Gravity Probe B data analysis status and potential for improved accuracy of scientific results

IOP PUBLISHING CLASSICAL AND QUANTUM GRAVITY

Class. Quantum Grav. **25** (2008) 114002 (11pp) doi:10.1088/0264-9381/25/11/114002

Gravity Probe B data analysis status and potential for improved accuracy of scientific results

C W F Everitt, M Adams, W Bencze, S Buchman, B Clarke, J Conklin, D B DeBra, M Dolphin, M Heifetz, D Hipkins, T Holmes, G M Keiser, J Kolodziejczak[1], J Li, J M Lockhart, B Muhlfelder, B W Parkinson, M Salomon, A Silbergleit, V Solomonik, K Stahl, J P Turneaure and P W Worden Jr

W W Hansen Experimental Physics Laboratory, Stanford University, Stanford, CA 94305-4085, USA

Received 28 February 2008, in final form 4 April 2008
Published 15 May 2008
Online at stacks.iop.org/CQG/25/114002

Abstract

Gravity Probe B (GP-B) is a landmark physics experiment in space designed to yield precise tests of two fundamental predictions of Einstein's theory of general relativity, the geodetic and frame-dragging effects, by means of cryogenic gyroscopes in Earth orbit. Launched on 20 April 2004, data collection began on 28 August 2004 and science operations were completed on 29 September 2005 upon liquid helium depletion. During the course of the experiment, two unexpected and mutually-reinforcing complications were discovered: (1) larger than expected 'misalignment' torques on the gyroscopes producing classical drifts larger than the relativity effects under study and (2) a damped polhode oscillation that complicated the calibration of the instrument's scale factor against the aberration of starlight. Steady progress through 2006 and 2007 established the methods for treating both problems; in particular, an extended effort from January 2007 on 'trapped flux mapping' led in August 2007 to a dramatic breakthrough, resulting in a factor of ∼20 reduction in data scatter. This paper reports results up to November 2007. Detailed investigation of a central 85-day segment of the data has yielded robust measurements of both relativity effects. Expansion to the complete science data set, along with anticipated improvements in modeling and in the treatment of systematic errors may be expected to yield a 3–6% determination of the frame-dragging effect.

PACS number: 04.80.Cc

[1] NASA Marshall Space Flight Center, Huntsville, AL 35812, USA.

1. The Gravity Probe B relativity mission

In 1960 L I Schiff showed that an ideal gyroscope acting as a local spacetime inertial reference would undergo two precessions with respect to a distant (undisturbed) inertial frame: (1) a geodetic precession in the orbital plane due to the space time curvature near the Earth and (2) a frame-dragging precession resulting from the angular momentum of the Earth which causes the gyroscope's spin axis to precess in the Earth's equatorial plane. For a polar orbit the two effects are at right angles to one another. At an altitude of 642 km, general relativity predicts a geodetic precession of -6606 mas yr^{-1} and a frame-dragging precession of -39 mas yr^{-1}. The GP-B experiment was designed to measure these precessions to a precision of <0.5 mas yr^{-1}.

At the heart of the experiment are four ultra-precise cryogenic mechanical gyroscopes. Each gyroscope rotor is a 3.81 cm diameter precision machined fused quartz sphere (63.5 g) coated with a uniform layer of niobium metal. They are spherical to 10 nm and are mass-balanced to 10 nm. Each is housed in a spherical quartz cavity with a rotor-to-housing gap of 32 μm and are electrostatically suspended via six housing electrodes. The gyroscopes operate at 2.5 K within a cryogenic vacuum probe inside a 2440 liter superfluid helium dewar (1.8 K). Below 9.3 K, the niobium coatings are superconducting. The orientation of the rotor's spin axis is determined by measuring the dipole magnetic moment, the London moment, generated by the spinning superconductor coating. Detection is by a highly sensitive SQUID magnetometer coupled to a 4-turn thing film superconducting coil patterned on the parting plane of the gyroscope housing.

