Implementation of the L5 SBAS MOPS

Todd Walter, Juan Blanch, and Per Enge
Stanford University

ABSTRACT

Several L1 and L5 capable satellites are already on orbit and many more are launching soon. The operators of the various Satellite Based Augmentation Systems (SBASs) are actively investigating incorporating the new constellations and new signals. We have recently proposed new messages for broadcast on L5. These messages are intended for use in the Minimum Operational Performance Standards (MOPS) and are designed to allow SBASs to communicate corrections and integrity information to the aircraft. Specifically, they allow the use of L1 and L5 in combination to achieve better performance than today's use of L1-only. The proposed L5 messages were designed to make better use of the limited data bandwidth between the SBAS geostationary satellites and the aircraft. In particular, they are intended to allow for the simultaneous correction of at least four constellations of navigation satellites.

This paper evaluates real satellite clock and orbit performance data and implements both the existing and proposed MOPS message formats. It then demonstrates the benefits of the proposed messages. It further demonstrates the new message structure's capability to handle satellite alerts. It provides a practical demonstration of functionality of the new message structure including the improved bandwidth capability to handle multiple constellations.

INTRODUCTION

The new messages proposed in 2012 [1] will require some changes in ground monitoring over what is done in today's Satellite Based Augmentation Systems (SBASs). In particular, we have proposed eliminating the fast correction message that was originally designed to mitigate selective availability. The fast correction messages are also often used to communicate alerts to the user when unexpected behavior is observed on the satellites.

The SBAS ground monitors will need to alter their clock estimation strategies as well as how they broadcast satellite alerts. This paper examines the behavior of the satellite clocks and proposes new approaches to estimating and broadcasting these corrections. We also examine historical alerts and how they have been communicated to the avionics. We then show how these same events can be handled by the new messages. The determination of what messages need to be sent, and at what time, is an important and often challenging algorithm for SBAS. We demonstrate a fairly simple approach that maximizes throughput of the desired messages, but also allows for alert messages to be sent as needed without significantly penalizing performance.

We also examine the parameters required to ensure that the integrity information is maintained over the lifetime of the message and determine values to strike a balance between preventing the need for alerts versus excessively raising confidence bounds. The analyses in this paper demonstrate that the proposed message structure works very well with observed and expected satellite performance. We further compare performance to the existing L1 MOPS structure and examine the relative merits of the two approaches. We will show that the proposed L5 approach improves both accuracy and availability while simplifying both the ground generation of the messages as well as their airborne application.

Finally, we examine the benefit of correcting four full core constellations of satellites. We show that there is significant benefit to going beyond the current 51-satellite limit that effectively prevents the use of more than two constellations by any service provider.

PROPOSED MOPS CHANGES

In a previous paper [1], we proposed several changes to the MOPS messaging approach. The proposal contained seven distinct components:

- An expanded PRN Mask to include Galileo and Beidou as well as other regional systems;
- Removal of the fast corrections;
- A new alert message that can update 91 satellites;

- Assembling all of the satellite corrections into a single message;
- Reduced dynamic range and quantization errors in the satellite corrections;
- The ability to broadcast the table that specifies the quantization of the confidence factors; and,
- An SBAS satellite orbit description message that can describe more than just geostationary orbits.

In this paper, we will specifically examine the effect of removing the fast corrections, the impact of reduced quantization error, and the use of the proposed alerting message.

The primary reason why the current MOPS [2] included fast corrections is that in previous years the GPS constellation deliberately degraded its accuracy by dithering its clocks. This action was known as Selective Availability (SA) [3]. It was implemented by GPS to prevent civilian users from obtaining the same accuracy as was obtained by military users. However, SA was easily defeated by differential corrections and it was removed from service in 2000. It is no longer part of GPS and it is not included as part of any of the other constellations.

Removing the fast corrections is important because they currently consume most of the available messaging bandwidth. Currently, the fast corrections are sent six seconds apart and each message updates up to 13 satellite clocks. Updating more than 39 satellites would require more than half of the available messaging capacity. Further, it is not possible to update more than 51 satellites concurrently under the existing MOPS. It is necessary to increase the time between clock updates in order to update much more than one constellation's worth of satellites.

However, extending the time between clock corrections allows greater uncertainty to accumulate at the user. The next section examines how quickly the satellite clock error grows with time. The quantization error in the corrections also contributes to this effect. We will show how the proposed changes mitigate some of the increased uncertainty caused by less frequent clock updates.

Finally, in order to update the integrity status of more than 51 satellites in a single message, we adopted an approach that reduces the amount of information sent. This change and its implications will be discussed in greater detail later in the paper.

SATELLITE CLOCK BEHAVIOR

The clocks used for satellite navigation are space qualified atomic clocks. GPS uses a mix of cesium and rubidium atomic clocks. GLONASS uses all cesium clocks, while Galileo will use hydrogen masers. These clocks all have extremely low frequency offsets and frequency drift rates. Figure 1 shows the change in GPS clock error over one week as estimated from the precise IGS GPS clock estimates [4] [5]. These estimates are provided for five-second intervals and are accurate to the centimeter level. The clock error estimated at the beginning of the week was subtracted from the subsequent error for each satellite in this figure. The changes are extremely linear and contain no obvious frequency drift.

Figure 2 shows a histogram of the variation in the

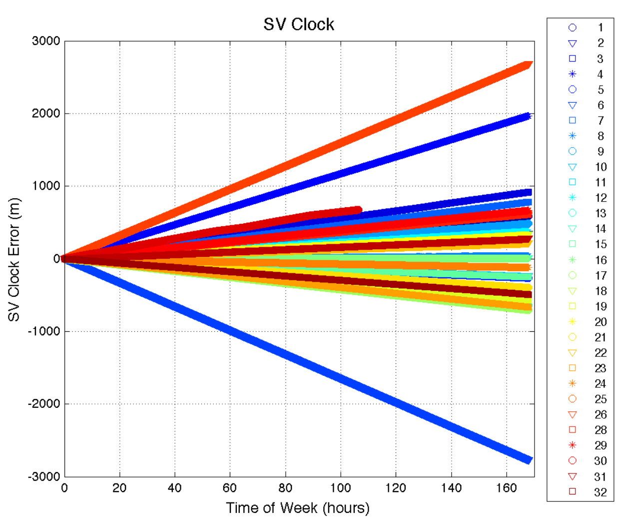

Figure 1. *Change in GPS clock error over March 24-30, 2013. The behavior is extremely linear.*

