Feasibility of Managing a Dynamic Constellation to a Fixed Constellation Definition

Bruce P. Ayati, *The Aerospace Corporation* Paul D. Massatt, *The Aerospace Corporation*

BIOGRAPHIES

Dr. Bruce P. Ayati is a Senior Member of the Technical Staff at The Aerospace Corporation in El Segundo, CA. He provides technical support for GPS constellation management. He received his Ph.D. through the Computational and Applied Mathematics Program of The University of Chicago.

Dr. Paul D. Massatt is a Senior Engineering Specialist at The Aerospace Corporation in El Segundo, CA. He designed the current GPS constellation and provides technical support for managing the constellation. He received his Ph.D. in Applied Mathematics from Brown University.

ABSTRACT

The GPS satellite constellation is dynamic. Earth gravitational forces cause orbits to vary, and solar and lunar gravitational forces produce out-of-plane perturbations. Launch dispersions often result in orbits that are not nominal. Even the number of satellites in the constellation can vary over time given anticipated and actual satellite failures and routine maintenance. The GPS satellite constellation is operationally managed to sustain the global-averaged availability of 3D-navigation coverage by routinely moving satellites within the constellation.

Long-term planners and builders of GPS augmentation systems often use fixed “orbital planes and slots” as defined in various GPS specifications. These reference orbits were meant to inform contractors that build GPS satellites, not users of the GPS, of what is needed in the way of satellite capability. These documents were not intended to convey that actual GPS satellites are in fixed orbits.

In this paper, we discuss the technical balance between operational flexibility and the demands of long-term planners. We show the sensitivity of positioning dilution of precision to perturbations in right ascension of ascending node, mean anomaly and inclination. Lastly, we give examples of how a flexible constellation structure can be adjusted to preserve and manage GPS coverage in a dynamic environment.

INTRODUCTION

The GPS constellation is designed to provide continuous global navigation coverage when the constellation is completely filled, to have little loss of coverage when one satellite is not operational, and to minimize the impact of multiple satellite failures. The initial 24 GPS satellites, built by Rockwell (since acquired by Boeing), performed beyond design and expectation with the result that for the past eight years the navigation community has enjoyed nearly continuous uninterrupted global coverage. Initial expectations that the constellation would be down at least one satellite 30% of the time and down at least two satellites 6% of the time did not materialize during the first eight years of the constellation since satellite failures were rare. However, these expectations were based on the longer-range assumption that GPS would be managed in a steady-state environment where the failure rate was equal to the launch rate. Hence, they were not applicable to the early years of the constellation when all of the satellites were relatively new and failures were unusual.

Current user expectations of GPS navigation service are predicated on the assumption that GPS will continue to provide the same navigation coverage that it has for the last eight years. Furthermore, various institutions, both military and civil, private and public, domestic and international, would like to build systems that rely upon the continued availability of GPS with at least the same performance that has been historically provided. In addition, organizations often prefer to compute their own coverage needs using different metrics than GPS was designed to support. Hence, many organizations would like GPS to provide guarantees that they will continue to operate the same constellation with the same frequency of having the constellation completely filled or down one satellite. Unfortunately, these expectations are not realistic and exceed the performance that GPS is designed and funded to maintain. It is simply unrealistic to expect that a constellation filled with older satellites that fail (unpredictably) as often as they are replaced will perform as well as a constellation filled with newer satellites that rarely fail.

In order to maintain the performance of the 24 satellite constellation when the failure rate is equal to the replenishment rate, the size of the constellation would need to be increased to at least 27 satellites with all 27 satellites managed to best support the navigation mission. Even with this increase, past performance cannot be guaranteed. Satellite failures and satellite longevity are unpredictable until enough failures of parts and satellites occur in a given block of satellites to enable predictions for other satellites. Since procurement schedules and launch planning must occur many years prior to the availability of such data, there is a risk as to how many resources will be available for launch replenishment. In 1996, Secretary of Air Force Widnall directed the Air Force to maintain the 24 satellite constellation. This has been interpreted from a procurement sense to maintain 24 on-orbit satellites with 95% availability. Such a management scheme gives a 50% expectation that there will be 27 satellites on orbit. Hence, it may be possible to maintain past performance if the 27 satellites are positioned wisely. However, it does not guarantee past performance since there is a 5% chance of falling below 24 satellites and since there is no guarantee that failures will occur in the same orbit planes where spare satellites are located.

In order to effectively utilize the satellite resources available to GPS to maintain nearly continuous global coverage, the constellation must be understood and managed for what it really is, a dynamic entity. While a new satellite constellation can be thought of and managed as a static constellation, an older constellation should not be viewed in this way. Satellites fail and replacement satellites must often be launched prior to actual failures. Hence the number of satellites available will vary over time according to both failure and replenishment rates. Launches generally cannot occur on demand as there is significant competition for launch pad availability. Instead, launches need to occur before actual failures, and additional satellites need to be incorporated into the constellation structure in order to increase the robustness of the constellation to potential unpredicted failures. Satellite orbits also degrade over time. Solar and lunar gravitational forces alter the inclination of satellites orbits resulting in non-uniform nodal regression rates that pull the right ascensions of orbit planes off by as much as six degrees. In addition to these effects, anomalies may occur that will cause a satellite's orbit to be impaired, yet the satellite will still be functional and useful.

