Section I
ADVANCES IN GUIDANCE, NAVIGATION AND CONTROL

SESSION I

Chairperson:	Bernard Kaufman Naval Research Laboratory
Co-Chairperson:	Lt. Tina-Marie D'Ercole SPAWAR
Local Chairperson:	David Eagle Martin Marietta Astronautics

The following paper numbers were not assigned:

AAS 90-017 to -019

PHASE-LOCK ROLL CONTROL FOR INERTIALLY-POINTING SPACECRAFT BY CORRELATIONS OF STAR INTENSITY PROFILES WITH A STORED REFERENCE

Bradford W. Parkinson* and Jeffrey Roger Crerie†

The spacecraft which houses the Stanford Relativity Gyroscope Experiment (GP-B) is designed to roll with a 10 minute period about its pointing axis. A precise measurement of roll phase is needed to demodulate the two relativity effects being measured by GP-B, and a steady roll rate is desired to effectively average disturbances to the experiment. Optimally, the spacecraft will be flown with no rotating machinery on board, so it is desired to control roll without the use of conventional rate gyros. A new technique has been devised to achieve ***highly accurate roll control without a rate gyro*** by employing one or more slit star sensors which rotate with the spacecraft, and correlate their output with a known reference to produce a measurement of roll offset.

The technique developed to control roll phase and rate ***mimics those used in pseudorandom noise telecommunication*** equipment. The algorithm regards the intensity pattern of the surrounding star field as pseudorandom noise which repeats itself every 360 degrees, and sets up a "phase-locked" loop to align the pattern with a stored reference pattern. Single-axis simulations confirm that such a device, when combined with a steady-state Kalman estimator, can control roll position to an accuracy of 25 arcsec RMS, and roll rate to an accuracy of 0.92 arcsec/sec RMS, even when the star sensor output and reference values are encoded with only one bit. This level of performance can be achieved using available thruster torques, and while staying within the mission-prescribed attitude acceleration limits. Algorithms for initial acquisition of roll phase have also been devised so that phase lock can be achieved within a fraction of the orbit period.

* Dr. Parkinson is a Research Professor in the Aeronautics and Astronautics Department and High Energy Physics Laboratory of Stanford University and the Program Manager of the Stanford Relativity Gyro Experiment, Stanford, California 94305.

† Mr. Crerie is a Ph.D. Student in the Aeronautics and Astronautics Department at Stanford University, Stanford, California 94305.

I. INTRODUCTION

The Gravity Probe B Experiment

Gravity Probe B (GP-B) is a space-based experiment to verify two previously untested predictions of Einstein's Theory of General Relativity. By placing near-perfect gyroscopes in a circular, polar orbit about the earth, GP-B intends to measure drifts of local inertial space due to the mass of the spinning earth. The theory predicts a gradual misalignment of a gyroscope's spin axis with respect to distant inertial space, the reference for which is a inertially-calibrated guide star (Rigel).

Einstein's theory predicts two orthogonal components of drift for a gyro in a 650 km circular, polar orbit that are unaccounted for by Newtonian gravitation. The geodetic drift, which occurs in the plane of the orbit due to the motion of the gyroscopes within the gravity field of the earth, would accrue at the rate of 6.6 arcsec/yr. The frame-dragging drift, a significantly smaller component (42 milliarcsec/yr) normal to the orbit plane, is caused by the massive earth's rotation dragging the inertial reference frame of the gyroscopes. These two relativistic effects are shown schematically in Fig. 1.[1,2,3]

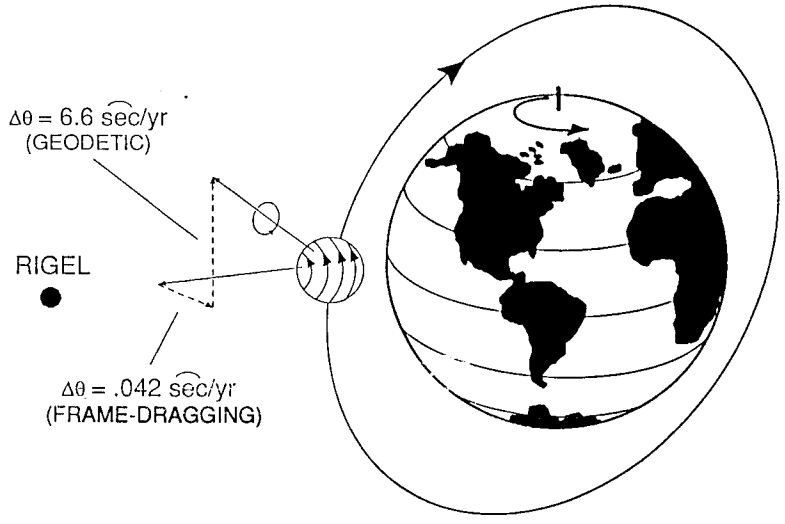

Figure 1: Relativistic Effects on the GP-B Gyroscope[1]

System Requirement

During the experiment, the GP-B spacecraft will roll about its pointing axis with a period of 10 minutes to help maintain thermal stability and average out disturbances to the gyroscopes. In order to demodulate the two orthogonal drifts being measured, roll phase must be known to within 100 arcsec RMS and must not drift by more than 20 arcsec in one year (random

walk). The available thruster torque about the roll axis is limited to 0.5×10^{-3} N-m.[6] Other system requirements include a maximum roll rate variation of $5 \times 10^{-4} * \omega_{roll}$ (5.08 arcsec/sec), and maximum attitude-control induced science gyro accelerations of 2×10^{-7} g, with less than 1×10^{-10} g at roll frequency.[15]

Motivation for Research

Conventional rate-integrating gyroscopes, due to their substantial power requirements, induced vibration and potential reliability problems, are undesirable for the GP-B spacecraft. One preference in the design of GP-B, in fact, is to fly the experiment with no rotating machinery other than the gyros which measure the relativistic effects. The research presented here involves the development of a roll control system that employs no rate gyros, with the means of measuring roll errors being pseudorandom intensity patterns found in the surrounding star field. Because the sensors used for this new algorithm are not required to image the star field, but rather just measure its intensity, and because the computational requirements are minimal (especially when the signals are given one-bit resolution), the phase-lock roll control system has a possible cost advantage over existing methods for attitude control. Even if we elect to include rate gyros, this technique can act as a valuable back up in the event of failure or high vibrations.

The Scope of this Paper

This paper will describe the design of the phase-lock roll control system, including the phase measurement algorithm and an optimal estimator and controller. Simulations of the concept are presented and discussed, with emphasis on satisfying the system requirement while maintaining a simple, inexpensive design. Section II discusses the overall phase-lock control concept, including a detailed description of the phase-measurement correlation scheme. Section III provides background on how a mission-specific star catalog was created for the purpose of simulation, and also explains the generation of the discretized model of the star intensity profile necessary for phase-lock control. This is followed (Section IV) by a study of the phase-lock discriminator characteristic, with attention to the inherent noise of the algorithm and it's color.

The optimal estimator and controller used to zero the roll phase error and maintain constant roll rate are presented in Section V, along with the results for a control system operating with full-precision signals from the sensor and stored model. There are significant advantages to encoding the sensor output and stored model with limited (one-bit) precision, and the results for a system with such simplifications are found in Section VI. Initial acquisition and reacquisition are discussed in Section VII, immediately followed by a summary of results and conclusion section.

