AD-A086 280 AEROSPACE CORP EL SEGUNDO CA GUIDANCE AND CONTROL DIV F/G 14/4

ORIENTATION OF MEASUREMENT SENSORS FOR OPTIMUM END-OF-LIFE PERF--ETC(U)

JUN 80 P K MAZAIKA F04701-79-C-0080

UNCLASSIFIED TR-0080-(5901-03)-1 SD-TR-80-41 NL

1 OF 1

AD A086280

END

DATE FILMED

8-80

DTIC

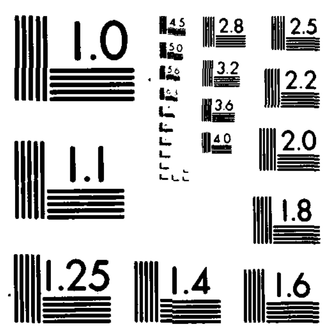

MICROCOPY RESOLUTION TEST CHART

15

REPORT SD-TR-80-41

12 LEVEL II

ADA086280

Orientation of Measurement Sensors for Optimum End-of-Life Performance

P. K. MAZAIKA
System Test and Evaluation
Guidance and Control Division
The Aerospace Corporation
El Segundo, Calif. 90245

25 June 1980

Final Report

DTIC
ELECTE
JUL 7 1980
S D
A

APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION UNLIMITED

DDC FILE COPY

Prepared for

SPACE DIVISION
AIR FORCE SYSTEMS COMMAND
Los Angeles Air Force Station
P.O. Box 92960, Worldway Postal Center
Los Angeles, Calif. 90009

80 7 3 008

This final report was submitted by The Aerospace Corporation, El Segundo, CA 90245, under Contract F04701-79-C-0080 with the Space Division, Deputy for Space Communications Systems, P.O. Box 92960, Worldway Postal Center, Los Angeles, CA 90009. It was reviewed and approved for The Aerospace Corporation by D. P. Duclos, Engineering Group. First Lieutenant J. C. Garcia, YLXT was the Deputy for Technology project engineer.

This report has been reviewed by the Public Affairs Office (PAS) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nations.

This technical report has been reviewed and is approved for publication. Publication of this report does not constitute Air Force approval of the report's findings or conclusions. It is published only for the exchange and stimulation of ideas.

Juan C. Garcia

J. C. Garcia, 1st Lt, USAF
Project Engineer

Joseph J. Cox Jr.

Joseph J. Cox, Jr., Lt Col, USAF
Chief, Advanced Technology Division

FOR THE COMMANDER

Burton H. Holaday

Burton H. Holaday, Col, USAF
Director of Space Systems Planning
Deputy for Technology

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: SD-TR-80-41	2. GOVT ACCESSION NO.: AD-A086280
3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle): ORIENTATION OF MEASUREMENT SENSORS FOR OPTIMUM END-OF-LIFE PERFORMANCE	5. TYPE OF REPORT & PERIOD COVERED: Final rept.
	6. PERFORMING ORG. REPORT NUMBER: TR-0080(5901-03)-1
7. AUTHOR(s): P. K. Mazaika	8. CONTRACT OR GRANT NUMBER(s): F04701-79-C-0080
9. PERFORMING ORGANIZATION NAME AND ADDRESS: The Aerospace Corporation, El Segundo, Calif. 90245	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS: 28
11. CONTROLLING OFFICE NAME AND ADDRESS: Space Division, Air Force Systems Command, Los Angeles, Calif. 90009	12. REPORT DATE: 25 June 1980
	13. NUMBER OF PAGES: 28
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	15. SECURITY CLASS. (of this report): UNCLASSIFIED
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Optimization, minimax, sensor configurations, gyro optimization, end-of-life

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

Relative orientations of on-board measurement sensors are found that ensure the best possible measurement accuracy in the end-of-life situation when all but three sensors have failed. For five and six sensors, the optimum orientations are, respectively, the single cone and dodecahedron configurations. For seven and eight sensors, new configurations are found that are, respectively, 25 and 20% more accurate than previously proposed configurations. In general, there is a tradeoff between performance at end-of-life and performance earlier in a configuration's lifetime.

CONTENTS

FIGURES

I. INTRODUCTION

Attitude and control performance depend upon the number, sensitivity, and orientation of the on-board measurement sensors. For a given number and sensitivity of the sensors, different orientations result in distinct performance characteristics. This report derives configurations of five through eight sensors that ensure the best possible measurement accuracy in the end-of-life situation when all but three sensors have failed. It is assumed the sensors are identical single degree-of-freedom sensors, such as gyros, accelerometers, or star-slit sensors, and that failure detection is accomplished by some other means than the sensors themselves.

A number of studies have approached the problem of optimizing the relative orientation of sensors[1,2] but not with the end-of-life criterion used here. In particular, Pejsa[1] gave optimum configurations for 3 through 10 sensors based on the combined measurement accuracy of all the sensors taken together. He found that relative sensor orientations pointed to the faces of regular polyhedra in the cases of 3, 4, 6, and 10 sensors, and that they were equally spaced around a cone with a half-angle of 54.74° for 3, 4, 5 and 7 sensors. For 3 or 4 sensors, the single-cone configurations are identical to the corresponding regular polyhedra (cube and octahedron) configurations.

The optimal end-of-life configurations may, in general, be different from the polyhedra and cone configurations examined by Pejsa, but it is difficult to visualize other good candidate configurations. Given that the sensor orientations lie in a single cone, it is possible to find the optimal cone half-angle, but this is an optimization over a small subclass of all possible sensor configurations. Following the examples of the regular polyhedra, it is more plausible that the optimum configurations resemble uniform distributions of sensor orientations over a hemisphere rather than concentration on a single cone.