Spin axis orientation is measured by observing the time-varying coupling between the dipole London moment (M_L) and the housing-fixed pickup loop. This is facilitated by rolling the spacecraft at a period of 77.5 s around the line of sight to the guide star (IM Pegasi); the instantaneous direction to the star is measured by a cryogenic star tracking telescope. In determining the two Schiff effects it is necessary to take into account the proper motion of the guide star with respect to the remote inertial frame, and also the 19.2 mas yr^{-1} additional geodetic effect in the plane of the ecliptic (de Sitter effect) due to the Earth's motion around the Sun.

It is important to notice that the output of the London moment readout is a voltage which, though proportional to angle, depends on several experimental parameters and hence is not directly in angular measure. A very beautiful feature of GP-B is that nature itself provides the calibration signal via the aberration of starlight. The apparent position of the guide star varies by exactly known amounts through the Earth's motion around the Sun (annual aberration) and the spacecraft's motion around the Earth (orbit aberration). Because the telescope is held pointing at the guide star, these aberration signals appear in the output of the gyroscope and between them provide exact and cross-checking scale factor calibrations.

2. Summary of current results

In April 2007, we presented the first public results of GP-B at the American Physical Society meeting in Jacksonville, FL. As figure 3 shows, the geodetic effect is immediately obvious in the north–south orbital plane in all four gyroscopes. The mean 1σ result then reported was -6638 ± 97 mas yr^{-1}, which yields after subtracting the requisite north corrections of +7 mas yr^{-1} for the solar geodetic effect and $+28 \pm 1$ mas yr^{-1} for the proper motion of the guide star, a geodetic value of -6673 ± 97 mas yr^{-1}, to be compared with the predicted -6606 mas yr^{-1}; this is consistent with the predictions of general relativity. Results in the west direction, while also consistent with GR, were less satisfactory. In some data runs the

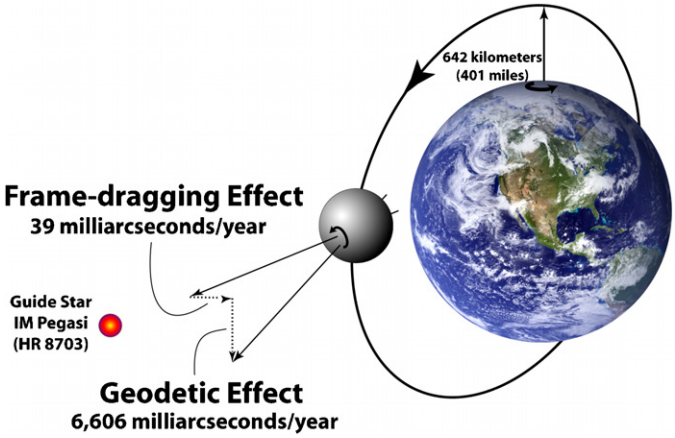

Figure 1. Predicted precessions on of the GP-B gyroscope (north and west are positive in this coordinate system).

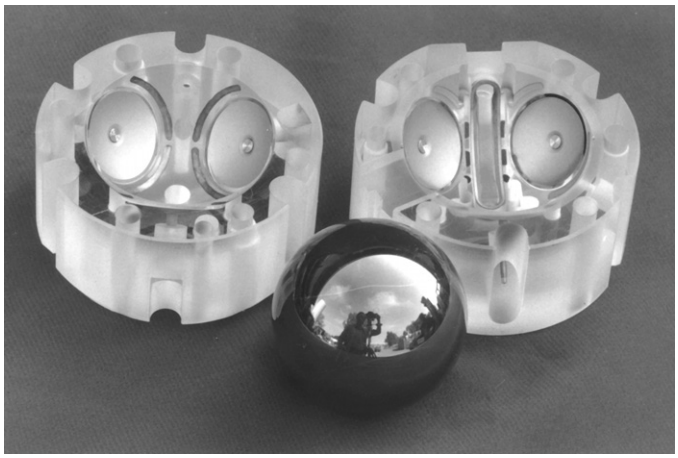

Figure 2. The GP-B gyroscope.

filtering algorithms appeared to yield authentic 'glimpses' of the frame-dragging effect but the results were far from stable; other runs under closely similar conditions either failed to converge or gave answers well outside the nominal 1σ limit.

