][3][4]. These messages take one second to broadcast and contain fewer than 250 bits of usable information. Every second, the SBAS provider must decide which message to send next. Should they update satellite information, ionospheric information, or bookkeeping information that is required to use the other messages correctly? Each SBAS must meet a strict six-second Time-To-Alert (TTA) and as a result many of these pieces of information time out relatively quickly and are in constant need of being refreshed. For the most stringent vertically guided aircraft operations, satellite integrity data may not be used if it is older than 12 seconds. Satellite clock and orbit corrections time out after four minutes, and ionospheric data times out after ten minutes. Each specific piece of information should be refreshed at least twice within that interval so that the user has at least two chances to obtain an update before the older data may no longer be used. Thus, integrity data should be broadcast at least every six seconds, satellite clock and orbit data at least every two minutes, and ionospheric data at least every five minutes. It is important to ensure that all data is broadcast within these time periods, otherwise users who miss a single copy of any message could lose the ability to use that information. In some cases, losing a message could lead to a loss of continuity as that data may be critical to maintaining sufficiently low protection levels.

If there is a sudden change in integrity status, such that the previously broadcast information is not guaranteed to be safe, the SBAS must immediately broadcast alert messages to inform the user of this change and to provide new information that is guaranteed to be safe. The alert is repeated three times such that four total alert messages are inserted into the regular message stream. This may interrupt previously planned messages and may cause some valuable information to time out.

The current WAAS message scheduler is dynamic. It decides at each second which message is most important to broadcast. However, because of the different update intervals, more than one message may be about to time out at the same time. The original scheduler was subsequently updated with a look ahead feature to identify messages that were about to time out and broadcast them earlier in the cycle to prevent needing to send them out at the last second. This update improved performance, but still does not guarantee that there are always two active (i.e., not timed out) copies of every message.

To address this issue, we have developed a rigid message scheduler that operates according to a preset pattern. Each message is scheduled to go out at a preset time that repeats over a fixed time interval. This rigid schedule prevents the possibility of collision of several simultaneously expiring messages. However, because there may still be integrity alerts, the scheduler cannot always follow the preset pattern. We also must allow alert message to preempt the regularly scheduled ones. Therefore, after such an alert, we need to recover back to the preset pattern. We have examined different options for this recovery algorithm that decides whether a preempted message should be skipped altogether, or whether it should itself preempt a subsequent message. Further, we show how we prevent these delayed messages from cascading over into subsequent intervals.

A benefit of this new scheduler is that it can easily accommodate new capabilities such as the ability to broadcast authentication messages that may need to be sent as often as every six seconds. We also show that the new scheduler has enough flexibility even to handle very large changes to the overall message set. We show how it can be implemented for the legacy WAAS L1 message stream and for the future WAAS L5 message stream and that it may be adapted for one or more constellations.

Table 1. - L1 SBAS Messages Regularly Used by WAAS

Message Type	Description	Update Interval	Time-Out Interval
1	Satellite Mask	120	600
2-5	Fast Correction/Integrity	6	12
7, 10	Degradation Information	120	12
25, 28	SV Correction/Covariance	120	240
18	Ionospheric Mask	300	600
26	Ionospheric Correction	300	600
9	GEO Ephemeris	120	240
17	GEO Almanac	120	-
63	Null	-	-

L1 Message Structure

The L1 message structure is relatively complicated and not all details will be fully described here. Instead readers may look to existing documentation for more complete descriptions of the message structure [3][4]. Table 1 lists the relevant message types regularly used by WAAS. The first column shows the message type number, the next column has a brief description of the contents, the third column shows the nominal update interval in seconds, and the final column shows the time-out interval (also in seconds). For simplicity we only list the time-out interval applicable to precision approach. Non-precision operations are given an extra chance to receive each message and therefore the time-out intervals for these operations are generally multiplied by 1.5. Again, for simplicity, we will only consider the precision approach time-out intervals.

Because of the integrity requirements, WAAS needs to broadcast one each of Message Types 2, 3, and 4 every six seconds. These messages alone consume half of the available bandwidth. Message Types 1, 7, and 10 are satellite mask and customizable parameters describing how the confidence in the corrections degrades with the time since the last received message. These messages contain information that only very infrequently changes. Each one needs to be sent at least once every 120 seconds in order to limit the overall time to first SBAS position fix for a receiver just starting up. Message Types 25 and 28 contain orbit/clock corrections and covariance matrix information respectively for up to two satellites (MT 25 has alternate configurations that allow for up to four satellites per message, but these configurations are not used by WAAS). Every corrected satellite needs to be updated at least every 120 seconds (e.g., if there are 20 correctable GPS satellites in view, then a minimum of ten MT 25s and ten MT 28s need to be sent).

Message Type 18 contains information about which ionospheric grid points will be corrected by the SBAS. There are 11 different global regions (bands) defined, but each SBAS only needs to send information about bands that are visible to their specific region. WAAS uses five bands (numbered 0, 1, 2, 3, and 9 in the SBAS Minimum Operational Performance Standards (MOPS) [3]) and therefore needs to send five different versions of MT 18 every 300 seconds. Message Type 26 sends ionospheric delay estimates and confidence bounds for up to 15 grid points within a specific band. WAAS needs to send 23 different versions of MT 26 to cover all 306 grid points in the full WAAS ionospheric mask. Each of these 23 messages must be sent at least every 300 seconds.