We developed a numerical iterative technique to perform the end-of-life optimization and applied it to the cases of five, six, seven, and eight sensors. The results for five and six sensors were, respectively, the single cone and dodecahedron configurations which have already been extensively studied.[3] For seven and eight sensors, new configurations were found that gave better end-of-life performance than the cone arrangements described by Pejsa.

We emphasize that in no case have we analytically proven the new configurations are true optima. In a strict mathematical sense, the configurations are only new lower bounds for the true optimal configurations. However, the numerical evidence, based on the convergence properties of the algorithm, indicates that the new configurations may well be the true optima. The new configurations are viable alternatives to the previous cone arrangements since they are certainly better at the end of life, and have different performance characteristics throughout their lifetimes. The relative merits of the different configurations are discussed in the results.

II. PERFORMANCE CRITERION

The linear filter equations used to do state estimation are well known.[2,4] The usual expression of these equations is

$$z = Hx + e \tag{1}$$

where z is an n-dimensional vector of sensor measurements, H is an $n \times 3$ geometry matrix, x is a 3-dimensional state vector, and e is an n-dimensional vector of measurement noise. The statistics of e are assumed to be

$$E(e) = 0, \; E(ee^T) = \sigma^2 I \tag{2}$$

where I is the $n \times n$ identity matrix. The least squares estimate $\hat{x}$ of x is given as

$$\hat{x} = (H^T H)^{-1} H^T z \tag{3}$$

and the estimation error covariance matrix is

$$C = E\,[(x-\hat{x})\,(x-\hat{x})^T] = (H^T H)^{-1}\,\sigma^2 \;. \tag{4}$$

The last equation indicates that the error covariance is a function of the sensor geometry H and the variance σ^2 of the measurement noise. Since we are only concerned with the sensor geometry, it will be assumed without loss of generality that $\sigma^2 = 1$.

One performance criterion of a sensor geometry matrix is the determinant, $|C|$, of the associated covariance matrix - good geometries yield small determinants. The square root of $|C|$ is proportional to the volume of the error ellipsoid associated with the measurements.[2] Rather than minimize $|C|$, it is often more convenient to maximize $|C^{-1}|$, or equivalently, maximize $|H^TH|$. H^TH is sometimes called the "information matrix" of the measurements because more information from better placed or additional sensors implies a smaller covariance in the estimate.

The matrix H is composed of n rows of 3-dimensional unit vectors where each vector describes the direction cosines of a sensor relative to the direction being measured, i.e.

$$H = \begin{bmatrix} - & h_1 & - \\ - & h_2 & - \\ & \vdots & \\ - & h_n & - \end{bmatrix} \; ; \; h_i h_i^T = 1 \text{ for } i = 1, \cdots, n. \tag{5}$$

The determinant $|H^TH|$ can be divided into the sum of triple products of each subset of three of the n vectors.[5] The resulting relation,

$$|H^TH| = \sum_{i=1}^{n-2} \sum_{j=i+1}^{n-1} \sum_{k=j+1}^{n} [(h_i \times h_j) \cdot h_k)]^2, \tag{6}$$

may be loosely interpreted to mean the total information in n vectors is the sum of the information contained in all subsets of three vectors. Note the special case where $n = 3$ implies that the information content of three sensors is simply the square of their triple product. In this case the volume of the error ellipsoid is inversely proportional to the magnitude of the triple product.

The above formula is useful for calculating the effect on estimation error when some of the original n sensors have failed; the remaining information is the sum over all possible triple products of the remaining functioning sensors. When all but three sensors have failed, such as might occur at the end of a long lifetime, the remaining information is the square of the triple product of the last three sensor direction vectors.

The value of a sensor configuration is defined as the minimum triple product squared within the configuration. This value corresponds to the information remaining after the worst case of $n - 3$ sensor failures. This worst case analysis gives a sure lower limit on end-of-life performance, and is a good figure of merit when a particular attitude specification must be met. The optimum configuration is the one with the highest value.

Another performance figure sometimes used is the minimum eigenvalue of the information matrix $H^T H$. The minimum eigenvalue yields the longest semiaxis of the error ellipsoid and is a measure of the uncertainty in the worst direction. For the case with all sensors taken together, maximizing the determinant of the information matrix gives the same result as maximizing the minimum eigenvalue.[6] The corresponding relation is not necessarily true for the end-of-life problem since the set of three vectors with the minimum eigenvalue may not even be the same as the set of three vectors with the minimum triple product squared. A calculation of the minimum eigenvalues within the optimized configurations indicated that although the optimization was only with respect to the triple product criterion, there was similar improvement in the minimum eigenvalues. The calculated worst direction uncertainties are included in the results.

III. NUMERICAL OPTIMIZATION

The goal of the optimization is to find configurations of maximum value, i.e., configurations with the maximum minimum triple product squared. The resulting nonlinear maximin optimization problem was solved using the iterative numerical method described in detail in the appendix. The general procedure is to start with some initial configuration and to continually perturb it in such a way that its value increases. This process is repeated until no further increase is possible, at which point the configuration should be optimized. In practice, this procedure does not always succeed.

Typical problems encountered in general nonlinear optimization are that the numerical scheme does not converge to a true maximum, or that the maximum to which it converges is only local rather than global. The local maximum problem can be attacked (for a convergent algorithm) by trying a sufficient variety of initial conditions. Different initial conditions may converge to different local maxima, and the analyst may then pick out the local maximum with the highest value as the solution. If the global maximum is the only local maximum, then all the initial configurations should converge to the same final configuration (allowing for rotations and reflections of that configuration). It is more difficult to handle a nonconvergent algorithm since the analyst cannot tell in general whether the algorithm stopped at a local maximum or merely failed to converge due to a pathological solution surface. However, if all the initial configurations converge to the same one or two maxima, one can be reasonably confident in the convergence of the algorithm.