Following the April 2007 announcement, we made steady improvements in the understanding and modeling of classical disturbance torques and other systematic error sources in the experiment, some of which were reported at the July 2007 GRG 18 meeting in Sydney, Australia. A critical turning point came a month later in August 2007 during an intensive study of a central 42 day period of data, subsequently extended to the 85 days from 10 December 2004 to 4 March 2005.

This interval was chosen because it had been studied extensively throughout the mission to isolate and understand some of the major disturbance sources on the measurements. The resultant 1σ error from a combined 3-gyro estimate[2] (gyros 1, 3, 4) obtained through these

[2] Three rather than four gyroscopes were combined here; resonance torques in gyro 2 during this period severely limited its contribution to the result. See section 3.7.

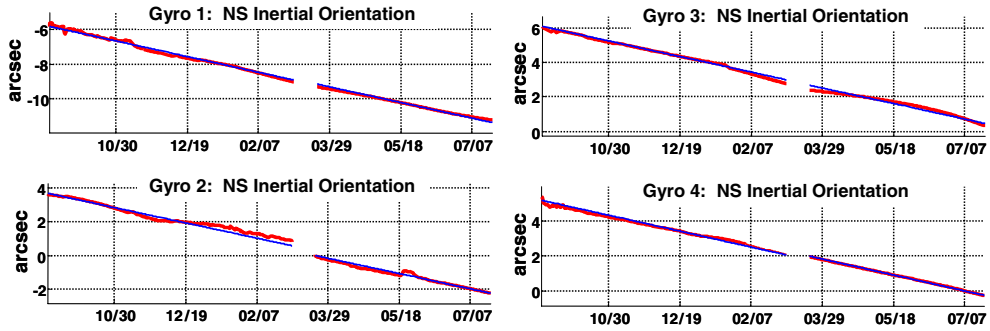

Figure 3. Direct measure of the north–south (NS) geodetic precession.

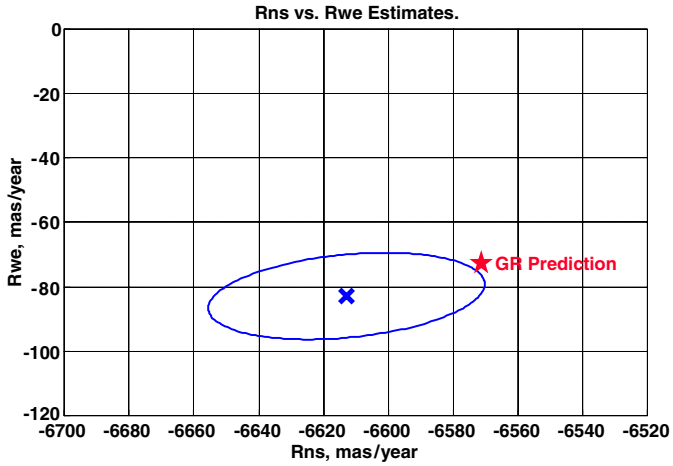

Figure 4. North and west drift measurements for an 85 day stretch of data (combined gyro estimates) with 1σ uncertainty ellipse.

advances, together with improved estimation codes, are presented in figure 4. They give -6617 ± 43 mas yr^{-1} for the north direction and -82 ± 13 for the west direction. Applying the requisite solar geodetic and proper motion corrections, these yield -6652 ± 43 mas yr^{-1} for the geodetic effect and -46 ± 13 mas yr^{-1} for frame-dragging.

The results presented in figure 4 are encouraging. Though limited to the central 85 day period, they exhibit an entirely new stability and robustness as compared with the April 2007 results. Nevertheless, caution is still essential. These are the outcome of 3-gyro averages; drift-rates for individual gyroscopes over the same period lie significantly outside the 1σ error curve of figure 4. Reasons exist for believing that a considerable part of the individual gyro spread may be the result of a 'separability' problem in the filtering machinery between classical torques and the relativity signals from the limitation to an 85 day period, and that with extension to the full 353 days of science data, the deviations will be substantially reduced. There are, however, also reasons for believing that some systematic errors requiring a separate and deeper treatment still remain in the signal.

In assessing how far the measurement may be improved, several factors enter. Consider first the ideal experiment with torque-free gyroscopes, limited only by SQUID noise in the

readout. Since the relativity signals increase linearly with time t and SQUID noise may be expected to average as $t^{1/2}$, the performance would improve as $t^{3/2}$. The limits in this case from measured SQUID noise range from 0.18 to 0.39 mas yr^{-1}, a 0.4% measurement of frame-dragging with 353 days of data.

