

POSITION COVARIANCE VISUALIZATION

J. Woodburn
Senior Member AIAA
Chief Orbital Scientist
Analytical Graphics, Inc.
Malvern, PA

S. Tanygin
Member AIAA
Lead Engineer, Attitude Dynamics and Control
Analytical Graphics, Inc.
Malvern, PA

ABSTRACT

The expanding role of positional covariance data in modern spacecraft operations leads to a need for better understanding of the time evolution of covariance by mission planners and operators. It is common to see positional covariance information presented as uncertainties in the radial, cross-track and in-track directions. While such a time history does provide some information, it tends to obscure the true directionality of the uncertainty. A methodology for improving the understanding of position covariance based on 3D visualization is presented. The methodology utilizes sequential estimation to provide insight into the evolution of the covariance in the presence of measurements and to incorporate the effects of dynamic modeling uncertainties between measurements. Display of the covariance as a triaxial ellipsoid representing a desired probability level is discussed. Various methods for interpolating positional covariance are examined resulting in a recommendation of a technique based on an eigenvalue-eigenvector decomposition of the covariance.

INTRODUCTION

An understanding of the uncertainty in the trajectories of spacecraft is quickly becoming a prerequisite for solving many of the pressing problems in the field of astrodynamics. The uncertainty in the trajectory is typi-

cally determined from the covariance resulting from orbit determination. The computation of the covariance during the estimation process depends on the contents of the estimation state, the available measurements, the uncertainty in the dynamical model, the uncertainty in the measurement model and the estimation algorithm. The position covariance, extracted as a 3x3 sub-matrix from the overall covariance, is useful both as a direct measure of orbit uncertainty and as input to subsequent analyses such as the computation of probability of collision with other satellites¹⁻⁷.

Visualization of the position covariance can be a powerful tool in the understanding its evolution. In this study, we use 3D visualization of the position covariance to obtain a better understanding of how the covariance evolves in the presence and absence of measurements. Observation of the behavior of the covariance resulting from sequential estimation provides insight into relationship between the dynamics of the problem and the observability of the state.

The geometrical representation of the position covariance in 3D space is an ellipsoid, the dimensions of which may be computed to represent a desired probability boundary. For visualization purposes, it is common to replay a scenario at different time steps or to select views at specific times. If the times requested for visualization are not on the grid of times associated with the ephemeris information, then an interpolation operation must be performed. We investigate several techniques for interpolating a time history of covariance. We also

examine how the principal axes of the covariance ellipsoid move relative to the Gaussian frame which is commonly used in the specification of orbit accuracy.

SURFACES OF EQUAL PROBABILITY DENSITY

One means of visualizing the position covariance is by constructing a surface of equal probability density. We restrict ourselves to a discussion of covariance matrices given in Cartesian coordinates. This is not a severe restriction since a covariance matrix, P_A , given in terms of a different element representation, A , can always be transformed to Cartesian coordinates, X , via the appropriate Jacobian matrix as shown in Equation 1,

$$P_x = \left[\frac{\delta X}{\delta A} \right] P_A \left[\frac{\delta X}{\delta A} \right]^T. \quad (1)$$

At least one prior study has demonstrated the benefits of appropriate coordinate selection and nonlinear mapping to produce more representative confidence surfaces for situations involving long prediction intervals and large uncertainties⁸. Our interest, however, is in situations where uncertainties are relatively small. Under such circumstances, the linear mapping of Equation 1 and the subsequent representation of the covariance as an ellipsoid in Cartesian space produces a very good approximation to the confidence surface. If we assume that the state estimate error, x , has a mean of zero and has a Gaussian probability density function, then such a surface will be an ellipsoid centered on the estimated position of the satellite. The shape of these surfaces can be determined from the probability density function,

$$p_x = \sqrt{\frac{1}{(2\pi)^n \text{Det}(P_x)}} \exp\left[-\frac{1}{2} x^T P_x^{-1} x\right] \quad (2)$$

where n is the dimension of the state and the mean of the state estimate error is zero⁹. The surfaces of equal probability density are therefore given by,

$$x^T P_x^{-1} x = k^2. \quad (3)$$

Equation 3 is the equation of an ellipsoid in n dimensions. It is important to note that the ellipsoid is defined not in terms of the the covariance matrix, but in terms of the inverse of the covariance matrix. For the purposes of visualization of the position covariance, we