Ten initial configurations were tried for n = 5, and all of them converged to the 54.74° single cone solution. Nine initial configurations were tried for n = 6, and all but one converged to either the dodecahedron or the 54.74° single cone. Thus, the algorithm converged to well-known configurations in 18 of the 19 test cases. This gave confidence that global optimization was possible by using the algorithm with a variety of initial configurations.

Seven initial configurations were chosen for the seven-sensor case. Five of the seven cases converged to configurations of nearly the same value. When these solutions were rotated into similar orientations, it was found that the respective vectors were within a neighborhood of a single solution. Additional sets of initial conditions were chosen in this neighborhood to further refine the solution. This process of progressively following the best family of solutions was repeated until the optimum configuration was determined to within 0.01 degrees. In both the six and seven-sensor cases, one of the initial configurations did not converge to either the global optimum or the single cone local optimum. Although their values were better than the single cone values, they were not as good as the global optimum values, and consequently, will not be included in the results.

The eight-sensor case was not examined as extensively as the previous cases. The initial configuration was chosen as the seven-sensor solution (see Table I) with an additional sensor direction at $\phi = 24.3°$, $\theta = 180°$ and the optimum eight-sensor configuration was derived numerically from it. Further optimization of this eight-sensor configuration may be possible, but a more efficient numerical algorithm is required.

Table I. Summary of Configurations and End-of-Life Accuracy

Number of Sensors	Configuration	Sensor Orientations (Polar Coordinates (ϕ,θ)) $0° \leq \phi \leq 180°, -180° \leq \theta \leq 180°$	Worst Case Performance Error Volume; Semi-axis Error
5	Single Cone	(54.75, 0.0), (54.75, 72.0), (54.75, -72.0), (54.75, 144.0), (54.75, -144.0)	1.98; 2.80
6	Dodecahedron	(37.38, 0.0), (37.38, 120.0), (37.38, -120.0) (79.19, 60.0), (79.19, -60.0), (79.19, 180.0)	2.13; 3.08
	Single Cone	(54.75, 0.0), (54,75, 60.0), (54.75, -60.0), (54.75, 120.0), (54.75, -120.0), (54.75, 180.0)	3.00; 4.18
7	End-of-Life	(61.16, 36.06), (61.16, -36.06), (61.07, 88.10), (61.07, -88.10), (61.16, 154.14), (61.16, -154.14), (24.30, 0.0)	3.31; 4.90
	Single Cone	(54.75, 0.0), (54.75, 51.43), (54.75, -51.43), (54.75, 102.86), (54.75, -102.86), (54.75, 154.29), (54.75, -154.29)	4.41; 5.85
8	End-of-Life	(64.39, 41.03), (64.39, -41.03), (64.39, 138.97), (64.39, -138.97), (29.91, 0.0), (29.91, 180.0), (70.79, 90.0), (70.79, -90.0)	3.92; 5.72
	Two Cone	(45.0, 0.0), (45.0, 90.0), (45.0, -90.0), (45.0, 180.0) (65.91, 45.0), (65.91, -45.0), (65.91, 135.0), (65.91, -135.0)	4.90; 7.11

IV. PERFORMANCE CHARACTERISTICS OF THE OPTIMUM CONFIGURATIONS

Relative sensor orientations as determined by the numerical scheme are given in Table I. These solutions, expressed in polar coordinates as indicated in Fig. 1, have been rigidly rotated to be aligned with the coordinate axes. The matrix H^TH was computed for all sets of three vectors within each configuration. The set of vectors with the minimum determinant of H^TH defines the configuration value, and the worst case error volume is the inverse of the square root of that determinant. Similarly, eigenvalues of H^TH were found for each set of three vectors. The maximum semiaxis error is the inverse of the square root of the minimum eigenvalue found among all sets of vectors. Although the optimization was only with respect to worst case error volume, the table shows there is similar improvement in the worst case semiaxis error.

The optimal configurations suggest some basic principles for good orientations. The table exhibits the dodecahedron expressed as a two-cone configuration indicating that cones are an integral part of good configurations. The seven-sensor case is displayed as six vectors lying on a cone with a seventh vector above it (Fig. 1). It can also be expressed as five sensors on a different cone by reversing the directions of the first two vectors and rotating coordinates. This leads to the vectors (76.0, 0.0), (73.9, 65.5), (73.9, -65.5), (75.1, 147.7), (75.1, -147.7), (23.4, 106.0), (23.4, -106.0) as an alternate description. Thus, the seven-sensor configuration is approximately described as the intersection of two cones, one with six sensors and the other with five sensors. The eight-sensor configuration is somewhat different. It has 4 vectors on a cone with 2 above and 2 below it,