In reality, three factors conspire to prevent our reaching that ideal limit: (1) various practical considerations probably limit the total useable data to between 307 and 323 days, depending on which gyroscope; (2) on-orbit interruptions to the mission resulted in the segmenting of the data into ten distinct subperiods; (3) the various not yet fully understood systematic errors just referred to.

Our present judgment is that with the extensions and improvements to the data analysis now to be discussed, a determination of frame-dragging in the 3–6% range is feasible.

3. Data analysis challenges and solutions

During the 353 days of science data, and a final 46 days of post-science calibration tests, we identified two unexpected, mutually reinforcing effects which significantly complicate the data analysis. One, observed in both, but only effectively sized during the calibration phase, was the presence of much larger than expected 'misalignment' torques proportional to the angle between the gyro spin vector and the spacecraft roll axis. The other, which complicates the process of calibrating the gyro scale factor, was that the polhode period of each gyroscope progressively changed over the course of the mission. The patterns observed were exactly those to be expected with damped polhode motion, including, in two gyroscopes (gyros 1 and 2), the passage across a separatrix as it transitions from spinning about the minimum to the intermediate to the maximum moment of inertia. The dissipation required to cause this effect was minute ($\sim 10^{-13}$ W) but pre-launch analyses appeared to have ruled out any such change. Particularly awkward was the fact that the settling times (30 to 70 days depending on the gyroscope) fell right within the range of the science data.

3.1. Simplified gyroscope dynamics and measurement model

To understand the challenges created by these two effects, and how they can be overcome, it is best to start with a simplified model of the gyroscope drift dynamics and orientation readout system:

$$\text{dynamics}: \quad \frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix} s_{\mathrm{NS}} \\ s_{\mathrm{WE}} \end{bmatrix}_i = \begin{bmatrix} R_{\mathrm{NS}} \\ R_{\mathrm{WE}} \end{bmatrix},$$

$$\text{measurement}: \quad z_i = C_g\, \hat{n}_i(\phi_r) \bullet \left(\begin{bmatrix} \tau_{\mathrm{NS}} \\ \tau_{\mathrm{WE}} \end{bmatrix} - \begin{bmatrix} s_{\mathrm{NS}} \\ s_{\mathrm{WE}} \end{bmatrix}_i \right) + \nu_{\mathrm{SQUID}}.$$

Dynamics. The individual gyro orientations, $\boldsymbol{s_i}$, with respect to inertial space are driven by the local inertial frame rotation rate, $\boldsymbol{R}$; for the GP-B configuration general relativity predicts the north and west components of $\boldsymbol{R}$ to be $R_{\mathrm{NS}} = -6606$ mas yr^{-1} and $R_{\mathrm{WE}} = -39$ mas yr^{-1}, other drift sources being at this stage taken to be negligible.

Measurement. The SQUID readout signal, z_i, is the dot product of the normal to the gyro readout loop, n, with the difference $(\boldsymbol{\tau} - \boldsymbol{s})$ between the satellite roll axis and the inertial orientation of the gyroscope multiplied by a scale factor, C_g. Because the vehicle is rolling about the line to the star, the orientation of n changes uniformly with the roll angle ϕ_r. When the guide star is visible the pointing direction $\boldsymbol{\tau}$ is measured to high accuracy by the cryogenic telescope. C_g comprises a very nearly constant London moment term M_{L}, plus a modulation at

polhode frequency due to the presence of trapped magnetic flux in the rotor, of amplitude 0.5–3%. M_L, depending on which gyro. The pointing angle includes also both the annual (20.4958 arcsec peak) and orbital (5.1856 arcsec peak) aberrations of starlight; these are computed respectively from JPL ephemerides data and precise orbit measurements from on board GPS receivers and laser reflectors, and used, as noted earlier, to provide absolute calibration signals for the gyro readout. The measurement noise, $\nu_{\,\mathrm{SQUID}}$, near roll frequency is $\sim$200 mas $\mathrm{Hz}^{-1/2}$, and this fact is crucial in the practical implementation of the scale factor calibration. C_g needs to be known to $\sim$1 part in 10^5 and that requires connecting unambiguously the data from many successive orbits, recalling that the guide star is only visible for just over half the orbit.

