

Harvard-Smithsonian Center for Astrophysics

Precision Astronomy Group

MEMORANDUM

Date: 20 April 1999 TM99-04
To: FAME Distribution via S.D. Horner and M.A. Murison
From: J.F. Chandler and R.D. Reasenberg
Subject: Draft paper on FAME mission covariance studies of the spiral reduction stage

The paper that follows is in draft form. It is slightly revised from the version that P.K. Seidelmann kindly distributed on our behalf a few weeks ago. The studies in the paper evolved over an extended period as the FAME mission evolved and our understanding of the problems developed. We have recently decided to redo the studies to add refinements and make the cases listed in the tables both more systematic and more applicable. We would welcome your suggestions in this regard as well as other comments or questions that you would be willing to offer us.

**A new approach to post-Hipparcos
astrometric surveys from space**

J. F. Chandler and R. D. Reasenberg

Smithsonian Astrophysical Observatory

Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138

Received _____; accepted _____

To appear in AJ — D R A F T — 1999 Apr 20

ABSTRACT

We consider a new class of small astrometric and photometric survey instruments in the Hipparcos tradition, which would support a rich and diverse ensemble of scientific investigations. In these instruments, there are two fields of view separated by the “basic angle,” and the star images fall on an array of CCD detectors operated in time-delay-integration mode. We have investigated the effect on mission accuracy of the instrument’s basic angle and of the frequency of interruptions of the spacecraft rotation. We conclude that it is advantageous for the basic angle to lie in a broad region centered on 100 deg. Further, we find a large advantage for uninterrupted rotation of the spacecraft. This is in contrast to the Hipparcos approach of frequent correction by gas jet impulses to effect the desired rate of precession. The improvement is about fourfold in changing from six to one rotation break per rotation, with significant further improvement obtained by having no breaks within a multi-rotation batch of data. In this last case, improvement also comes from increasing the length of the batch, with saturation around 36 rotations. Finally, we note that increasing the width of the detector array conveys an advantage greater than merely increasing the observation rate, since it serves to connect a series of rotations through the repeated observation of individual stars.

Subject headings: astrometry — simulations — space vehicles: instruments — methods: numerical

1. Introduction and science objectives

Following the successful ESA mission Hipparcos (Perryman et al. 1989; Perryman et al. 1997), there has been interest in a follow-on astrometric survey mission that is even more capable. Several proposed missions have been studied, including ROEMER (Bastian et al. 1993; Høg 1993; Høg 1995), DIVA (Röser et al. 1997), FAME (Seidemann et al. 1998), GAIA (Lindgren, Perryman, & Loiseau 1995), and FAME-98¹ (Reasenber & Phillips 1998 and Phillips & Reasenber 1998). Here we report on one aspect of a study of such instruments.

All of these proposed instruments take from Hipparcos its two essential geometric characteristics. First, each instrument has two distinct look directions (separated by the “basic angle”), often with the corresponding fields combined on a single detection plane. The repeated observation of many stars in both fields of view provides rigidity in the spacecraft rotation model. Second, each instrument follows a scan pattern that involves both a nominal spin axis orthogonal to the look directions and the precession of that spin axis around the Sun direction. As the spin axis precesses, its angular separation from the Sun direction, ξ , remains approximately constant.

The five parameters of primary interest to any astrometric mission are position (2), proper motion (2), and parallax (1). In addition, photometric data and, in some cases,

¹FAME-98, Full-sky Astrometric Mapping Explorer, should not be confused with the Fizeau Astrometric Mapping Explorer (FAME, hereafter FAME-95), which was the subject of an unsuccessful MIDEX proposal in 1995. FAME-98 has been a collaboration between the Smithsonian Astrophysical Observatory (SAO) and the US Naval Observatory (USNO). More recently, the collaboration has expanded to include Lockheed Martin Missiles and Space Co., the Naval Research Laboratory, IPAC, and Omitron.

crude spectra may be obtained. For a given star, the mission precision is derived from three factors: (a) the single measurement precision, (b) the number of measurements over the mission, and (c) the observing geometry. Optimizing the mission design includes trades between the first and second of these, with due consideration of the subtler issue of systematic error. The last of these may figure importantly, by way of parameter correlations and the condition number of the normal equations, when the astrometric parameters of the stars are estimated.

Among the subjects to be addressed by a modern astrometric survey mission are: calibration of the standard candles (Cepheids and RR Lyrae stars), leading to a better cosmic distance scale; determination of parallax and thus absolute stellar magnitudes for a wide variety of stellar types, including both population I and population II stars; measurement of the space velocities of stars in the solar neighborhood; investigation of the dynamics of open clusters and OB associations; detection and characterization of companions, including stars, brown dwarfs, and planets; determination of the masses of stars in binary systems for a wide class of stars; improvement of the ephemerides of the outer planets and some of their brighter satellites; improvement of the ephemerides of minor planets and the determination of the masses of some by means of mutual perturbations; and a light-deflection test of general relativity. From the star data will come: a better understanding of the internal structure of stars; the ages of clusters; an estimate of the density of dark matter in the Galaxy; and constraints on the mechanism responsible for the spiral arms.

The focus of the present study has been FAME-98, a full-sky astrometric and photometric survey instrument with a nominal mission length of 2.5 years (and a possible extension to five years). The present nominal FAME-98 mission will measure 4×10^7 stars with an astrometric accuracy better than 0.05 mas for $V=9$ and 0.3 mas for $V=15$,

and will perform multi-color photometry with an uncertainty under 0.05 mag for most targets. Central to the instrument is a three-mirror telescope with a square primary mirror (nominally 55 cm across in the current design). This telescope is preceded by a “complex mirror” that combines the two fields of view and is followed by an array of CCD detectors operated in time-delay-integration mode. The instrument and mission are described by Reasenberg and Phillips (1998), and the optical system and detector-plane considerations are described by Phillips and Reasenberg (1998).

In addition to measuring the five astrometric parameters, the instrument would estimate the brightness (stellar magnitude) and temperature of the brighter stars. For a Sun-line, $V=9$ star, the statistical uncertainties would be under 1 milli-magnitude and 10 K, respectively, but the systematic aspects of these uncertainties have not been addressed. On-board multi-band photometry would provide an independent estimate of star temperature. The determination of brightness and temperature by the instrument will not be further addressed here, but are addressed briefly by Reasenberg and Phillips (1998). For further discussion of the scientific applications of missions of this class, see Seidelmann et al. (1998).

In Section 2, we define the nominal data analysis system that provides the conceptual framework for the mission simulations described in Section 3. The three main purposes of these simulations are (1) to address the question of whether frequent rotation corrections (e.g., firing of a gas jet or any other torsional impulse) of the astrometric spacecraft would seriously degrade the astrometric output of a mission, (2) to understand the constraints imposed by this architecture upon the instrument’s basic angle, and (3) to demonstrate the feasibility of the data analysis system. In Section 4, we present the results of the simulations. Finally, in Section 5, we discuss the implications of these studies. The most important of these is that it is strongly beneficial to avoid frequent gas jet firings for

controlling the spacecraft precession. The best alternative is one which avoids gas jets entirely, such as providing the necessary torque through solar radiation pressure on a large shield. Should that technique prove infeasible, the precession could be provided by pairs of firings separated by about 0.25 rotation, at intervals of 1.5 rotation. With this mode, the fuel use is only slightly greater than in a series of frequent firings throughout each rotation, and a far better knowledge of the spacecraft can be obtained.

2. Data analysis approach

To understand the potential of any astrometric mission, one needs to consider the complete package that includes the instrument, the mission (i.e., data-collecting particulars), and the analysis. All three will affect not only the nominal uncertainty in the positions of the observed stars, but also the systematic errors and the subtler matter of connecting the measurements to a coordinate frame.

The modern instruments differ from Hipparcos in many ways. One important difference is in the means of detecting star positions. Hipparcos used a modulating grid to define the metric. Light passing through the grid was relayed by a lens system to an image dissector tube. Thus, in principle, only one star could be observed at a time. However, the image dissector was capable of jumping quickly from one star to another. By this means, Hipparcos created “frames” of data: rapid sequences of measurements that could be treated as simultaneous. The instantaneous Hipparcos field of view was 0.9×0.9 deg (Perryman et al. 1989).

The modern instruments use a set of CCD array detectors, which provide both the metric (pixels) and the detection. The CCDs are operated in time delay and integration (TDI) mode. As the instrument rotates, the star image moves across the CCD surface in

the “scan” direction. The CCD is clocked to keep the growing charge packet under the moving image. For this system, the primary observable is the epoch of an event, namely, the passage of the (centroid of a) star image over the edge of the CCD detector. The secondary observable is the position of the star in the “cross-scan” direction. The position accuracy is likely to be an order of magnitude worse in the cross-scan direction than in the scan direction. The CCD detector offers both higher quantum efficiency than Hipparcos’ detector and relay system and, in principle, the possibility to collect photons from more than one star at a time. FAME-98 will collect photons for astrometry in about 20 $2K \times 4K$ CCD chips, each viewing 0.22×0.44 deg, for a total area of 1.9 sq. deg, more than twice the area of the Hipparcos detector.

In order to provide a context for the studies discussed in the next section, we describe here the nominal data analysis method implicit in our studies. The Hipparcos Mission provided examples of data processing methods that we used as the starting point for our approach, which has three stages following the determination of the event parameters.² The approach combines a hierarchical scheme with numerical techniques to reduce the problem in both complexity and computational demand.

(Stage A) The “observing-spiral” reduction (Hipparcos: great-circle reduction) would address the target events collected during a “batch interval” of a few hours to a few days and would yield a rotation model for the instrument during that period. In our studies, we have assumed the analysis would use an average of three stars per sq. deg as the spiral fiducial stars. The Stage-A analysis could be performed as soon as the target events are available from the first observing spiral. (Even a partial set could be used as part of the

²A more complete description of the data processing approach for FAME-98, including the centroiding of the star light data and the detection of complex targets (e.g., binaries), will be published elsewhere.

instrument check-out.)