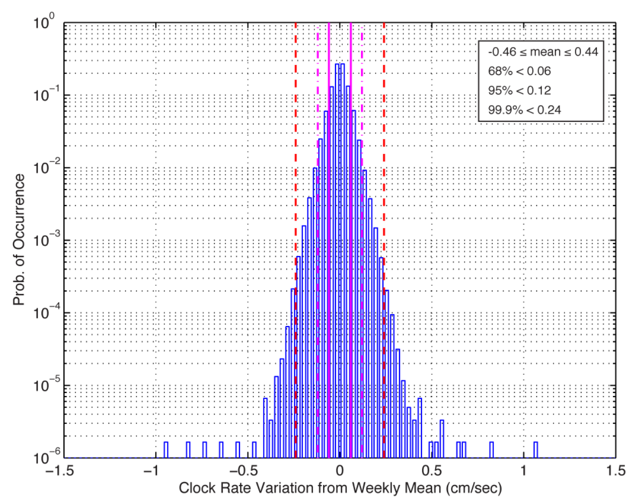

Figure 2. *Histogram of estimated short-term (six-minute) clock rate minus the weekly average.*

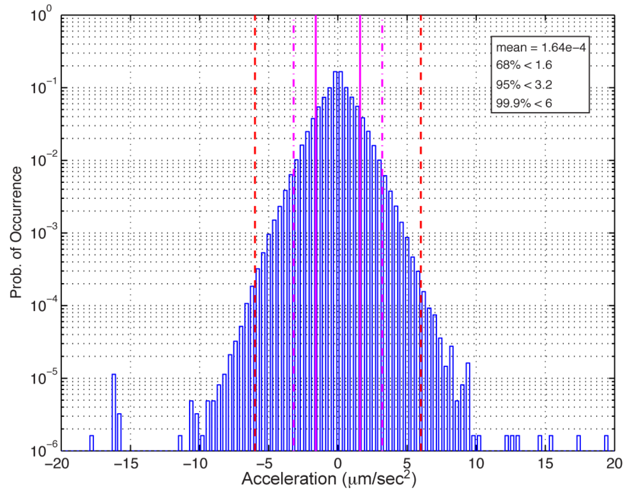

Figure 3. *Histogram of estimated short-term (six-minute) clock acceleration.*

estimated clock frequency obtained by subtracting a weekly mean frequency from the frequency estimate obtained from a sliding six-minute linear fit. The frequencies are all nearly constant over the week. The drifts do not appear to change significantly over the course of hours or days. Some satellites have small frequency offsets, but all have mean values within 0.5 cm/second. The variations about the mean have a standard deviation of 0.06 cm/sec and nearly all of the variations are contained within 0.5 cm/sec of their respective mean values.

Figure 3 shows a histogram of the clock acceleration estimates as formed by using a sliding six-minute second order fit to the raw clock data. All satellites are essentially zero-mean and all are well under 0.1 mm/sec^2. In fact, these estimates are likely dominated by noise from the clock estimation process and the actual clock acceleration values are even smaller still.

Therefore, it is obvious that GPS clock errors contain no significant second order term over timeframes that correction messages need to remain accurate. A simple linear prediction should suffice to provide clock correction information to the users. In the next section we will look at a simple estimator and its accuracy as a function of age of the correction.

SATELLITE CLOCK PREDICTION ACCURACY

In the current MOPS, the long-term corrections are broadcast at least every two minutes. We will assume that for the next MOPS satellite corrections should be sent at least this often. Further, because a user may miss

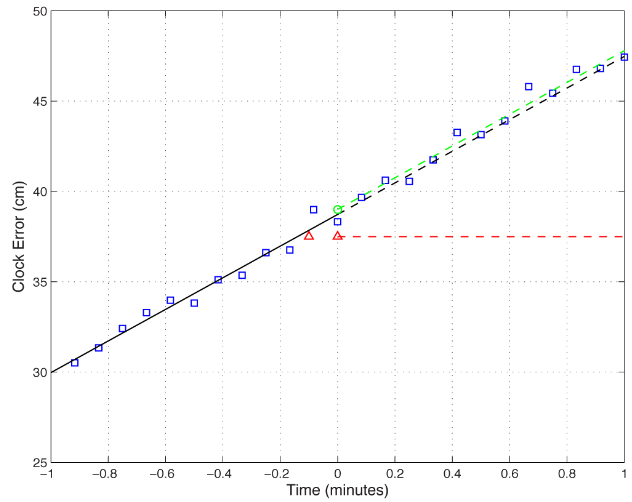

Figure 4. *The process of predicting the data is shown. The raw data is shown by the blue squares. The internal fit by the black line. The quantized value for broadcast is illustrated by the green circle and dashed green line. Finally, the red triangles show the current MOPS* fast corrections *and extrapolated correction.*

messages, they are allowed to use the previous message should they fail to receive the latest update. In less demanding phases of flight they may hold onto an even older correction in case they have missed the two most recent updates. Therefore, the expected lifetime of the message is up to 360 seconds.

Because the message should remain valid for up to 360 seconds, we will use a similar length of data to estimate the rate of change of the clock. Using much shorter intervals of data could lead to increased uncertainty in the rate of change estimate. Using significantly longer data sets may fail to respond to changes in clock slope should such a change occur. Therefore, we used the previous six minutes of IGS precise five-second satellite clock estimates to perform a linear fit and then estimated the current clock error and the error rate.