From measurements of pointing, roll angle, and SQUID output, together with knowledge of the aberration signals, a time history of the orientation of the individual gyroscopes can be formed. Estimates of the drift terms can be formed by a fit to the linear gyro dynamics equation.

3.2. Patch effect torques and their complications to the data analysis effort.

During the design of the GP-B experiment, we recognized that patch effect fields on the gyro rotors might cause disturbance torques, and invested considerable effort to eliminate them. Nevertheless, we found, as already stated, that larger than expected disturbance torques were present, and, following an extensive error tree analysis combined with new laboratory tests on gyro rotors, established that patch charges consistent with a $\sim$100 mV cm dipole (and higher terms) were present and were in fact the root cause of three distinct effects.

(1) Misalignment torques proportional to the angle between the gyro spin vector and the spacecraft roll axis, producing a gyro drift perpendicular to the plane formed by these two vectors.
(2) The polhode damping, already described, which causes the spin axis as viewed in the body frame to transition toward the maximum moment of inertia. This has the immediate effect of smearing the previously stable spectrum of polhode-driven scale factor variations, and hence intermingling variations in C_g with the changes in gyro orientation.
(3) Resonant-effect torques that appear when a harmonic of the polhode period is in zero-beat with the roll frequency of the spacecraft.

These three effects complicate the data analysis, but once understood can be modeled accurately and separated from the general relativity rate estimates. Our primary effort following the completion of the science data gathering phase of the mission has been to identify, isolate and remove their disturbing consequences.

3.3. Modified gyroscope dynamics and measurement equations

Two additional terms in the dynamics equation and one in the measurement equation are sufficient to take the three identified effects in to account. In the dynamics equation, the first of the new terms is the classical drift due to the misalignment ($\boldsymbol{\tau} - \boldsymbol{s}$) multiplied by a time-dependent torque coefficient $K_\mu\,(t)$; the second represents the so-called 'resonance' torques. In the measurement equation, C_g. becomes time dependent not in the simple repetitive way originally conceived but evolving continuously with the changing polhode period:

$$\text{dynamics}: \quad \frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix} s_{\mathrm{NS}} \\ s_{\mathrm{WE}} \end{bmatrix}_i = \begin{bmatrix} R_{\mathrm{NS}} \\ R_{\mathrm{WE}} \end{bmatrix} + K_\mu(t)\left(\begin{bmatrix} \tau_{\mathrm{NS}} \\ \tau_{\mathrm{WE}} \end{bmatrix} \times \begin{bmatrix} s_{\mathrm{NS}} \\ s_{\mathrm{WE}} \end{bmatrix}_i\right) + \begin{bmatrix} f^{\mathrm{RES}}_{\mathrm{NS}}(\phi_r, \phi_P) \\ f^{\mathrm{RES}}_{\mathrm{WE}}(\phi_r, \phi_P) \end{bmatrix}$$

$$\text{measurement}: \quad z_n = C_g(t, \phi_P)\hat{n}(\phi_r) \bullet \left(\begin{bmatrix} \tau_{\mathrm{NS}} \\ \tau_{\mathrm{WE}} \end{bmatrix} - \begin{bmatrix} s_{\mathrm{NS}} \\ s_{\mathrm{WE}} \end{bmatrix}_i\right) + \nu_{\mathrm{SQUID}}.$$