The WAAS L1 signal includes data and ranging information on its Geostationary Earth Orbiting (GEO) satellites. Each GEO is required to broadcast a single Message Type 9 to provide the ephemeris for itself and an MT 17 to provide almanac data for all three WAAS GEOs. These messages need to be sent at least every 120 seconds. Further, the GEO satellites require that MT 28 values be broadcast for them as well, adding three more satellites that are always in view and are correctable for WAAS. These do not require MT 25 corrections as the MT 9s already contain the internal WAAS orbit and clock estimates for these satellites.

In order to accommodate all of these requirements and to meet the minimum update intervals with margin, we have created a sequence that repeats every 540 seconds built around five very similar blocks each of 108 seconds length. This sequence gets all of the messages out with margin and includes many instances of an empty or null message, MT 63, that may be used to allow a graceful recovery from alerts. Message Type 63 is an empty message, it has no content and is sent when no other information needs to go out. Tables 2 and A.1 - 4 show these five successive 108 second blocks. This 540 second sequence contains 90 instances each of MTs 2, 3, and 4 (every six seconds); five instances each of MTs 1, 7, 9, 10, and 17, as well as each version of 25 and 28 every 108 seconds; two instances of each type of MTs 18 and 26 every 270 seconds; and 62 MT 63 messages. If there are fewer than 23 GPS satellite to be corrected, then one or more of the MT 25 and MT 28 messages may also be replaced with MT 63s.

This scheduler meets all of the required time out intervals. The fast corrections and UDREIs are sent exactly every 6 seconds. Messages with 120-second required update periods are sent every 108 seconds (leaving 12 seconds of margin) and messages with 300-second required update periods are sent every 270 seconds (leaving 30 seconds of margin). This margin will be used below in combination with the MT 63s to accommodate alerts.

Table 2. - First 108 second block for L1 WAAS

	$t_0 + 0$	$t_0 + 1$	$t_0 + 2$	$t_0 + 3$	$t_0 + 4$	$t_0 + 5$
$t_0 = 0$	MT2	MT3	MT4	MT1	MT25$_{01}$	MT26$_{01}$
$t_0 = 6$	MT2	MT3	MT4	MT63	MT25$_{02}$	MT28$_{01}$
$t_0 = 12$	MT2	MT3	MT4	MT63	MT25$_{03}$	MT28$_{02}$
$t_0 = 18$	MT2	MT3	MT4	MT63	MT26$_{02}$	MT28$_{03}$
$t_0 = 24$	MT2	MT3	MT4	MT7	MT25$_{04}$	MT26$_{03}$
$t_0 = 30$	MT2	MT3	MT4	MT63	MT25$_{05}$	MT28$_{04}$
$t_0 = 36$	MT2	MT3	MT4	MT10	MT25$_{06}$	MT28$_{05}$
$t_0 = 42$	MT2	MT3	MT4	MT63	MT26$_{04}$	MT28$_{06}$
$t_0 = 48$	MT2	MT3	MT4	MT18$_{1}$	MT25$_{07}$	MT26$_{05}$
$t_0 = 54$	MT2	MT3	MT4	MT63	MT25$_{08}$	MT28$_{07}$
$t_0 = 60$	MT2	MT3	MT4	MT9	MT25$_{09}$	MT28$_{08}$
$t_0 = 66$	MT2	MT3	MT4	MT63	MT26$_{06}$	MT28$_{09}$
$t_0 = 72$	MT2	MT3	MT4	MT63	MT25$_{10}$	MT26$_{07}$
$t_0 = 78$	MT2	MT3	MT4	MT63	MT25$_{11}$	MT28$_{10}$
$t_0 = 84$	MT2	MT3	MT4	MT17	MT25$_{12}$	MT28$_{11}$
$t_0 = 90$	MT2	MT3	MT4	MT63	MT26$_{08}$	MT28$_{12}$
$t_0 = 96$	MT2	MT3	MT4	MT18$_{2}$	MT63	MT26$_{09}$
$t_0 = 102$	MT2	MT3	MT4	MT63	MT26$_{10}$	MT28$_{13}$

Alerting Logic

Next, we describe the logic for handling alerts. Alerts occur rarely on WAAS. On average about three or four per day are seen, but in rare instances as many as 19 alerts in one day have occurred [5]. If the integrity monitors identify the necessity of an alert, the schedule must be interrupted. Internal flags indicate the type of alert and which satellites are affected. Some alerts may affect more than one satellite. These alerts may be broadcast via MTs 2, 3, or 4 if all affected satellites are in one of those messages. If the affected satellites span multiple fast correction messages, then a special alerting message called MT 6 [3] must be sent.

If the alert affects only a single fast correction message and it is already scheduled, then it is sent as scheduled and no message needs to be deferred. Otherwise an MT 2, 3, or 4 must be sent and if the previously scheduled message is not an MT 63, the scheduled message is added to a deferred message list. If the alert affects multiple fast correction messages, an MT 6 will be sent instead. If the scheduled message is not an MT 63, the scheduled message is added to the deferred message list. If the alert is not satellite specific then another special alerting message, a Zero-Filled MT 0 (ZFTZ) [3], is sent. This message tells a receiver to clear their history of previously received messages and that the GEO is currently not usable for aviation guidance. No messages are added to the deferred list if this ZFTZ message is sent and the deferred list is cleared of any prior messages.

Alerts must be sent four times in a row upon initiation. The scheduler keeps track of the alert status through a counter. When the alert is initiated, a first message is sent and the alert status counter is decremented. If the alert is still present the next epoch, another alert message is sent. This process is repeated until four alert messages have gone out in a row and the alert status counter is back to zero, or the initial alert condition is removed (i.e., a satellite is initially declared unusable, but before four seconds have elapsed, it is declared usable again with the original confidence level; in this case the alert counter will be set to zero outside of the message scheduler).