Figure 4 illustrates the process used. The blue squares represent the precise IGS data. At each time epoch, the previous six-minutes of data was fit to a first order model as shown by the solid black line. The estimate at time zero is quantized to fit within the proposed satellite correction message (3 cm quantization steps) [1]. The predicted slope was also quantized for the message (0.12 mm/sec quantization steps). The unquantized predicted fit (dashed black line) and the quantized prediction (dashed green line) are then compared against the raw data. For comparison, we also quantized the raw value at time zero using the current L1 MOPS message format (12.5 cm quantization steps). The prediction from six

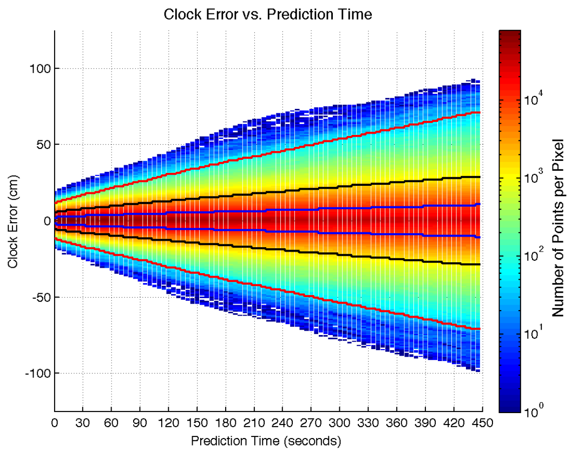

Figure 5. *Two-dimensional histogram of unquantized clock prediction errors vs. time since prediction. The blue, black, and red lines show the error growth for 68%, 95%, and 99.9% of the data, respectively.*

seconds prior was also quantized and then their difference was used to form the range rate correction and extrapolate forward in time (dashed red line).

Because the quantization steps for the current MOPS are so large and the current clock rates are so low, the extrapolated line almost always has zero slope. Fortunately, it is intended to correct for only a short time (six seconds or up to 18 in the event of lost messages). However, even over these short prediction times it is obvious that the current approach is much less accurate.

Figure 5 shows a two-dimensional histogram of results for the unquantized prediction. The x-axis shows the time of prediction from zero to 450 seconds. The y-axis shows the magnitude of the prediction error and the color indicates the frequency of occurrence. Smaller prediction errors were much more common than larger prediction errors. The solid blue line shows the value that contained 68% of the data at the time indicated. The solid black line shows containment of 95% of the data and the solid red line shows containment of 99.9% of the error values. The histogram for the proposed message format including quantization looks virtually identical. However, the corresponding histogram for the existing L1 MOPS approach is significantly worse.

Table 1 shows the average one-sigma clock and clock rate errors for the three prediction methods. Also shown is the overbounding sigma value for the worst performing satellite. Note that there is almost no difference between the unquantized predictions and the proposed approach. However, the current L1 MOPS has a substantially larger error growth rate for the reasons mentioned above.

1-σ Values	No Quant.	Proposed	L1 MOPS
Average	2.9 cm + 0.03 cm/s	3.0 cm + 0.03 cm/s	3.6 cm + 0.53 cm/s
Worst SV overbound	5.3 cm + 0.075 cm/s	5.3 cm + 0.078 cm/s	4 cm + 1.4 cm/s

Table 1. *The average one-sigma clock and clock rate errors for the three prediction methods*

It should be noted that the prediction method used here is based on very low noise, precise, IGS clock estimates. The operational SBASs will have noticeably larger uncertainty in the raw data used to generate the fit. However, using a somewhat longer data set to perform the linear fit may offset this effect. The raw clock estimates will be a strong function of the number of observing reference stations and the amount of carrier smoothing that has been performed. Well observed satellites, should not have more than about 10 cm of error. Perhaps using 1000 instead 360 seconds will offset the effect of the increased noise. This needs to be investigated using actual SBAS measurement data.

The results indicated that there was a difference in performance between the different satellites. This is a well known result [5] that is due to the different clock technologies implemented on the satellites. Figure 6 shows the 95% prediction accuracy after 30 and 120 seconds for each of the GPS satellites. The satellites are divided by generation and those that employ cesium clocks vs. rubidium clocks are highlighted. It is obvious that within each block there is similar performance for similar clock technology. The rubidium clocks perform much better than the cesium clocks. It is promising to

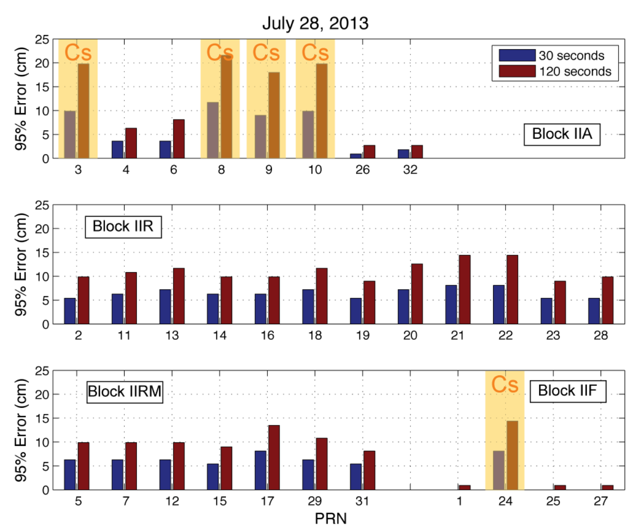

Figure 6. *95% prediction accuracy for the GPS satellites after 30 and 120 seconds.*

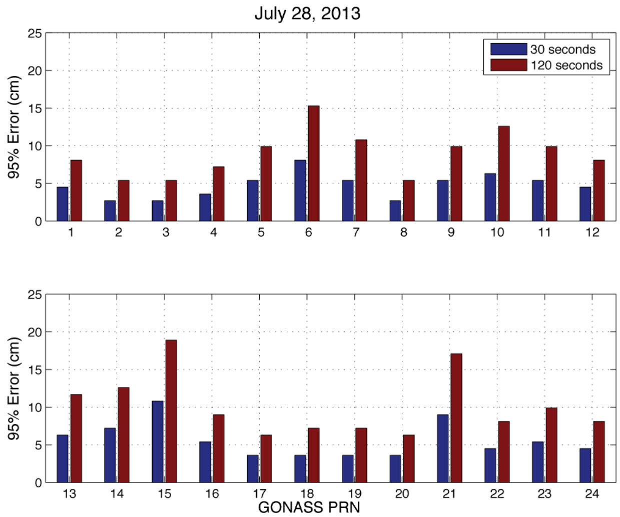

Figure 7. *95% prediction accuracy for the GLONASS satellites after 30 and 120 seconds.*

note that the newest Block IIF satellites, which are the first to include the new L5 signal, have substantially smaller errors and error growth when using rubidium clocks.