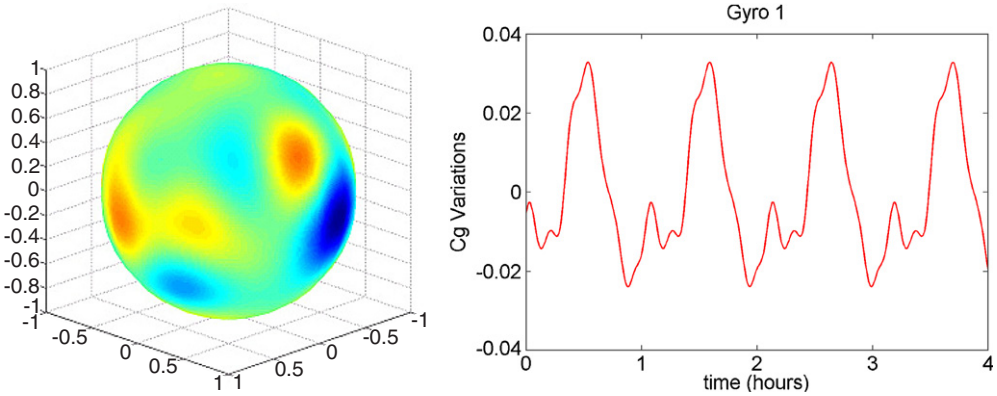

Figure 5. Gyro 1 trapped flux map (left) and resulting C_g variations (right).

Fixing the equations is simple; applying them to the actual experimental situation is not. To do so requires exact knowledge through the entire science phase of two key parameters: (1) for torque computation, a continuous history of the vehicle pointing angle τ, not only when the guide star is visible, but also during the times amounting to almost half of each orbit when it is occulted by the Earth; (2) for C_g determination, a similar continuous history of the polhode phase ϕ_P to connect successive separate aberration measurements. Neither was immediately available. While it had been our intention before launch to refer vehicle pointing during the guide star invalid periods to the science gyroscopes, the exigencies of on-orbit operations restrained us from doing so; instead, pointing was referred to the much less accurate external rate-gyros. (Likewise for the polhode phase.) The gyro readout had not been designed to measure ϕ_P; the only important pre-launch consideration with respect to polhode motion was to have the frequency determination stable enough to allow trapped flux terms to be spectrally separated from the science band around space vehicle roll frequency.

We now discuss how both difficulties have been surmounted, starting with a process of trapped flux mapping to determine ϕ_P.

3.4. Treating polhode phase and scale factor variations

Though the readout was not designed to measure ϕ_P, data was available from engineering telemetry channels in the form of 2 s long 2.2 kHz bursts, which could be applied for this purpose. From it we obtained 'snapshots' of the trapped flux signal over a number of gyroscope spin cycles, and from the evolving flux pattern thus seen were able to track with remarkable precision the motion of the spin axis in the body frame over time.

This trapped flux mapping (TFM) technique produces two critical results: (1) a history of the polhode phase angle good to 1° over the entire 353 days of science data and (2) a direct estimate of polhode-driven variations in the C_g scale factor. The mapping is accomplished by fitting a constant-coefficient rotor-fixed flux map to a dynamic model of the polhode motion of the rotor. Figure 5 shows (left) the flux map for Gyro 1 (which, as would be expected with trapped flux in a superconductor, remained perfectly constant within the measurement limit throughout the mission), and (right) the periodic variation in scale factor over a 4 h time period as the polhode motion modulates the direction of the flux pattern. The results for gyros 2, 3 and 4 had a similar form with successively lower flux levels.

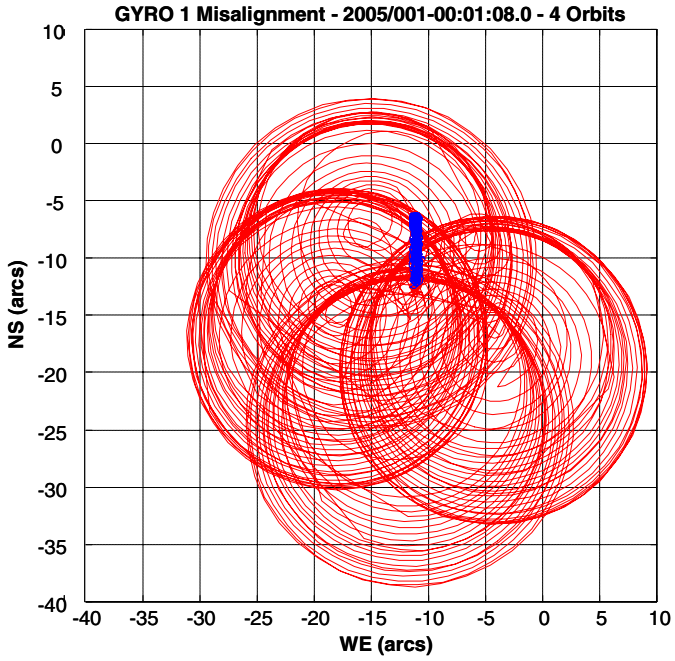

Figure 6. Representative in-view (blue) and occulted (red) pointing.