With funding limited and both launch and satellite procurements completed for at least the next eight years, it is better to manage the constellation dynamically rather than waste satellite resources. This means that people who rely upon GPS will need to know with greater fidelity what to expect from GPS as it undergoes a transition from a relatively new constellation where failures and launches are rare to a steady-state performance where failures and launches are more frequent.

The nominal constellation is useful as a basis for user perception of GPS coverage. The user should be aware that satellite drift affects these results and that spare satellites are routinely used to improve the coverage and robustness so that the effects of this degradation are rarely seen. If the constellation is managed by slots, these effects could be exaggerated. Moreover, unpredictable failures make maintaining constellation slots impractical. The first portion of the paper addresses the coverage users might experience from a fixed constellation definition and from typical orbit degradations. The second section addresses some of the options available to the Air Force for managing the constellation in order to provide as much coverage and robustness to failure as possible for all classes of users.

ORBIT PERTURBATION ANALYSES

GPS users need to know what service to expect and they have varying requirements. GPS cannot be optimized or managed for all users simultaneously. One way to address the coverage needs of a myriad of users is to define GPS as a constellation of satellites that provide signals. By defining the location of the satellites and their availability, users can themselves assess the level of service they are provided. The primary difficulty of this approach is that GPS orbits are not static. Tolerances for orbit deviations must be specified before users can determine their coverage. Typical deviations are $\pm 6^\circ$ in right ascension of the ascending node (RAAN) and $\pm 2^\circ$ in inclination. However, anomalies exist that can exceed these variations. The slot definition for each satellite must be lax.

We conduct an orbit perturbation analysis of the designed 24-satellite constellation to determine the utility of this concept for users. This analysis is based on a full constellation. Both average and worst case constellation values are demonstrated. A second analysis provides a glimpse at what can be expected in terms of actual orbit drift due to gravitational forces.

In the first analysis, we randomly perturb the RAAN of each satellite, according to a uniform distribution over each ”slot”, and maintain station keeping in mean anomaly (MA). The simulations are done for various values of the mask elevation angle. Inclination perturbations are ignored because they are small and insignificant.

In the second analysis, we perturb the RAAN and inclination of satellites according to an orbit deviation model. In this model, the deviations are functions of the time since launch of each satellite. Time since launch is determined randomly via a joint cumulative distribution function of the probability of satellite failure due to wear-out and the probability of electronic failure. Start epochs of the simulations are also determined randomly. Station keeping adjusts the reference longitude of the ascending node (LAN) to the RAAN dispersion so that relative differences of satellite MA are preserved.

SLOT SIMULATIONS

Deviations from an orbital position can vary for several reasons, such as launch injection and gravitational forces. We define a fixed slot about an orbital position and deviate the satellite within this slot. Because satellite RAAN may drift as much as $\pm 8°$ due to gravitational forces and an additional $\pm 3°$ from launch injection, a slot definition in RAAN may need to be as large as $\pm 11°$. If the constellation is managed to slot definitions, satellites positions could be at the extremes of a slot more frequently than they currently are.

RAAN is varied according to a uniform probability distribution within the slot while the variation of LAN follows a parabolic distribution. The uniform variance of RAAN simulates the users perspective of RAAN deviations within a fixed constellation definition. The results presented provide an expectation of what may be typical in terms of coverage as well as a view of how poorly coverage can degrade at an extreme.

We present Monte Carlo simulations of PDOP<6 availability of the nominal 24 satellite GPS constellation over a 24 hour period. We incorporate dispersions in RAAN and station keeping in MA. The RAAN dispersions are chosen from a uniform distribution in intervals of $\pm 2°, \pm 6°, \pm 12°$. The station keeping in MA is done within bounds of $\pm 4°, \pm 8°$. We consider PDOP<6 availability for mask elevation angles of $2°, 5°, 10°, 15°$ using best four satellite combinations.

The quantities we investigate are the global average coverage (the average over all grid points and all Monte Carlo runs), the worst case location coverage (worst coverage over all grid points and all Monte Carlo runs), and the maximum gap (longest continuous outage over all runs and all spatial grid points.) The data in each case is taken from 100 Monte Carlo runs. The grid used is $2° \times 2°$ in space and 10 seconds in time.

Figures 1-3 give the results of the simulations. A baseline simulation (BASE) of no dispersion in RAAN and MA is given in each table for each value of the mask elevation angle.

For the case of global average coverage (Figure 1), a loss of coverage of 10^{-4} corresponds roughly to a combined outage of 20 minutes over 1 million square miles. As a point of reference, the size of the continental United States is approximately 3 million square miles.

In all but the $2°$ mask elevation angle case, increasing the tolerances strongly degrades coverage. Slot definitions in a fixed constellation that are large enough to allow for all likely orbit deviations will be too large to guarantee effective coverage.