The presence of an alert will likely cause some messages to end up on the deferred list. Normally, an alert requires four messages to be sent in a row and could lead to four messages being added to the deferred list. However, it is possible that back-to-back alerts lead to more than four messages on the deferred list. Message Types 0, 6, and 63 are never added to the deferred list. MT 0 is special type of alert that is never deferred. Similarly, MT 6 is an alerting message that can only be preempted by MT 0. Finally, MT 63 is an empty message that does not need to be sent if preempted.

The only WAAS Message Types that can be on the deferred list are: MTs 1, 2, 3, 4, 7, 9, 10, 17, 18, 25, 26, or 28. No more than one version of a specific message can be on the deferred list. MTs 1, 2, 3, 4, 7, 9, 10, and 17 are each unique and each may only be on the list one time apiece. There are up to five versions of MT 18, and up to 23 versions of MT 26, and each one may only be on the deferred list one time apiece. There are up to 12 versions of MTs 25 and 13 versions of 28, and each one may only be on the list one time apiece.

The list is sorted by importance where MTs 2, 3, and 4 messages are the most important, MTs 7 and 10 have second priority, MTs 9, 25 and 28 have third priority, MT 26 has fourth priority, MT 1 has fifth priority, MT 18 has sixth priority, and MT 17 has the lowest priority. Within each priority grouping the messages are sorted from oldest to newest, with older being given higher priority.

Each deferred message may only preempt a message of the same or lesser priority:

- MT 2 may preempt MTs 1, 3, 4, 7, 9, 10, 17, 18, 25, 26, 28, or 63
- MT 3 may preempt MTs 1, 2, 4, 7, 9, 10, 17, 18, 25, 26, 28, or 63
- MT 4 may preempt MTs 1, 2, 3, 7, 9, 10, 17, 18, 25, 26, 28, or 63
- MT 7 may preempt MTs 1, 9, 10, 17, 18, 25, 26, 28, or 63
- MT 10 may preempt MTs 1, 9, 7, 17, 18, 25, 26, 28, or 63
- MT 9 may preempt MTs 1, 17, 18, 25, 26, 28, or 63
- MT 25 may preempt MTs 1, 9, 17, 18, other 25s, 26, 28, or 63
- MT 28 may preempt MTs 1, 9, 17, 18, 25, 26, other 28s, or 63
- MT 26 may preempt MTs 1, 17, 18, other 26s, or 63
- MT 1 may preempt MTs 17, 18, or 63
- MT 18 may preempt MTs 17, other 18s, or 63
- MT 17 may preempt MT 63

If the deferred message is selected to preempt a lower or equal ranking message, the previously deferred message is removed from the deferred list and the newly preempted message is added to the deferred list (unless it is an MT 63).

If a message is broadcast that has a corresponding match on the deferred list, that message is removed from the deferred list. For example, if MT 25 subscript number 5 is on the deferred list and the scheduled message is MT 25 subscript number 5, then that message is sent as scheduled and the corresponding version is also removed from the deferred list.

Figure 1 contains a flowchart of the above described algorithm. If there have been no recent alerts, then the deferred list is likely empty, and the schedule proceeds nominally as described in the previous section. When alerts occur, messages will be bumped, and the deferred list will accumulate a backlog of messages that can be sent in the strategically placed empty slots. Nearly all alerts only affect a single satellite [5] and therefore MT 6 and MT 0 are very rarely used.

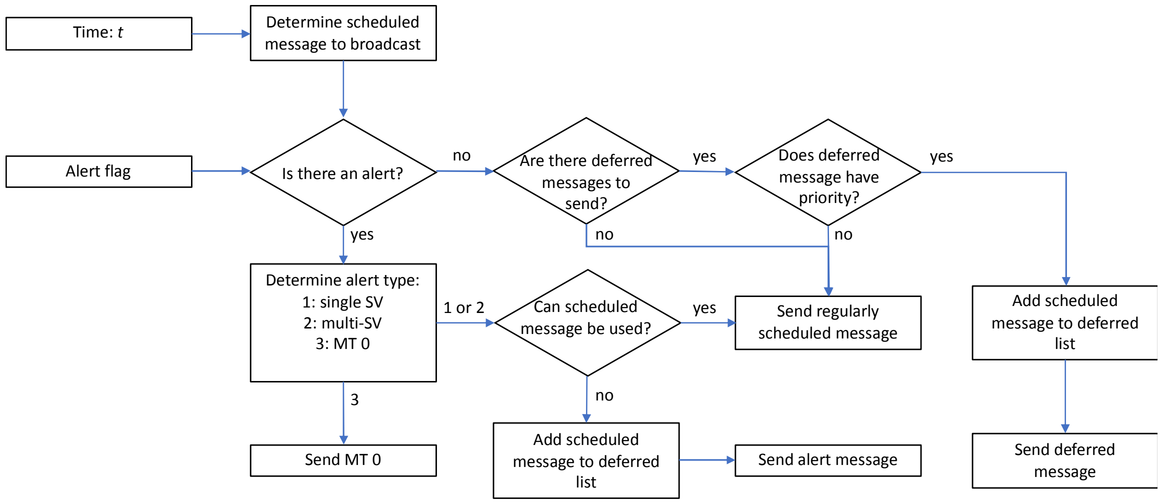

Figure 1. - Flowchart for WAAS L1 Rigid Message Scheduler

Evaluating Nominal Performance

We used our MATLAB Algorithm Availability Simulation Tool (MAAST) [6] to simulate a day's worth of 1 Hz data using a GPS almanac with 31 satellites (October 28, 2018), 3 GEO satellites, and the current WAAS ionospheric mask [7]. MAAST produced 86,400 estimates of User Differential Range Error Indices (UDREIs) for each satellite and Grid Ionospheric Vertical Error Indices (GIVEIs) for each grid point. These internal integrity values are used to determine which satellites are in view and what messages need to be sent. UDREI values below 14 indicate specific numerical gaussian overbounds for the post-correction residual satellite error, while a value of 14 indicates that the satellite has no current correction (i.e., it is out of view). Similarly, GIVEI values below 15 indicate specific numerical gaussian overbounds for the post-correction residual ionospheric error, while a value of 15 indicates that the grid point has no current correction (i.e., there are an insufficient number of nearby measurements).