With this, data from multiple orbits could be chained together coherently in order to calibrate C_g to enough precision to resolve mas changes in the gyro orientation. Absolute polhode phase is now known to $\sim 1^\circ$ for all four gyroscopes over the entire mission, spin speeds to $\sim$1 nHz and spin-down rates to $\sim$1 pHz s^{-1}.

3.5. Estimating a continuous pointing history using science gyroscope SQUID data.

With the addition of the torque terms to the dynamical model, we need, as already remarked, a continuous history of the vehicle pointing angle τ not only when the guide star is in view but also when it is occulted by the Earth. Though the pointing in the guide star invalid phase was referred to the external rate gyros, its magnitude and phase could be determined rather accurately from the SQUID readout of science gyroscopes. Figure 6 illustrates results for four orbits. As expected, pointing is much poorer during occultation with a coning motion having a characteristic range of $\sim$25 arcs as compared with the 50 mas (plus orbital aberration) when the star is visible, but the combined telescope/science gyro measurements have allowed us to determine τ to better than 10 mas RMS over the mission, good enough to meet the requirements of the full misalignment torque model.

3.6. Misalignment torque modeling and removal

During the science phase of the mission, we observed some surprising shifts in gyroscope spin direction during anomalous periods when the spacecraft was temporarily disabled. Definitive proof of the existence of much larger than expected misalignment torques came in tests during the final calibration phase where the vehicle was deliberately pointed away from IM Pegasi to a succession of other nearby real and virtual stars. A combination of on-orbit observations and

ground-based tests on flight-like rotors established that these torques arise from patch effect charge interactions between the housing and rotor. The interaction is complex because the patches, like the magnetic trapped flux, are rotor-fixed and consequently their location with respect to the housing is continually changing because of the damped polhode motion of the rotor.

We now distinguish two methods of separating relativity from classical drifts, which are conveniently referred to as the *geometric* and *algebraic* methods.

First, while the magnitude of the torque (as parametrized by K_μ in the model) may be uncertain, its direction, given by the cross product of vehicle pointing and gyro spin vectors is known. Also, the vehicle pointing changes through the year in an exactly known way due to the annual aberration of starlight. The geometric method exploits this precise geometrical relationship by making a change of variables in the model so as only to observe the component of relativistic drift perpendicular to the torque direction. By plotting the sinusoidal variation of this term against the phase of the torque direction over the year, the need for an explicit model of K_μ is eliminated, greatly simplifying the analysis. The measurements of the geodetic effect reported at the APS meeting in April 2007 were obtained with this elegant method. Beautiful and illuminating as it is, however, it was in its original form limited by the requirement for short batch lengths (5 days), by being restricted to data from times when the guide star is in view, and by biases resulting from estimating drift-rates from short data intervals. The limit to it from SQUID noise using the full science data set is $\sim$20 mas yr^{-1}, a factor of 50 higher than the $\sim$0.4 mas yr^{-1} with year-long integration quoted earlier. We are exploring ways in which this limitation can be overcome, partly in combination with the algebraic method now to be described.

The algebraic method embodies the same geometrical relationships but explicitly models K_μ and thus allows a complete integration of the dynamic model. In addition, it can process all four gyros together to form a joint estimate of the general relativity effects under study. The relativity drift rates $\boldsymbol{R}$ are manifestly identical for all gyros, and by virtue of being on the same space vehicle, vehicle pointing is identical for all as well. This is the method by which the 85 days of data in December 2004 to March 2005 were processed; extensions to more segments are underway.

3.7. Resonance torques—underlying physics understood and effects mitigated.

In early 2007 we identified an additional torque from interactions of the patch effect charges when a harmonic of the polhode frequency zero-beats with the space vehicle roll frequency. A non-roll averaged torque is produced which causes step change in orientation. The physical underpinning of this torque is now well understood and when combined with the precise polhode angle measurements from trapped flux mapping to identify intervals where these torques are active it become possible either to excise (geometric method) or physically model (algebraic method) the effect. This surprising effect turns out to be much easier to treat than would first appear. The resonance intervals are gyro specific and non-overlapping. When they occur in any one gyro, good data from the other three is used to measure the magnitude and direction of the step and make an accurately calibrated correction.