(Stage B) The global fit (Hipparcos: sphere solution) would interconnect the observing-spiral rotation models to yield a single model of spacecraft rotation over the whole mission. It would be sufficient to have as few as 1000 global fiducial stars, which would naturally be a subset of the spiral fiducial stars, chosen for their even distribution over the sphere, their presence in the Hipparcos catalog, and their stable astrometric properties. There would be no harm in changing global fiducial stars after the analysis had started.³ The global fit could first be performed well (although probably for position only) after about three months of collecting data, when full sky coverage is first available. We suspect that the quality of the global fit will improve rapidly with increasing data span both before and after the first availability of full sky coverage. After about a year, the analysis could be extended to include proper motion and parallax.

(Stage C) The application of the spacecraft rotation model, with the parameter values estimated from the first two stages, to the determination of the astrometric parameters of the program stars would create the catalog (Hipparcos: astrometric parameter determination). Each star would be fit separately and analyzed for peculiarities such as motion due to a dark companion.

During Stages A and B, the goal is to develop a rotation model for the instrument, not to determine the astrometric parameters of the fiducial stars. These stars serve two functions. First, they provide stable celestial points to be repeatedly observed at high precision through both instrument ports. The resulting sets of repeated observations of each fiducial star provide the internal connections to form the “rigid models” discussed

³The global fit would likely be repeated with a few additional sets of global fiducial stars as a check for systematic errors.

in Section 3.4. Second, through the application of a mild *a priori* constraint on their positions, the fiducial stars provide approximately correct orientations for the observing spirals and the global fit, thus ensuring that the orientation of the mission catalog will be closely aligned to the nominal catalog (for example, Hipparcos) before any final correction is applied.

It is only in Stage C that stellar astrometric parameters are estimated and preserved. During Stages A and B, the astrometric parameters of the fiducial stars could be estimated and discarded. However, a numerical technique for speeding up the analysis would bypass the estimates of the fiducial-star positions. (See Appendix A.)

To clean up the solution iteratively, the Stage-C procedure would be applied only to the combined set of fiducial stars. With improved estimates of positions, proper motions, and parallaxes for these stars, the observing-spiral reductions and global fit would be repeated. We have not yet demonstrated the convergence rate of this three-stage procedure, but we believe it to be rapid for the expected range of corrections to the *a priori* star parameters. Unless this iteration uncovers and precipitates the removal of bad data (blunders) or shows some of the fiducial stars to be unsuitable, it is unlikely that it would *need* to be repeated. However, any modern mission will need a catalog for various near-real-time analyses, and each extension of the catalog would naturally include an iteration. (Thus, the question of convergence rate is largely moot.)

For many purposes, the orientation of the reference frame is not important. However, intercomparisons of the new star catalog with other data would be facilitated by providing a standard orientation connected to the Earth ecliptic and equator. The ecliptic has only a little meaning in a spaceborne astrometric study, where it enters through stellar aberration, and the equator has none. However, a reasonable frame orientation can be assured by using standard catalog positions for the *a priori* position estimates for the fiducial stars.

If these were given modest uncertainties, say 10 to 100 mas, they would anchor the frame well without disturbing the estimated astrometric parameters of the stars or the spacecraft rotation model. Alternatively, a few stable objects like quasars could be included in the solution with *a priori* estimates of their astrometric parameters based on a standard reference frame.

3. Mission simulations

In this section, we describe simulations in which we produce pseudo-observations and subject them to the Stage-A analysis. We start with descriptions of the spacecraft motion and the observed stars. Then, we discuss the analysis of the pseudo-observations, including the numerical techniques that make the analysis tractable. Finally, we introduce the figure of merit that we use to characterize the simulation results.

3.1. Spacecraft operational description

To keep the instrument model as simple as possible for the purpose of simulations, we have distilled whole classes of detector designs down to a model with two parameters: the extents in the scan and cross-scan directions. The detector is represented as two rows of CCDs running in the cross-scan direction and separated by a constant offset in the scan direction. (This differs in detail from the patterns in the proposed instruments.) The first parameter of our “distilled” model represents the base for obtaining the (nearly) instantaneous rate of rotation of the spacecraft. The second determines the width of the observing spiral band on the sky. In all, the model includes four adjustable parameters that describe the observing geometry: (1) the basic angle between the two fields of view in the scan direction, (2) the scan-direction angular offset between the look angles of the two

rows of detectors, and (3 and 4) the cross-scan offsets of the two fields of view from the body-fixed equator. These angles are treated as constants to be estimated from the data.

In a real mission, we expect that there would be a period of latency after each rotation correction, when no useful observations could be made. The action of attitude control gas jets is likely to have a component that changes the spin rate, thus requiring the data immediately after each attitude correction event (ACE) to be used to re-estimate the spin rate so that the CCD arrays can be synchronously clocked. However, for simplicity, our simulations make no allowance for such a latency time. Thus, each star is observed four times per rotation (once in each field of view by each of the two rows of detectors) as long as it lies close enough to the nominal observing plane (normal to the spacecraft’s rotation axis). We have assumed a uniform observation uncertainty of 0.35 milliarcsec (mas) in the scan direction and 3.5 mas in the cross-scan direction for each observation.

The spacecraft rotation is modeled as a hierarchy of three motions: (1) the slow, annual revolution of the Sun-spacecraft line about the normal to the ecliptic, (2) the considerably faster precession of the spacecraft’s nominal spin axis about the Sun-spacecraft line, and (3) the relatively rapid rotation about the spin axis. The first of these is treated as a known, constant, circular motion, but the other two are treated together in terms of the three Euler angles relating the spacecraft frame to the slowly rotating Sun-oriented frame. These three angles are (1) the rotational phase ϕ , (2) the azimuth ν (about the Sun direction) of the nominal spin axis, and (3) the Sun-axis angle ξ . Each angle is decomposed into a sum of Legendre polynomials with adjustable coefficients over each rotation span.

In the absence of nutations, and in the presence of only a steady torque due to (Sun) light pressure⁴, five rotation coefficients would be estimated (initial phase, phase rate, initial precession angle, precession rate, and constant Sun-axis angle). However, in a system

⁴In this study, we ignore the variations in solar luminosity, as well as the highly variable,

undergoing occasional ACE, a more realistic description would include some higher-order coefficients, as well as a fresh determination of all the coefficients after each ACE. For this reason, we use the term “rotation break” in referring to an ACE or to any other event requiring a new set of coefficients.

We consider two modes of controlling the precession of the spacecraft: (1) periodic ACE by firing gas jets and (2) a smooth torque from light pressure on the Sun shield. To simulate the first mode, we choose a set of span lengths randomly in a Gaussian distribution with selectable mean and standard deviation. In the second mode, the entire batch interval is a single rotation span. In order to operate in the second mode, it is preferable for the spacecraft both to spin and precess quickly. This is, in part, due to the relatively large torque developed by the solar radiation pressure on the Sun shield in most spacecraft configurations. (See Reasenberg 1997, 1999a, 1999b.) Since the tracks of the star images on the CCD detectors must be aligned moderately well with the CCD columns, fast precession requires fast rotation as well. These considerations drive the choices of parameter values in the studies that follow.

3.2. Star observations

In our simulations, the sky is populated by stars thrown randomly with a uniform probability density of three per square degree, corresponding roughly to the mean density of stars between $V=8$ and $V=10$ (Allen 1976, p 243). (In so doing, we ignore the five-fold increase in the star density from the galactic pole to the galactic plane.) The sky model in our simulations includes two adjustable position parameters for each star observed. The number of observations made during a batch interval depends on the length of the batch

but small, torque due to solar wind pressure.

interval, the average density of stars, the rotation rate, and the angular width of sky swept out by the detector array. However, the number of distinct stars seen (and thus the total number of parameters to estimate) depends also on the precession rate.

3.3. Data analysis

A brute-force approach to estimating all these parameters via least-squares techniques would require enormous amounts of computer time. The case we are considering involves the observation of a few thousands to a few tens of thousands of stars in each batch interval, thus requiring many thousand parameters for the star coordinates, in addition to a few hundred to a few thousand for the rotation model. The direct solution of the resulting system of normal equations – even the smallest of those we are considering – would require about three days of computing on an ordinary desktop computer. There is a severe time penalty due to solving for the star coordinates, which we do not need at Stage A. Indeed, the coefficient matrix is only sparsely filled, and the processing time is much shortened by pre-reducing the matrix in batches of observations from one star at a time. (See Appendix A for an introduction to pre-reduction of the normal equations.) This technique eliminates the uninteresting parameters (the star coordinates in this case) and produces normal equations for the remaining parameters that give results, covariances, and postfit statistics that are algebraically identical to those that would have been obtained from the full normal equations. Although the pre-reduction effort is proportional to the (large) number of stars observed, it runs about two orders of magnitude faster than the direct inversion in this scenario.

To avoid numerical problems in certain low-probability situations, three kinds of constraints have been applied. First, each rotation coefficient is assumed to be known *a priori* with an uncertainty that gives a peak amplitude of 1000 arcsec over a span.

This constraint makes little difference in the results, since the solutions generally give uncertainties many orders smaller. Second, each randomly chosen span length is forced to lie within three standard deviations of the mean by discarding any value outside that range and replacing it by a newly chosen random value. The very last span may, of course, be shorter than this limit, but the potential numerical problems of a very short span are avoided by the *a priori* constraints. (In processing real data, such short spans would not be created artificially, as we have done here to make different sets of results intercomparable.) Third, the star positions are assumed to be known *a priori* with some uncertainty, as discussed in Section 2. Throughout this study, except for a few test cases presented in Table 4, the *a priori* position uncertainty is 10 mas.

3.4. Figure of merit

A successful Stage-A analysis should yield a “rigid” model of the instrument rotation. To understand this concept, consider the pointing directions of the instrument at two separate epochs within the batch interval, and the angle $\Delta\phi$ between those pointing directions. Consider also the nominal star-position uncertainty (in the scan direction) based on averaging over all measurements of a star during a batch interval, $\sigma_{batch} \approx \sigma_0/\sqrt{N_{batch}}$, where σ_0 is the single-measurement uncertainty, and N_{batch} is the number of observations of a star during the batch interval. By a rigid model we mean that $\sigma(\Delta\phi) \ll \sigma_{batch}$. In the present case, there can be 4 measurements of a given star during a single instrument rotation, and a (minimum) average of 6.2 successive rotations in which a given star would be measured.⁵ Thus, $N_{batch} \geq 25$, and the nominal star-position determinations would have an uncertainty $\approx \sigma_0/\sqrt{25}$ or better.