ORBIT DEVIATION MODEL

In this analysis we investigate the effects that gravitational forces alone have on orbit deviations and the resulting loss of coverage. The Earth's gravitational bulge causes a nodal regression of RAAN as a function of inclination. The Sun's gravity and the Moon's gravity cause small changes in inclination and RAAN. Small changes in inclination result in nodal regression rate variations that can produce RAAN dispersions as large as eight degrees over a satellites life. Variations in the gravitational field of fixed ground tracks produce satellite in-orbit drift that must be corrected through the occasional use of satellite thrusters. This produces a parabolic movement in time of satellite LAN.

The model indicates that satellites may have RAAN deviations as large as $8°$. This would require a broadening of any slot definition. However, because the proportion of these outliers are small, coverage is not strongly degraded.

We estimate the effects of RAAN and inclination drift of satellites on PDOP< 6 availability. Our approach is to utilize a nominal cumulative distribution function (cdf) for a Block IIR satellite to get a distribution of time since launch for each satellite in the constellation. The start epoch for the simulation is randomly generated so that the distribution will not be tied to the current satellites. We then feed this distribution, combined with analytical representations of RAAN drift and inclination over time, into a Monte Carlo simulation to get snapshots of the constellation configuration. We discuss the derivation of the model for determining RAAN and inclination dispersions. The model will then be used to determine steady-state coverage variability as a function of satellite life. A variation of this model that assumes no satellite is younger than eight years is used to assess the potential impact of having all the satellites launched close together. This was the case when the first GPS Block II satellites were launched.

To approximate RAAN drift of each satellite, we use a first order approximation in the oblateness parameter of RAAN as a function of inclination [3],

$$\frac{d\Omega}{dt} = K(\cos i)°/\text{day},$$

where Ω is RAAN and i is inclination. The constant

$$K = \frac{-9.9639}{(1-\epsilon^2)^2}\left(\frac{R}{R+\bar{h}}\right)^{3.5}$$

is dependent on $\bar{h} = \frac{h_a+h_p}{2}$ and $\epsilon = \frac{h_a-h_p}{h_a+h_p+2R}$. The equatorial radius of the earth is given by R=6378 km. The values h_a and h_p are the height above Earth at apogee and perigee of the orbit, respectively. Assuming a circular orbit with $\bar{h} = 20182$ km gives $K = -0.06761$. For RAAN drift we have the equation

$$\frac{d\Delta\Omega}{dt} = K(\cos(i) - \cos(55°))°/\text{day}.$$

A first order Taylor expansion about $55°$ gives

$$\frac{d\Delta\Omega}{dt} = -K\frac{\pi \sin(55°)°/\text{day}}{180°}(i - 55°).$$

Integrating gives

$$\Delta\Omega = K_1 \int \Delta i \, dt,$$

where $\Delta i = i - 55°$ and $K_1 = -K\frac{\pi\sin(55°)°/\text{day}}{180°} = 9.666 \times 10^{-4}\text{day}^{-1}$.

We assume that the drift in inclination for each orbit plane

		RAAN ±2°		RAAN ±6°		RAAN ±12°	
	BASE	MA ±4°	MA ±8°	MA ±4°	MA ±8°	MA ±4°	MA ±8°
2° elev	1	1	0.99998	0.99999	0.99996	0.99991	0.99981
5° elev	1	0.99997	0.99975	0.99989	0.99964	0.99945	0.99903
10° elev	0.9982	0.99731	0.9952	0.99621	0.99412	0.99279	0.99032
15° elev	0.9711	0.96709	0.96032	0.96439	0.95733	0.95413	0.94664

Figure 1: Global Average Coverage. A value of 1 indicates a coverage of 0.999995 or better.

		RAAN ±2°		RAAN ±6°		RAAN ±12°	
	BASE	MA ±4°	MA ±8°	MA ±4°	MA ±8°	MA ±4°	MA ±8°
2° elev	1	0.9885	0.9817	0.9806	0.9802	0.9753	0.9693
5° elev	0.9986	0.9876	0.975	0.9805	0.971	0.9614	0.9469
10° elev	0.9692	0.9563	0.9413	0.9384	0.9318	0.9058	0.8891
15° elev	0.9048	0.8687	0.852	0.8509	0.8372	0.8087	0.7725

Figure 2: Worst Case Location. A value of 1 indicates a coverage of 0.999995 or better.

		RAAN ±2°		RAAN ±6°		RAAN ±12°	
	BASE	MA ±4°	MA ±8°	MA ±4°	MA ±8°	MA ±4°	MA ±8°
2° elev	0	16.6	26.4	27.9	28.6	35.5	44.2
5° elev	1.95	17.8	34.1	28.1	41.8	50.3	54.8
10° elev	29.4	44.8	61	56.9	71.6	81.6	96.9
15° elev	55.8	71.3	88.5	90.2	101	131	167

Figure 3: Maximum Gap. Time is measured in minutes.