are specifically interested in the 3x3 submatrix that corresponds to the positional uncertainty. Gura and Gersten describe two possible methods for reducing the n dimensional space to the 3 dimensional subspace of interest¹⁰. In the first method, the inverse of the entire n dimensional covariance matrix is computed and the desired 3x3 submatrix is extracted. This method can be thought of in geometrical terms as creating the intersection of the n dimensional ellipsoid with the three dimensional subspace. The second method consists of extracting the 3x3 submatrix first, then performing the inversion operation. In this method, the behavior of the other elements of the state is ignored. Geometrically, the second method can be thought of as computing the projection of the n dimensional space onto the three dimensional subspace. Since it is only the position subspace that is of interest in this analysis, we will follow the second method.

Gura and Gersten make several other points that are important enough to reiterate here. It is common for the covariance matrix to be rotated into a set of axes that are familiar to the person working with the data. In the case of satellite position covariance this often means rotating the covariance into a frame defined by the radial, in-track and cross-track directions. This frame is often referred to as the RIC frame, the UVW frame or the Gaussian frame. It is then a common practice to treat the square roots of the diagonal elements of the resulting matrix as the dimensions of a sigma surface oriented along those axes. In fact, it is the reciprocals of the square roots of inverse of the covariance matrix that define the intersection of the axes with the one sigma ellipsoid. The dimensions of the ellipsoid are only directly related to the diagonal elements when the axes of interest coincide with the principle axes of the ellipsoid. In this special case, the dimensions of the ellipsoid can be determined from either the covariance or its inverse, both of which will be diagonal matrices.

To obtain a convenient representation of the covariance for extracting the ellipsoid dimensions, an eigenvalue-eigenvector decomposition is performed on the covariance matrix. The original covariance matrix is related to its factors as

$$P_x = U^T D U. \quad (4)$$

The diagonal matrix, D , contains the eigenvalues of P_x while the columns of U contain the corresponding eigenvectors. If we only deal with the 3x3 position submatrix, then the orthogonal transformation matrix, U , is a three dimensional rotation matrix which may be used to convert from the reference axes relative to

which the covariance is defined and the principle axes of the ellipsoid. The elements of D correspond to the squared dimensions of a one sigma surface of the covariance ellipsoid.

COVARIANCE GENERATION

A prerequisite for the display of covariance information is the computation of such information. For the purpose of this paper, an extended sequential filter was used in conjunction with a fixed interval smoother. The extended filter algorithm consists of a sequence of time updates, which move the state estimate and covariance forward in time, and measurement updates which fold in measurements to update the estimate of the state and covariance at the time of the measurement. The filter process flow is shown in Figure 1. The filter algorithm employs local linearization during the time and measurement updates. The time update portion of the algorithm also includes the effects of dynamical modeling errors for the geopotential, atmospheric density and solar radiation pressure.

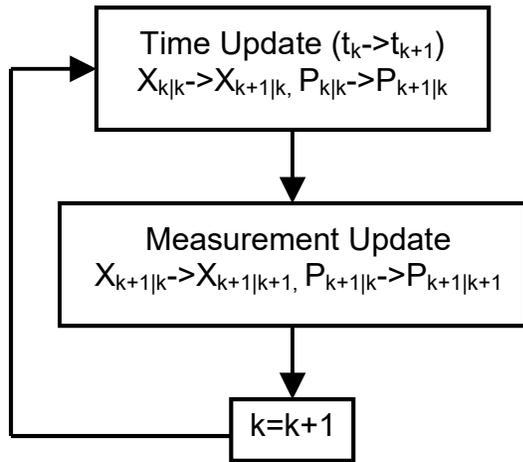


Figure 1. Filter process flow

The subscripts on the state and covariance in the filter flow diagram, Figure 1, give the time associated with the state and the time associated with the last processed measurement. For example, $X_{k+1|k}$ represents the estimated state at time, t_{k+1} , having processed all measurements through time t_k .

The fixed interval smoother is initialized using an estimate from the filter and moves the state estimate and

covariance backwards in time to the epoch of each prior measurement update and time update produced by the filter. The smoothed estimate and the covariance of the smoothed estimate are then updated using information computed during the execution of the filter. The smoother process flow is shown in Figure 2. The state estimates and covariance resulting from the filter contain discontinuities at the measurement update times, but the state estimates and covariance resulting from the smoother are continuous and smooth.