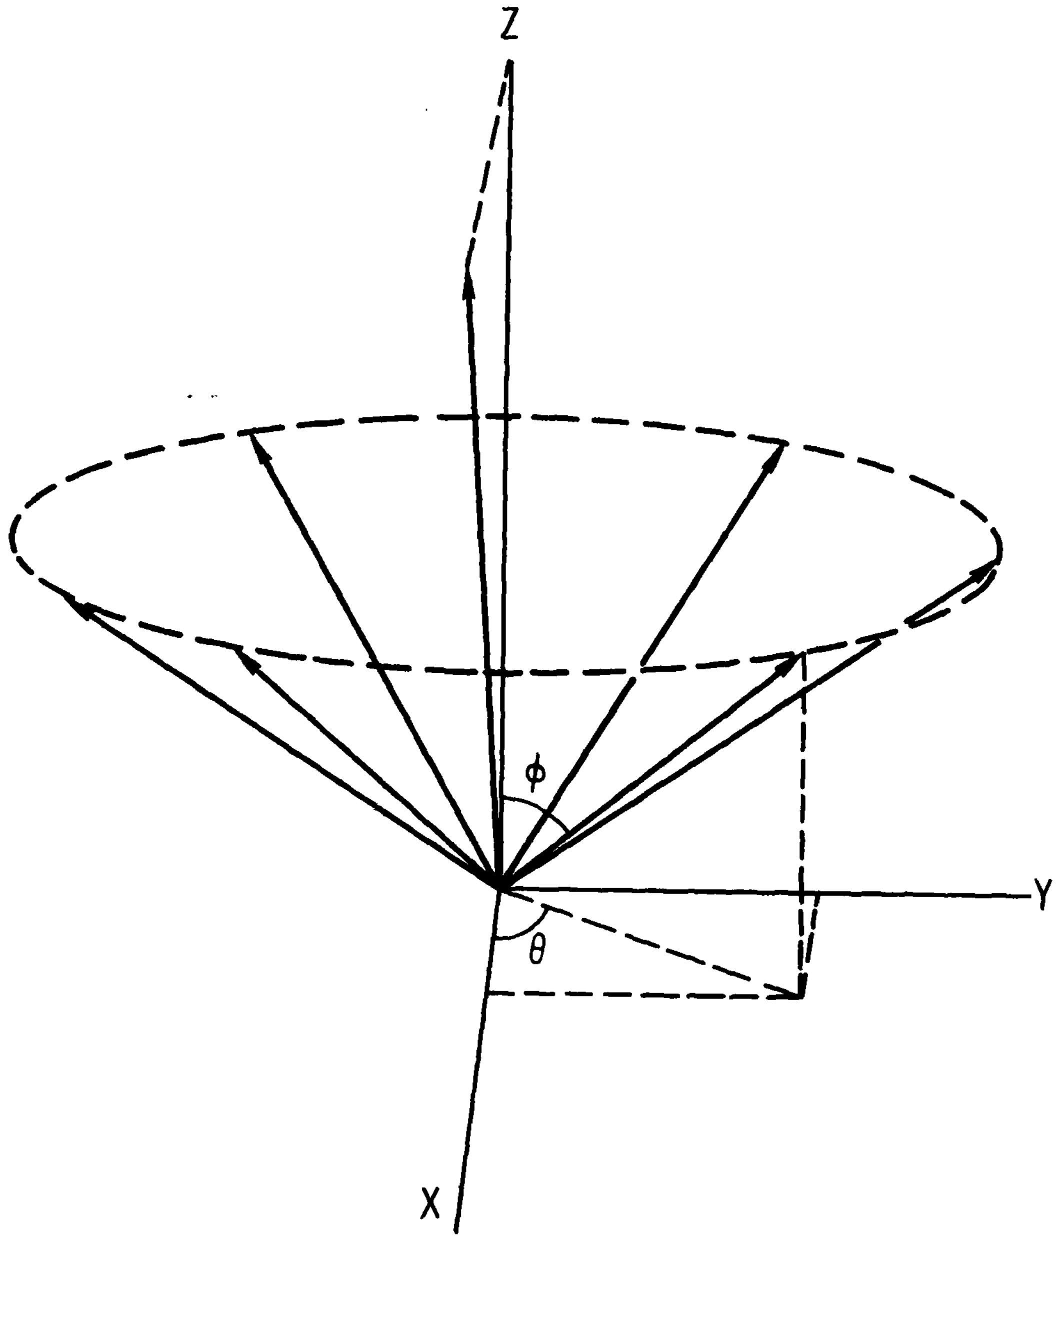

Fig. 1. Optimum End-of-Life Configuration for Seven Sensors

and resembles a 10-sensor icosahedron arrangement with two sensors removed. (Adding two sensors at (90, 30) and (90, -30) yields the approximate icosahedron). One implication of this result is that the optimal end-of-life nine-sensor configuration may closely resemble an icosahedron with only one sensor removed.

The optimum arrays can be compared to the cone arrays by their performance as individual sensors fail out of the original full array. This gives the short term characteristics of the configuration before the end-of-life situation is reached. We assume a situation where thermal and power constraints require that only three sensors be operating at any one time; when more than three sensors are operable, such as immediately after launch, the user chooses the three sensors that are closest to an orthogonal triad and turns the remaining ones off. Thus, the user is assumed to maximize the array accuracy given that particular sensors have failed. For any fixed number of failures, the resulting accuracy depends on which of the particular sensors failed. Performance will be measured as the volume of the error ellipsoid of the operating sensors relative to that for an orthogonal triad. An orthogonal triad has a relative volume of 1; less favorable configurations have larger relative volumes.

Two performance curves are plotted for each full configuration. One gives the mean performance after a given number of failures, and the other gives the worst possible performance after a given number of failures. It is always assumed the user has chosen the best available set of sensors. Averages were computed as if the failure rates were the same for all sensors, whether active or inactive. In some applications, failure rates for standby sensors are significantly lower than failure rates for active sensors, or there is a significant probability of failure when a sensor is initially turned on. These effects have been omitted from this study.

The performance of the single cone configuration of five sensors is exhibited in Fig. 2. There are no orthogonal triads within the configuration, so the performance figure is greater than one, even with all the sensors available. There is no loss of performance when one sensor fails - this is typical of the single cone configurations. When two sensors fail, there is a range of possible performances depending on the remaining sensors, with a worst case performance of just under 2.00.