This information was fed into the message scheduler and we examined the resulting output. All messages went out on time as expected and the schedule was followed exactly. Figure 2 shows the results for PRN 08. It is initially out of view. The UDREI values in the MT 3 message are set to 14. As it comes into view, the UDREI gradually decreases down to a floor value of 5. As it begins to go out of view of the WAAS network, the UDREI increases until PRN 08 is no longer observed and the UDREI is again set to 14. This pattern repeats twice a day. MT 25 and 28 messages are only broadcast for this satellite while it is in view. The right side of Figure 2 shows the first 30 seconds when the satellite comes into view. Note that there is a lag between when the system first has an internal estimate for a numerical UDRE value and when the next MT 3 message is scheduled and can broadcast the UDREI of 13. There is an even greater lag for when the corresponding MT 25 and 28 messages are broadcast.

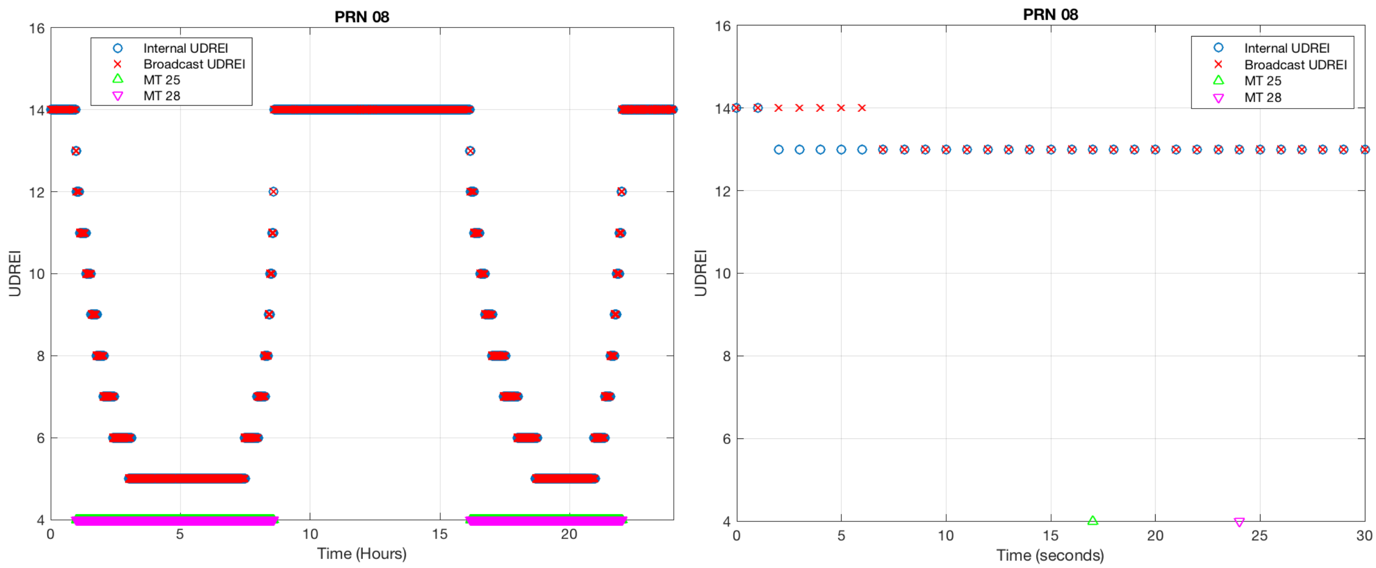

Figure 2. – The internal and broadcast values for UDREI for PRN 08 (a) over the course of the day and (b) as it first comes into view.

In order to minimize this lag on MTs 25 and 28, we added logic to our scheduler to insert an unscheduled copy of each of these messages in the first available empty slot after the satellite first has valid corrections. Otherwise, the regularly scheduled message may not be until up to 107 seconds later, which would delay the ability to make use of the satellite. This logic reduces the mean waiting time for MT 25 from 53 down to 9 seconds and for MT 28 from 51 down to 16 seconds. Thus, a user gains an extra 37 seconds of utility out of the satellite on average.

Testing the Alert Recovery Approach

The full day nominal schedule from the previous section was rerun, but now with added alerts. The alerts were applied every 53 seconds to a random satellite. Fifty-three seconds was chosen as it does not match any of the message update intervals. It is extremely rare to have back to back alerts within 53 seconds of each other. By having so many alerts we could test the effect of an alert at all different points in the scheduler cycle.