⁵This applies to Versions 4, 5, and 6 in Table 1. Stars near the region of maximum observing-band overlap may be observed a much larger number of times during a batch

If the rigid-model criterion were not met, then there would be a large correlation among the estimated positions of stars nearby on the sky, and the average position of a group of such stars would have an uncertainty not much smaller than the uncertainty in the position of one. If we require that we be able to obtain an improved average position of, say, a dozen stars in a small patch of sky, then we must have $\sigma(\Delta\phi) < \sigma_0/20$. This requirement also helps to guard against regional biases.

The figure of merit $\bar{\sigma}$ used in this study is based on the rigid-model concept. It is formed by selecting an evenly spaced grid of epochs covering the entire batch interval, with 100 such epochs for every six rotations, plus one at the end (thus dividing each six rotations into 100 intervals). We then calculate the variance for the modeled difference ($\Delta\phi$) in rotation phase ϕ between pairs of epochs, using standard error propagation from the least-squares solution. The logarithms of these variances are averaged for all possible pairs corresponding to each lag, and, further, the results of 16 separate runs, with independently chosen sets of rotation breaks (but the same set of stars), are averaged together to smooth out any dependence on the exact lengths of the rotation spans. Examination of plots of unaveraged cases (not shown), verifies that averaging over 16 runs is helpful for this purpose. Finally, the values are averaged for all lags from 0.5 to 3 rotations, and the result is expressed in the form of $\bar{\sigma}$, the rms angular uncertainty for that set of runs.

4. Results of the simulations

In this section, we examine four kinds of results. First, plots of $\sigma(\Delta\phi)$ provide some understanding of the detailed cohesion of the rotation model. Second, the figure of merit $\bar{\sigma}$ as a function of the basic angle shows a broad region of acceptable values with a few

interval.

isolated bad points. Third, the figure of merit for a variety of operating modes shows that it is important to avoid rotation breaks and that, in the absence of rotation breaks, it is advantageous to have long batches of data, up to a saturation point around 36 rotations. Finally, we discuss some tests that support the validity of our analysis.

4.1. $\sigma(\Delta\phi)$ versus $\Delta\phi$

Figure 1 shows typical examples of the average uncertainty as a function of lag for frequent rotation breaks due to ACE and for smooth precession with only a single rotation span per batch interval. Because of the large number of independent rotation models in the former case, we have included two examples of the latter: one with the same degree of complexity in the single rotation model (3 ϕ coefficients) as in each of the models in the frequent-ACE case and another with a much more complex model (108 ϕ coefficients) resulting in the same number of ϕ coefficients overall as in the frequent-ACE case. There is some structure in all three examples. The first feature is a relatively deep dip near zero lag, which is much deeper in the case of a single rotation span with a simple model. Not surprisingly, in the multi-span case, there are also dips in the uncertainty at lags corresponding to multiples of the rotation period, but not quite so deep. The same star being observed on successive passes will help to tie together the coefficients of the rotation models for spans separated by a whole number of rotations. The amplitude of these dips in uncertainty is always less than a factor of 3 in all the cases studied. In the two examples with simple rotation models, there are only faint traces of dips at intervals corresponding to the basic angle (70 deg in this figure), but such dips show plainly in the example with a complex rotation model. For simple models, the effect is largely suppressed by the smoothing implicit in the limited numbers of coefficients.

It was after examining plots of these kinds that we selected the figure of merit $\bar{\sigma}$. The

motivation for averaging lags from 0.5 to 3 rotations is threefold: (1) to avoid the deep dip near zero lag, (2) to include as much in the average as possible (even for the smallest batches studied), and (3) to include only lags that permit sampling the entire batch interval. The remainder of this section deals with $\bar{\sigma}$ in a variety of cases.

4.2. Figure of merit versus basic angle

In a preliminary series of tests, we held the basic angle fixed at 45 deg and varied the number of ϕ coefficients over a wide range, from as few as two to as many as fifty per span. The case with the most complex ϕ models, but with no rotation breaks, included about the same number of ϕ parameters overall (50) as the cases with the least complex ϕ models and frequent breaks (72). We found that the figure of merit was quite insensitive to the number of ϕ coefficients estimated, changing by less than 25% from best to worst with all other conditions held fixed. In contrast, we found that the runs with no rotation breaks were about a factor of 9 better than the corresponding runs with six breaks per rotation, as long as the total number of parameters was roughly comparable. These preliminary tests are not discussed further.

In order to learn the effect of the basic angle on mission performance, we have done our simulations in sets with values of the basic angle in steps of 10 deg, covering the range from 0 to 180 deg (even though the 0 deg case is degenerate). In practice, the simulation software uses algorithms that require the basic angle to differ from both extremes (0 and 180 deg) by at least the angular offset between the two rows of detector elements (taken to be 0.1 deg in this study). Therefore, we have offset the two endpoints correspondingly.

Using this scheme, we have examined three different versions of spacecraft operation, one with slow rotation (period of 2 hr) and two with an intermediate rotation (period of 30

min). Three additional versions, with the current nominal fast rotation (period of 20 min), are included in the analysis described in Section 4.3. The slow rotation is a nominal from the FAME-95 concept, while the intermediate rotation rate was a guess at the rate that might be used for the FAME-98 design. As mentioned above, the faster rotations go with correspondingly scaled-up precession rates.

Because of a redesign of the optics to have a shorter focal length, which is needed to maintain star dwell time on the detector with a faster rotation (Phillips & Reasenberg 1998), the three groups of versions also differ in their fields of view (0.75 deg, 1.6 deg, and 2.2 deg, respectively).⁶ These increases lead to a larger minimum overlap of the observing spiral band between successive instrument rotations (50%, 75%, and 83%, respectively). Thus, in the three groups of versions, each star seen by the instrument is observed on a minimum of two, four, or six successive rotations, respectively (except for some stars seen near the beginning or end of the batch interval). The remaining distinction between versions is the duration of a batch interval. Because we do not yet know the limits on this duration, we have explored values varying by up to a factor of 24 in number of rotations. To make comparisons easier, we arranged for both intermediate versions to match the slow version: one in the number of rotations and the other in the number of hours. Also, two of the fast versions have batch intervals of the same lengths as the two intermediate versions. Indeed, when we began this study, we expected that the batch duration in hours might be an important factor in understanding the results. However, since we have scaled the rotation and precession rates together, while the relatively slow annual motion of the Sun has little impact on the results of a spiral reduction, and since we have not adjusted the assumed measurement uncertainty as we changed the dwell time of a star image on a single CCD, the best measure of duration is the rotation phase ϕ .

⁶The present nominal field of view for FAME-98 is 2.2 deg in diameter.

For each of these versions, we consider three patterns of precession, one with short spans between rotation breaks (6 per rotation), one with long spans (1 per rotation), and one with no rotation breaks. For the two patterns with rotation breaks, we use the same degree of complexity in the rotation model: for each span, there are 9 coefficients (5 for ϕ , 2 for ν , and 2 for ξ). For the pattern with no rotation breaks, we similarly include 4 coefficients for the spin vector orientation (2 for ν and 2 for ξ), plus either 16 or 4 ϕ coefficients per rotation in the span — high and low resolution cases.

The results of all these runs are shown in Figures 2, 3 and 4. Each part of each figure shows a summary of a full set of runs at different basic angles. Though they differ in detail, they present a remarkably consistent picture. In each, there is a broad minimum in the uncertainty, punctuated perhaps by a few “bad” angles, and the angles near 0 and 180 deg are particularly bad. The differences among the figures can be characterized largely by (1) the depth of the minimum, (2) the relative badness of the bad angles, and (3) the degree of asymmetry about 90 degrees. Not all the bad angles appear in each figure. Interestingly, Figure 2c has a small bump at 70 deg; a separate run (not plotted in the figure) with a 72 deg basic angle shows a (slightly) higher uncertainty than 70 deg. Other runs for exactly $1/7$ and $1/8$ of 360 deg show no hint of bumps at those angles.

It is easy to understand the existence of intrinsically bad basic angles. By observing the same star through both fields of view, the instrument ties together the rotation model with many links offset by the basic angle. When the angle is commensurate with a complete rotation, these many links repeatedly tie the observing spiral together into nearly disjoint subsets, but without linking the batch as a whole, i.e., without including links between subsets. (Because of each star is seen on several successive rotations, there are already many such links at 360 deg.)

Since the rotation breaks are not spaced uniformly, the “poisonous” effect of

commensurability is smeared out by somewhat more than the average width of a rotation span (60 deg in Figures 2a and 2d, 360 deg in Figure 2b). This explains why the bump at 60 deg is scarcely noticeable and why it is more noticeable in Figure 2d (with more nearly uniform span lengths) than in Figure 2a. It also explains why 120 and especially 90 are much less severe than 180 and why the bump at 180 is so much reduced in Figure 2b. A corresponding mechanism operates in Figure 2c to suppress the badness at the special angles smaller than 90 deg — the effective angular resolution (in inverse cycles) of the rotation model is given approximately by the number of ϕ coefficients divided by four times the number of rotations. (We have verified that reducing the number of coefficients from 95 to 59 suppresses the bumps at 90 deg and 120 deg. The results are otherwise very similar to Figure 2c and are not shown here.)

Figure 3 (the first intermediate version) is especially similar to Figure 2 (the slow version), since the durations (spin period, batch interval, etc.) have all been scaled alike. The difference comes from the wider field of view in the latter, which results in more observations being made (about 41,000 vs about 19,000). Quantitatively, we expect the uncertainties to be smaller with more observations, and indeed they are. The statistical advantage of the larger number of observations explains a decrease in the mean uncertainty by a factor of 1.47. The actual factors of decrease range from 1.80 to 2.04, suggesting an advantage from the enhanced cross linking.