In addition to evaluating GPS performance, we obtained 30-second precise clock information from IGS for the GLONASS constellation. We found that the clock performance was comparable to the cesium clocks in the GPS constellation. The 95% values after 30 seconds were all below 11 cm and after 120 seconds were all below 20 cm (see Figure 7). All of the GLONASS satellites had smaller prediction errors than the worst performing GPS satellite.

Thus, a starting value for the bound on the prediction accuracy would be an initial value of 5.3 cm with an increase in the uncertainty of 0.08 cm/second over the time of validity of the message. The optimal estimation strategy for the operational systems remains to be determined. However, it is likely that SBASs will have higher noise on their raw clock estimates and therefore a larger increase in uncertainty over time. In the next section, we will assume that a one-sigma error growth term of 0.2 cm/second is achievable.

EFFECT OF CLOCK UNCERTAINTY

The proposed approach for the L5 MOPS [1] eliminates the fast corrections. Consequently, there is a longer period of time between clock correction updates to the user. The confidence bound sent to the user needs to reflect this increased uncertainty. This can be done in one of several ways. The uncertainty at the end of life of the message can be incorporated into the dual frequency confidence bound, Dual Frequency Range Error (DFRE) from the outset. The uncertainty at the end of the update period can be included in the initial value and the degradation parameters can be set to increase the sigma term after each update period. Finally, the DFRE can be set to a value corresponding to the initial broadcast time and a linear uncertainty can be included in the degradation parameters to continuously increase the user overbounding sigma.

We recommend the last approach, as it would provide the minimum value to the user at all times. However, the second approach may be slightly simpler to analyze, as the DFRE values are constant for the user who has not missed messages. We will now look at the effect of the growing clock uncertainty at the end of the update period compared to the current L1-only MOPS.

The current MOPS has clock updates every six seconds. Table 1 shows that the overbound growth term is 1.4 cm/sec. Therefore, after six seconds the uncertainty will need to be increased by 8.4 cm. The current MOPS enforces an even larger growth in the error as a six second update mandates an acceleration growth term of either 0.46 cm/sec^2 or 0.58 cm/sec^2. A linear growth term is also included, but in most situations this term is currently set to zero. The fast correction degradation term is described in Appendix A of the MOPS as [2]:

$$\varepsilon_{fc} = a\left(t - t_u + t_{lat}\right)^2 / 2 \quad (1)$$

where a is the acceleration degradation factor, t is the current time, t_u is the time of the update, and t_{lat} is the latency in the system. Using the minimum degradation value and a latency of four seconds results in a degradation value six seconds after the update of 23 cm.

Our previous analysis [1] showed that by using a fixed message scheduler for the L5 messages in combination with not broadcasting ionospheric corrections, we could send satellite corrections every 36 seconds for one constellation, every 60 seconds for two constellations, every 90 seconds for three constellations, and every 120 seconds for four constellations. Using the more conservative growth term for the proposed approach of 0.2 cm/seconds results in increased uncertainty values of 7.2 cm, 12 cm, 18 cm, and 24 cm at the end of the update periods, respectively, for one to four constellations. Thus, we can see that for one or two constellations there is a noticeable reduction in this uncertainty term at the end of the update interval. The current MOPS is not capable of

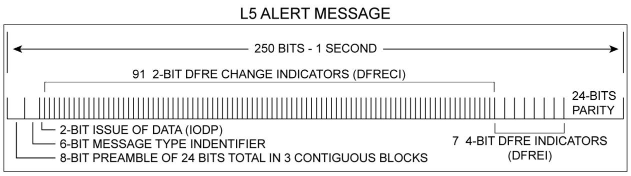

Figure 8. *Proposed L1/L5 Service alert / integrity update message.*

adapting to more than two corrected constellations without making substantive changes, so it is difficult to compare it to the proposal for the three and four constellation cases.

USE OF THE ALERT MESSAGE

One of the key aspects of our proposal is a single alert message that can update the integrity status of 91 satellites simultaneously. As previously described [1], the advantage is that this information, which must be broadcast every six seconds, can be sent using only a single message. The trade-off is that this message cannot update the status of every satellite to any arbitrary state. The updated information is more constrained. Figure 8 shows the content of the message. It includes 91 DFRE Change Indicators (DFRECIs). The two-bit DFRECI indicates one of four states:

- DFREI is unchanged,
- DFREI is changed to a numerical value,
- The satellite is not monitored (NM), or
- The satellite should not be used (DNU).

If the DFREI has changed to a numerical value (and not to NM or DNU), the new numerical value is indicated in one of the seven four-bit DFREI values at the end of the message.

If more than seven satellites should be changed to a new numerical value, there could be a loss of performance with this new message. There is no problem regarding integrity because every satellite can be indicated NM or DNU if required. However, there could be a loss of performance. Satellites whose DFRE has decreased may have a delay before that information is sent to the user. Alternatively, additional satellites whose DFREs need to be increased may instead have to be set to NM or DNU instead of a higher numerical value that could otherwise support the intended operations.

There are two situations that lead to changes to the required DFRE values: nominal changes due to satellites rising and setting within the reference network; and unexpected changes due to a sudden loss of observability or failure in either the SBAS or the core constellations. We will examine both in more detail.

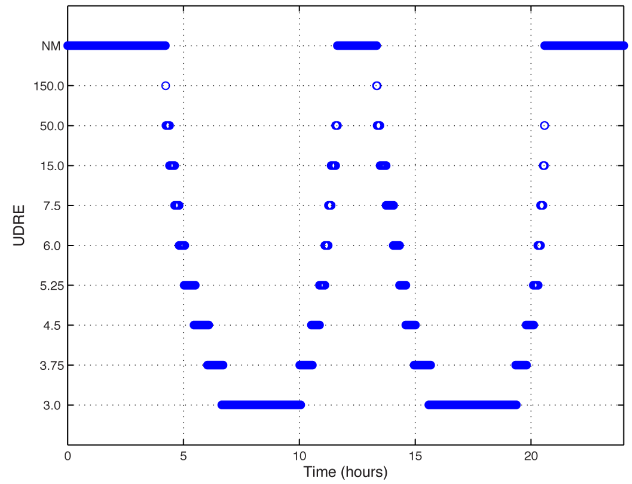

Figure 9. *Time history UDRE changes for a particular satellite observed by WAAS on a nominal day.*

Nominally the satellites come into view of the observation network and are gradually seen by more and more stations. After some time at maximum observability, they begin to set within the system, seeing fewer and fewer stations until they no longer meet a minimum observability threshold. This situation is repeats itself day after day, subject to variations of the satellite's orbit.