4. Conclusions

Our assessment is that the underlying physics of the major disturbances to the science measurement is now well understood and that methods to remove them credibly from the measurement are maturing rapidly. In particular, precise knowledge of the polhode phase is

Table 1. Estimated relative improvement in statistical error versus the number of segments analyzed..

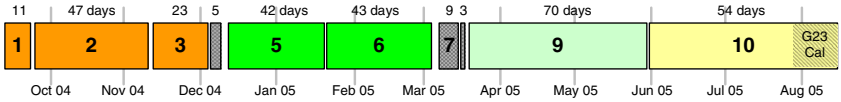

Segments	**5, 6**	**5, 6, 9**	**5, 6, 9, 10**	**3, 5, 6, 9, 10**	**2, 3, 5, 6, 9, 10**	**All**
Total duration (days)	85	155	209	232	279	307
Relative improvement ($t^{3/2}$)	1	2.5×	3.9×	4.5×	5.9×	6.9×

Table 2. Overall estimate of improvement potential.

Data analysis category	Improvement potential
Extension of analysis to all major segments	3.9– 5.9×
Full trapped flux mapping (TFM) scale factor knowledge	2×
Modeling improvements for identified systematic effects (thermal, etc)	1.5– 2×
Final refinement of data grading (use more good and less bad data)	1.2– 1.5×
Overall improvement potential	14– 35×

the key synchronizing element needed to tie together the entire mission data set. The gyroscope drift estimation machinery can accurately model polhode-modulated mission-length variations of the readout scale factor and separate them from estimates of gyro drift.

Analysis through the fall of 2007 has produced a robust result based on 85 days of data from 12 December 2004 to 5 March 2005. This time period was chosen because of its representative character—fairly active polhoding, an interesting solar flare anomaly, distinct resonances torque intervals—making it an excellent period to validate the data analysis codes and methods. The results shown in section 2 are encouraging and exhibit the expected increase in precision of the estimate when going from an initial 42 day period to the longer 85 day period. A gyroscope drift consistent with frame-dragging is clearly seen. This must, however, be viewed with proper caution as a suggestive rather than established result since it is the average from three gyros whose individual outputs show wider scatter.

The accuracy of the relativity drift estimates and their separability from disturbance torque effects is directly and strongly affected by the quantity of data used. In principle, it is straightforward to extend the analysis to longer periods, but a fair amount of preparation and internal consistency checking is needed to ensure that the results derived from longer data stretches are valid. Table 1 presents the expected relative improvement as more segments are included. No fundamental 'breakthroughs' are needed to extend the analysis to these other segments. An intense effort is now in process to carry this work to completion.

The ten segments shown in table 1 are punctuated by various space vehicle anomalies where the science data was either corrupted or not available. Events include large solar flares that blind the telescope, radiation-induced soft computer failures, attitude control issues, etc. Practically speaking, only about 307 days of the total 353 days of science data are of high quality and are readily analyzed. With the large gains expected from the inclusion of additional segments, results from other refinements in the analysis will further improve the overall result; these are summarized in table 2. The gains come from the finalization of the C_g determination from trapped flux mapping, as well as from improved data grading and preprocessing and elimination of other systematic effects (thermal, resonance torques, etc).

The improvements listed in table 2 are from independent sources and should in principle multiply. If the full effect of these improvements is realized, the uncertainty in the GR measurement can be reduced to $\sim$0.5 mas yr^{-1}, equivalent to a $\sim$1% measurement of frame-dragging. In a real system, however, complete independence is unlikely; a conservative assessment of the likely accuracy of the frame-dragging measurement is 3–6% using the present analysis approach and toolsets.

Acknowledgements

In addition to the authors listed, significant contributions to GP-B data analysis were provided by P Boretsky, J Kozaczuk, D Santiago, P Shestople and S Wang. This work was supported by NASA under contract NAS8-39225 through the George C Marshall Space Flight Center (MSFC).