II. PHASE-LOCK CONTROL

Phase-Lock Loops: Foundations in Telecommunication

Phase-lock loops have proven very useful in the field of telecommunications, and more particularly in the processing of information contained in binary pseudorandom noise (PRN) sequences. The satellites of the Global Positioning System (GPS) transmit known PRN sequences that allow receivers and ground stations to accurately determine their position. These sequences also provide a clock calibration so that precise time can be inferred from position measurements. In fact, GP-B will use an on board GPS receiver to maintain time synchronization. GPS receivers rely on phase-lock control loops to synchronize and average the measurement of the satellite signals.

The Application to Spacecraft Attitude Control

When considering the problem of how to determine a spacecraft's angular orientation and rate, engineers can reference celestial objects with known positions or trajectories. Devices that use the star field for attitude determination (star trackers, slit detectors) share the common attribute that they model the known position and magnitude of selected stars or star clusters and then match reality against these models. For the phase-lock attitude control algorithm being presented, individual stars are not relevant except in their contribution to the overall intensity pattern of the celestial sphere. This pattern is discretized and considered as a pseudorandom noise information source. This unique feature has many advantages: a more continuous flow of control information, greater resiliency to model errors, and in most cases, drastically diminished storage and computational effort for a system which achieves comparable accuracy.

The Phase-Lock Roll Control Concept

A single axis simulation of the phase-lock roll control system for the inertially-pointing GP-B spacecraft serves as the first verification of this concept. Because GP-B maintains the angular orientation of its telescope axis with 60-milliarcsecond accuracy, and rolls (spins) about this axis with a nominal 10 minute period, an optical sensor which points normal to the telescope axis will continuously scan a select band of the star field which will be referred to as the **relevant band** (see Fig. 2). Of course this band is constant for the whole mission. The output of the slit detector is sampled at discrete intervals and is proportional to the integrated intensity of incoming starlight through the aperture, which has dimensions 0.1 deg by 10 deg for the simulations discussed in this report.

An on board model of the intensity would contain the expected sensor output at discrete roll positions (every 0.1 deg) around the entire 360 degree band. These reference values are correlated with the actual sensor output using the simplified phase-lock algorithm described below to produce a

measurement proportional to the roll phase error. This error signal feeds into a steady-state Kalman estimator to generate commands to the actuators, which are continuously-firing proportional helium thrusters in the case of GP-B.[5]

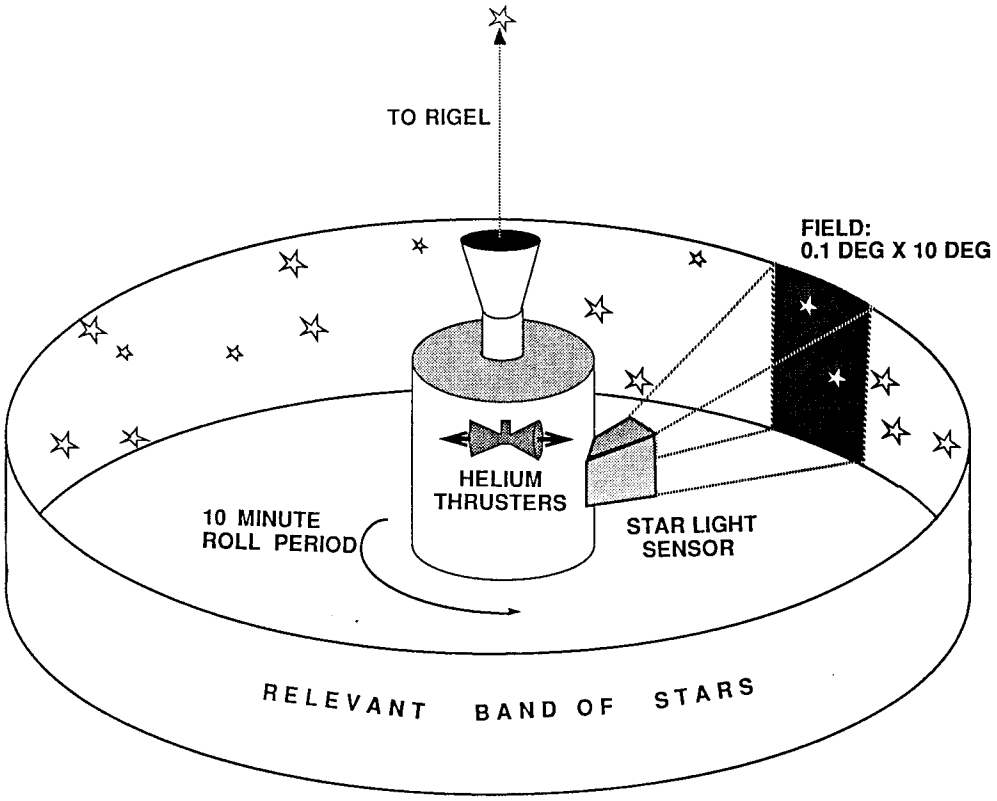

Figure 2: The Phase-Lock Roll Control Concept

A Simplified Phase Measurement Algorithm

Fig. 3 depicts schematically the phase-lock algorithm used to derive a measurement of roll phase error. At discrete time intervals, the products of a star sensor output (logarithm of an integrated intensity) and two interpolated reference values, one early and one late, are taken. The early and late reference values represent the expected sensor output 0.05 deg ahead of and behind the commanded roll position, respectively. This choice of phase shift between early and late reference values will be justified by the results of their correlations with the sensor output, which will be presented in Section IV.

A lock indicator signal verifies whether the roll phase error is within the operable range of the measurement algorithm. It monitors the product of the sensor output and an "on time" reference value, which is the expected sensor output at the commanded roll position. This product is averaged over time, using a fading memory filter, such that the resulting lock indicator

signal has a mean of 1.0 if the system is phase-locked, and has zero mean if the system strays outside the hold-in range. In the event that the system leaves the hold-in range, a reacquisition procedure would be initiated.

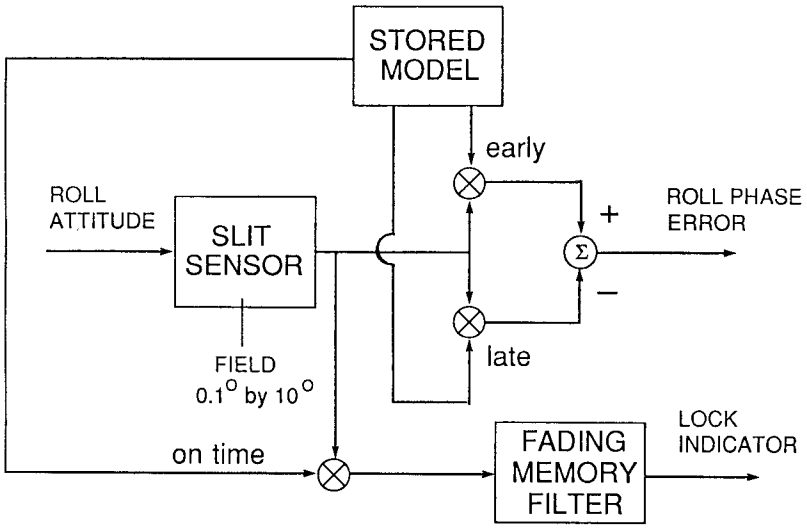

Figure 3: Roll Phase Measurement Algorithm

III. GENERATING AND MODELLING THE CELESTIAL REFERENCE

The phase-lock roll control system simulation requires an accurate re-creation of the star field for both the generation of the stored model and the determination of what the sensor actually "sees" during operation. Star positions and visual magnitudes were extracted from an extensive star catalog provided by ETLON Software of Colorado, which contains astronomical data for all known stars down to 7th magnitude in brightness. An Euler angle transformation was used to redefine star positions in terms of an inertial reference frame with one axis aligned with the GP-B's telescope pointing axis, and a new database, stripped of all stars that lie outside of the relevant band, was generated.