Figure 9 shows an example of typical satellite behavior for the current L1 service under WAAS. The satellite will be out of view, and then come into view with a large UDRE, reach a minimum value and then have its UDRE increase until it goes out of view again. This example has two visible passes in one day.

We ran a simulation using the existing WAAS UDRE algorithm. We first simulated having a single constellation of 31 satellites, and then added additional constellations of 27 satellites each until we ended up with four constellations with a total of 102 satellites in orbit. We neglected geostationary satellites, as nominally their DFRE values should not change over time. Figure 10 shows histograms for each of the four cases describing the percentage of time within a 30-second window that a certain number of DFRE values changed to a new numerical value (we did not count transitions to NM). As can be seen, the vast majority of the time, a single constellation will have no more than three expected changes in 30 seconds. At most we saw seven nominal changes with four constellations in operation. However, this only occurred less than 0.07% of the time.

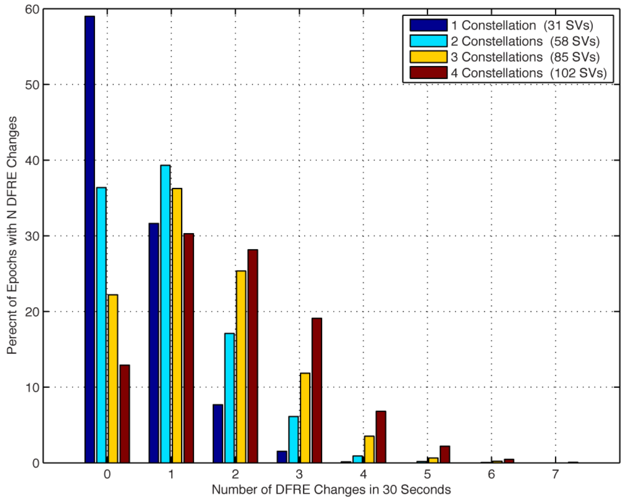

Figure 10. *Histogram of the number of DFRE changes in a 30-second window for one to four constellations.*

The thirty-second window was chosen because the DFRE needs to be broadcast every six seconds. This allows five opportunities to update the DFRE. The specified message loss rate is less than one in a thousand under the worst operating conditions. Assuming that the loss for messages spaced apart by six seconds is independent, there should be below a 10^{-15} probability of losing five consecutive updates. There may be operational conditions, such as banking away from the SBAS satellite that could lead to a prolonged outage (at least 25 seconds to drop five consecutive updates). If necessary an even longer period may be chosen.

This period should be evaluated further to make sure that it is sufficiently long to support operational requirements. This is particularly important because if the user does miss five consecutive DFRE updates they will no longer be able to tell if a DFRECI indication of no change refers to a DFRE value that they have received. They may need to stop using that satellite until they receive a full update of the DFRE. For four constellations this wait could be up to 120 seconds to receive the full DFRE for every satellite (assuming no more missed messages). A 25 second outage is a fairly long outage and one should not necessarily expect to immediately resume service after such. However, this should be discussed with the community at large to ensure that such a limitation is acceptable.

The nominal transitions to higher DFREs should not require any alerting. Appropriate selection of the degradation parameters will cause the applied confidence sigmas to increase as required over time. Therefore, there should be no need to repeat the DFRE increases in the six-second alert message updates. In the event of a nominal DFRE increase, the no change indication will signify that the most recent satellite correction with the increased DFRE is valid and the older correction with the smaller DFRE that has been increased by the degradation parameters is also valid. The error growth terms identified in the previous section can be set to handle more than 99.9% of nominal DFRE increases. Therefore, half of the transitions identified in Figure 10 do not need to be indicated with a full DFRE. Instead of a maximum of seven nominal transitions in a 30 second window, we only have a maximum of five that decrease within that window. Further indicating decreases in DFRE is a matter of performance not integrity. Once a decrease has been indicated once or twice, it is safe to resume an indication of no change. The user either has the newer smaller DFRE or the older larger DFRE that is still safe. The vast majority of users will have received the update after one or two alert messages. Therefore nominal transitions should easily work with the seven available slots to update the full DFRE.

In addition to nominal changes, there may also be true alerts due to unexpected conditions. These could involve loss of data or errors either in the SBAS operation or in the core constellations. A loss of all or nearly all data is well within the capabilities of the proposed message as all satellites can be set to NM via the DFRECIs. Similarly, a raft of integrity failures on multiple satellites can also be accommodated as any or all satellites can be set to DNU. What cannot be handled as gracefully is a change of conditions that leads to more than seven satellites needing to have their DFREs increased, but not to the point of being NM or DNU.

Such a situation would be a very exceptional event and has not been observed in the operational lifetime of WAAS. Observed GPS failures to date have only affected one satellite at a time. Potential threats that affect multiple satellites primarily come through updates to the broadcast navigation messages. Since the use of navigation messages is controlled by the IODE contained in the SBAS correction message, these could either be delayed to avoid overlap or avoided altogether. However, it may be better to discontinue use of a constellation that is experiencing multiple simultaneous faults.

There is the potential that the SBAS system itself may be the source of multiple simultaneous increases. WAAS has been designed to avoid such a situation. The reliability of its individual components is kept very high and each component is backed up by redundant elements. Thus, the loss of data from an individual receiver is already uncommon, but the loss from two co-located receivers is extremely rare. However, even in this rare event, the redundancy of multiple reference stations prevents the loss of any one from creating a sudden change in many satellites. Over its history WAAS has not had cases of multiple satellite alerts [6]. However, it is understood that other systems may operate differently and the possibility of reducing the number of simultaneous corrections to create additional DFRE slots in the alert message should be internationally coordinated.