Figure 3 compares the performance of a dodecahedron to a single cone. The single cone gives better initial performance because it contains orthogonal triads, but the dodecahedron is 29% better in worst case end-of-life performance.

The new seven-sensor configuration is compared with the single cone configuration in Fig. 4. The average performance of the configurations is very similar, with the new configuration slightly better. For worst case performance, the new configuration is 25% better at the end-of-life but the single cone is better in the intermediate lifetime when two or three sensors have failed. A comparison of the eight-sensor configurations (Fig. 5) indicates similar behavior. Average performances of the configurations are nearly the same, however, for worst case performance, the new configuration is 20% better at the end-of-life while the two-cone configuration is better in the intermediate lifetime of two and three failures.

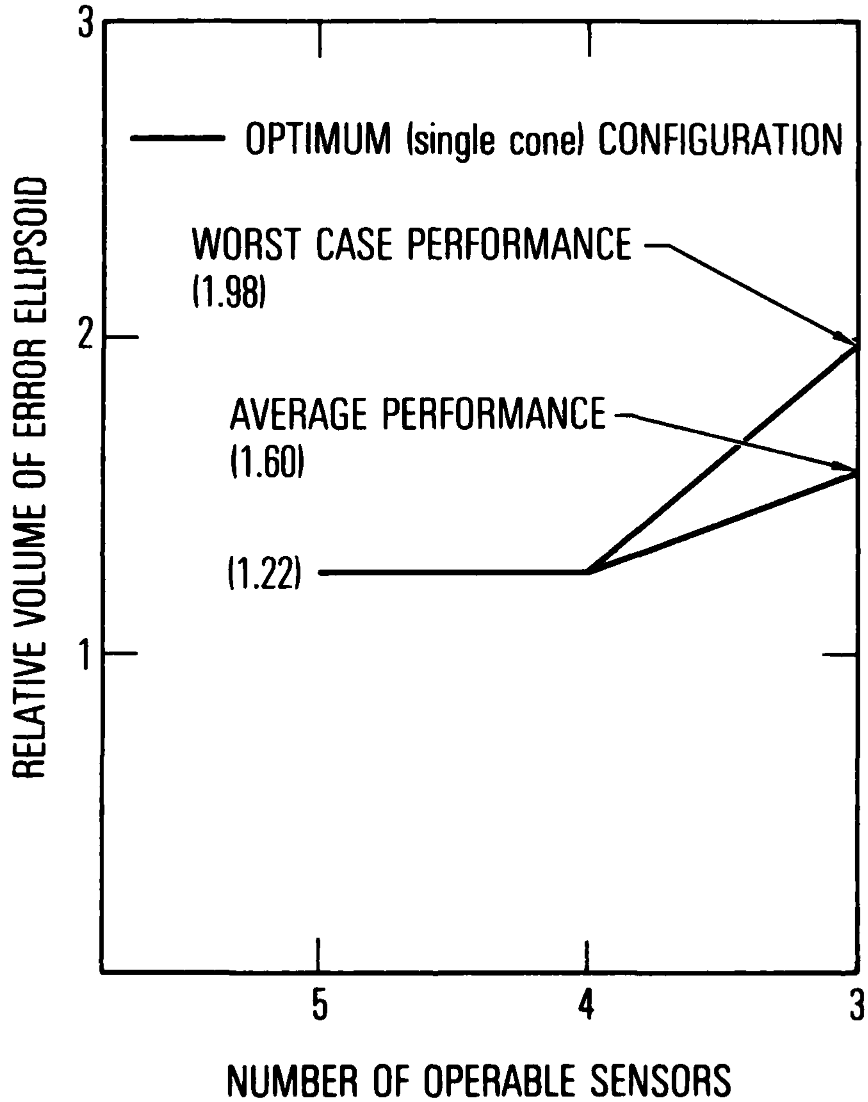

Fig. 2. Performance of Optimum Five-Sensor Configuration

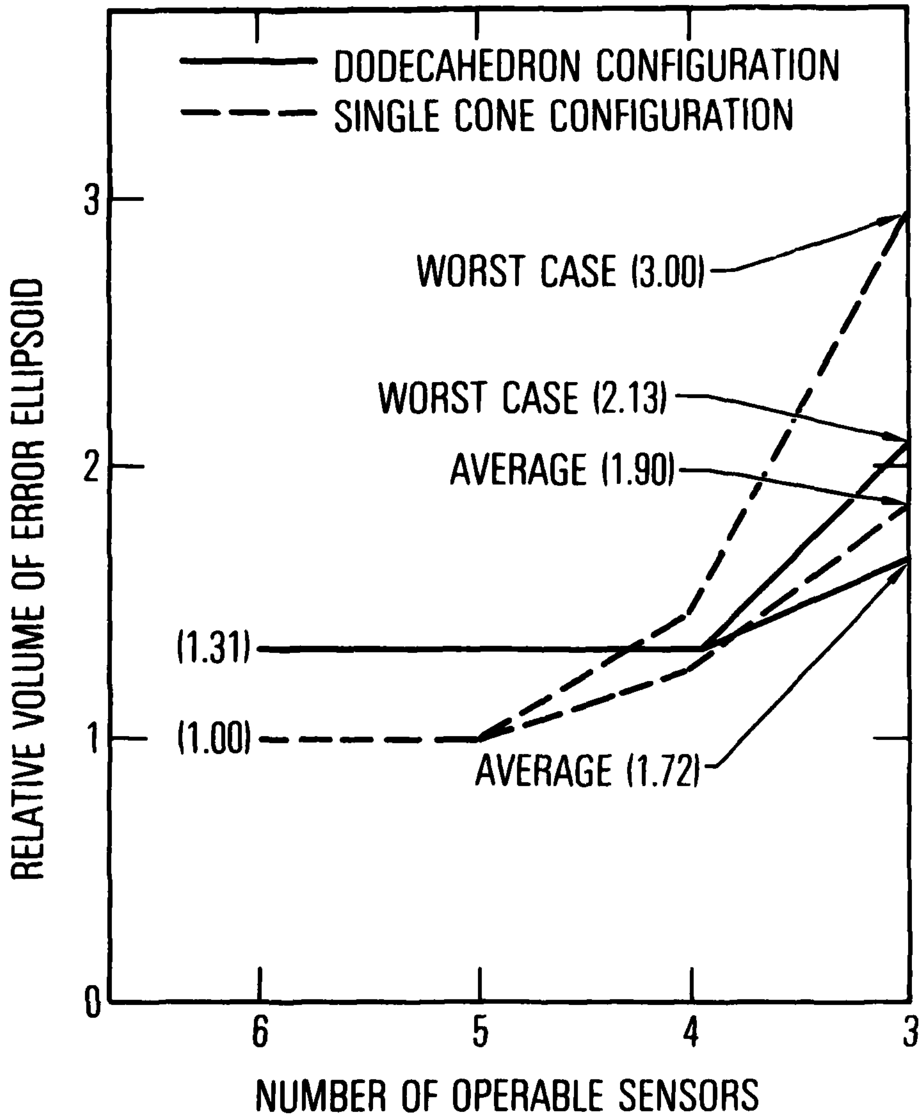

Fig. 3. Performance of Optimum Six-Sensor Configurations

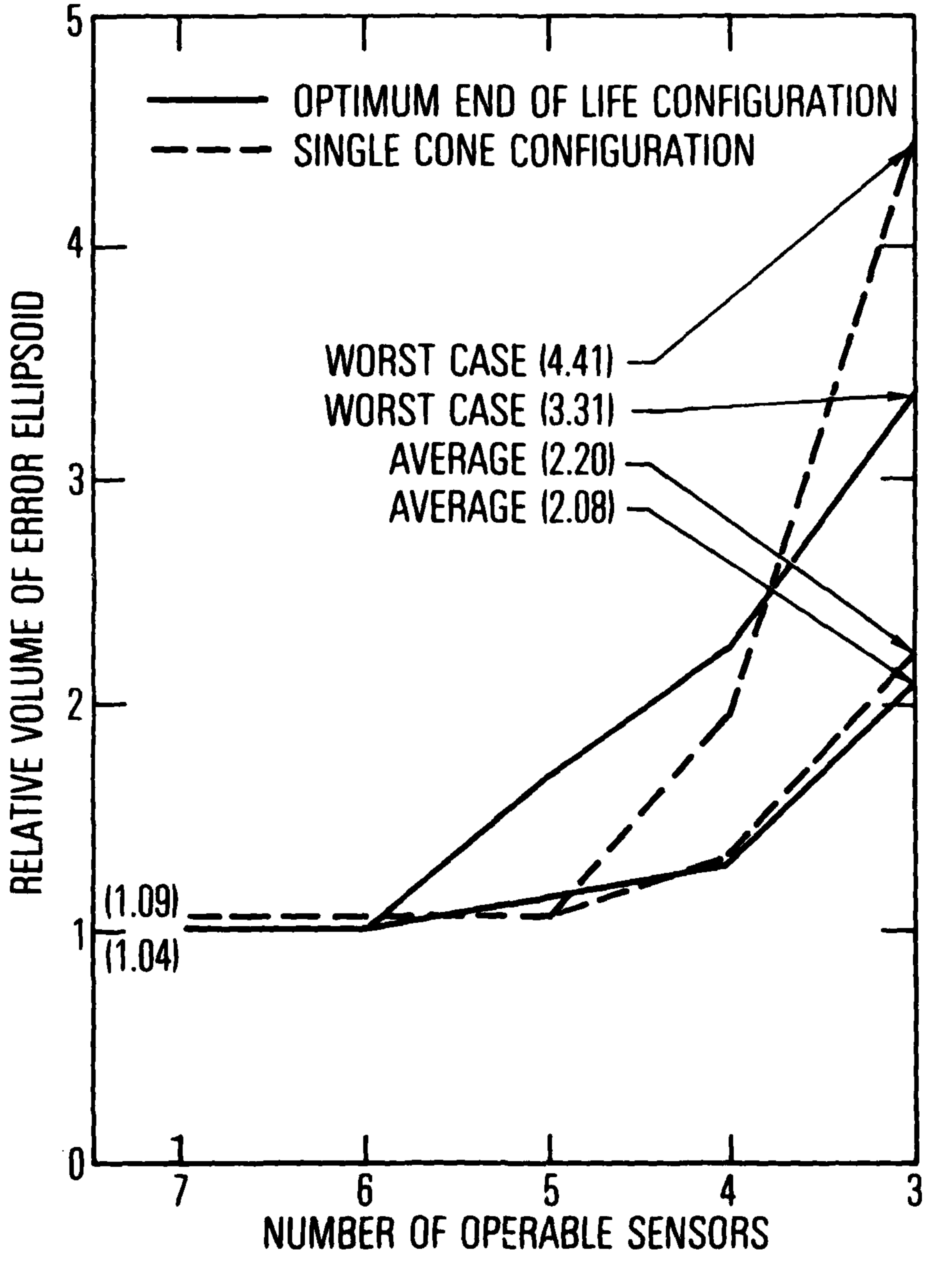

Fig. 4. Performance of Optimum Seven-Sensor Configurations

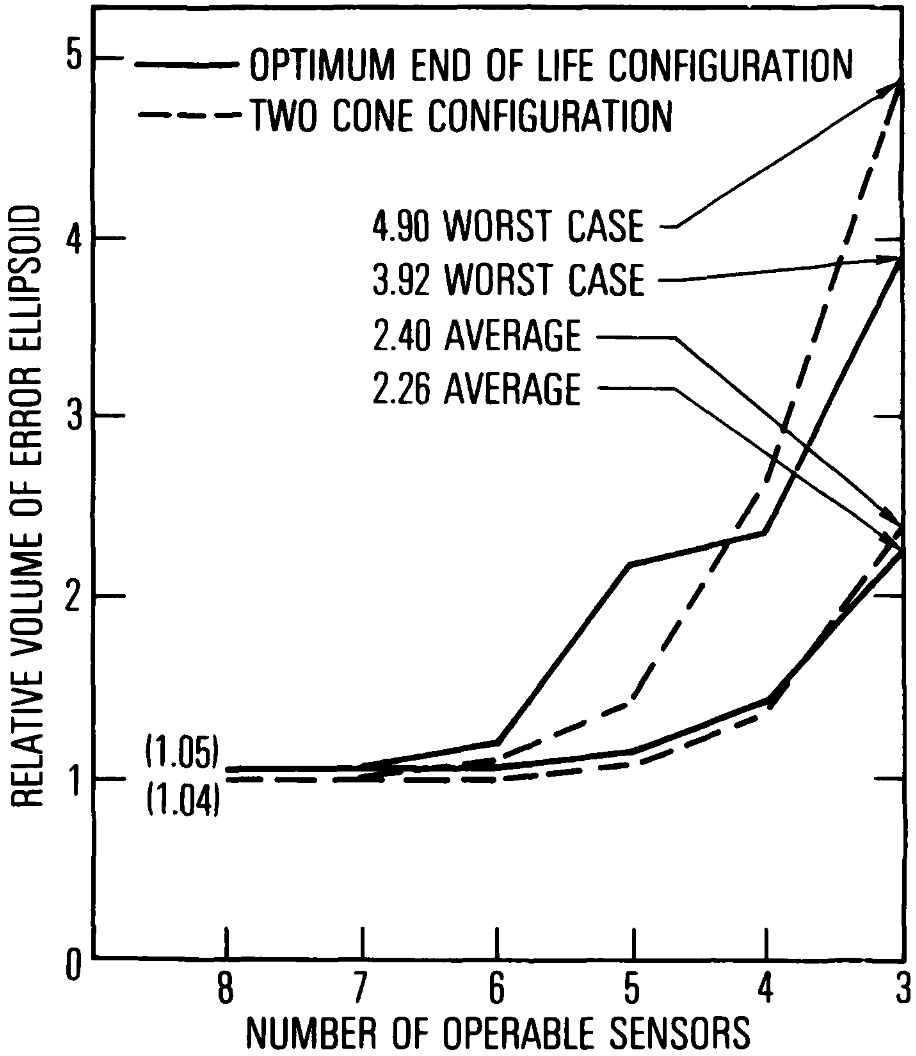

Fig. 5. Performance of Optimum Eight-Sensor Configurations

V. SUMMARY