A nominal full day will have 86,400 total messages that will include 14,400 instances each of MTs 2, 3, and 4; 800 instances of each type of MTs 1, 7, 9, 10, 17, 25, and 28; 320 instances of each type of MTs 18 and 26; and 9,120 instances of MT 63. With the added alerts every 53 seconds we found that all expected non-empty messages were still broadcast. An additional 5,439 fast corrections (MTs 2, 3, 4) were sent as the alerts, ultimately replacing 5,439 MT 63s. The largest time delay per message type was:

- 4 seconds for MTs 2, 3, and 4
- 7 seconds for MTs 7, 9 and 10
- 11 seconds for MTs 25 and 28
- 24 seconds for MT 26
- 30 seconds for MTs 1 and 18
- 36 seconds for MT 17

Thus, none of MTs 7, 9, 10, 25, or 28 were delayed past their required update interval as the planned schedule has 12 seconds of margin for those messages. Similarly, neither MT 18 nor 26 were delayed past their required update interval as the planned schedule has 30 seconds of margin for those messages. Figure 3 shows the histogram of all delayed messages. There were 13,027 messages that were delayed. Most were only delayed by one second and 76% were delayed by 6 seconds or less. The maximum delay was 36 seconds and that is for a message that is of lowest priority.

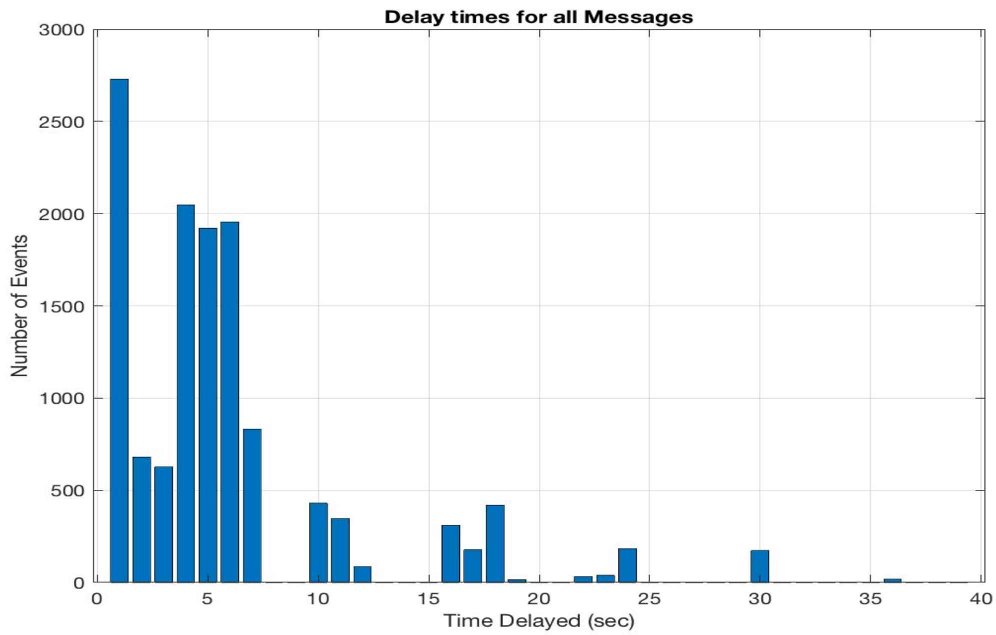

Figure 3. – Histogram of the delay time for all non-empty scheduled messages

The alert recovery mechanism was always able to return to the preplanned schedule within 36 seconds of the alert. Therefore, as long as alerts are at least that far apart, the rigid scheduler should be able to quickly recover from any alert. If there are alerts that occur back to back, four seconds followed immediately by another four seconds, then the scheduler will be more significantly stressed and that may lead to a situation where some data times out. For example, there was an event on the operational system using the existing message scheduler in November of 2018 where successive alerts on MT 3 prevented an MT 4 from being broadcast for a 12 second window, causing its UDREIs to all time out for one second. In this case the sequence was …2, 3, 4, 28, 28, 25, 2, 3, 3, 3, 3, 3, 3, 3, 3, 4,… There is little that any message scheduler could have done to avoid this situation as those alerts had to be sent at those times. Instead one could examine the integrity algorithm that declared the alert. In this case, it bumped it to a higher numerical UDREI and four seconds later decided it had to be bumped again to not monitored. Had it been bumped to not monitored at the beginning, the second alert could have been avoided. However, tuning this algorithm is beyond the scope of this paper.

L5 Message Structure

We now turn to the new Dual-Frequency Multi-Constellation (DFMC) SBAS messages that will be broadcast on L5 [2]. Its message structure is considerably simpler than the L1 message structure [3]. Table 3 shows the subset of message types that are likely to be implemented by WAAS.

Table 3. - DFMC SBAS Messages Planned for WAAS Usage

Message Type	Description	Update Interval	Time-Out Interval
31	Satellite Mask	120	600
32	Satellite Corrections	15 - 204	30 - 408
35	Integrity Information	6	12
37	Customization Data	120	240
47	SBAS Satellite Almanacs	120	240

Message Type 31 is a satellite mask; it is used to identify which piece of integrity information in Message Type 35 goes with which satellite (i.e., the first value goes with PRN 1, the second with PRN 2, etc.). Its information rarely changes over time and it is crucial to have a valid MT 31 in order to use MT 35. Therefore, it has an unusually long time-out giving a receiver at least five chances to successfully decode one. Message Type 32 contains all of the correction information for one satellite. There is one MT 32 for each satellite that is corrected by WAAS. If WAAS can see and correct 18 different GPS satellites at one time, then there are 18 different versions of MT 32 that need to be broadcast. The update and time-out intervals are variable because different SBAS service providers may make different choices as to which constellations to correct and how many active satellites to maintain. Obviously with 18 satellites to correct, each message cannot go out more often than once every 18 seconds. In reality there are many more messages that need to go out, so updating every 30 seconds would be more realistic. If WAAS chose to also correct Galileo, then there may be more than 30 satellites in view, and they all could not be updated more than every 45-60 seconds. Another SBAS may opt to correct all four constellations leading to update intervals of 90 seconds or longer. This update interval is specified in Message Type 37 and can be varied between 15 and 204 seconds at the discretion of the SBAS service provider.