		BASE	$\mu = 114$	$\mu = 144$	$\mu = 168$	OLD
Average Availability	MA 0°	1	1	1	0.99999	0.99997
	MA ±4°	0.99996	0.99996	0.99996	0.99995	0.99992
Worst Case Availability	MA 0°	0.9998	0.9916	0.9912	0.9854	0.9876
	MA ±4°	0.9848	0.9842	0.9832	0.9829	0.9796
Maximum Gap (minutes)	MA 0°	0.29	9.23	11.3	21	17.9
	MA ±4°	21.9	22.8	24.2	24.6	29.3

Figure 4: PDOP < 6 Availability. Results are shown as a function of station keeping in mean anomaly (MA) and mean (μ months) of the normal distribution for failure due to wearout. BASE gives the availability of the nominal 24 satellite constellation with no orbit plane deviations. OLD gives the availability for the case of $\mu = 168$ with the added restriction that all satellites are at least 8 years old. A value of 1 represents availability of 0.999995 or better.

follows a sine curve with period 8980 days[1]. We take

$$\Delta i = a(\cos(\phi) - \cos(\pi t/(4490 \text{ days}) + \phi)).$$

A constant ϕ_0 is randomly chosen from a uniform distribution. We set $\phi = \phi_0$ to determine the nominal RAAN of the A plane. The nominal RAAN of orbit planes B through F are determined by a series of $60°$ phase shifts from the A plane. Then, assuming $\Delta\Omega(0) = 0$,

$$\Delta\Omega = K_2 \cos(\phi)\, t - K_3 \sin(\pi t/(4490 \text{ days}) + \phi) + K_3 \sin(\phi),$$

where $K_2 = aK_1$ and $K_3 = a\frac{4490 \text{ days}}{\pi}K_1$.

Although the model in this section does not incorporate moon effects, if we choose $a = 1.1$, then it does maintain a qualitative match with a model that does account for the moon [1].

For the Monte Carlo simulations, we use a Weibull distribution to model probability of failure of electronic systems and a normal distribution to model probability of satellite failure due to wearout. The Weibull cdf has the form $\Phi(t) = 1 - \exp(-(t/\beta)^\alpha)$. The normal distribution, Ψ, has mean μ and standard deviation σ. The function $\Upsilon = (1 - \Phi)(1 - \Psi)$ gives the probability that a satellite is still alive. We normalize Υ and invert its cumulative integral to get the distribution of satellite ages for use in a Monte Carlo procedure[2]. Time since launch is fed into the curves for $\Delta\Omega$ and Δi to obtain RAAN and inclination drift for a single satellite. This is then used with a program that computes the coverage of the entire constellation.

Sources of error and sensitivities in this analysis include the approximations used in determining RAAN drift as a function of inclination and on the use of cosine curves for relating inclination and RAAN. The assumptions made on the cdf's for satellite failure in the Monte Carlo simulation are also areas of sensitivity in this analysis. This is particularly the case with the assumptions made on satellite longevity.

The results show that global average PDOP < 6 availability is not strongly degraded by orbit plane deviations. Worst case availability and maximum gap are affected somewhat more, particularly in the case of strict station keeping in mean anomaly. However, the distribution of orbital deviations will have outliers that require a broadening of any slot definition to a point where it will not be informative. The slot definitions will need to be further broadened when the effects of launch injections are considered.

The bulk of the simulations discussed have satellite ages that are somewhat broadly distributed between young and old. The current constellation has satellites that are generally older due to having been launched near the initial deployment of the constellation. A simulation based on using only older satellites shows some further degradation of availability. In time we expect the age distribution of the constellation to broaden to include a larger proportion of younger satellites. This should result in an improvement of constellation availability.

We look at the cases of station keeping in mean MA of $\pm 4°$ and $0°$, the latter case representing very strict station keeping.

Longevity predictions for Block IIR satellites are not able to utilize prior performance of those satellites and may prove to be conservative. For the Weibull distribution, we use the specification parameters for Block IIR satellites, $\alpha = 1.52$ and $\beta = 159$ months [4]. For the normal distribution, we use a standard deviation of $\sigma = 12$ months and means of $\mu = 114$, 144 and 168 months.

The quantities we are interested in are the global average coverage (the average over all grid points and all Monte Carlo runs), the worst case location (worst coverage over all grid points and all Monte Carlo runs), and the maximum gap (longest continuous outage over all runs and all spatial grid points.) The data in each case is taken from 100 Monte Carlo runs. The grid used is $2° \times 2°$ in space and 10 seconds in time. The table in Figure 4 shows the PDOP<6 availabilities given by the simulations as a function of station keeping in mean anomaly and mean (μ) of the normal distribution for failure due to satellite wearout. We show the results for station keeping of $0°, \pm 4°$ and for five cases of orbit deviations: the cases of $\mu = 114, 144, 168$ months, the case of no orbit deviations, and the case of $\mu = 168$ months with the added restriction that no satellite is less than 8 years old.