4.3. Tabulations of $\bar{\sigma}$

A summary of the results from the previous subsection is shown in Tables 1, 2, and 3, which reflect the co-evolution of this study and the instrument concept. To make the comparisons plain, we have condensed these results by including only the values for basic angles of 70 deg (a possible selection suggested by some aspects of the instrument’s optical

design) and 100 deg (approximately the minimum of each curve in the figures). The results for 70 deg are slightly worse than for 100 deg in all cases, but never by more than 20%, and typically by much less. In addition, we include here the analysis of the three fast-rotation versions, for which only the two representative basic angles are considered, not the full range.

Table 2 also shows the results of some simulations designed to evaluate the possible benefits of combining batch intervals for analysis. The lines labeled “4 spans” represent three ways of making such combinations; an entry here under Version 3 is rather like a concatenation of 4 low-resolution batches of Version 2, for example. The first of these lines shows cohesions that are actually worse than the corresponding values for the single-span, next-to-the-left versions that were, in effect, combined to produce this line. This apparent worsening of the cohesion is an artifact of the definition of the figure of merit, which is an average over all $\Delta\phi$'s within our chosen range of 0.5 to 3 rotations, including those that step across rotation breaks. The second 4-span line shows that excluding the step-across regions from the figure of merit does result in an “improvement,” but still does not come close to the 1-span case for the same version. Some improvement was expected, since there are only 60% as many independent stars, but the same number of star observations, in the 4-span composite. The third 4-span line uses the same modified figure of merit and also imposes fixed values on the coefficients for the ν and ξ models. In comparison with the last line, which has the same simplified models, this third 4-span line shows that the need to estimate the spacecraft axis direction is what prevents the 4-span analysis from doing as well as the corresponding 1-span case.

4.4. Verification of technique

Where possible, as stated above, the simulations with rotation breaks are repeated with 16 different randomly chosen sets of span lengths, and the results are averaged to suppress the statistical noise from the span lengths. However, one set of simulations (48 hr batch interval with long rotation spans) requires so much computing time for each case that we have dispensed with the 16-case averaging for that set. With the longest batch intervals, the need for averaging is reduced, since there are so many spans per batch (144, compared with only 6 for Version 1 with long spans). We also dispensed with 16-case averaging in the detailed simulations covering the whole range of possible basic angles for short spans in Version 3 (though not in the data shown in Tables 2 and 3). In so doing, we verified (for a single value of the basic angle) that the difference between the 16-case average and the single case is acceptably small. The 16 individual cases show a mean of -2.617 for the log of the uncertainty and a root-mean-square of 0.010 deviation about that mean.

We considered the possibility that our Monte Carlo sky might skew the results of the simulations. To test that hypothesis, we made pairs of runs for many cases with two different Monte Carlo skies, created by using two different random number seeds for the star selection. We found that the largest difference due to switching Monte Carlo skies was less than 3% in the figure of merit. We conclude that the sets of stars are so large (> 2000 stars) that a great deal of smoothing must occur naturally; the behavior we see is intrinsic and insensitive to the detailed distribution of star positions.

We also considered the effect of varying the strength of the *a priori* constraints on the positions of the fiducial stars and of varying the complexity of the model of spin axis orientation (parameters ν and ξ). Table 4 shows the results of these tests for one particular case: Version 2 with no rotation breaks and a low-resolution ϕ model.

As shown in the table, there is a striking degradation of the figure of merit whenever

the precession rate $\dot{\nu}$ is estimated and a further degradation at the weakest *a priori* level when the Sun-axis angle ξ is estimated. This effect of precession rate stems from the geometry: the precession and spin are not orthogonal, and the cross-scan measurements, which break the degeneracy between precession and spin, have an uncertainty ten times as large as the scan-direction measurements. The resulting greater uncertainty in the rotation rate gives a progressively greater uncertainty in $\Delta\phi$ with greater lags and, thus, a greater average uncertainty in $\Delta\phi$. (There is a corresponding degradation in the uncertainty of the initial rotation phase that comes from estimating the initial precession angle, but that has no effect on the figure of merit, which is computed on angle differences.) The effect diminishes slightly with stronger *a priori* constraints on the star positions, but does not vanish. (To confirm this behavior, we have examined one further tenfold decrease in the *a priori* uncertainty of star positions, but we have excluded these numbers from the table as “unrealistic.”) By the same token, when ξ is not assumed to be known, and the star positions are not constrained tightly enough, ξ must be estimated from the same (cross-scan) information that disambiguates the rotation and precession rates, even if the precession rate itself is not being estimated. We see this effect in the leftmost column of Table 4, where the star *a priori* position uncertainties are 100 mas. We have also condensed the results in Table 4 to the cases with equal numbers of ξ and ν coefficients and extended them up to 10 each in order to see the effect of modeling more complex motions of the spacecraft axis. See Figure 5. Unlike irregularities in the spin rate, these higher-order terms in the precession and nutation appear to degrade the model cohesion quickly.

Note that, in analyzing real data, there will not be an option to estimate fewer than two parameters each for ν and ξ . The experiments summarized in Table 4 were intended only to give insight into the problem.

5. Discussion

The similarities of the plots in Figures 2, 3, and 4 lead to several conclusions about the design of the instrument. Most obviously, there is very little difference in performance over a wide range of the basic angle, so that other considerations can govern the choice. Moreover, the “bad” basic angles should be avoided if at all possible. Indeed, the bad angles are so few that avoiding them should be relatively easy.

The clearest point revealed by these studies is that avoiding frequent rotation breaks should be a high priority. The worst single-span case shown in Table 3 (fewest observations, high-resolution rotation model) gives a better result than the best case with multiple rotation breaks, even though the latter represents almost two orders of magnitude more observations (albeit only about three times as many per revolution). A more even-handed comparison would entail equal numbers of parameters and observations on both sides, and we can approximate such a comparison by considering three cases within any given version of operation: multiple breaks with short and long spans vs a single span with high resolution. These three have, respectively, 54, 9, and 16 orientation parameters per rotation. In comparing the first two cases against the third, we see that the single-span case is better by factors of about 12 and about 3, respectively. The “fair” comparison must lie somewhere in between.

Not surprisingly, the case with a mean of 1 break per rotation period (long spans) displays a cohesion that is intermediate between the two extremes represented by short spans and no breaks at all. Thus, any reduction in the frequency of rotation breaks can be expected to improve the instrument performance. This is explicitly shown for the range from 6 breaks per rotation to 1 break per 36 rotations. We see no limit to the extension to more frequent breaks. For less frequent breaks, the limited examples of the Version 6 results (when compared to the Version 5 results) suggest that the effect saturates around

36 rotations. However, this should be investigated as the other aspects of the instrument become better understood, permitting realistic modeling of effects that may limit the improved performance with increasing number of rotations per break.

The severe penalty for having frequent ACE or other rotation breaks is explained in part by examination of Table 4 and Figure 5. After each break, the new precession rate must be estimated, and the table reveals the significant degradation in the figure of merit resulting from the need to estimate the precession rate.

Interestingly, the improvement seen in Tables 2 and 3 from Version 1 to Version 2 is more than the simple “root-N gain” from the increased number of observations. In all four comparable cases in each table, the improvement is about a factor of 2, almost as much as the increase in the number of observations. We take this to mean that the improvement is partly due to breaking degeneracy in the solutions; observing each star more times during the batch interval leads to a greater degree of interconnection of the rotation model. The improvement is less from Version 2 to Version 3, and even less from Version 4 to 5 and 6, to the point of saturation, as noted above.

In order to realize the advantages of avoiding frequent ACE, while precessing at the necessary rate, we prefer to use the torque from solar radiation pressure, rather than from gas jets. Should further study show that this is not a workable solution, there is a way of using gas jets that is less disruptive than having frequent firings. Precession would be achieved by a pair of firings separated by about 0.25 rotation. Between the firings, the spacecraft would experience a large Eulerian nutation and might not be able to take data. The second firing would end the nutation and leave the spin vector shifted. Such pairs would be repeated every 1.5 spacecraft rotation. Assuming $I_3/I_1 = 0.8$, this approach would use only 20% more fuel than the optimal (fuel minimizing) approach of having many firings per rotation and thus having the spacecraft in a continual string of Eulerian nutation

states (Reasenber 1998). If observations are not made during the quarter rotation between torsional impulses, then there is a 17% loss of observing time, but, while observations are being made, the star images move along a fixed path in the detector plane, thus increasing the precision of the cross-scan measurements. (Further analysis of this approach must include the variance in the jet pointing and thrust, as well as the requirements for thruster placement.)

Because of the hierarchical nature of the data analysis, each stage can be viewed in isolation from the rest, taking inputs from the preceding stage and providing outputs for the next. The overall return from the mission, then, depends upon the successful operation of each stage. From Tables 2 and 3, we see that some of the options considered give cohesions for Stage A that are much better than the design goal of 0.05 mas for the minimum uncertainty of bright-star positions. Assuming that Stage B works as well at knitting the individual batches into a global model of the spacecraft orientation, then the only further requirement is that the mission last long enough to yield the desired uncertainties with respect to the spacecraft reference.

The work described in this paper was supported in part by the USNO and by the Smithsonian Institution.

A. Partial pre-reduction

Here, we set forth the algebraic basis for the technique of partial pre-reduction (PPR) of the normal equations for weighted least-squares parameter estimation. The technique has two parts: the reduction of the normal equations *per se*, which removes a subset of the variables from the problem, and the reconstruction of estimates and variances of the reduced-away parameters. (See Reasenber 1975 and Chandler 1989, respectively, for

further details.) The PPR technique is applicable when the parameter set can be partitioned into two subsets (which we will call the “interesting” and “uninteresting” parameters), and there is no (immediate) need for the covariances and solutions for the uninteresting parameters. Since the effort required to invert a square matrix goes as the cube of its linear dimension, there is a potentially large computational advantage in reducing the size of the matrix to be inverted, even if that entails some extra manipulation. One is motivated to use PPR when the uninteresting parameters are numerous and either multiple inversions are required for the interesting parameters (e.g., to look at solutions for subsets of these) or the portion of the coefficient matrix for the uninteresting parameters is block diagonal.