Message Type 35 contains the confidence values for up to 53 satellites. WAAS would use this message exclusively if only correcting one or two constellations. Message Types 34 and/or 36 could be used to handle up to 92 simultaneously corrected satellites, but further discussion of the use of these messages is beyond the scope of this paper. Message Type 37 contains bookkeeping information to allow the user to interpret the information contained in the other messages. It allows flexibility in how the bits are interpreted (e.g., changing the time-out intervals for MT 32). Its information should change only very rarely, but different SBAS providers may choose different options. Finally, MT 47 contains almanac information for up to two SBAS satellites. This information is necessary to positively identify the source of the message. WAAS typically operates three geostationary satellites to broadcast its data although it briefly had four during the summer of 2019. Thus, two different versions of MT 47 are required to transmit almanacs for all of WAAS's GEO satellites.

Therefore, WAAS needs to broadcast one MT 31 at least every 120 seconds, up to 24 MT 32 messages, an MT 35 every six seconds, an MT 37 every 120 seconds, and two MT 47 messages every 120 seconds. The degree of freedom is the update rate for the MT 32 messages. Given the size of the GPS constellation and of the WAAS network, a maximum of 24 should only ever be visible over an approximately ten-minute time-frame.

Because of the six second update interval of the integrity data, it is best to have message repeat intervals be multiples of 6. 30 seconds is the minimum possible to fit in 24 MT 32s along with everything else. Table 4 shows one possible implementation. Each row shows a six-second sequence of messages. The row starts with an MT 35 then the first four rows follow it up with five MT 32s. The last row has four MT 32s and then either an MT 31, MT 37, or one of two MT 47s. These would cycle through in rotation. The first 30-second block would put MT 31 in this last spot, the next 30 second block would send MT 37, the third block would send the first MT 47, and the fourth block would send the second MT 47. Then the whole process would repeat every 120 seconds.

Table 4. - A 30 second/120 second repeat cycle for L5 WAAS

	$t_0 + 0$	$t_0 + 1$	$t_0 + 2$	$t_0 + 3$	$t_0 + 4$	$t_0 + 5$
$t_0 = 0$	MT35	MT32-SV01	MT32-SV02	MT32-SV03	MT32-SV04	MT32-SV05
$t_0 = 6$	MT35	MT32-SV06	MT32-SV07	MT32-SV08	MT32-SV09	MT32-SV10
$t_0 = 12$	MT35	MT32-SV11	MT32-SV12	MT32-SV13	MT32-SV14	MT32-SV01
$t_0 = 18$	MT35	MT32-SV16	MT32-SV17	MT32-SV18	MT32-SV19	MT32-SV20
$t_0 = 24$	MT35	MT32-SV21	MT32-SV22	MT32-SV23	MT32-SV24	MT31/37/47

This approach minimizes the time between MT 32 updates, but it allows very little flexibility and every message only just meets its update interval. This scheme does not handle alerts well due to its lack of margin against the time out intervals. A more robust implementation is shown below in Table 5

Table 5. - A 36 second/108 second repeat cycle for L5 WAAS

	$t_0 + 0$	$t_0 + 1$	$t_0 + 2$	$t_0 + 3$	$t_0 + 4$	$t_0 + 5$
$t_0 = 0$	MT35	MT32-SV01	MT32-SV02	MT32-SV03	MT32-SV04	MT31/63
$t_0 = 6$	MT35	MT32-SV05	MT32-SV06	MT32-SV07	MT32-SV08	MT37/63
$t_0 = 12$	MT35	MT32-SV09	MT32-SV10	MT32-SV11	MT32-SV12	MT63
$t_0 = 18$	MT35	MT32-SV13	MT32-SV14	MT32-SV15	MT32-SV16	MT47-1/63
$t_0 = 24$	MT35	MT32-SV17	MT32-SV18	MT32-SV19	MT32-SV20	MT47-2/63
$t_0 = 30$	MT35	MT32-SV21	MT32-SV22	MT32-SV23	MT32-SV24	MT63

If WAAS had fewer than 24 correctable satellites, then the unneeded MT 32 messages would be replaced with MT 63. Here the 36-second block has three different versions to form a complete 108-second cycle. In the first version, perhaps MT 31 and the first MT 47 are sent while MT 63s are sent in the last slot of rows corresponding to $t_0 = 6$ and 24. In the second block, MT 37 is sent and all other messages in the last column are sent as MT 63. In the final block, the second MT 47 is sent and again all other messages in the last column are sent as MT 63. After 108 seconds, the entire pattern is repeated.

At first glance this seems wasteful as we are sending many empty messages and adding an unnecessary six seconds between MT 32 updates. However, the empty slots are essential to recovering from alerts. When an alert occurs, it will preempt the regularly scheduled message. We then have two choices: we can lose the message forever or we can send it out somewhat delayed. It is far better to send it out with minimal delay. With Table 4, our only real option is to lose it forever. With Table 5, we have several MT 63 slots that can be used to send the displaced messages. As another benefit we are sending the MT 31, 37, and 47 messages slightly more often than necessary (every 108 seconds instead of 120) so that we can afford to delay their transmission by up to 12 seconds without risking having the user time-out critical information.

Conclusions

The rigid message scheduler works as designed. It sends all of the messages out on time when there are no alerts present. We have minimized the delay time from when a satellite first becomes observed and the required message set is sent to the user such that they can incorporate it into their position solution and protection levels. Alerts cause delay for some of the messages, however, these are typically less than six seconds and the schedule is fully recovered within ~40 seconds at the longest. This message scheduler design has been delivered to the FAA and they are planning to incorporate it into future versions of WAAS. It is a substantial improvement over the existing message scheduler which can sometimes send out delayed messages even in the absence of alerts.