Altogether, 737 stars were extracted from the original ETLON star catalog. Then, dimmer stars down to 9th magnitude were scattered at random positions within the band in densities approximating those found in the true star field,[12] resulting in a simulated star field with a total of 6713 stars. These were used to generate a stored model of the intensity pattern by calculating the expected sensor output at 3600 discrete roll positions spaced 0.1 degrees apart. Fig. 4(a) depicts the resulting full-precision model for two roll periods (720 degrees). Notice that the model profile repeats itself every 360 degrees. For the later studies, a one-bit version of the model was used, with +1's representing positions that provide intensities above a threshold value, and -1's for those below the threshold. This threshold was set so as to

produce an approximately equal number of +1's and -1's over one revolution, thus maximizing their information content. A portion of the one-bit model, with interpolation between entries, can be compared to it's original full-precision representation in Fig. 4(b).

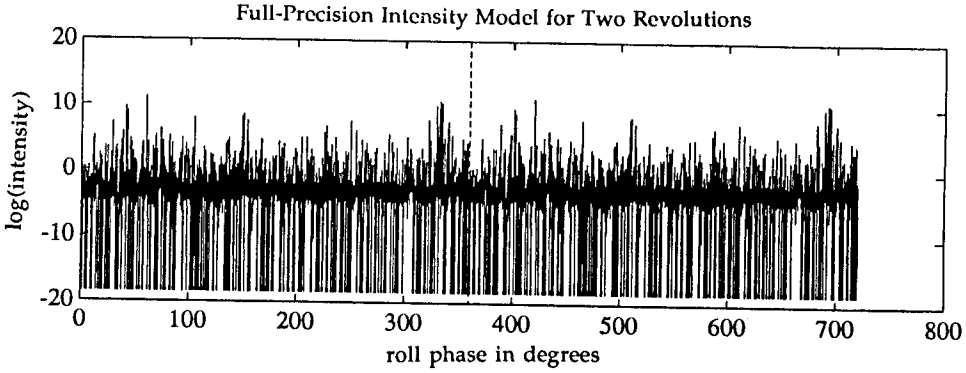

Figure 4(a): Full-Precision Intensity Model

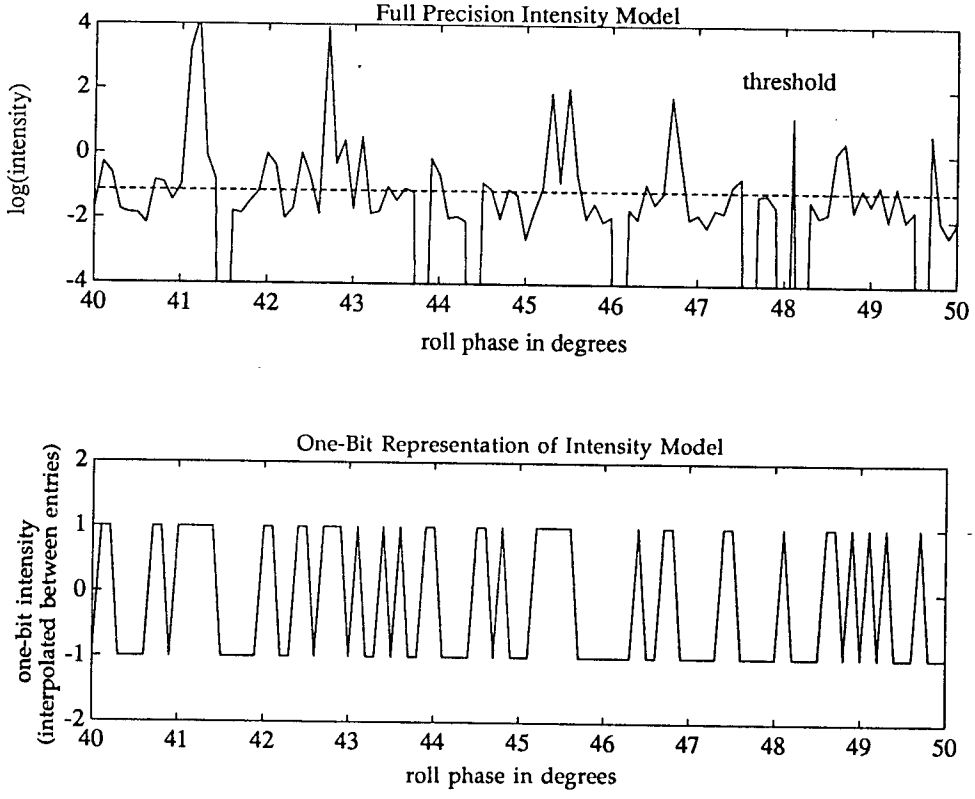

Figure 4(b): Full-Precision and Corresponding Binary Representation of the Intensity Model (shown for a fraction of a roll period).

Simulations have shown that there are, in fact, very few instances where the actual sensed intensity is very near the threshold intensity. The histogram in Fig. 5 counts how often the output values differed from the threshold value by various amounts. Note how rarely the sensed intensities are in the vicinity of the threshold value, which is equivalent to the intensity of a single star with an apparent magnitude of 7.1. The location of the threshold suggests that the sensor hardware must be capable of determining the presence of a 7th magnitude star within its window.

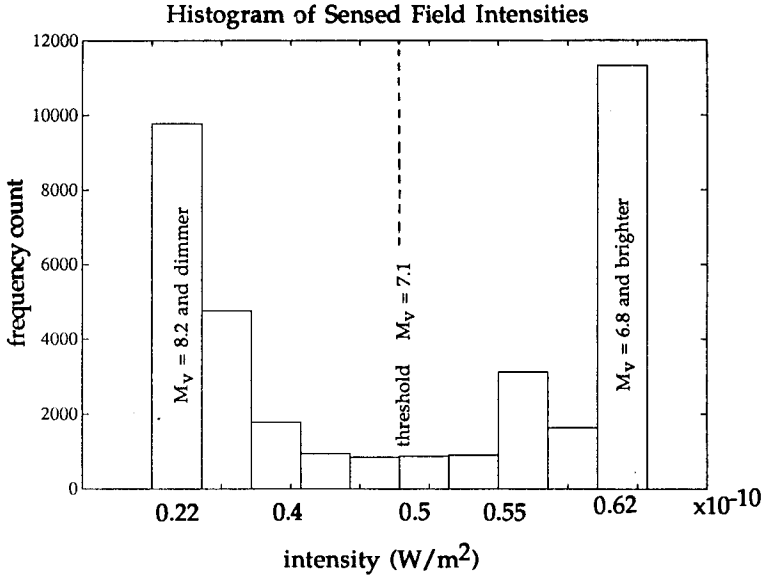

Figure 5: Histogram of Sensor Output Values Prior to One-Bit Encoding

IV. PHASE-LOCKED CORRELATOR: NOISE AND SENSITIVITY STUDIES

The Difference of Two Correlations: The Phase-Lock Gain Curve

Fig. 6(a) shows the average result of the early and late correlations over one revolution of the spacecraft (one 360 degree sweep of the relevant band) as a function of a constant difference between the actual and commanded roll positions (the control error). When the difference between these two correlation functions is taken (early minus late), it is clear that the result is a signal that, when averaged over time, is proportional to the roll offset over a narrow hold-in range, as shown in Fig. 6(b). Keeping in mind that the gain curve plotted in Fig. 6(b) represents a time average of output values, we must now concern ourselves with the nature of the noise inherent to this special correlation scheme.