When the parameter set has been partitioned into two subsets, the normal equations can be rewritten from the ordinary form

$$BX = U \tag{A1}$$

to the partitioned form

$$\begin{pmatrix} C & F \\ F^T & D \end{pmatrix} \begin{pmatrix} Y \\ Z \end{pmatrix} = \begin{pmatrix} V \\ W \end{pmatrix} \tag{A2}$$

where the symmetric coefficient matrix B has been partitioned into symmetric pieces C and D and rectangular sub-matrices F and F^T , the vector of unknown parameter adjustments X has been similarly partitioned into Y and Z , and the right-hand side vector U into V and W . The sub-matrices C , Y and V all pertain to the interesting parameters while D , Z and W all pertain to the uninteresting ones; F provides the connection between the two subsets. Equation A2 can be reduced to a purely “interesting” set of normal equations by the PPR process, which starts with the following additional quantities

$$\begin{aligned} \bar{C} &= C - FD^{-1}F^T \\ \bar{V} &= V - FD^{-1}W \\ \bar{F} &= FD^{-1} \end{aligned} \tag{A3}$$

$$\begin{aligned}\bar{Z} &= D^{-1}W \\ \bar{S}_0 &= S_0 - \bar{Z}^T W = S_0 - W^T D^{-1}W\end{aligned}$$

where S_0 is the weighted prefit sum-squared residual. By combining these definitions with equation A2, we obtain the reduced matrix equation for Y :

$$\bar{C}Y = \bar{V} \tag{A4}$$

Not only does the inversion of \bar{C} in equation A4 yield the desired solution vector Y , but it can also be shown that \bar{C}^{-1} is the upper left corner of B^{-1} , the inverse for the whole set of normal equations. In terms of the quantities defined in equations A3, it is easily shown that

$$B^{-1} = \begin{pmatrix} \bar{C}^{-1} & -\bar{C}^{-1}\bar{F} \\ -\bar{F}^T\bar{C}^{-1} & D^{-1} + \bar{F}^T\bar{C}^{-1}\bar{F} \end{pmatrix} \tag{A5}$$

gives the inverse of the partitioned matrix in equation A2. Thus, the covariances for the interesting parameters follow from the solution of the pre-reduced equations in the usual way. Moreover, the reduced prefit sum-squared residual \bar{S}_0 serves in the same way, since the postfit sum-square is given by

$$S = S_0 - X^T U = \bar{S}_0 - Y^T \bar{V} \tag{A6}$$

It follows that the pre-reduction can be applied in successive stages, with only a few parameters being reduced away at each stage. Such a procedure may avoid unnecessarily inverting any large matrices at all.

We also find that the “uninteresting” solutions and variances can be obtained at this point with little extra effort from the quantities already computed. The solutions are given by

$$Z = \bar{Z} - \bar{F}^T Y = D^{-1}(W - F^T Y) \tag{A7}$$

and the variances by the diagonal elements of B^{-1} shown in equation A5. In essence, the time savings of partial pre-reduction stem from the many elements of B^{-1} that are

never calculated. If the strict reduction is relaxed merely to the extent of calculating the “uninteresting” variances, the added computation time is kept to a minimum.

REFERENCES

- Allen, C. W. 1976, *Astrophysical Quantities* (London: University of London Athlone Press)
- Bastian, U., Gilmore, G., Halbwegs, J. L., Høg, E., Knude, J., Kovalevsky, J., Labeyrie, A., van Leeuwen, F., Lindegren, L., Pel, P. W., Schrijver, H., Stabell, R., & Thejll, P. 1993, “ROEMER, Proposal for the Third Medium Size ESA Mission (M3),” Lund (available on request from L. Lindegren, email: lennart@astro.lu.se)
- Chandler, J. F. 1989, “Restoration of Prereduced Normal Equations,” unpublished memorandum TM89-01
- Høg, E. 1993, in *IAU Symp. 156, Developments in astrometry and their impact on astrophysics and geodynamics*, ed. I. I. Mueller & B. Kołaczek (Dordrecht: Kluwer), 37
- Høg, E. 1995, in *IAU Symp. 166, Astronomical and astrophysical objectives of sub-milliarcsecond optical astrometry*, ed. E. Høg & P. K. Seidelmann (Dordrecht: Kluwer), 317
- Lindegren, L., Perryman, M. A. C., & Loiseau, S. 1995, in *SPIE Proceedings 2477, Spaceborne interferometry II*, ed. R. D. Reasenberg (Bellingham: SPIE), 91
- Perryman, M. A. C. et al. 1989, *The Hipparcos Mission*, SP-1111 (Paris: ESA)
- Perryman, M. A. C. et al. 1997, *The Hipparcos and Tycho Catalogues*, SP-1200 (Noordwijk: ESA)
- Phillips, J. D., & Reasenberg, R. D. 1998, in *SPIE Proceedings 3356, Space Telescopes and Instruments V*, ed. P. Y. Bely and J. B. Breckenridge (Bellingham: SPIE), 832
- Reasenberg, R. D. 1975, “Partial Prereduction of the Normal Equations,” unpublished memorandum TM75-03

- Reasenber, R. D. 1997, “Effects of radiation pressure on the rotation of FAME,” unpublished memorandum TM97-03
- Reasenber, R. D. 1998, “Spacecraft rotation, precession by means of gas jet impulses,” unpublished memorandum TM98-05
- Reasenber, R. D. 1999a, “Precession of the FAME spacecraft,” unpublished memorandum TM99-03
- Reasenber, R. D. 1999b, “A new approach to post-Hipparcos astrometric surveys from space,” BAAS, in preparation
- Reasenber, R. D., & Phillips, J. D. 1998, in SPIE Proceedings 3356, Space Telescopes and Instruments V, ed. P. Y. Bely and J. B. Breckenridge (Bellingham: SPIE), 622
- Röser, S., Bastian, U., de Boer, K. S., Høg, E., Röser, H. P., Schalinski, C., Schilbach, E., de Vegt, C., Wagner, S. 1997, in Hipparcos, Venice '97: presentation of the Hipparcos and Tycho catalogues and first astrophysical results of the Hipparcos space astrometry mission (Noordwijk: ESA), 777 (<http://astro.estec.esa.nl/Hipparcos/venice.html#session12>)
- Seidelmann, P. K., Johnston, K. J., Germain, M., Monet, D. G., Murison, M., Urban, S., Davinic, N., Rickard, L. J., Weiler, K., Shao, M., Vaughan, A., Fanson, J., Norris, D., Bahcall, J., DeVegt, C., Gatewood, G., Gould, A., Jefferys, W., Lasker, B., Miyamoto, M., Sandage, A., Van Altena, W., & York, D. 1998, in Advances in Astronautical Sciences, 98-154, in press

Fig. 1.— Geometric mean uncertainty in the difference in ϕ between pairs of points on the evenly-spaced grid, as a function of the separation (or “lag”) between the points. The upper curve, with the “X” symbol, represents an average of 16 runs, each with an independent set of Gaussian-distributed rotation spans of 20 min average length. Each run is set up identically, aside from the span lengths. Each span has a separate rotation model consisting of 3 coefficients for ϕ , 2 for ν , and 1 for ξ . The lower curve, with the “Y” symbol, represents a single run with no rotation breaks. The overall rotation model for this case has 3 coefficients for ϕ , 2 for ν , and 1 for ξ . The middle curve, with the reversed “N” symbol, is like the lower curve, except that the model has 108 ϕ coefficients. For all three curves, the basic angle is 70 deg, and the spacecraft precession rate is 6 deg/d.

Fig. 2.— (a) Geometric mean uncertainty in the difference in ϕ between pairs of points on the evenly-spaced grid, averaged over all lags from 0.5 to 3 rotations. The rotation period is $P = 2$ hr, and the batch interval is $B = 12$ hr. Each point represents an average of 16 runs, each with an independent set of Gaussian-distributed rotation spans of $S = 20$ min average length ($P/6$) and standard deviation of 5 min ($S/4$). Each span has a separate rotation model consisting of 5 coefficients for ϕ , 2 for ν , and 2 for ξ .

(b) Same as (a), except that the distribution of rotation spans has a mean of $S = 2$ hr (P) and a standard deviation of 30 min ($S/4$).

(c) Similar to (a), but each point represents a single run with no rotation breaks. The overall rotation model has 95 coefficients for ϕ , 2 for ν , and 2 for ξ .

(d) Same as (a), except that the distribution of rotation spans has a standard deviation of 1 min ($S/20$).

Fig. 3.— (a) Geometric mean uncertainty in the difference in ϕ between pairs of points on the evenly-spaced grid, averaged over all lags from 0.5 to 3 rotations. The rotation period is $P = 0.5$ hr, and the batch interval is $B = 3$ hr. Each point represents an average of 16 runs,

each with an independent set of Gaussian-distributed rotation spans of $S = 5$ min average length ($P/6$) and standard deviation of 1.25 min ($S/4$). Each span has a separate rotation model consisting of 5 coefficients for ϕ , 2 for ν , and 2 for ξ .

(b) Same as (a), except that the distribution of rotation spans has a mean of $S = 30$ min (P) and a standard deviation of 7.5 min ($S/4$).

(c) Similar to (a), but each point represents a single run with no rotation breaks. The overall rotation model has 95 coefficients for ϕ , 2 for ν , and 2 for ξ .

(d) Same as (c), except that the overall rotation model has 24 coefficients for ϕ (4 coefficients per spacecraft rotation).

Fig. 4.— (a) Same as (3a), except that the batch interval is $B = 12$ hr, and each point represents only one run.

(b) Same as (3b), except that the batch interval is 12 hr.

(c) Same as (3c), except that the batch interval is 12 hr. This is qualitatively similar to (3d), both in appearance and in that there are only 4 ϕ coefficients per spacecraft rotation for both.

Fig. 5.— The effects of *a priori* constraints on star coordinates and of model complexity in the orientation of the spacecraft spin axis. The four lines represent four different levels of *a priori* constraints: (no symbols) 100 mas, (X) 10 mas, (O) 1 mas, and (T) 0.1 mas. This figure corresponds to Version 2 with no rotation breaks, 24 coefficients for ϕ , and the stated numbers of coefficients for each of ξ and ν .