We see that the degradation of global average availability is not substantial for either case of MA station keeping. For worst case availability and maximum gap, the degradation due to orbit deviations is more substantial in the case of strict station keeping (MA $0°$). The case of all satellites being at least 8 years old corresponds more closely to the current constellation. As expected, this case provides worse global average availability than the others. In time the age distribution of the constellation should broaden, resulting in an expected improvement in coverage.

It should be noted that the worst case availability and maximum gap are highly dependent on the number of iterations of the Monte Carlo simulation. These data will tend to vary in simulations that consist of only 100 Monte Carlo runs. For example, rerunning the $\mu = 168$ and OLD case for MA station keeping of $0°$ (Figure 4) gave worst case coverage values of 0.9891 and 0.9888, respectively, and maximum gap values of 15.7 min and 16.1 min, respectively.

Figure 5 shows the analytical distribution of RAAN and inclination deviations (assuming that time since launch is uni-

1. The figure of 8980 days is obtained by an approximation using the average daily change in RAAN, $d\bar{\Omega}/dt = 0.04009°/\text{day}$ to get $\bar{\Omega} = 0.04009°/\text{day}\, t$. Using the relationship $di/dt = a\sin(\bar{\Omega}) = a\sin((2\pi/360°)0.04009°/\text{day}\, t) = a\sin((2\pi/8980 \text{ days})\, t)$ gives the period. The cosine form of Δi is obtained from integrating $d\Delta i/dt = A\sin(\pi t/4490 \text{ days} + \phi)$ and assuming $\Delta i(0) = 0$.

2. See Appendix.

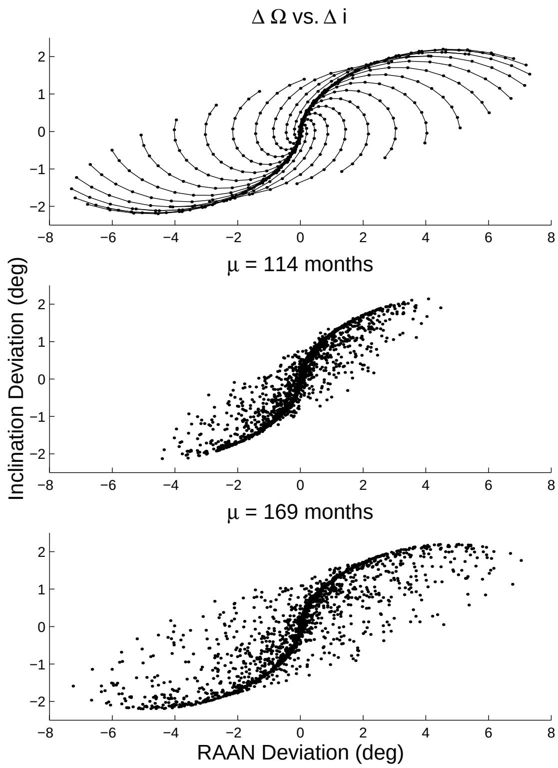

Figure 5: Scatter plot of RAAN vs. Inclination Deviation. The top graph shows parametric plots of $\Delta\Omega$ vs. Δi for different values of RAAN ($\phi = \pi/16, \pi/32, \ldots, 2\pi$). Time goes from 0 to 15 years. Each dot in the top graph signifies a one year separation from the previous dot. The middle and bottom graphs show RAAN vs. inclination deviations from Monte Carlo simulations with time since launch values determined by the joint distribution of a Weibull distribution with exponent $\alpha = 1.52$ and divisor $\beta = 159$ months, and a normal distribution with means of $\mu = 114, 168$ months and a standard deviation of 12 months.

formly distributed) in the top plot. The curves represent different values of RAAN. In the middle and bottom plots, data are shown for two cases: when time since launch was taken from the joint distribution of the Weibull and normal with mean $\mu = 114$ and $\mu = 168$. There are 2400 points total on each plot (24 satellites times 100 Monte Carlo runs.) As expected, the RAAN vs. inclination data is concentrated along the "S" curve that appears in the analytical distribution.

Simulations show that increasing the mean time to satellite wearout does not substantially increase the relative number of large orbit plane deviations. This is reflected in the minimal degradation in global average availability. The additional number of large orbit plane deviations is reflected somewhat more in the worst case availability and maximum gap data.

Figure 6 shows the data for the constellation at epoch May 23, 2000 at midnight. We see that the data fit into the envelope described by the orbit deviation model used in this paper. However, the actual satellite deviations have a broader and