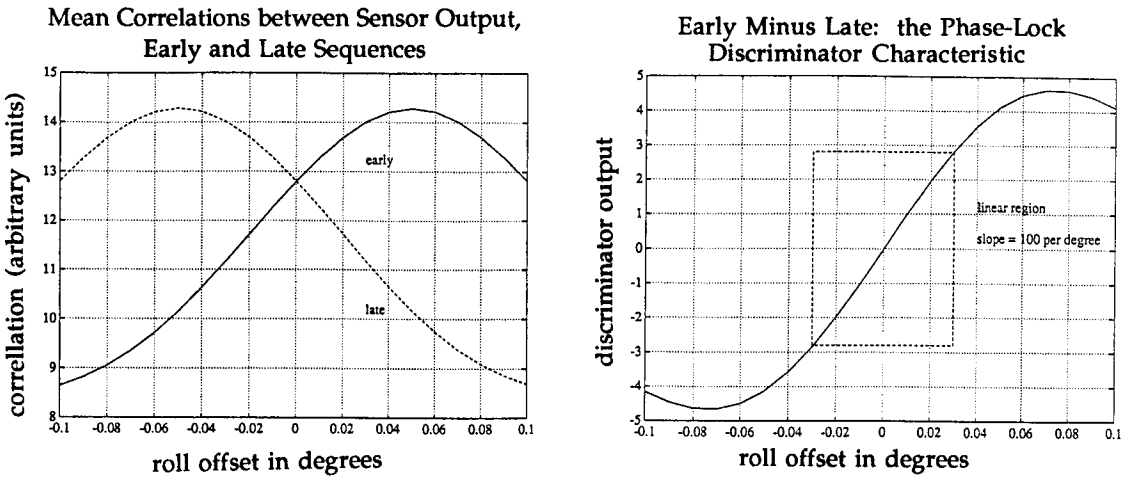

Figure 6(a): Early and Late Correlations, 6(b): (Early-Late): The Phase-Lock Gain Curve

Inherent Noise of the Phase-Locked Correlator

Even if we considered an ideal system, in which there was zero measurement error in reading the intensity pattern of the relevant band and no mismatches between discrete values from reality and the stored model, the output of the phase lock algorithm would, for two reasons, still be corrupted by noise. First, the stored model is discretized and values are interpolated from it, thus creating reference values that deviate from reality except at the discrete positions of the stored values. The noise resulting from this effect is quite small compared to that due to the fact that the product of two instantaneous values taken from random, cyclostationary processes which differ only in phase varies wildly as a function of time, and it is only ***in the mean*** that a predictable function of the roll phase difference exists. Fig. 7 plots the covariance of the measurement error as a function of roll offset.

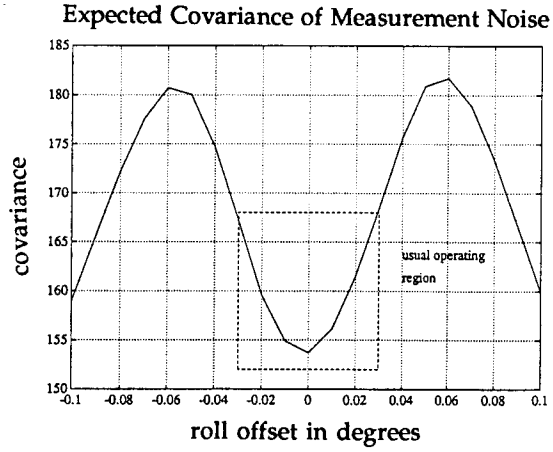

Figure 7: Covariance of Measurement Error vs. Roll Offset

Investigation of Roll Phase Measurement Noise Color

One of the initial assumptions of the phase-lock roll control algorithm was that the noise of the discriminator just described would approximate white noise within the bandwidth of the controller. A series of FFT's for several fixed roll offsets in the hold-in range were performed on the discriminator noise sequences to determine their frequency content, and they were found to be virtually indistinguishable. Fig. 8(a) plots the power spectral density (PSD) of a noise sequence recorded at zero offset, which approaches a flat spectrum over the entire frequency range. The spike at 6 Hz is associated with the time required for stars to traverse the sensor field width, and corresponds exactly with the spatial sampling frequency of the model with a nominal roll rate of 0.1 rpm. Fig. 8(b) affirms the uniformity of the spectral energy near the bandwidth of the controller, which is near 0.04 Hz.

Because the noise of the phase-lock roll offset measurement displays the general characteristics of a zero-mean, stationary white noise process, an optimal steady-state Kalman estimator can be used to estimate the states of the system: roll phase and roll rate. These estimated states are then fed back with constant gains to generate commands to the thrusters, which drive the roll offset toward zero.

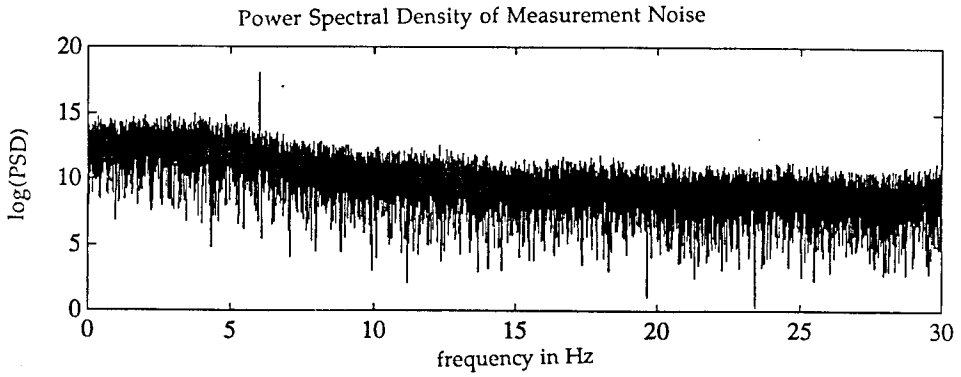

Figure 8(a): Frequency Content of Measurement Noise

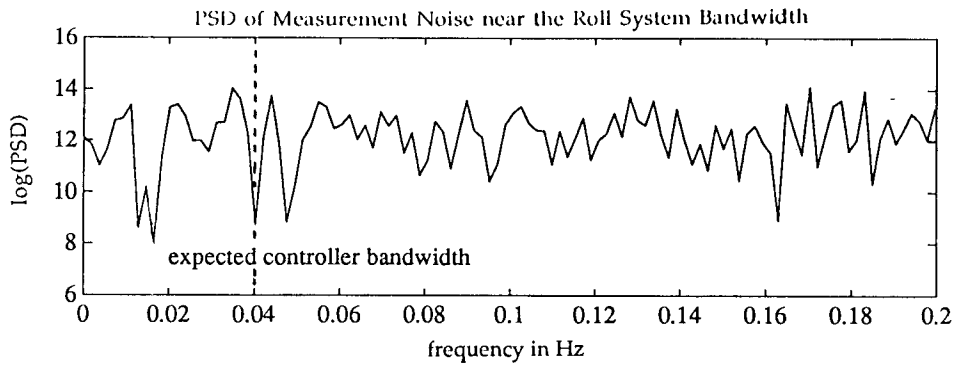

Figure 8(b): Frequency Content of Measurement Noise Near Roll System Bandwidth

V. CONTROLLER SIMULATIONS

Closed-Loop Control using a Steady-State Kalman Estimator

A block diagram of the control law for the phase-lock roll control system is shown in Figure 9. The $1/Js^2$ plant represents the spinning inertia of the spacecraft about its pointing axis. A four-state Kalman estimator was designed to optimally abate noise from three sources: (i) the roll phase measurement error (v_n) discussed in Section IV, (ii) plant disturbances caused by zero-mean, random thruster noise with a uniform distribution (w_n), and (iii) a slowly varying thruster bias (random walk) with a period of approximately one hour. The resulting estimates of roll phase error ($\Delta\theta$), roll rate error ($\Delta\omega$), and two states of the thruster bias (r,s) are fed back through an optimal controller to produce the command thruster torque (u) that will zero the roll phase error.