Relative sensor orientations were derived that optimized end-of-life measurement accuracy in the worst case situation when all but three sensors had failed. For five and six sensors, the results were the well-known single cone and dodecahedron configurations. For seven and eight sensors, new configurations were found. The seven-sensor configuration is roughly described as six vectors lying unequally spaced on a cone with a half-angle of 61.2° while the seventh points 24.3° away from the cone axis. The eight-sensor configuration somewhat resembles a 10-sensor icosahedron arrangement with two sensors removed. The new configurations are at least 20% more accurate in end-of-life performance than previously proposed configurations, but are less accurate in the intermediate lifetimes when only two or three sensors have failed. The optimization was performed by numerical iterative methods and further improvement may be possible, particularly in the case of eight sensors.

APPENDIX

Maximin Optimization Using Linear Programming

The maximin optimization problem is to find the configuration which has the maximum minimum triple product squared. We describe the configuration by a vector y with 2n components that specifies the 2n independent variables in the n sensor problem, and we denote each triple product squared by $f_i(y)$ where each index value i corresponds to a different set of three vectors. The objective of the optimization is to maximize C(y) where

$$C(y) = \min_{1 \leq i \leq N} \{f_1(y), f_2(y), \ldots, f_N(y)\} \tag{7}$$

and N is the total number of distinct sets of three vectors.