Table 1. Defining characteristics of the six versions of the simulations. (Results of the simulations are given in Tables 2 and 3.)

| Parameter | Version 1 | Version 2 | Version 3 | Version 4 | Version 5 | Version 6 |
|--------------------------------------|-----------|-----------|-----------|-----------|-----------|-----------|
| Field of view (deg) | 0.75 | 1.60 | | | 2.20 | |
| Rotation period (min) ^a | 120 | 30 | | | 20 | |
| Precession rate (deg/d) ^b | 6 | 24 | | | 36 | |
| Batch interval (hr) | 12 | 3 | 12 | 3 | 12 | 48 |
| Batch interval (rot'n) | 6 | 6 | 24 | 9 | 36 | 144 |
| Observations (1000's) | 19 | 41 | 165 | 85 | 343 | 1377 |
| Obs/rot'n (1000's) | 3.2 | 6.9 | 6.9 | 9.5 | 9.5 | 9.6 |
| Stars observed (1000's) | 2.1 | 3.0 | 7.4 | 4.4 | 10.9 | 36.0 |
| Stars observed/rot'n | 358 | 502 | 307 | 489 | 302 | 250 |
| Display location | Figure 2 | Figure 3 | Figure 4 | (none) | (none) | (none) |

^aThe measurement precision assumed for this study is independent of the time a star image spends on the detector.

^bIn all cases, the precession rate is 0.5 deg per rotation.

Table 2. Mean uncertainty in $\Delta\phi$ (mas) for a basic angle of 70 deg in six design/operation versions.

| Case | Version 1 | Version 2 | Version 3 | Version 4 | Version 5 | Version 6 |
|---|-----------|-----------|-----------|-----------|-----------|-----------|
| Short spans (6 breaks/rot'n) | 1.059 | 0.586 | 0.394 | 0.394 | 0.298 | |
| Long spans (1 breaks/rot'n) | 0.238 | 0.130 | 0.125 | 0.108 | 0.104 | 0.102 |
| No breaks, high resolution ^a | 0.096 | 0.047 | 0.032 | 0.032 | 0.025 | |
| No breaks, low resolution ^a | 0.083 | 0.037 | 0.018 | 0.021 | 0.012 | 0.010 |
| 4 spans, low resolution | | | 0.040 | | 0.022 | 0.012 |
| 4 spans, low resolution ^b | | | 0.034 | | 0.019 | 0.011 |
| 4 spans, low resolution ^{b,c} | | | 0.012 | | 0.010 | 0.010 |
| No breaks, low resolution ^c | | 0.013 | 0.012 | 0.010 | 0.010 | 0.010 |

^aThe terms “high” and “low” resolution refer to the ϕ model. The high-resolution cases have 16 times as many ϕ coefficients as there are rotations, while the low-resolution cases have only 4 times as many.

^bThe mean uncertainty in $\Delta\phi$ is computed using only the pairs of grid points that fall in the same rotation span.

^cThe rotation model coefficients for ν and ξ are assumed to be known, and are not included in the solutions.

Table 3. Mean uncertainty in $\Delta\phi$ (mas) for a basic angle of 100 deg in six design/operation versions.

| Case | Version 1 | Version 2 | Version 3 | Version 4 | Version 5 | Version 6 |
|---|-----------|-----------|-----------|-----------|-----------|-----------|
| Short spans (6 breaks/rot'n) | 1.005 | 0.537 | 0.376 | 0.365 | 0.285 | |
| Long spans (1 breaks/rot'n) | 0.230 | 0.128 | 0.124 | 0.106 | 0.103 | 0.101 |
| No breaks, high resolution ^a | 0.092 | 0.045 | 0.028 | 0.028 | 0.021 | |
| No breaks, low resolution ^a | 0.082 | 0.036 | 0.017 | 0.021 | 0.012 | 0.010 |

^aThe terms “high” and “low” resolution refer to the ϕ model for the single rotation span. The high-resolution cases have 16 times as many ϕ coefficients as there are rotations in the batch interval, while the low-resolution cases have only 4 times as many.

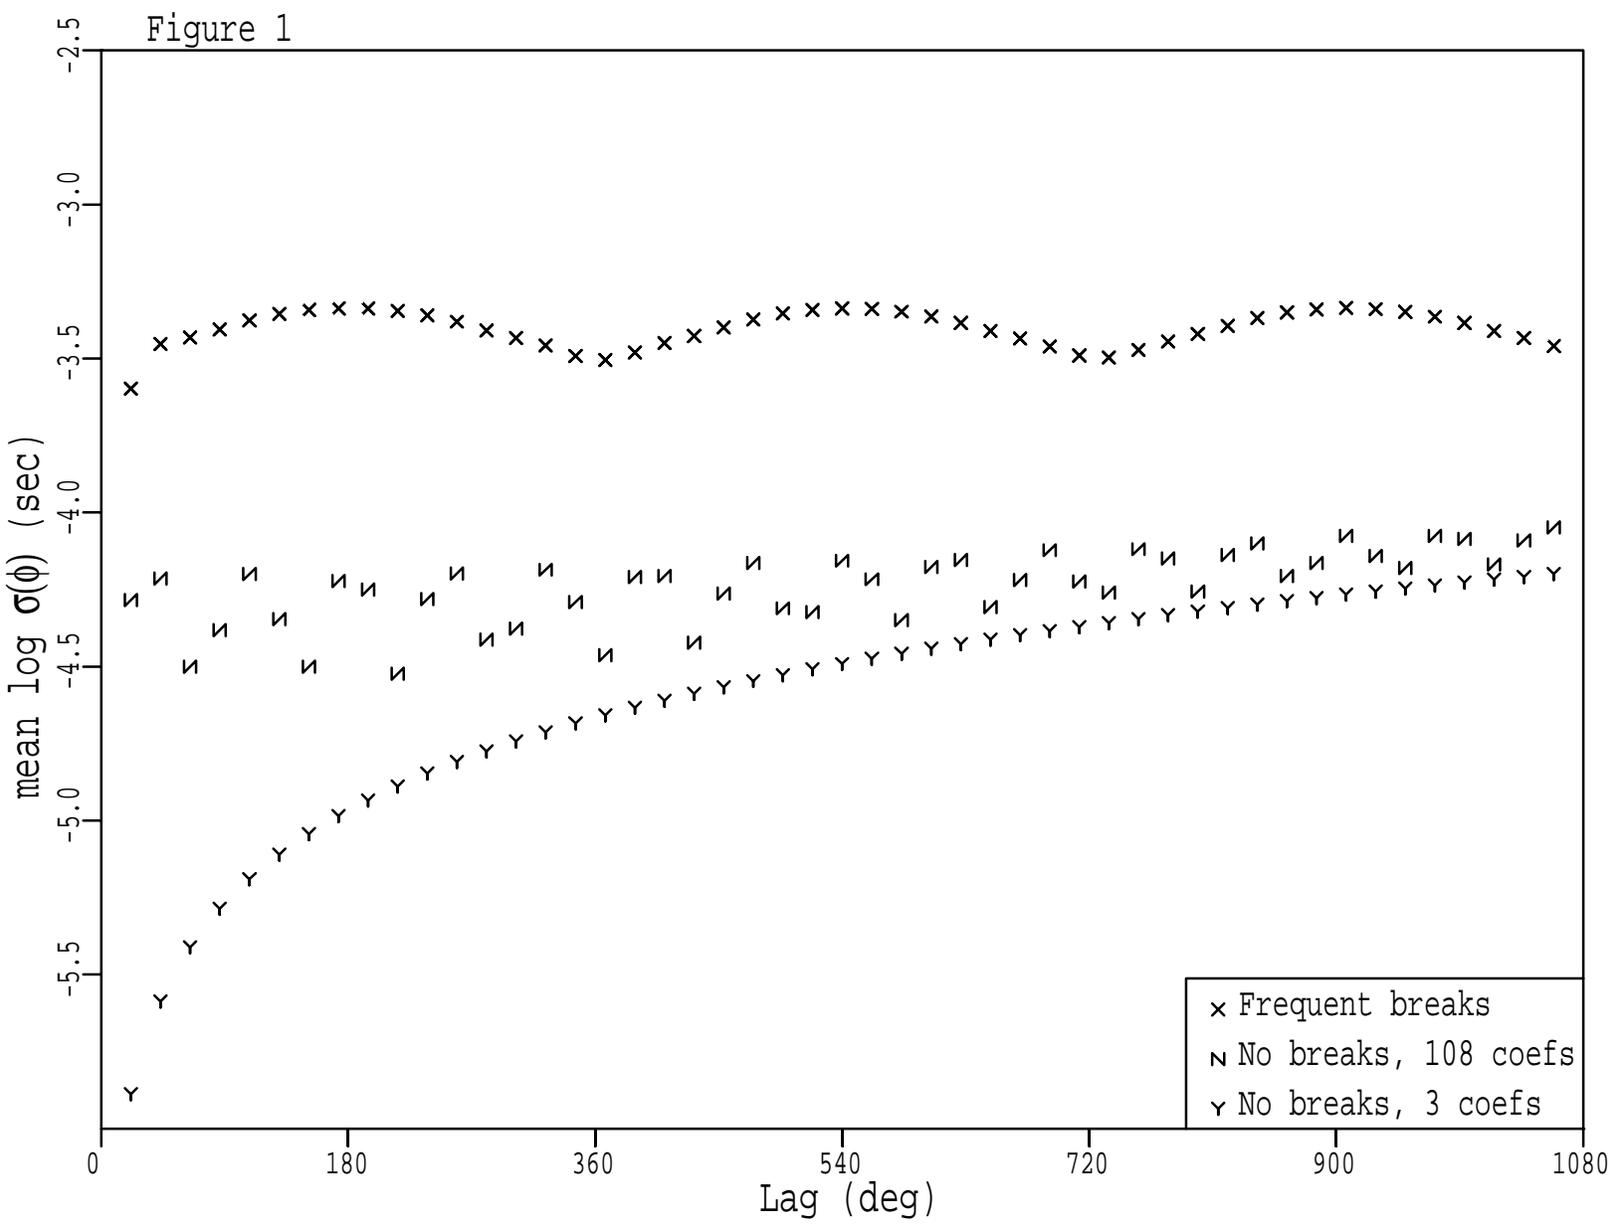
Table 4. Mean uncertainty in $\Delta\phi$ (mas) for a basic angle of 70 deg with four levels of *a priori* star position constraint, using Version 2^a with no breaks and low resolution.