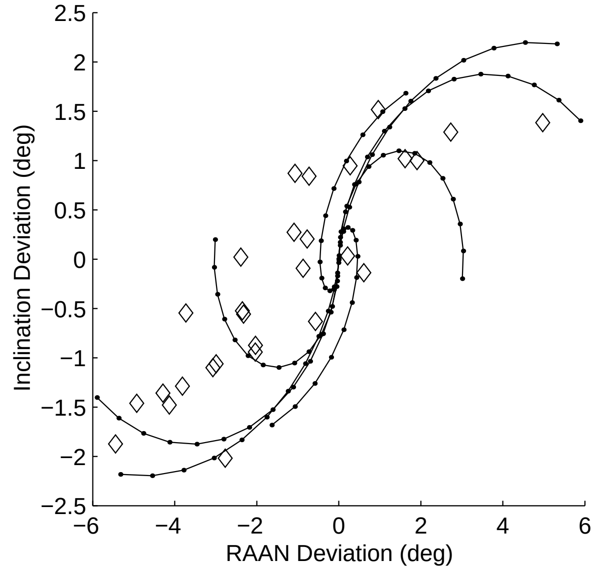

Figure 6: RAAN and Inclination Deviations from Constellation Data. Epoch is May 23, 2000 at midnight. The $\Delta\Omega$ vs. Δi data for each satellite is shown as a diamond. Curves representing the analytical forms of $\Delta\Omega$ and Δi for different values of RAAN are superimposed as a reference ($\phi = \pi/4, \pi/2, \ldots, 2\pi$). Time goes from 0 to 13 years. Each dot in the top graph signifies a one year separation from the previous dot.

more skewed distribution than those from the simulations.

The increased number of large deviations is at least partially explained by the fact that most of the satellites were launched near the initial deployment of the constellation, giving an older age distribution of satellites than than the one used in the simulations. Figure 7 shows the RAAN and inclination deviations from Monte Carlo simulations where the mean of the normal distribution for satellite wearout was taken to be $\mu = 168$ months with the added constraint that each satellite is at least 8 years old. The constellation data shown in Figure 6 is more closely aligned with this simulation.

The skew towards more negative RAAN deviations is partially due to the choice of mean nodal rate of the reference GPS Block II constellation [2]. The Rockwell nodal rate of $-0.04009°$/day was used to compute the reference RAAN values. Use of an Aerospace estimate of the nodal rate of $-0.04034°$/day adds between $0.5°$ and $1°$ to these RAAN values.

LAN CORRECTION

LAN correction is needed to mitigate the effects of RAAN dispersion on MA. Without LAN correction, MA drift can be twice as large as RAAN deviations. An assumption made in the previous analyses was that LAN is corrected. We extend the orbit deviation analysis to investigate the loss of GPS PDOP < 6 availability when the LAN reference is not adjusted. The results show that the degradation of availabil-

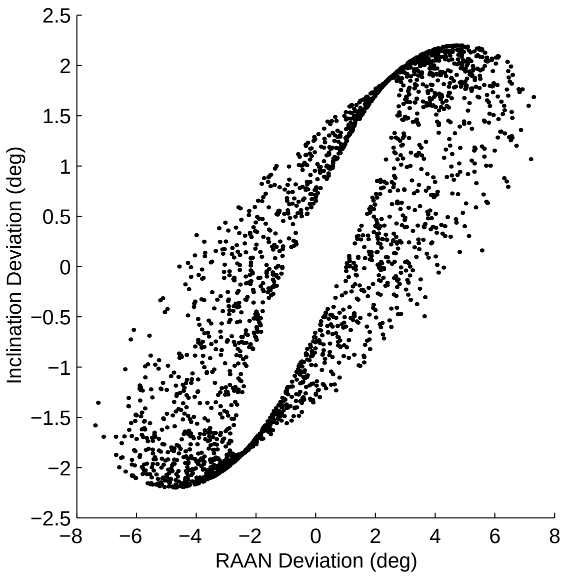

Figure 7: RAAN and Inclination Deviations for a Constellation Consisting of Older Satellites. The $\Delta\Omega$ vs. Δi data from 100 Monte Carlo runs for the nominal 24 satellite constellation where time since launch is taken from the joint distribution with mean of the normal given by $\mu = 168$ months with the added constraint that each satellite is at least 8 years old.

ity is significant when LAN is not adjusted for deviations in RAAN.

We look at the cases of station keeping in mean anomaly (MA) of $\pm 4^\circ$ with and without LAN correction. The quantities we are interested in are the global average coverage (the average over all grid points and all Monte Carlo runs), the worst case location (worst coverage over all grid points and all Monte Carlo runs), and the maximum gap (longest continuous outage over all runs and all spatial grid points.) The data in each case is taken from 100 Monte Carlo runs. The grid used is $2^\circ \times 2^\circ$ in space and 10 seconds in time. We use a mask elevation angle of 5°. The table in Figure 8 shows the PDOP$<$6 availabilities given by the simulations.

We see that the loss of global average availability is significant when we do not adjust the LAN reference to account for dispersions. Although the numbers are not as stable for worst case and maximum gap, we see a discernible degradation for those values as well.

CONSTELLATION MANAGEMENT

GPS satellites are built with internal redundancies to increase the lifetime and reliability of each satellite. Without such redundancies, the design life and mean mission duration would be diminished considerably. Consequently, it is not unusual for a single satellite to have a significant portion of its lifetime vulnerable to a single component failure. It is also not unusual to see more than half of the operational satellites in the constellation vulnerable to a single component failure. Hence, it is extremely difficult to try to predict which satellite in the constellation will fail next. Based on the histories or specifications of various component satellite parts, one can determine probabilities for each satellite, but the end result often indicates that there are perhaps a dozen suspect satellites, each with low probability of failing soon, but for which the collective probability that one of them will fail is high.