BENEFITS OF MULTIPLE CONSTELLATIONS

The current L1 MOPS already has mechanisms to support 51 simultaneous corrections that could be made to apply to two constellations. If there were no intention to ever broadcast corrections to more than two constellations, it would make sense to more closely follow the current structure. However, there is desire on the part of the countries fielding their core constellations to have those satellites used in their own airspace. It is quite likely that future aviation receivers will need to be able to process each available constellation. Thus, there will be tracking channels capable of processing any of the four core constellation's satellites. As we will show in this section, there is continued benefit to processing more and more satellites. Conventional wisdom to date has indicated that the benefit diminishes as the number of received satellites increases. While this is true, it is also true that there still is significant benefit to simultaneously using three or four constellations.

We implemented a set of simulations using our Matlab Availability Analysis Simulation Toolset (MAAST) [7]. MAAST accurately emulates the integrity algorithms in WAAS and can predict availability under different conditions. We used the current WAAS reference network and implemented a recently proposed covariance based DFRE algorithm approach [8] to determine the DFRE values to be broadcast given a set of almanacs for the different constellations. For the protection calculation formulation, we set the individual line of sight confidences by the following formulation:

$$\sigma_i^2 = \sigma_{clk\_orb,i}^2 + \sigma_{trop,i}^2 + \sigma_{air\_iono\_free,i}^2 \qquad (2)$$

where σ_i^2 is the total variance for the line of sight; $\sigma_{clk\_orb,i}^2$ is the combination of the DFRE, the δDFRE from the MT28-like parameters, and the degradation terms; $\sigma_{trop,i}^2$ is the tropospheric bounding variance as in Appendix A of the MOPS [2]; and $\sigma_{air\_iono\_free,i}^2$ is the iono-free combination of the airborne noise and multipath bounding variance. Due to the potentially large number of satellites in the sky, we added off-diagonal covariance terms to the weighting matrix **W** to account for tropospheric correlations. These took the form of $\sigma_{TVE}^2 \cdot m(E_i) \cdot m(E_j)$ where σ_{TVE}^2 is the bounding vertical tropospheric value, and $m(E)$ is the elevation angle based mapping function, both specified in the MOPS [2]. We also added a separate clock state to the observation matrix, **G**, for each constellation. Therefore, the position-domain covariance matrix, $\left(\mathbf{G}^T \cdot \mathbf{W} \cdot \mathbf{G}\right)^{-1}$, is 4x4 for one constellation, 5x5 for two constellations, continuing up to 7x7 for four constellations.

We also applied a degradation parameter of 0.2 cm/sec to the $\sigma_{clk\_orb,i}^2$ term. The product of the update interval and this one-sigma growth bound was added to the product of DFRE and δDFRE. The sum of these two products is then squared to find $\sigma_{clk\_orb,i}^2$. The update intervals used were 36 seconds for one constellation, 60 seconds for two, 90 seconds for three, and 120 seconds for four. The square root of the third diagonal element (in an East, North, Up frame) is the overbounding variance on the vertical position error. The VPL was calculated by multiplying the square root of this variance by 5.33. We did not apply any additional nominal bias terms as have been described in earlier papers. Instead, we assumed that as in today's systems, the nominal biases could be minimized, and bounded by inflating the sigma terms. The HPL was similarly calculated using an analogous equation as can be found in Appendix J of the MOPS [2].

To simulate the constellations, we used almanacs containing 24 satellites each for GPS, GLONASS, Galileo, and Beidou. We simulated users on a 2° x 2° global grid and ran 300 time steps over 10 sidereal days. At each time step and for each user, the geometry matrix, **G**, weighting matrix, **W**, VPL, and HPL were calculated. These were used to determine availability at each grid point on Earth and then to determine coverage maps.

Figures 11 – 14 show the VPL maps for one, two, three, and four constellations, respectively. As can be seen, the VPLs drop most significantly going from one to two constellations. The figures show 99.5% VPL, that is, the value of VPL at each location that the user was at or below 99.5% of the time. This measure is very sensitive to poor geometries. Going from one to two constellations ensures that there never are any poor geometries. However, as Figures 13 and 14 demonstrate, there continue to be benefits going to three and four constellations. Although there are no large gaps in the sky to lead to bad DOPs, the extra satellites provide averaging against satellite and local multipath errors. Thus, doubling the number of satellites in the sky can lead to almost a square root of two reduction in the overall VPLs and HPLs.