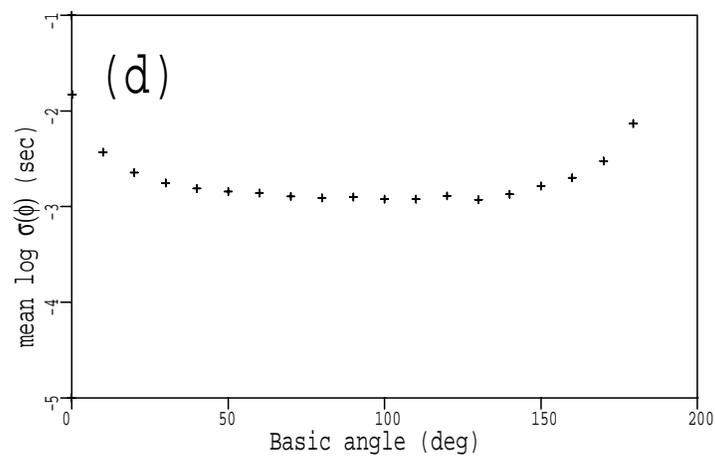
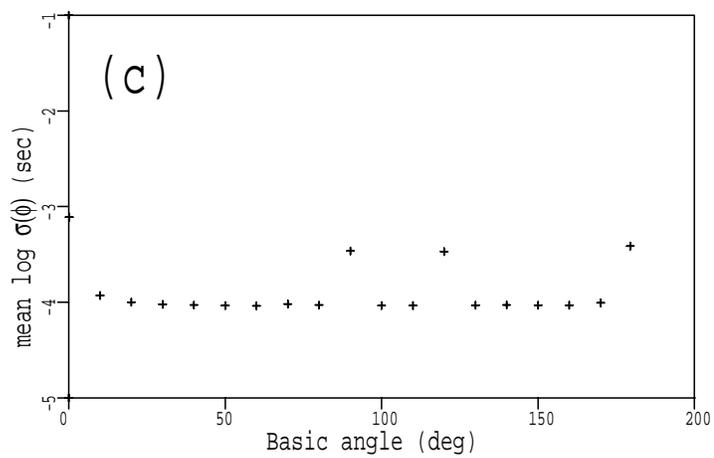
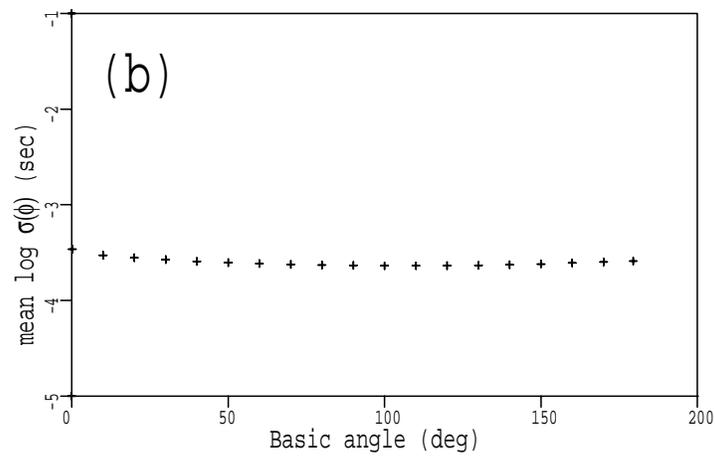
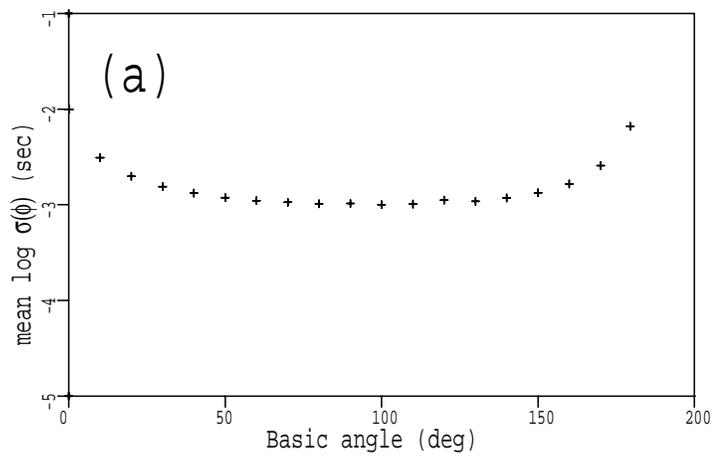
| ν terms ^b | ξ terms ^b | 100 mas | 10 mas | 1 mas | 0.1 mas |
|--------------------------|--------------------------|---------|--------|-------|---------|
| 0 | 0 | 0.013 | 0.013 | 0.012 | 0.012 |
| 0 | 1 | 0.025 | 0.013 | 0.012 | 0.012 |
| 0 | 2 | 0.025 | 0.013 | 0.012 | 0.012 |
| 1 | 0 | 0.013 | 0.013 | 0.012 | 0.012 |
| 1 | 1 | 0.026 | 0.013 | 0.012 | 0.012 |
| 1 | 2 | 0.026 | 0.013 | 0.012 | 0.012 |
| 2 | 0 | 0.037 | 0.036 | 0.029 | 0.026 |
| 2 | 1 | 0.065 | 0.037 | 0.029 | 0.026 |
| 2 | 2 | 0.065 | 0.037 | 0.029 | 0.026 |

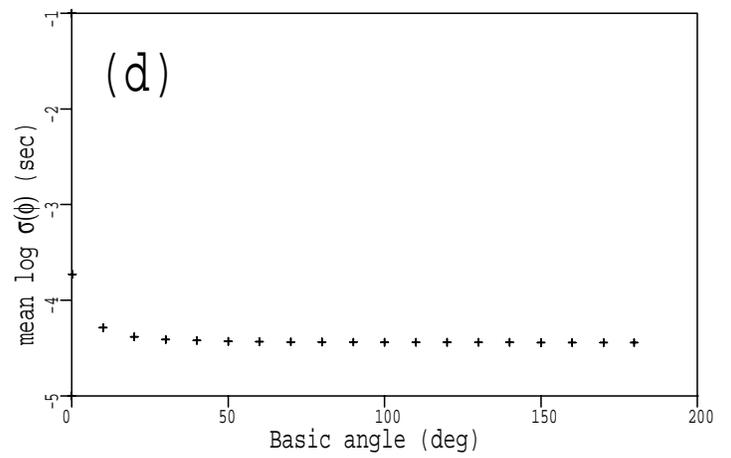
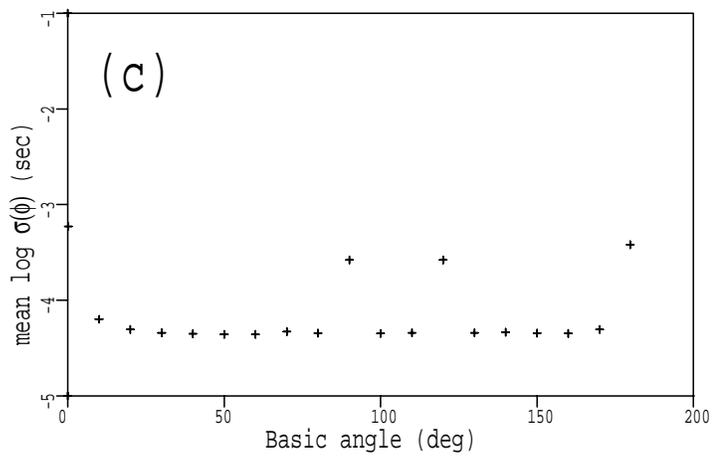
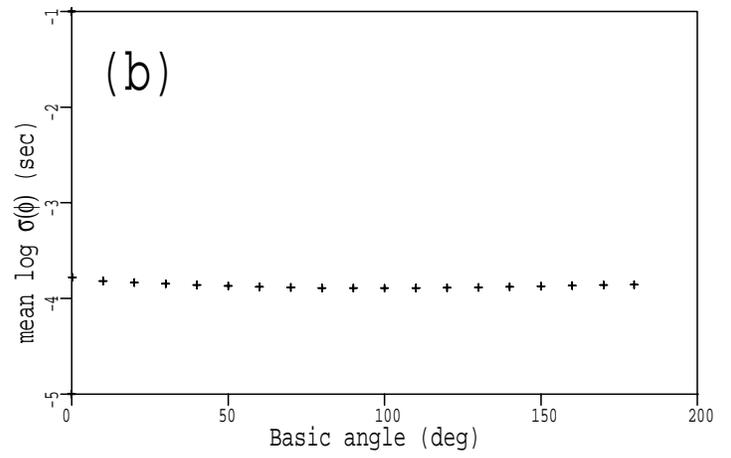
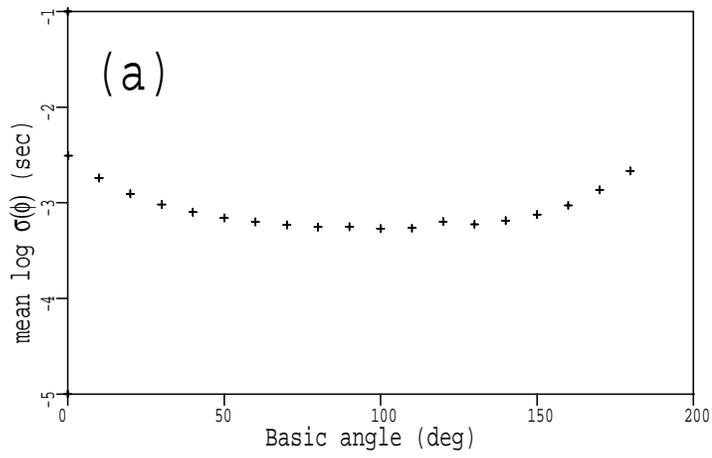
^aSee Table 1 for a description of Version 2.