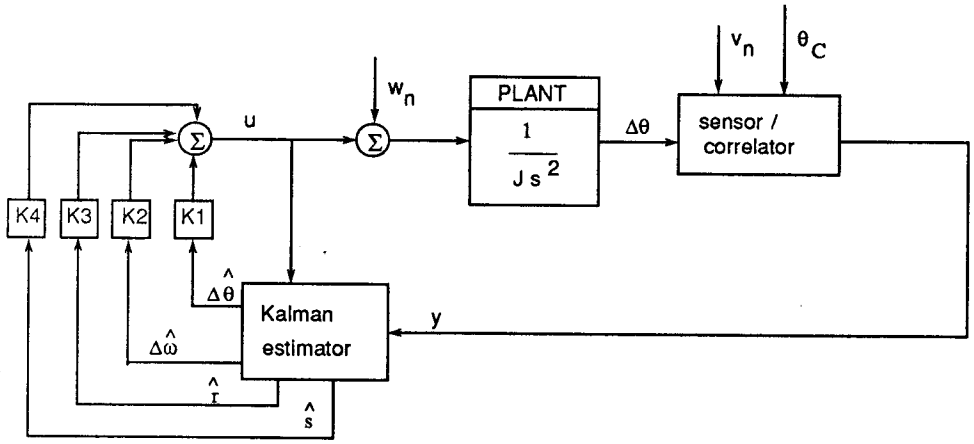

Figure 9: Block Diagram of Roll Control System

Choice of Sample Rate

A 60 Hz sample rate was chosen for the simulations discussed in this report. The sample rate cannot be made too fast, since the starlight sensor relies on the integrated intensity of incoming photons, and shorter integration times would cause the shot noise of the sensor (the noise due to the random arrival of photons) to dominate the signal. Sample rates slower than 20 Hz cause the system to repeatedly drift outside the hold-in range of the discriminator, leading to dynamical instability for the current sensor/model correlation scheme (which only requires two multiplications and a subtraction per sample). More elaborate correlations of larger sequences of the stored model and the sensor output could resolve this instability and allow slower sample rates.

Full-Precision Control System Parameters

A DLQG* estimator was designed assuming a phase measurement noise covariance of 2.01×10^5 arcsec^2, and RMS white-noise thruster accelerations of 0.01 arcsec/sec^2, along with 0.01 arcsec/sec^2 RMS accelerations due to a slowly-varying thruster bias (random walk, T = 1 hour). The RMS magnitudes of both thruster noise components represent 10% of the available thruster torque about the roll axis. The limits imposed on the available thruster torque given in the system requirement (enabling nominal angular accelerations of 0.1 arcsec/sec^2 RMS) most influenced the choice of regulator gains. The regulator, designed by DLQR*, had weighting factors on the phase error and rate error states of 1, with the commanded accelerations weighted by 1×10^5. Table 1 provides the results of a modern controller design for the phase-lock control algorithm with full-precision (to the resolution of the computer) sensor output and model.

	gains	z-plane poles
estimator	$l1 = 8.2113 \times 10^{-6}$	$0.99979 +/- j3.594 \times 10^{-4}$
	$l2 = 2.0236 \times 10^{-7}$	0.99959
	$l3 = 2.4715 \times 10^{-9}$	0.99171
	$l4 = 1.0417 \times 10^{-12}$	
regulator	$k1 = 3.1602 \times 10^{-3}$	$0.99934 +/- j6.678 \times 10^{-4}$
	$k2 = 7.9563 \times 10^{-2}$	0.9999918
	k3 = 1.0008	0.99171
	k4 = -1.7098	

Table 1: Parameters of the Modern Controller Design--Full Precision

Full-Precision Control System Simulation Results

Simulations were performed on the full-precision modern controller design to verify that all of the system requirements are met. The results of these simulations are provided in Fig. 10. These time responses are for five roll periods (unless otherwise specified) and consider both the 10% white noise and 10% random walk thruster disturbances. Note in Fig. 10(a) that the steady state roll phase error response satisfies the system requirement of 100 arcsec RMS by keeping phase error to 21 arcsec RMS, with a mean phase error over five roll periods of 0.51 arcsec. Figure 10(b) shows that tracking the estimate of the roll phase error provides only slightly more information,

* DLQG/DLQR controller design, or modern control theory, involves selecting full state feedback gains (DLQR) by solving a discrete algebraic Riccati equation found from minimizing a quadratic cost function of the states and controls. The dual estimator problem (DLQG) is used, with the separation principal, for compensator design. See reference 7 for details.

since the RMS difference between the estimated phase error and the true phase error is 15 arcsec. Even smaller excursions in roll phase error could be achieved by weighting that state by more than 1 in the regulator design, but the penalty would be larger roll rate variations and an increase in the required thruster torque.

Figs. 10(c) and (d) show the magnitudes of the roll rate error and required control accelerations, respectively, for the full-precision controller design. While the instantaneous roll rate error has a magnitude of 0.57 arcsec/sec RMS, the mean rate error over five roll periods is 4×10^{-3} arcsec/sec. The required thruster accelerations to achieve this level of control have an RMS magnitude of 0.045 arcsec/sec^2, which corresponds to 50% of the available thruster torque in roll. Fig. 10(e) shows the recovery of the phase-locked controller to initial errors in roll phase and roll rate with magnitudes that could be encountered after an initial phase acquisition.