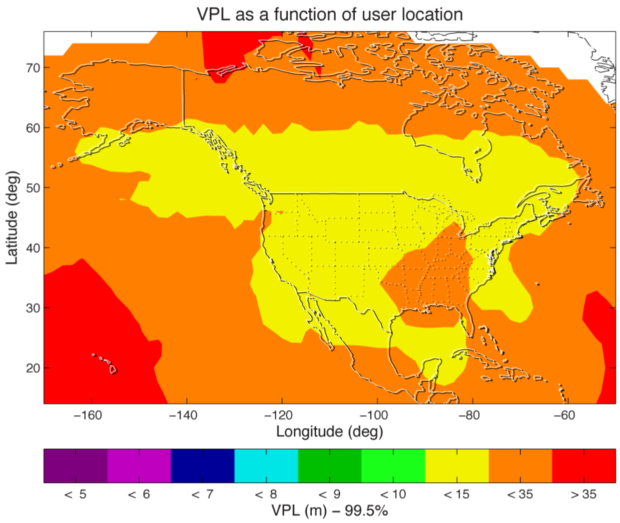

Figure 11. *99.5% VPL for a single constellation of 24 GPS satellites*

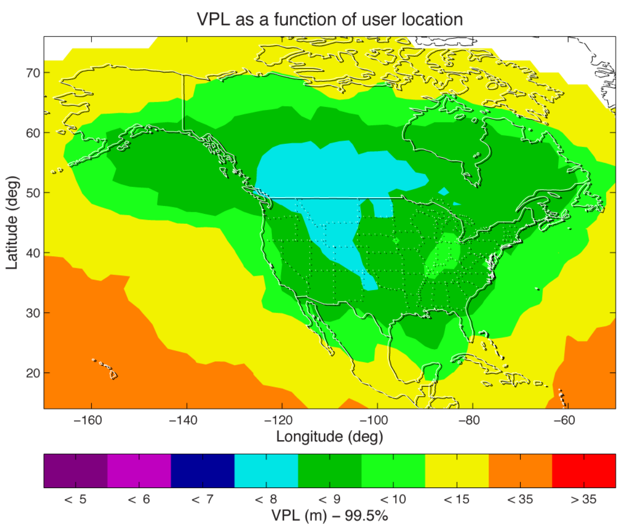

Figure 12. *99.5% VPL for two constellations of 24 satellites each*

Currently, LPV-200 [9] is the most demanding level of service provided by SBAS. However, this level of service had not been envisioned when WAAS was first commissioned. At that time, LPV was the most demanding service planned. It was observed that WAAS was capable of providing lower VPLs, and therefore a new and better service level was created. We anticipate that a similar path may be taken with dual-frequency SBAS. Figures 11-14 clearly show that much lower VPLs can be achieved if additional constellations are corrected. Therefore, new operations and procedures may be created to exploit lower VALs and HALs. A 10 m VAL easily supports an autoland procedure. Even more

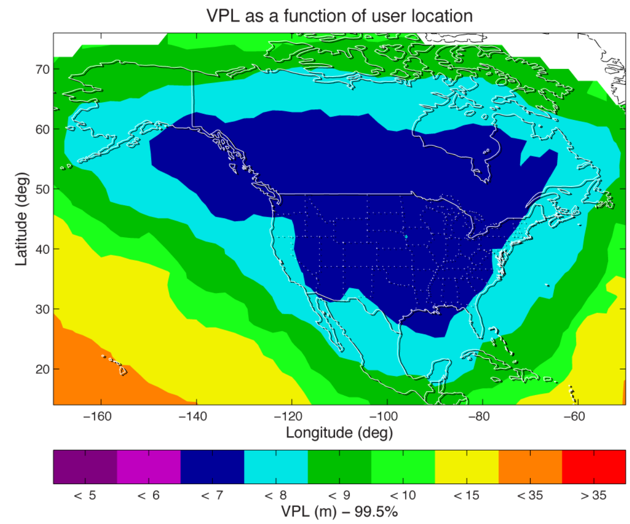

Figure 13. *99.5% VPL for three constellations of 24 satellites each*

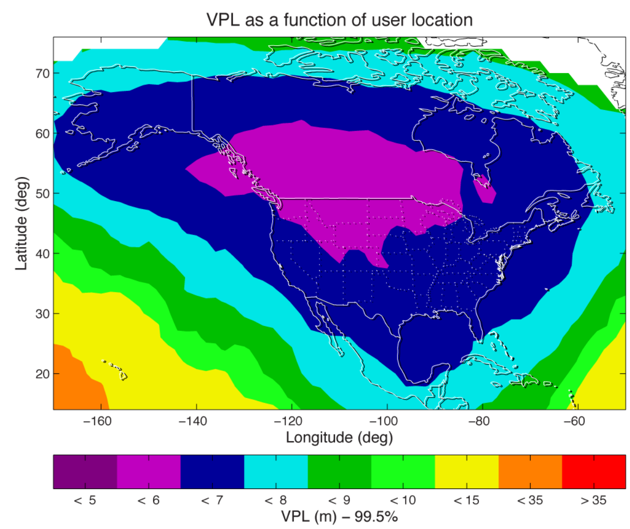

Figure 14. *99.5% VPL for four constellations of 24 satellites each*

demanding operations could be supported with a 7 m VAL.

Figures 15-18 show the corresponding 99.5% HPLs for the same four constellation scenarios. Again, the largest improvement comes from the transition from one to two constellations. However, there are still significant improvements going to three and then four constellations. Perhaps there will be very demanding surface operations that could be met with a four constellation SBAS. What has been simulated here assumes four constellations of 24 satellites each. GPS currently has 31 active satellites. Galileo and Beidou are planning at least 27 satellites on orbit. Therefore it is possible that service will be even better than what is projected here.