Given that the number of satellites that could fail is generally high, it is preferable to use spare satellites to increase the constellation's robustness to unpredictable failures rather than try to guess which satellites will fail next. There are several different strategies that can be employed in order to do this. All of these methods, and more, could be utilized by the Air Force to manage the constellation. The effectiveness of one method over another will vary depending upon the performance histories of the satellites, funding commitments and the available satellite procurement levels.

In this section, we will list a few of the potential management stratagems that the Air Force may choose to utilize in order to both maintain coverage and minimize program costs. This description does not commit the Air Force to following any one strategy over another, but simply to inform users of the system of various strategies that could be employed so that they can consider the effect these strategies might have on their ability to rely upon GPS to meet their own particular navigation needs.

One method that has been frequently employed is to launch new satellites into the orbit plane that has the greatest probability of falling below four satellites, and place the new satellite into the orbit position that most supports the robustness of the constellation to random, unpredicted failures. These spare satellite locations can then be thought of as crisis management centers where the operator is prepared to move the satellite into any location within the orbit plane upon first indication of failure. One problem that has occurred with this method is that sometimes the operator has been able to bring a failed satellite back into operation with the result that two satellites then occupy the same constellation slot for a period of time before the less reliable satellite is moved to a new location. The second problem with this method is that there is a down-time due to the satellite failure, coverage can further be degraded as the spare satellite is located to the new position, and the spare satellite is removed from a position where it was contributing to the robustness of the constellation, thus opening up new vulnerabilities to satellite failure.

A second possible management scheme is to launch into the plane with the most risk of falling below four satellites, replace the satellite that has the greatest risk of failure, and move the replaced satellite to the orbital position that contributes most to the robustness of the constellation. This strategy reduces the likelihood of unexpected failures within the constellation structure by filling the constellation structure with as many healthy satellites as possible. But there are some negative consequences to this strategy that need to be considered. There are potential communication-related prob-

		$\mu = 114$	$\mu = 144$	$\mu = 168$	OLD
Average Availability	LAN corrected	0.99996	0.99996	0.99996	0.99995
	LAN not corrected	0.99994	0.99992	0.9999	0.99983
Worst Case Availability	LAN corrected	0.9848	0.9842	0.9832	0.9829
	LAN not corrected	0.9791	0.9808	0.9821	0.9747
Maximum Gap (minutes)	LAN corrected	21.9	22.8	24.2	24.6
	LAN not corrected	28.1	25.5	25.8	26.6

Figure 8: PDOP < 6 Availability. Results are shown as a function of LAN correction and mean (μ months) of the normal distribution for failure due to wearout. OLD gives the availability for the case of $\mu = 168$ with the added restriction that all satellites are at least 8 years old.

lems to having two operational satellites located too close to one another in the same orbit plane. Hence, the satellite being replaced may need to be moved out of its slot before the new satellite can be brought into the slot. One way to bring this about without coverage loss is to bring the satellite on-line in an orbit location near the constellation slot prior to moving the other satellite out of its slot, and then move the new one in while the other is departing from its slot. This also allows testing of the new satellite to be completed before it is used as a replacement. A complication with this strategy, however, is that older satellites cannot be moved at all times during the year. Hence, there can be delays involved in replacing the satellite and in moving it to a spare location. The spare satellite, being both older and immovable during certain times of the year, is less effective as a replacement for a failed satellite within the same orbit plane. It also takes longer to fill the spare slot and provide the robustness needed to protect against failures that may occur in different orbit planes.

There is one other significant difficulty with the above sparing strategies. Locating the spare satellites in "optimal slots" to improve the robustness of the constellation is nontrivial since the constellation is dynamic. If the spare satellite is located in a position that maximizes the robustness relative to the 24 satellite constellation, then the coverage of multiple spares may end up providing redundant protection. On the other hand, if the spare is placed to optimize robustness relative to all of the current satellites in the constellation, then whenever a spare satellite fails or a spare satellite is moved to replace a different satellite failure, the spare satellites in the other orbit planes would no longer be optimally located and a vulnerability could be opened up in the constellation. Since the real constellation is dynamic, the locations of optimal slots for spares is also dynamic.

Another alternative strategy that is receiving serious consideration is to adopt a larger, more robust, constellation structure that is compatible with the 24 satellite constellation structure. Each orbit plane of the 24 satellite constellation consists of two isolated satellites and two satellites spaced about 30 to 40 degrees apart in their common orbit plane. It is possible to replace three of the isolated satellites in the constellation with pairs selected to maximize the robustness of the constellation.