Numerical optimization was accomplished by a direct search method that for any point y finds a neighboring point $y + \Delta y$ such that $C(y + \Delta y) > C(y)$. In general, there are many possible directions for Δy that increase the configuration value. To obtain the largest possible increase for a given length of Δy, the problem was rephrased at each point y as a constrained optimization with the objective of maximizing one of the f_i's subject to the constraints that it was less than or equal to all the other f_i's. The minimum function f_i at y was chosen as the objective function.

The constrained problem may be solved by locally linearizing all the functions $f_i(y)$ and solving the resulting linear programming problem. This approach to the constrained problem is similar to linear approximation programming,[7,8] but there is a distinct difference because of the context of the constrained problem within the overall maximin problem. Ordinarily, the solution to the linearized problem may violate some constraints of the full problem which necessitates a restoration move to satisfy the constraints. In the maximin problem, however, the linearized move is an improvement as long as

$$\min_{1<i<N} \{f_i(y + \Delta y)\} > \min_{1<i<N} \{f_i(y)\} \cdot \tag{8}$$

In particular, this means that a different f_i may be the minimum at $y + \Delta y$ (indicating that a constraint of the previous full problem was violated) and yet a restoration move may not be required. Thus, restoration moves are not necessary since a different objective function can be chosen at each step. The algorithm is described in detail below.