This message scheduler also offers substantial flexibility to create preplanned patterns that can accommodate significant changes such as adding authentication messages every six seconds [8] or adding another constellation (e.g., Galileo) to the operation. The current WAAS scheduler would require a significant re-tuning exercise to accommodate either of these modifications whereas the rigid scheduler simply requires an update to the predefined pattern. Thus, we feel that the new rigid scheduler creates much more flexibility for WAAS design changes as we look into incorporating some of these significant changes.

References

[1] Walter, T., "Chapter 12: Satellite Based Augmentation Systems" in Springer Handbook of Global Navigation Satellite Systems, Published 2017 by Springer International Publishing Switzerland, DOI: 10.1007/978-3-319-42928-1.

[2] EUROCAE WG-62, "Minimum Operational Performance Standard for Galileo / Global Positioning System / Satellite-Based Augmentation System Airborne Equipment," ED-259A v0.03 draft publication October 2019.

[3] RTCA SC-159 WG2, "Minimum Operational Performance Standards for Global Positioning System/Satellite Based Augmentation System Airborne Equipment," RTCA publication DO-229E, December 2016.

[4] Walter, T., "WAAS MOPS: Practical Examples," in proceedings of the ION National Technical Meeting, San Diego, California, January 1999.

[5] Bill Wanner, FAA Technical Center, private communication.

[6] Jan, S.S., Chan, W., and Walter, T., "MATLAB Algorithm Availability Simulation Tool," Published in GPS Solutions, Vol. 13., No. 4, September 2009.

[7] Walter, T., Shallberg, K., Altshuler, E., Wanner, W., Harris, C., and Stimmler, R., "WAAS at 15," Published in NAVIGATION , Vol. 65, No. 4, Winter 2018, pp. 581-600, DOI 10.1002/navi.252.

[8] Neish, A., Walter, T., and Powell, J. D., "SBAS Data Authentication: A Concept of Operations," Published in Proceedings of the 32nd International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS+ 2019), Miami, Florida, September 2019.

Appendix

This appendix includes the last four blocks of the proposed WAAS L1 message. These blocks are nearly identical to the first block, but they vary in their ionospheric information (MT 18 and 26).

Table A1. - Second 108 second block for L1 WAAS

	$t_0 + 0$	$t_0 + 1$	$t_0 + 2$	$t_0 + 3$	$t_0 + 4$	$t_0 + 5$
$t_0 = 108$	MT2	MT3	MT4	MT1	$MT25_{01}$	$MT26_{11}$
$t_0 = 114$	MT2	MT3	MT4	MT63	$MT25_{02}$	$MT28_{01}$
$t_0 = 120$	MT2	MT3	MT4	MT63	$MT25_{03}$	$MT28_{02}$
$t_0 = 126$	MT2	MT3	MT4	MT63	$MT26_{12}$	$MT28_{03}$
$t_0 = 132$	MT2	MT3	MT4	MT7	$MT25_{04}$	$MT26_{13}$
$t_0 = 138$	MT2	MT3	MT4	MT63	$MT25_{05}$	$MT28_{04}$
$t_0 = 144$	MT2	MT3	MT4	MT10	$MT25_{06}$	$MT28_{05}$
$t_0 = 150$	MT2	MT3	MT4	MT63	$MT26_{14}$	$MT28_{06}$
$t_0 = 156$	MT2	MT3	MT4	$MT18_{3}$	$MT25_{07}$	$MT26_{15}$
$t_0 = 162$	MT2	MT3	MT4	MT63	$MT25_{08}$	$MT28_{07}$
$t_0 = 168$	MT2	MT3	MT4	MT9	$MT25_{09}$	$MT28_{08}$
$t_0 = 174$	MT2	MT3	MT4	MT63	$MT26_{16}$	$MT28_{09}$
$t_0 = 180$	MT2	MT3	MT4	MT63	$MT25_{10}$	$MT26_{17}$
$t_0 = 186$	MT2	MT3	MT4	MT63	$MT25_{11}$	$MT28_{10}$
$t_0 = 192$	MT2	MT3	MT4	MT17	$MT25_{12}$	$MT28_{11}$
$t_0 = 198$	MT2	MT3	MT4	MT63	$MT26_{18}$	$MT28_{12}$
$t_0 = 204$	MT2	MT3	MT4	$MT18_{4}$	MT63	$MT26_{19}$
$t_0 = 210$	MT2	MT3	MT4	MT63	$MT26_{20}$	$MT28_{13}$