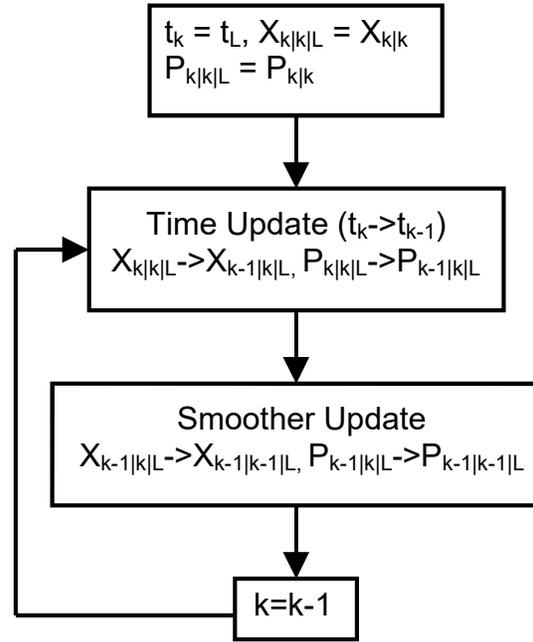


Figure 2. Smoother process flow

The subscripts on the state and covariance in the smoother flow diagram, Figure 2, give the time associated with the state, the time associated with the last processed update and the time that the smoothing process was initiated. For example, $X_{k-1|k|L}$ represents the smoothed state at time, t_{k-1} , having processed all filter output from the start of the smoothing interval, t_L , to time t_k .

Time histories of the angular separations between the principle axes of the covariance ellipsoid and the Gaussian frame are given in the following figures. Figure 3 demonstrates the large variations that occur during the filtering of measurements. Figure 4 shows that significant motion also exists between the frames in the absence of measurements. The two histories shown in Figures 3 and 4 are the angle between the long axis of

the ellipsoid and the in-track direction and the angle of the smallest eigenaxis rotation that is required to align the long axis of the ellipsoid with the in-track direction and orient the other axes of the ellipsoid along the radial and cross-track directions. In this second measure, there is no preference given to aligning the minor axis in either the radial or cross-track directions. The combination requiring the smallest rotation angle was chosen. The results shown here are based on the processing of two way ground ranging on a satellite in Low Earth Orbit (LEO).

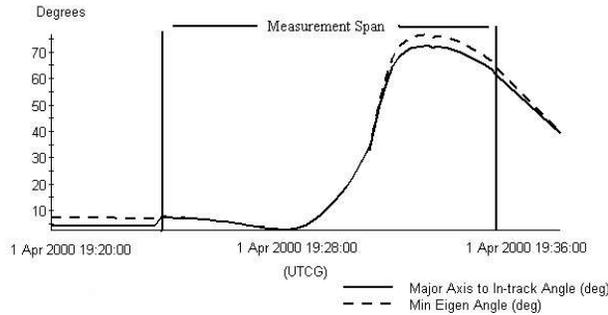


Figure 3. Ellipsoid Orientation during Measurement Processing

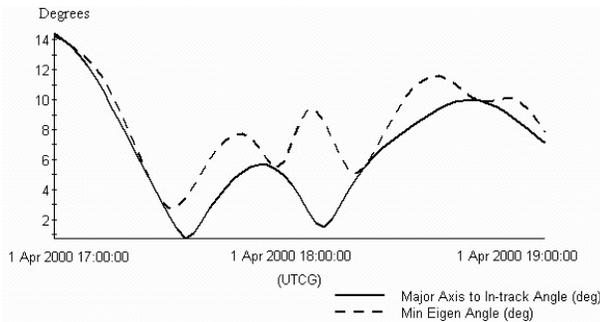


Figure 4. Ellipsoid Orientation during Propagation

COVARIANCE INTERPOLATION

The output of the sequential filter and fixed interval smoother algorithms are time histories of state estimates and associated covariance matrices. Each time history provides all the information needed to produce a visualization of the covariance at the times included in the history. If we desire to produce visualizations at

any arbitrarily selected time with the span of the generated data, we need methods for interpolating the state and covariance information. There are a number of widely used methods for ephemeris interpolation including Lagrange interpolation and Hermitian interpolation. The interpolation of the covariance information presents a slightly more difficult problem.