Each iteration starts by calculating the triple product squared of each set of 3 vectors and denotes them $f_i(y) = 1, 2, \ldots N$ in order of increasing magnitude.

For nearby points, the function values are approximated by

$$f_i(y + \Delta y) = f_i(y) + \sum_{j=1}^{2n} \frac{\partial f_i}{\partial y_i} \cdot \Delta y_j, \; i = 1, \cdots, N, \tag{9}$$

where $y = (y_1, y_2, \ldots, y_{2n})$.

The partial derivatives are approximated locally using a difference approximation

$$\frac{\partial f_i}{\partial y_j} = \frac{f_i(y + h_j) - f_i(y)}{|h_j|} \tag{10}$$

where h_j is a small vector in the y_j direction. The coordinate directions y_j are redefined at each step of the iteration as two mutually perpendicular coordinates that are perpendicular to each vector. They may be visualized as the allowable directions of motion of the vector tip in a local area on the surface of a unit sphere.

Given the rate of change of the functions with respect to each of the local coordinates, we seek the direction that maximizes the minimum of the functions. Choosing f_1 to remain as the minimum, the problem and constraints may be expressed as:

$$\underset{\Delta y}{\text{maximize}} \left\{f_1(y + \Delta y) - f_1(y)\right\} \tag{11}$$

subject to:

$$f_i(y + \Delta y) > f_1(y + \Delta y), \quad i = 2, \ldots, N.$$

This formulation avoids ambiguities arising from strictly rotational modes of changes. In general, there are three possible directions Δy that give no change in any of the functions since they correspond to a rigid rotation of the entire configuration. The maximization formulation avoids these possibilities because they yield no change in configuration value.

Using the linear function approximations, the maximization is expressed as a linear programming problem. An additional constraint is added to bound Δy to the region of validity of the linear approximations. This constraint also bounds the feasibility region for the solution. The resulting equations are:

$$\underset{\Delta y}{\text{maximize}} \sum_{j=1}^{2n} \frac{\partial f_1}{\partial y_j} \Delta y_j$$

subject to:

$$\sum_{j=1}^{2n} \left(\frac{\partial f_i}{\partial y_j} - \frac{\partial f_1}{\partial y_j} \right) \Delta y_j \geq f_i(y) - f_i(y), \quad i = 2, \ldots, N. \qquad (12)$$

$$|\Delta y| \leq \text{constant}$$

There is no constraint on the signs of Δy_j in this formulation, so it was converted into the usual non-negative format by the change of variables

$$\Delta y_j = \Delta y_j' - \Delta y_j'' \qquad (13)$$

where

$$\Delta y_j' \geq 0, \quad \Delta y_j'' \geq 0.$$

The magnitude constraint on Δy was expressed by the linear relation

$$\sum_{j=1}^{2n} \Delta y_j' + \Delta y_j'' < 100. \tag{14}$$

This allowed the solution to be found easily using a standard linear programming routine. The resulting optimal vector Δy was scaled into a smaller size reflecting the local validity of the linear approximation as defined by the length h_j of the vectors used to approximate the partial derivatives.

The configuration value $C(y + \Delta y)$ at the new point is tested against $C(y)$, and the iteration is repeated from the beginning if there is an increase in value. In the case of no increase, the accuracy of the linear approximation is improved by reducing the magnitudes of Δy and h until the solution of the linearized problem is a valid solution of the full problem. The algorithm is terminated when five of the above reductions still do not yield an improved solution.

The linear programming formulation is a first-order technique and, consequently, had difficulties near local maxima and in regions of high curvature. In practice, it was very helpful in bringing arbitrary configurations into near optimal configurations, but not as successful in actually finding the optimum from a point near the optimum. Convergence was very good for the five-sensor case but became progressively worse as the number of sensors increased. For the seven and eight sensor cases, the optimization process needed to be assisted by adding initial conditions in the neighborhood of previous suboptimal solutions in order for the true optimum to be obtained. With this supplement, the algorithm succeeded in finding the optimum configurations to the desired accuracy.

REFERENCES

1. Pejsa, A. J., "Optimum Orientation and Accuracy of Redundant Sensor Arrays," AIAA Paper No. 71-59, presented at AIAA 9th Aerospace Sciences Meeting, N.Y., Jan. 25-27, 1971.

2. Harrison, J. V. and Gai, E. G., "Evaluating Sensor Orientations for Navigation Performance and Failure Detection," IEEE Transactions on Aero. and Elect. Systems, Vol. AES-13, No. 6, Nov. 1977.

3. Gai, E., Harrison, J. V. and Daly, K. C., "FDI Performance of Two Redundant Sensor Configurations," IEEE Transactions on Aero. and Elect. Systems, Vol. AES-15, No. 3, pp 405-413 (1979).

4. Gelb, A., ed., Applied Optimal Estimation, The M. I. T. Press, 1974.

5. Franklin, J. N., Matrix Theory, Prentice-Hall, Inc., 1968.

6. Uyeminami, R. T., Dixon, M. V., Fuhrman, T. A., Melzer, S. M., and Potter, J. E., High Altitude Attitude Determination Study, TR-0078(3526)-1, The Aerospace Corp., El Segundo, CA 90245.

7. Griffith, R. E. and Stewart, R. A., "A Nonlinear Programming Technique for the Optimization of Continuous Processing Systems," Management Science 7, pp 379-392 (1961).

8. Glass, H. and Cooper, L., "Sequential Search: A Method for Solving Constrained Optimization Problems," Journal of Assoc. for Computing Machinery 12, pp 71-82 (1965).