The position of each of the individual satellites can be fine-tuned to optimize the 27-satellite constellation. In addition, the satellite locations most critical to the maintenance of the constellation can be proactively maintained to ensure that they are always filled with reasonably healthy satellites. Incorporating a much more robust constellation and a proactive management style relative to the most critical constellation slots can mitigate much of the risk inherent in maintaining a constellation of older satellites. The major drawback of this approach is the possibility that satellites will fail at a higher rate than we can replace them. In this case, the constellation structure is versatile enough to allow for selective replenishment and contraction by replacing paired satellite locations with a single satellite until the launch schedule can once again support the full 27-satellite constellation. One must think of the constellation management enterprise as a versatile dynamic management program that both proactively anticipates and mitigates potential failures and reactively responds and corrects actual failures.

CONCLUSIONS

Satellite failure and orbit perturbations render the GPS constellation dynamic. Although users need to know coverage for varying missions, rigid slot definitions, based on the notion of a static constellation, neither adequately inform users nor permit operators flexibility.

Active and flexible management is needed to mitigate the effects of multiple satellite failures, to make effective use of spares and to account for satellite anomalies. In the potential management schemes described in this paper, none are likely to hurt performance of any specific mission, even though all are currently tuned to maximizing coverage and robustness at PDOP$<$6 and a 5° mask elevation angle.

Management flexibility is required if the operator is expected to provide and maintain good coverage given existing resource limitations.

APPENDIX

In this appendix we discuss how to obtain the age distribution of satellites in a constellation from the distributions for individual satellite failure. The two failure distributions considered are a Weibull cdf of the form $\Phi(t) = 1 - \exp(-(t/\beta)^\alpha)$ and a normal distribution, Ψ, with mean μ and standard devi-

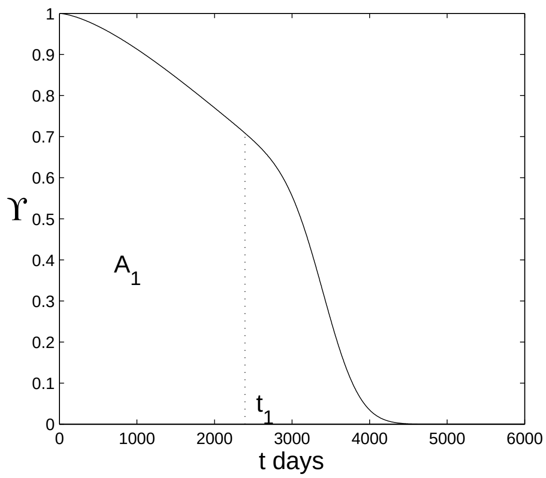

Figure 9: Plot of a probability curve that a satellite is still alive, Υ. A_1 is the integral of Υ from 0 to t_1.

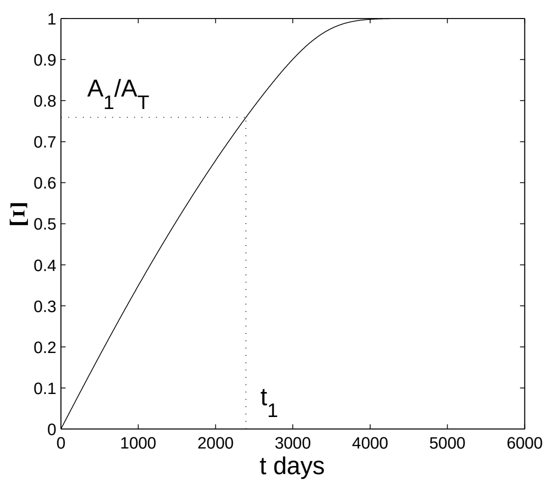

Figure 10: Plot of a cumulative distribution function for satellite ages, Ξ. The cdf is the normalized cumulative integral of Υ. A_1 is the integral of Υ from 0 to t_1 and A_T is the total area under Υ.

ation σ.

The function $\Upsilon = (1-\Phi)(1-\Psi)$ gives the probability that a satellite is still alive. Figure 9 shows Υ for $\alpha = 1.52$, $\beta = 159$ months, $\sigma = 12$ months and $\mu = 114$ months. If A_1 is the area under Υ from time 0 to t_1 and A_T is the total area under the curve, then A_1/A_T is the proportion of satellites that are less that t_1 days old. Thus, normalizing Υ and taking its cumulative integral gives the cumulative distribution function of satellite ages, Ξ. The cdf is shown in Figure 10.

To use the cdf in a Monte Carlo procedure, we take a random variable, X, that is uniformly distributed on the interval $[0, 1]$. Then $Y = \Xi^{-1}(X)$ is a random variable representing satellite age that is distributed according to Ξ.

References

[1] C. C. CHAO, *An analytical integration of the averaged equations of variation due to sun-moon perturbations and its applications*, Tech. Rep. SD-TR-80-12, The Aerospace Corporation, 1979.

[2] C. C. CHAO AND C. J. JOHNSON, *GPS reference nodal rate investigation*. Internal Aerospace Corporation ATM NO. 95(5473-31)-1, 1994.

[3] L. J. GINSBERG, R. D. LUDERS, J. S. DENNER, K. E. GORDON, AND L. E. SPECKMAN, *Orbit planner's handbook*. Internal Aerospace Corporation Publication, 1997.

[4] J. WOMACK. Personal Communication.