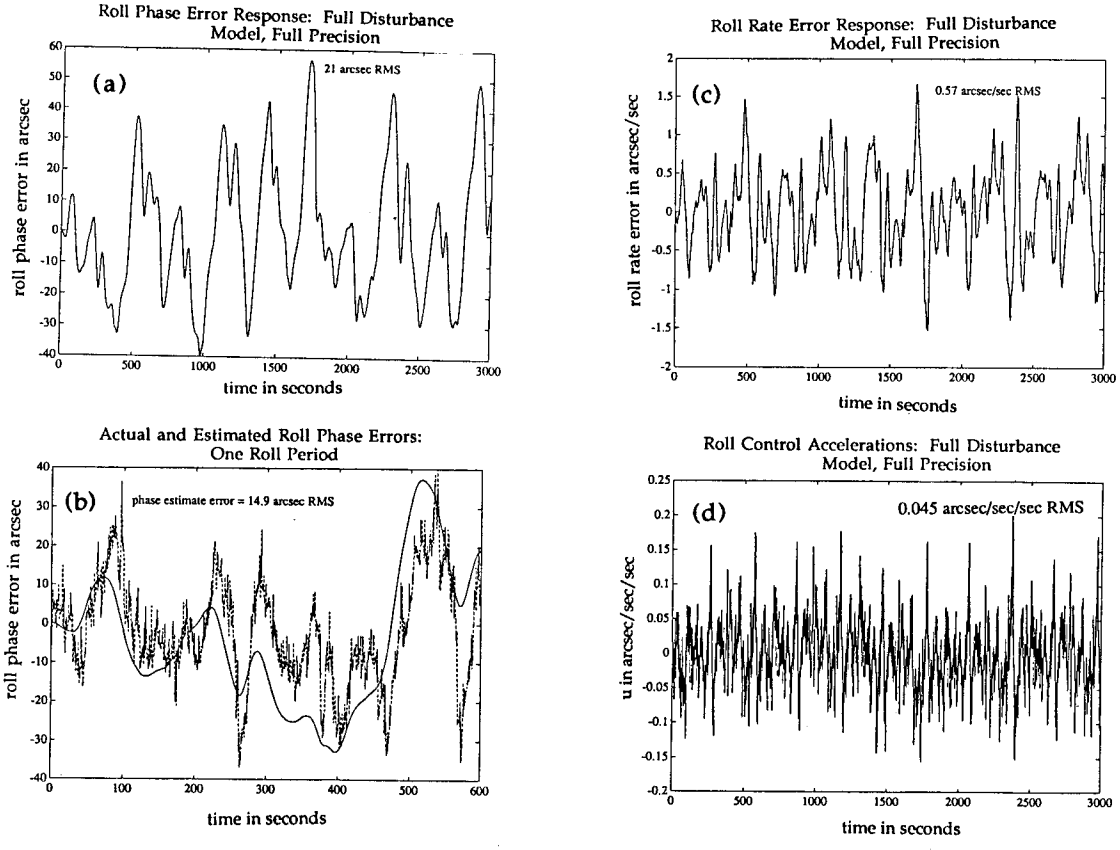

Figure 10: (a) Roll Phase, 5 Roll Periods, (b) Comparison of Estimated and Actual Phase Errors, 1 Roll Period, (c) Roll Rate Error, and (d) Commanded Accelerations, 5 Roll Periods (all using full disturbance model).

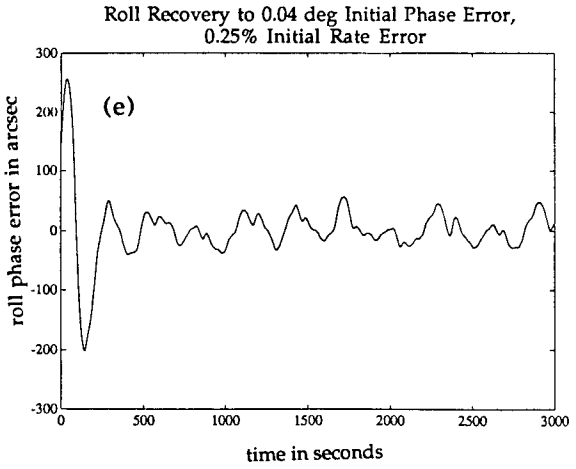

Figure 10(e): Response to Initial Errors in Roll Phase and Rate

VI. CONTROLLER SIMPLIFICATION: ONE-BIT ENCODING

The Advantages of Limited-Precision Sensor Output and Stored Model

One of the complications of the phase-lock algorithm is how to handle gain and bias changes in the measurement, whether due to long-term thermal effects or imprecise modelling of the reference. For the former, we again appeal to the techniques used in pseudorandom noise telecommunication equipment: why not encode the stored model and the output of the sensor with one bit resolution? This would have many positive impacts on the design. First, if bright areas of the star field produce a +1 and dim areas produce a -1, there would tend to be fewer discrepancies between the stored model and reality. Also, the introduction of bright, star-like objects such as the distant planets traversing the relevant band could only have the effect of introducing a short sequence of positive bits in the sensor output where a series of negative and/or positive bits exists in the model. Computational effort on the part of the discriminator is much relieved when the values are binary. The early and late values, however, are not binary since they result from interpolations between binary values in the stored model. (Not interpolating the model results in higher measurement error covariances, causing a degradation in overall controller performance.)

Scale factor changes in the sensor characteristic can be eliminated by use of a simple counting procedure to ensure that the output of the starlight sensor produces +1's and -1's in equal proportions during each revolution of the spacecraft. The original stored model would be pre-calculated so as to contain the same number of positive and negative bits, since this maximizes the information content of the model. If, during operation, it is determined that the sensor is producing too many +1's (and thus too few -1's) or vice

versa, the threshold that determines the sign placed on the bits is adjusted slightly in the appropriate direction to equalize the counts. This also fixes the scale factor of the discriminator, since now the model and the actual sensor output will both have a mean value of zero, and the correlations between them at various offsets will tend to remain constant.

Limited-Precision Control System Parameters

When the sensor output and stored model were given one-bit resolution before correlation, the shape of the phase-lock discriminator was unchanged, but the new scale factor became 11.0, and the measurement noise covariance became 1.01×10^5 arcsec^2 (compared to 2.01×10^5 arcsec^2 for the full-precision signals). A distinct feature of the one-bit algorithm is the shape of its measurement covariance as a function of roll offset. Note in Fig. 11 that the measurement error covariance reaches a maximum when the system has zero phase offset when limited precision signals are used, whereas the same algorithm with full-precision signals shows a minimum at zero offset (see Fig. 7). To explain this, one must consider that the binary values being correlated in the one case are generally in direct agreement (2 bits with the same sign) or direct disagreement. The advantage of full-precision signals is that when they are not the same, they can still be approximately equal, leading to a smoother correlation near zero roll offset.

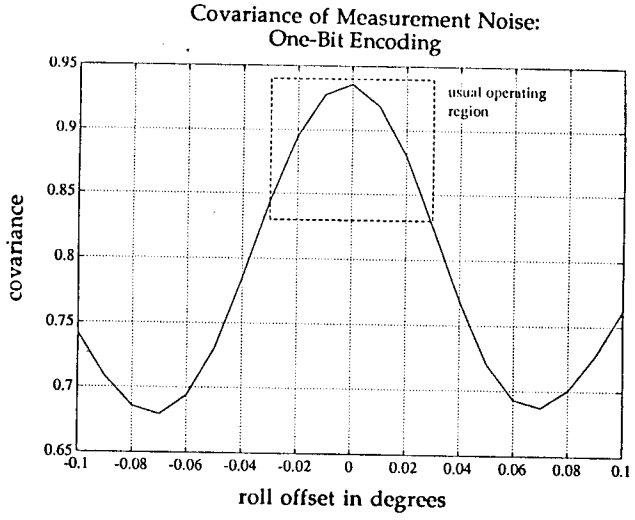

Figure 11: Covariance of Measurement Error vs. Roll Offset--One-Bit Encoding of Sensor Output and Stored Model.