The extended filter algorithm produces updates to the state and covariance at each measurement time. These updates result in discontinuities in the time histories. Each discontinuity is characterized by having two entries at the same time; one entry representing the estimate and uncertainty in the estimate prior to the incorporation of the measurement and another entry representing the estimate and uncertainty in the estimate considering the measurement. Using information that spans a discontinuity in an interpolation procedure can result in large interpolation errors. We therefore restrict ourselves to interpolation methods requiring information at only the two time points in the history which bound the time at which the covariance is desired.

Element by element

One means of interpolating covariance matrices is to interpolate each element of the covariance matrix, or some matrix derived from the covariance matrix, independently. All methods of this type have the common flaws that they treat matrix entries as being independent of one another when they are not and that an artificial dependence on the coordinate system used in the representation of the covariance is introduced.

The simplest method of interpolating the position covariance matrix is to perform independent interpolation on each element of the covariance matrix itself. There is an obvious problem with this method, the interpolated position covariance matrix is not guaranteed to be positive definite. Another disadvantage of this method is that dimensional change will occur at a higher rate in the vicinity of the smaller ellipsoid. This can be demonstrated via a simple example where interpolation is desired at three intermediary times between two entries in the covariance time history. The entries from the time history are chosen to be,

$$P(1) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad P(2) = \begin{bmatrix} 9 & 0 & 0 \\ 0 & 9 & 0 \\ 0 & 0 & 9 \end{bmatrix}.$$

We then interpolate the covariance matrices using direct element by element interpolation to construct ellipsoids at times 1.25, 1.5 and 1.75. The results of the interpolation are given in Table 1. The resulting variation of the rate of growth of the ellipsoid can be easily computed based on a constant rate of σ^2 as,

$$\dot{\sigma} = \frac{1}{2\sigma} \frac{d(\sigma^2)}{dt}. \quad (5)$$

Table 1. Direct Interpolation of P

Time	σ^2	σ	$\dot{\sigma}$
1.0	1.0	1.0	4.0
1.25	3.0	1.73	2.31
1.5	5.0	2.23	1.79
1.75	7.0	2.65	1.51
2.0	9.0	3.0	1.33

One solution to the problem of maintaining a positive definite matrix through the interpolation process is to perform element by element interpolation on a square root form of the covariance matrix. One square root form of the covariance matrix can be obtained through the use of Cholesky decomposition. The relationship between the covariance matrix and its Cholesky factorization is given by

$$P = L L^T, \quad (6)$$

where L is a lower triangular matrix¹¹. The process therefore involves first converting the time history of covariance to a time history of lower triangular square root matrices. Interpolation is then performed to produce a lower triangular matrix at the requested interpolation time and the covariance is reconstructed based on Equation 5.

A third possible method of performing element by element interpolation is to interpolate the entries of corresponding sigma-correlation matrices. The sigma-correlation matrix is constructed from the covariance matrix by taking the square roots of the diagonal elements and dividing the off diagonal elements by the square roots of the diagonal elements in the same row and column. This process may be represented as,

$$P_{\sigma\rho} = \begin{bmatrix} \sigma_{11} & \rho_{12} & \rho_{13} \\ \rho_{21} & \sigma_{22} & \rho_{23} \\ \rho_{31} & \rho_{32} & \sigma_{33} \end{bmatrix}, \quad (7)$$

where,

$$\rho_{ij} = \frac{\sigma_{ij}^2}{\sigma_{ii}\sigma_{jj}}.$$

The method requires first transforming the covariance matrices into sigma-correlation matrices, then interpolating. After the interpolation is performed, the covariance matrix is reconstructed.

Finally, since the ellipsoid is really defined by the inverse of the position covariance, see Equation 3, we examine the possibility of performing interpolation on the square root of the inverse. In this case, the time history of covariance is converted to a time history of inverse lower triangular matrices, L^{-1} , the interpolation is followed by an inversion of the lower triangular matrix and the interpolated covariance is reconstructed based on Equation 5.