Table A.2. - Third 108 second block for L1 WAAS

	$t_0 + 0$	$t_0 + 1$	$t_0 + 2$	$t_0 + 3$	$t_0 + 4$	$t_0 + 5$
$t_0 = 216$	MT2	MT3	MT4	MT1	$MT25_{01}$	$MT26_{21}$
$t_0 = 222$	MT2	MT3	MT4	MT63	$MT25_{02}$	$MT28_{01}$
$t_0 = 228$	MT2	MT3	MT4	MT63	$MT25_{03}$	$MT28_{02}$
$t_0 = 234$	MT2	MT3	MT4	MT63	$MT26_{22}$	$MT28_{03}$
$t_0 = 240$	MT2	MT3	MT4	MT7	$MT25_{04}$	$MT26_{23}$
$t_0 = 246$	MT2	MT3	MT4	MT63	$MT25_{05}$	$MT28_{04}$
$t_0 = 252$	MT2	MT3	MT4	MT10	$MT25_{06}$	$MT28_{05}$
$t_0 = 258$	MT2	MT3	MT4	MT63	MT63	$MT28_{06}$
$t_0 = 264$	MT2	MT3	MT4	$MT18_{5}$	$MT25_{07}$	MT63
$t_0 = 270$	MT2	MT3	MT4	MT63	$MT25_{08}$	$MT28_{07}$
$t_0 = 276$	MT2	MT3	MT4	MT9	$MT25_{09}$	$MT28_{08}$
$t_0 = 282$	MT2	MT3	MT4	MT63	$MT26_{01}$	$MT28_{09}$
$t_0 = 288$	MT2	MT3	MT4	MT63	$MT25_{10}$	$MT26_{02}$
$t_0 = 294$	MT2	MT3	MT4	MT63	$MT25_{11}$	$MT28_{10}$
$t_0 = 300$	MT2	MT3	MT4	MT17	$MT25_{12}$	$MT28_{11}$
$t_0 = 306$	MT2	MT3	MT4	MT63	$MT26_{03}$	$MT28_{12}$
$t_0 = 312$	MT2	MT3	MT4	$MT18_{1}$	MT63	$MT26_{04}$
$t_0 = 318$	MT2	MT3	MT4	MT63	$MT26_{05}$	$MT28_{13}$

Table A.3. - Fourth 108 second block for L1 WAAS

	t_0 + 0	**t_0 + 1**	**t_0 + 2**	**t_0 + 3**	**t_0 + 4**	**t_0 + 5**
t_0 = 324	MT2	MT3	MT4	MT1	$MT25_{01}$	$MT26_{06}$
t_0 = 330	MT2	MT3	MT4	MT63	$MT25_{02}$	$MT28_{01}$
t_0 = 336	MT2	MT3	MT4	MT63	$MT25_{03}$	$MT28_{02}$
t_0 = 342	MT2	MT3	MT4	MT63	$MT26_{07}$	$MT28_{03}$
t_0 = 348	MT2	MT3	MT4	MT7	$MT25_{04}$	$MT26_{08}$
t_0 = 354	MT2	MT3	MT4	MT63	$MT25_{05}$	$MT28_{04}$
t_0 = 360	MT2	MT3	MT4	MT10	$MT25_{06}$	$MT28_{05}$
t_0 = 366	MT2	MT3	MT4	MT63	$MT26_{09}$	$MT28_{06}$
t_0 = 372	MT2	MT3	MT4	$MT18_{2}$	$MT25_{07}$	$MT26_{10}$
t_0 = 378	MT2	MT3	MT4	MT63	$MT25_{08}$	$MT28_{07}$
t_0 = 384	MT2	MT3	MT4	MT9	$MT25_{09}$	$MT28_{08}$
t_0 = 390	MT2	MT3	MT4	MT63	$MT26_{11}$	$MT28_{09}$
t_0 = 396	MT2	MT3	MT4	MT63	$MT25_{10}$	$MT26_{12}$
t_0 = 402	MT2	MT3	MT4	MT63	$MT25_{11}$	$MT28_{10}$
t_0 = 408	MT2	MT3	MT4	MT17	$MT25_{12}$	$MT28_{11}$
t_0 = 414	MT2	MT3	MT4	MT63	$MT26_{13}$	$MT28_{12}$
t_0 = 420	MT2	MT3	MT4	$MT18_{3}$	MT63	$MT26_{14}$
t_0 = 426	MT2	MT3	MT4	MT63	$MT26_{15}$	$MT28_{13}$

Table A.4. - Fifth 108 second block for L1 WAAS

	t_0 + 0	**t_0 + 1**	**t_0 + 2**	**t_0 + 3**	**t_0 + 4**	**t_0 + 5**
t_0 = 432	MT2	MT3	MT4	MT1	$MT25_{01}$	$MT26_{16}$
t_0 = 438	MT2	MT3	MT4	MT63	$MT25_{02}$	$MT28_{01}$
t_0 = 444	MT2	MT3	MT4	MT63	$MT25_{03}$	$MT28_{02}$
t_0 = 450	MT2	MT3	MT4	MT63	$MT26_{17}$	$MT28_{03}$
t_0 = 456	MT2	MT3	MT4	MT7	$MT25_{04}$	$MT26_{18}$
t_0 = 462	MT2	MT3	MT4	MT63	$MT25_{05}$	$MT28_{04}$
t_0 = 468	MT2	MT3	MT4	MT10	$MT25_{06}$	$MT28_{05}$
t_0 = 474	MT2	MT3	MT4	MT63	$MT26_{19}$	$MT28_{06}$
t_0 = 480	MT2	MT3	MT4	$MT18_{4}$	$MT25_{07}$	$MT26_{20}$
t_0 = 486	MT2	MT3	MT4	MT63	$MT25_{08}$	$MT28_{07}$
t_0 = 492	MT2	MT3	MT4	MT9	$MT25_{09}$	$MT28_{08}$
t_0 = 498	MT2	MT3	MT4	MT63	$MT26_{21}$	$MT28_{09}$
t_0 = 504	MT2	MT3	MT4	MT63	$MT25_{10}$	$MT26_{22}$
t_0 = 510	MT2	MT3	MT4	MT63	$MT25_{11}$	$MT28_{10}$
t_0 = 516	MT2	MT3	MT4	MT17	$MT25_{12}$	$MT28_{11}$
t_0 = 522	MT2	MT3	MT4	MT63	$MT26_{23}$	$MT28_{12}$
t_0 = 528	MT2	MT3	MT4	$MT18_{5}$	MT63	MT63
t_0 = 534	MT2	MT3	MT4	MT63	MT63	$MT28_{13}$