The change in the phase-lock measurement characteristic led to a refiguring of the estimator gains using the same thruster noise covariances but a new measurement scale factor and noise covariance. The regulator gains were the same as for the full-precision controller design, since weighting factors identical to those used in the full-precision case produced a dynamic response which had phase error, rate error and acceleration covariances with similar proportions. Table 2 summarizes the new control system parameters for the limited-precision design.

	gains	z-plane poles
estimator	$l1 = 8.3985 \times 10^{-5}$	$0.99977 +/- j4.0794 \times 10^{-4}$
	$l2 = 2.3287 \times 10^{-6}$	0.99953
	$l3 = 3.2043 \times 10^{-8}$	0.99171
	$l4 = 1.2585 \times 10^{-11}$	
regulator	(same as for the full-precision controller design)	

Table 2: Parameters of the Modern Controller Design--One-Bit Encoding of Sensor Output and Intensity Model

Limited-Precision Control System Results

Simulations using one-bit encoding of the sensor output and stored model produced responses similar to those using full-precision signals but with increased error state and control covariances. This seems to be in contradiction with the fact that limited-precision phase measurement error covariance was half that of the full-precision algorithm. The steady-state roll error response for five roll periods, shown in Fig. 12, has an RMS value of 25 arcsec when subjected to the torque disturbance model discussed in Section V. The rate error and necessary control accelerations had responses like those in Figs. 10(c) and (d), except with RMS magnitudes of 0.92 arcsec/sec and 0.066 arcsec/sec^2, respectively. Simulations were also performed with no torque disturbances. As with the full-precision controller, these resulted in only slightly improved performance, with a 23 arcsec RMS phase error.

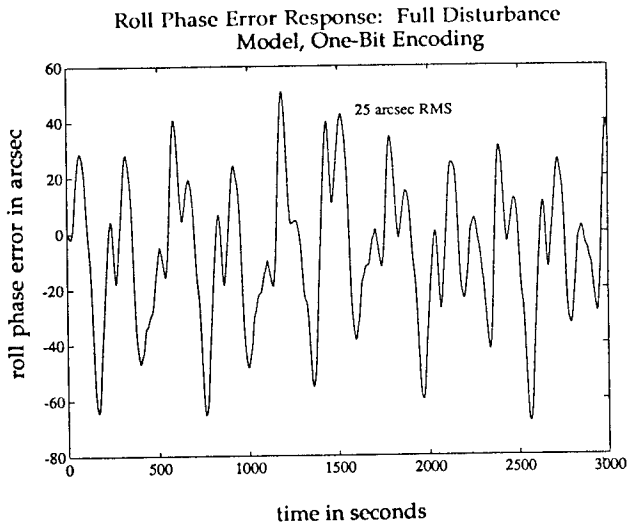

Figure 12: Roll Phase Error Response with Full Disturbance Model

VII. INITIAL ACQUISITION AND REACQUISITION

Description of Acquisition Procedure

An acquisition scheme has been devised to determine the roll phase of the GP-B spacecraft from zero knowledge, using a sliding correlation of 32 consecutive stored reference values with the actual sensor outputs. The procedure is analogous to tossing a fair coin 32 times and predicting the sequence of outcomes in order. Probability theory tells us that the odds of getting more than 24 in the sequence correct are less than 1 in 10^5. As a result, there will tend to be only one strong maximum in the entire correlation, and it will occur at or near the point where the real intensity pattern and the stored model align. A successful initial acquisition occurs when the absolute phase difference between the predicted (acquired) roll angle, and the actual roll angle is within the hold-in range of the phase-lock algorithm (+ or - 0.05 degrees) so that phase-lock control can be successfully initiated.

It is important to realize, however, that this acquisition scheme is only effective, meaning a strong maximum will only occur, if the actual roll rate of the spacecraft is within one or two percent the nominal rate of 0.1 rpm (the rate on which the stored model is based). Start-up of the GP-B spacecraft upon insertion into the desired orbit will probably not provide an initial roll rate that is accurate enough to perform a successful initial acquisition. A more precise initial rate, however, can be achieved by autocorrelating long sequences of the the actual sensor output over several roll periods, and inferring rate from the time difference between neighboring maxima of the autocorrelation (which would be separated by 2π radians).

Initial Acquisition Results

An ideal acquisition is depicted in Figure 13. Note that a single maximum occurs at 10.0 degrees, which was the origin of a 32-bit sequence extracted from the stored model for correlation with actual sensor outputs. There are four significant error sources, however, that must be modeled to simulate a true acquisition scenario: (i) An initial condition error caused by the random choice of phase origin when starting the correlation. It may be that none of the phase locations sampled exactly align with those of the stored values, but they can differ by at most by 0.005 degrees (assuming a 3600-element model, 0.1 rpm nominal roll rate, and a 60 Hz sample rate). (ii) Discrepancies between reality and stored model due to sensor noise. The frequency of these errors is dependent on the sensor, but is assumed to be 5% for the acquisition trials to be presented. (iii) Errors due to misalignment in the sensor mounting. Here we will assume that the sensor is "twisted" about its optical axis by a random angle with zero mean and uniform distribution from -0.5 to +0.5 degrees. (iv) Roll rate uncertainty after rate determination using an autocorrelation of sensor outputs in frequency and time was estimated to be 0.03% of the nominal roll rate. A histogram of successful acquisition trials with all of the above error sources included is shown in Fig. 14. Out of 1000 trials, 13 were unsuccessful, having absolute maxima outside the hold-in range of the phase-lock device (+/- 0.05 degrees), suggesting a 98.7% success rate for acquisition on a single 360 degree roll.

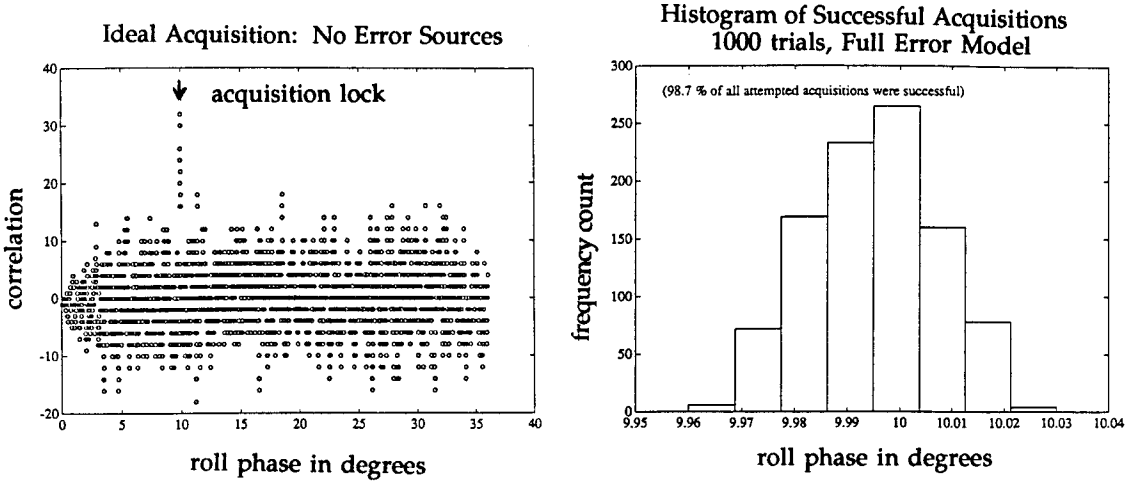

(a) (b)

Figure 13(a): Absolute Maximum Resulting from Correlation of 32-bit Sequence from Stored Model with Sensor Output
Figure 13(b): Histogram of Simulated Acquisition Runs (intended lock position is at 10.0 degrees)

Reacquisition and the Lock Indicator

In the event that the phase error strays outside the hold-in range of the phase-lock algorithm, the discriminator will produce measurement outputs that are zero-mean (since the stored model and reality become uncorrelated), and which resemble the outputs produced at zero roll phase error (!). For this reason, a lock indicator must be included (see Fig. 3). The lock indicator correlates the sensor output with the **current** (as opposed to early or late) value in the stored model, and then filters these results through a fading memory averager with a time constant of 30 seconds. The average value of the lock indicator will be 1 as long as the phase-lock algorithm is functioning properly, but quickly goes to zero in the event that control is lost. Fig. 14 shows the lock indicator's response to a roll phase error that drifts outside the hold-in range of +/- 0.05 degrees. Such an event would necessitate reacquisition, which follows a procedure identical to that of acquisition, but the sliding 32-bit correlation could be confined to a much smaller window of the star field.