All of the element by element interpolation methods are, of course, able to correctly reconstruct the covariance matrices on the nodes of the time history. While interpolating between the nodes, however, element by element methods have the potential to exhibit quadratic behavior in the dimensions of the ellipsoid. This quadratic behavior, indicated by a change in sign of the rate of the dimension, is due to interpolation of the off-diagonal elements and the subsequent squaring operation used to reconstruct the covariance. We note that the dependence upon the coordinate system with respect to which the covariance is specified is also related to the existence and therefore interpolation of the off-diagonal elements. This also means that the result of the interpolation is dependent upon the coordinate system used in the specification of the covariance. The existence of such dependence is purely an artifact of the interpolation method and is a strong argument against the use of any of the element by element methods. The following example will serve to illustrate some of the problems that occur with these methods. Let the following two covariance matrices represent matrices from the time history,

$$P(1) = \begin{bmatrix} 6.2975 & -19.6421 & 61.9912 \\ -19.6421 & 72.3975 & -205.0296 \\ 61.9912 & -205.0296 & 673.5415 \end{bmatrix},$$

$$P(2) = \begin{bmatrix} 0.7771 & 2.3897 & -9.4083 \\ 2.3897 & 62.2937 & -186.5467 \\ -9.4083 & -186.5467 & 656.2044 \end{bmatrix}.$$

The dimensions of the ellipsoids at the nodes and at three intermediary times are shown in Table 2. The dimensions are ordered such that σ_x is always the largest dimension and σ_z is always the smallest. The method designations are: P_{op} refers to interpolation of the sigma-correlation matrices, \sqrt{P} refers to interpolation of the Cholesky decomposition matrices and \sqrt{P}^{-1} refers to interpolation of the inverses of the Cholesky decomposition matrices. The extreme quadratic behavior of the square root formulations is not typical. Most of the time, all three formulations appear to perform as expected, but occasionally strange behavior is observed.

Table 2. Element by Element Interpolation

Method	Time	σ_x	σ_y	σ_z
P_{op}	1.0	861.716	95.685	22.950
P_{op}	1.25	855.368	95.085	52.062
P_{op}	1.5	850.557	94.335	51.220
P_{op}	1.75	846.499	93.482	40.337
P_{op}	2.0	842.656	92.590	25.149
\sqrt{P}	1.0	861.716	95.685	22.950
\sqrt{P}	1.25	644.011	92.538	37.580

Method	Time	σ_x	σ_y	σ_z
\sqrt{P}	1.5	547.172	91.492	47.897
\sqrt{P}	1.75	631.013	92.037	39.960
\sqrt{P}	2.0	842.656	92.590	25.149
\sqrt{P}^{-1}	1.0	861.716	95.685	22.950
\sqrt{P}^{-1}	1.25	541.810	96.483	30.604
\sqrt{P}^{-1}	1.5	432.254	97.388	36.783
\sqrt{P}^{-1}	1.75	523.645	95.515	33.146
\sqrt{P}^{-1}	2.0	842.656	92.590	25.149

Factorization

Interpolation results free of the undesired higher order behavior of the element by element interpolation method can be obtained by interpolating the eigenvalue-eigenvector factorization of the covariance matrices. The relationship between the covariance matrix and its factors is given in Equation 4 and repeated here for reference,

$$P_x = U^T D U.$$

To support interpolation using this factorization, we first convert the covariance time history to a time history of ellipsoid dimensions and quaternions. The ellipsoid dimensions are simply obtained by taking the square roots of the diagonal elements of D . The quaternions representing the orientation of the ellipsoid are constructed from the orthogonal transformation matrix, U . Interpolation of the ellipsoid axis lengths and quaternions is then performed independently. At the end of the interpolation step, the ellipsoid dimensions and orientation needed to produce a 3D visualization of the covariance is known. If the covariance is required for other purposes, it can be reconstructed by squaring the ellipsoid dimensions and entering them into D , con-

verting the interpolated quaternion into a transformation matrix and applying Equation 4.