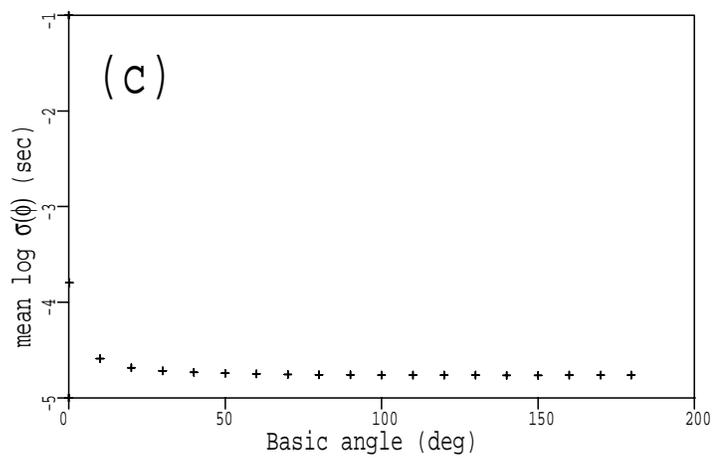
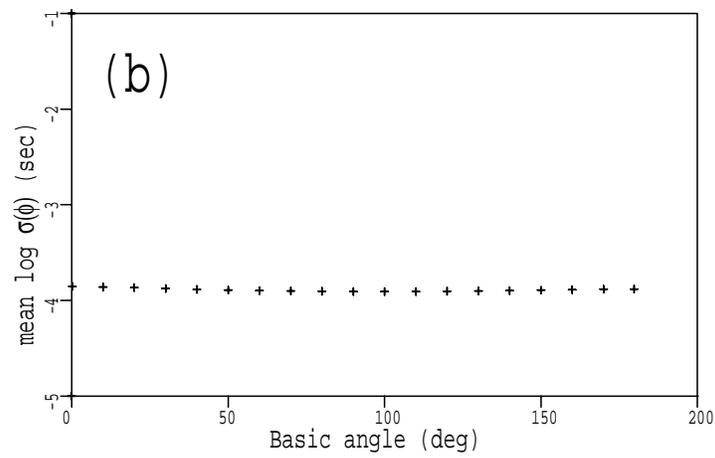
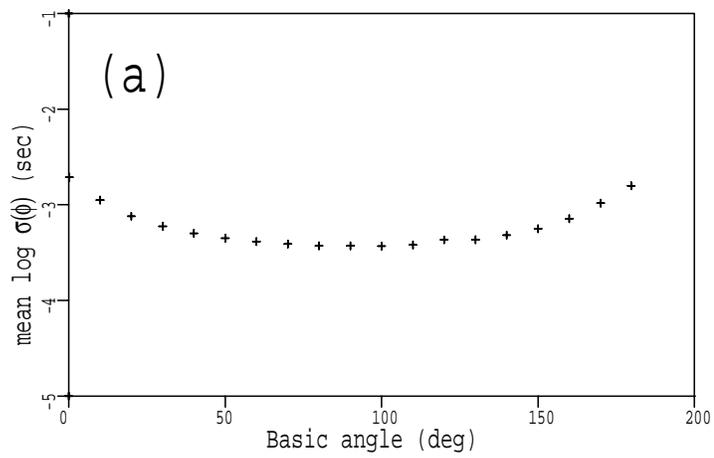
^bThe rotation models are all “low resolution” as described in Table 2, i.e., 4 ϕ terms per rotation, but the numbers of ν and ξ terms estimated are different in each case. “1” means just an initial angle while “2” means an initial angle and a rate.

Figure 1









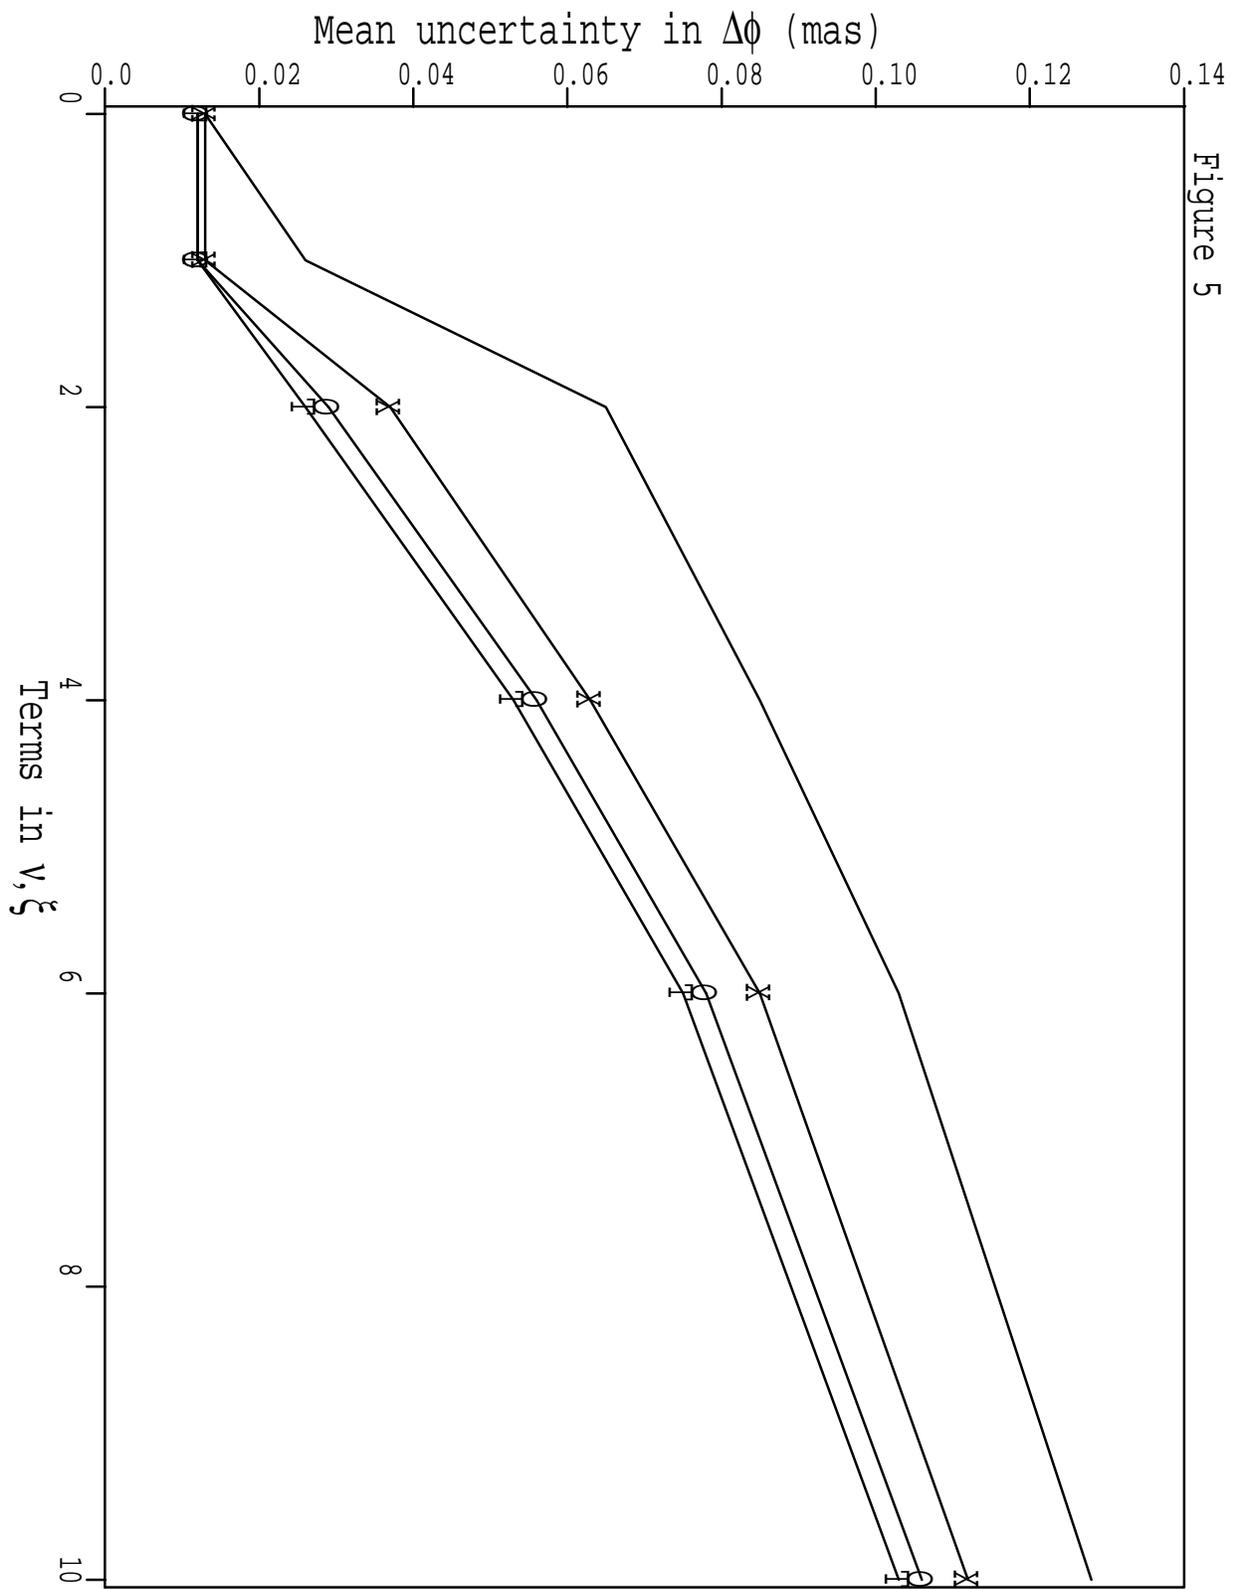


Figure 5