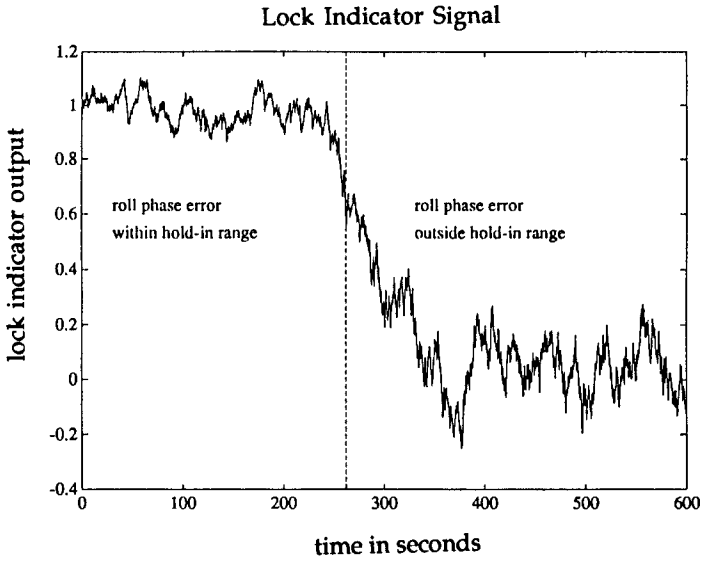

Figure 14: Response of the Lock Indicator to a Loss of Phase-Lock

VIII. SUMMARY OF RESULTS & CONCLUSION

Summary of Results

Table 3 highlights the results of the full-precision and one-bit controller simulations. Note that in simulation, both controllers satisfy the system requirement for 100 arcsec RMS roll phase error despite the presence of conservatively large thruster disturbances. This can be accomplished with available control torques, as indicated by the RMS magnitude of the required control accelerations. In general, the performance of the controller in response to zero disturbance is only slightly better than that with the full disturbance model. This is due to the comparatively large measurement error covariances that are inherent to the phase-lock correlation scheme.

	zero disturbance		10% disturbance	
controller accuracy*	full precision	1-bit	full precision	1-bit
position (arcsec)	15	23	21	25
velocity (arcsec/sec)	0.53	0.89	0.57	0.92
acceleration (arcsec/sec^2)	0.044	0.065	0.045	0.066

Table 3: Summary of Roll Control System Performance
* (RMS quantities)

Conclusion

Simulations of the phase-lock roll control algorithm have shown its promise as a method of control, free of rate gyroscopes, for inertially-pointing spacecraft such as GP-B. More specifically, initial studies of the phase-lock algorithm and controller simulations support the following conclusions:

- By treating the output of slit sensors which continuously scan a band of the star field as a PRN sequence, and comparing this to a model of the expected sensor output at discrete roll positions, roll control can be maintained to within 25 arcsec RMS in phase and 0.92 arcsec/sec RMS in rate.

- Encoding the stored model and sensor output with one-bit resolution led to a 19 percent increase in the RMS roll error over the simulations which used similar components but with full, floating-point accuracy.

- Acquisition can be accomplished by correlating larger sections of the model (32 consecutive values) with the actual sensor output to determine roll phase with an accuracy better than 0.05 degrees, which is the hold-in range for the current phase-lock algorithm.

- The phase-lock roll control scheme is a feasible means of maintaining accurate roll phase, without rate gyroscopes, for inertially pointing spacecraft with spin rates on the order 1 rpm. It is also inherently simple, and inexpensive.

- Existing sensor hardware is not optimized for application to the phase-lock roll control algorithm. The combination of a precision optical system with a high-sensitivity, low-noise photon detector without imaging capability would most suit this type of control.

Plans for Future Work

There are many areas that need refining and still more that remain unexplored. One priority is to research current sensor technology so that a more precise model of the starlight sensor can be constructed. Our initial impression is that silicon photodiodes appear well-suited for application to this project because of their swift rise times and potentially low noise levels. Also, the associated optics must be defined: whether or not to focus the image on the sensor, focal lengths and aperture sizes, shutter mechanisms to protect the sensors from bright, occulting bodies. All of these decisions will eventually lead to a laboratory prototype.

Validation of the concept through theoretical analysis is currently under way, and involves the study of correlations between binary-valued cyclostationary random processes. Immersion into the statistical properties of the problem should also uncover guidelines to the optimization of the device through variations of key parameters, such as model resolution, sensor field width and sample rate. The simulation of the roll control system as a part of all control systems on the GP-B spacecraft has begun, overseen by Dr. Ben Lange at Stanford University.

REFERENCES

1. Parkinson, Bradford W., C. W. Francis Everitt, and Daniel DeBra. The Stanford Reltivity Gyro Experiment. *Guidance and Control*, Vol. 63, *Advances in the Astronautical Sciences*, AAS, etc.
2. Parkinson, B. W., C. W. F. Everitt, J. P. Turneaure, and R. T. Parmley. The Prototype Design of the Stanford Relativity Gyroscope Experiment. *38th Congress of the International Astronautical Federation*, October 10-17, 1987, Brighton, United Kingdom. IAF-87-458.
3. Everitt, C. W. F. Preface, Papers on the Stanford Relativity Gyroscope Experiment. *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 619, 23-24 January, 1986.
4. Bardas, D., et. al., Hardware Development for Gravity Probe B. *Proceedings of SPIE - The International Society for Optical Engineering*. Vol. 619, January, 1986.
5. Chen, Jeng-Heng. Helium Thruster Propulsion System for Precise Attitude Control and Drag Compensation of the Gravity Probe-B Satellite. December 1983. Stanford Ph.D. Thesis.
6. Wiktor, Peter. Thruster Control for GP-B. January, 1990. Unpublished Report.
7. Bryson, Arthur E., Jr., and Yu-Chi Ho, Applied Optimal Control. John Wiley & Sons, 1986.
8. Egan, William F., Frequency Synthesis by Phase Lock. John Wiley & Sons, 1981.
9. Fairbank, W. M., C. W. F. Everitt, D. B. DeBra, et. al. Report on a Program to Develop a Gyro Test of General Relativity in a Satellite and Associated Control Technology. June, 1980.
10. Franklin, Gene F. , Powell, David J., Digital Control of Dynamic Systems. Addison-Wesley Publishing Company, 1980.
11. Parkinson, Bradford W., Kasdin, N. Jeremy, Twenty Milliarcsecond Pointing System for the Rolling GP-B Spacecraft. *Guidance and Control*, Vol. 66, *Advances in the Astronautical Sciences*, AAS, etc.
12. Robinson, J. Hedley, Muirden, James., Astronomy Data Book, 2nd Ed. John Wiley & Sons, 1972.
13. Spilker, James J. Jr., Digital Communications by Satellite. Prentice-Hall Inc., 1977.
14. Tapley, Mark, frequent communication during Spring 1988.
15. Vassar, R.H., Dougherty, H.J. Gravity Probe B Control System Concept Document. March 1989. Unpublished Report.
16. Wertz, James R., Ed., Spacecraft Attitude Determination and Control. Reidell Publishing Corp., 1978.