At this point, the eigenvalue-eigenvector decomposition appears to be a perfect solution. There are two subtle issues, however, that complicate this process. The first issue is of the order of the eigenvalues. If we simply follow a standard algorithm to compute the desired factorization, there is no guarantee that the first eigenvalue at time t_k will correspond to the first eigenvalue at time t_{k+1} . We compensate for this problem by choosing a consistent ordering of the eigenvalues based on their size. Noting that all eigenvalues must be positive or zero, we place the largest eigenvalue first and the smallest eigenvalue last. When the order of the eigenvalues is modified, the associated coordinate transformation must be updated. This update is easily achieved after conversion to quaternion form by applying a 90 degree rotation about the axis which is not involved in the swap. This rotation is performed by multiplying the existing quaternion by a quaternion evaluated as

$$q_i = 0, q_j = 0, q_k = \frac{1}{\sqrt{2}}, q_s = \frac{1}{\sqrt{2}}, \quad (8)$$

where q_i and q_j are the vector components of the quaternion associated with the axes being interchanged and q_s is the scalar component of the quaternion. Note that achieving the correct order of the eigenvalues may require two axis interchange operations.

The second issue is the lack of uniqueness in the rotation between the reference frame in which the covariance is provided and the principle axes of the covariance ellipsoid. This issue is a result of the squared, symmetric nature of the covariance. If we say that we will define the X axis in the ellipsoid frame to be along the major axis of the ellipsoid, the selection of the +X direction is arbitrary. This arbitrary selection of axes occurs during the factorization of each of the provided covariance matrices. If the selection is made in different directions on two subsequent decompositions, interpolation between those points will produce a 180 degree flip of the ellipsoid. To avoid this problem, the rotations between subsequent quaternions must be minimized. The symmetry of the ellipsoid combined with the constraint of preserving the order of the eigenvalues allows for four potential orientations of the axes as shown in Figure 5.

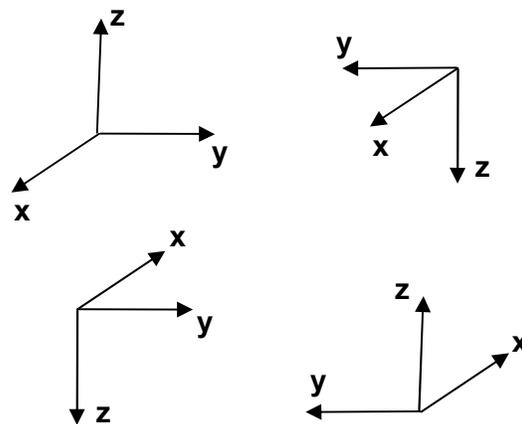


Figure 5. Equivalent Ellipsoid Axes

If we assume that the set of axes in the upper left corner of Figure 5 are the axes resulting from the eigenvalue ordering procedure, then moving clockwise around the figure, the other sets of axes are constructed via 180 degree rotations about the x, y and z axes respectively. Transformations between the original axes and each of the equivalent axes can be easily performed by multiplying the original quaternion by

$$q_i = 1, q_j = 0, q_k = 0, q_s = 0 \quad (9)$$

where q_i is the vector component of the quaternion about which the rotation is being applied. Each of the possible orientations can then be checked to determine which one yields the smallest rotation from the prior quaternion.

VISUALIZATION

Visualization of the covariance ellipsoid requires that the orientation and size of the ellipsoid be known. The orientation of the ellipsoid is given by the transformation between the ellipsoid principle axes and the reference axes of the covariance data. Additional coordinate transformations will have to be added if the reference frame for the graphics window is not the same as the reference axes of the covariance information. The dimensions of the ellipsoid may be specified in one of two ways: by a number of standard deviations at which to draw the surface or via a probability that the true location of the satellite is inside the ellipsoid. If the ellipsoid is specified via a probability level, then this value must be converted to a number of standard devia-

tions prior to visualization. The relationship between the probability level and the number of standard deviations, k , for a three dimensional random vector is given by¹⁰,

$$P_3(k) = \text{erf}\left(\frac{k}{\sqrt{2}}\right) - \sqrt{\frac{2}{\pi}}k e^{-(k^2/2)}. \quad (10)$$

The number of standard deviations corresponding to a 95% probability level for a three dimensional ellipsoid is 2.796 based on Equation 9. Note that this is significantly different than the value of 1.960 which applies to a one dimensional problem. Reference 10 provides tables of $P_n(k)$ for spaces of various dimensions (n).

Once the orientation and dimensions of the ellipsoid have been determined, the ellipsoid can be rendered. Figure 6 shows an example of such a rendering. The ellipsoid is shown as a mesh and the principle axes of the ellipsoid are depicted in relation to the Gaussian frame. This frame was captured as measurements were being processed. Note that the difference between the major axis of the ellipsoid and the in-track direction which is depicted by the RIC basis vector pointing to the left of the page.