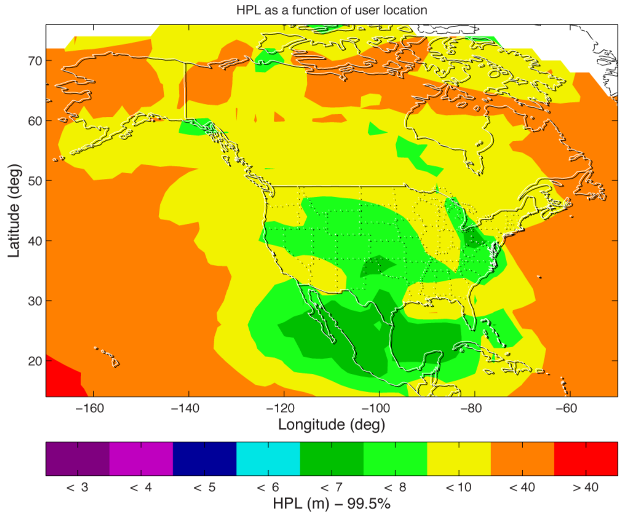

Figure 15. *99.5% HPL for a single constellation of 24 GPS satellites*

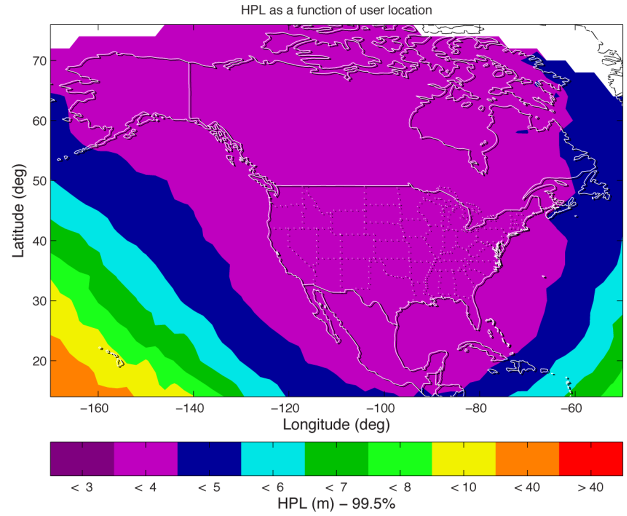

Figure 17. *99.5% HPL for three constellations of 24 satellites each*

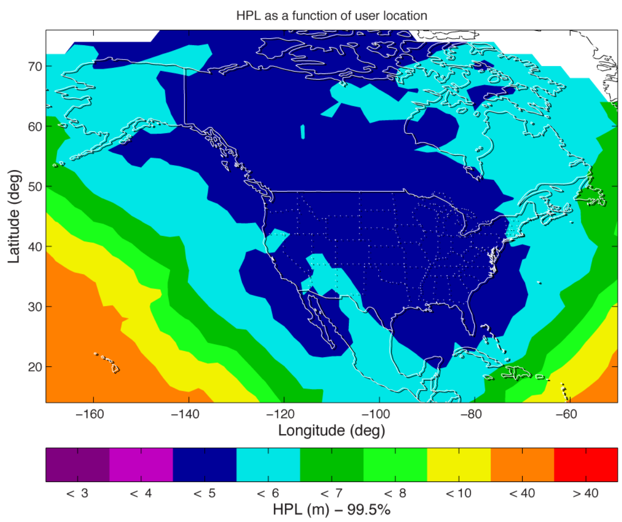

Figure 16. *99.5% HPL for two constellations of 24 satellites each*

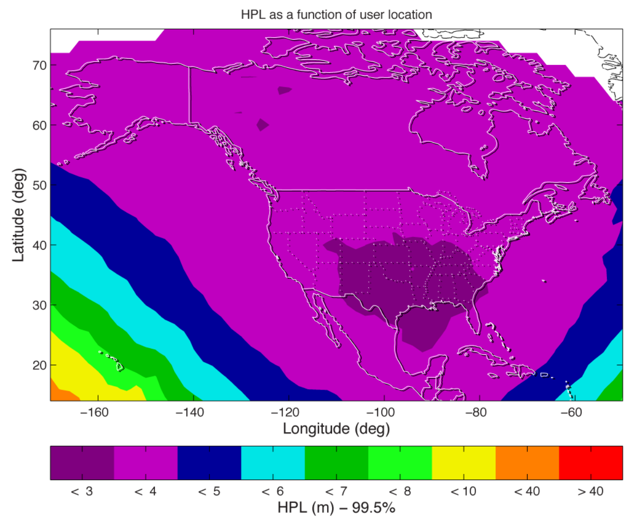

Figure 18. *99.5% HPL for four constellations of 24 satellites each*

At the moment, this performance is speculative. It will need to be verified experimentally that the actual accuracies reduce by comparable levels as are simulated here. However, it is clear that with many more satellites the dependency on each individual ranging source decreases. It will take larger individual errors or many combined errors to create positioning errors that threaten these protection levels. It is also possible that with so many ranging sources there will be correlation among the multipath errors that will also need to be taken into account. GPS and GLONASS are already operational. Galileo and Beidou have partial constellations. We are rapidly approaching the point where these algorithms can be prototyped and operationally evaluated.

CONCLUSIONS

We previously proposed a new message structure that can be expanded to provide 91simultaneous corrections allowing SBAS use of four full constellations and supporting geostationary and other SBAS satellites [1]. We demonstrate in this paper that some of the concerns raised by the proposed methodology do not create significant limitations to performance. Specifically, we investigated satellite clock performance and verified that it is extremely predictable over the several minutes that a user may be applying a correction message. We saw no significant penalty from eliminating fast corrections and having longer time intervals between successive updates. The reduced quantization steps more than offset this error term, leading to better accuracy and availability than is offered by the current L1-only method.

Next, we examined limitations in the proposed alerting message. We found that it achieved its target of updating 90 integrity values and worked with all situations that have been encountered by WAAS over its ten years of operation. It is possible that other systems may have more difficulty working with the proposed approach and so it is still under evaluation.

Finally, we demonstrated the benefit of creating a message structure that can support four full constellations. Although it remains to be seen whether SBAS service providers elect to augment more than two constellations, it has been decided that there is value in creating a message structure that can support such an approach. We have further demonstrated that correcting more than two of the core GNSS constellations can significantly reduce the protection levels and open the doors to improved levels of service.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the FAA Satellite Product Team for supporting this work under Cooperative Agreement 2012-G-003. However, the opinions expressed in this paper are the author's and this paper does not represent a government position on the future evolution of SBAS.

REFERENCES

[1] Walter, T., Blanch, J., and Enge, P., "L1/L5 SBAS MOPS to Support Multiple Constellations," Proceedings of the 25th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2012), Nashville, TN, September 2012, pp. 1287-1297.

[2] RTCA, "Minimum Operational Performance Standards for Global Positioning System/Satellite-Based Augmentation System Airborne Equipment," RTCA publication DO-229D Change 1, 2013.

[3] van Graas, F. and Braasch, M., "Selective Availability," Chapter 17 of Global Positioning System: Theory and Applications Vol I, edited by Parkinson and Spilker, AIAA Blue Book, 1996.

[4] Dow, J.M., Neilan, R. E., and Rizos, C., The International GNSS Service in a changing landscape of Global Navigation Satellite Systems, Journal of Geodesy (2009) 83:191–198.

[5] Senior, K. L., Ray, J. R., and Beard R. L., "Characterization of periodic variations in the GPS satellite clocks," GPS Solutions, (2008) 12:211–225.

[6] Tim Schempp, private communication.

[7] Jan, S.S., Chan, W., Walter, T., and Enge, P., "Matlab Simulation Toolset for SBAS Availability Analysis," Proceedings of the 14th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GPS 2001), Salt Lake City, UT, September 2001, pp. 2366-2375.

[8] Blanch, J., Walter, T., and Enge, P., "A Clock and Ephemeris Algorithm for Dual Frequency SBAS," Proceedings of the 24th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2011), Portland, OR, September 2011.

[9] Lawrence, D., Bunce, D., Mathur, N. G., and Sigler, C. E., "Wide Area Augmentation System (WAAS) - Program Status," Proceedings of the 20th International Technical Meeting of The Satellite Division of the Institute of Navigation (ION GNSS 2007), Fort Worth, TX, September 2007.