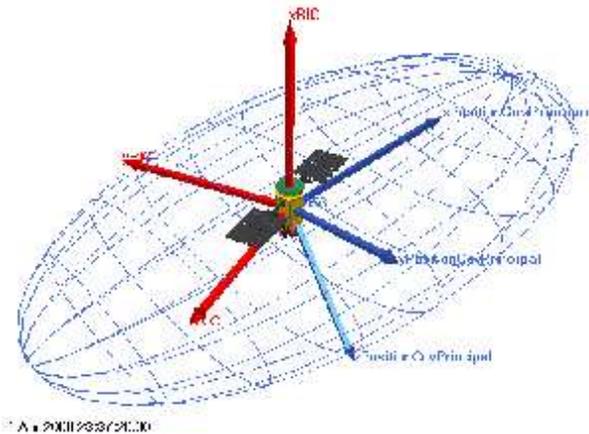


Figure 6. Oriented Ellipsoid

CONCLUSIONS

A method for converting the position covariance sub-matrix to the dimensions and orientation of an ellipsoid representing a surface of constant probability density has been presented. These quantities, which are re-

quired to support 3D visualization of the covariance, are obtained via an eigenvalue-eigenvector decomposition of the position covariance. Several methods of performing covariance interpolation based on the use of covariance at only two time points have been analyzed. A method based on interpolation of the eigenvalue-eigenvector decomposition of the position covariance has been shown to be preferable to any of the examined methods based on independently interpolating matrix entries. The recommended interpolation method has the additional benefit that the results of interpolation can be used for the visualization of the covariance without need for further transformations. It should be noted that we did not examine the use of any interpolation methods which use derivative information, such as Hermitian interpolation. Such methods may prove to be useful in appropriately capturing higher order behavior of the covariance during interpolation if the appropriate derivative information is available.

REFERENCES

1. Alfriend, K.T., Akella, M.R., Lee, D., Wilkins, M., Frisbee, J., Foster, J.L., Probability of Collision Error Analysis, AIAA Paper 98-4279, Presented at the AIAA/AAS Astrodynamics Specialist Conference, Boston, MA, August 1998.
2. Coppola, V., Woodburn, J., Determination of Close Approaches Based on Ellipsoidal Threat Volumes, AAS Paper 99-170, Presented at the AAS/AIAA Space Flight Mechanics Meeting, Breckenridge, CO, February 1999.
3. Wright, J.R., Close Approach Spacecraft Maneuver Criterion, AAS Paper 99-393, Presented at the AAS/AIAA Astrodynamics Specialist Conference, Girdwood, Alaska, August 1999.
4. Lee, D., Alfriend, K.T., Effect of Atmospheric Density on Collision Probability, AAS Paper 00-181, Presented at the AAS/AIAA Space Flight Mechanics Meeting, Clearwater, FL, January 2000.
5. Patera, R.P., A General Method for Calculating Satellite Collision Probability, AAS Paper 00-182, Presented at the AAS/AIAA Space Flight Mechanics Meeting, Clearwater, FL, January 2000.
6. Chan, K., Analytical Expressions for Computing Spacecraft Collision Probabilities, AAS Paper 01-119, Presented at the AAS/AIAA Space Flight Me-

chanics Meeting, Santa Barbara, CA, February 2001.

7. Gottlieb, R., Sponaugle, S.J., Gaylor, D.E., Orbit Determination Accuracy Requirements for Collision Avoidance, AAS Paper 01-181, Presented at the AAS/AIAA Space Flight Mechanics Meeting, Santa Barbara, CA, February 2001.
8. Junkins, J.L., Adventures on the Interface of Dynamics and Control, AIAA Paper 97-0002, Presented at the 35th Aerospace Sciences Meeting and Exhibit, Reno, NV, January 1997.
9. Papoulis, A., Probability, Random Variables and Stochastic Processes. New York: McGraw-Hill, 1965.
10. Gura, I.A., Gersten, R.H., On Analysis of n-Dimensional Normal Probabilities, Aerospace Report No. TR-0066(5129-01)-2.
11. Golub, G.H., Van Loan, C.F., Matrix Computations, Third Edition. Baltimore: Johns Hopkins University Press, 1996.