

Gravity Probe B data analysis: II. Science data and their handling prior to the final analysis

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2015 Class. Quantum Grav. 32 224019

(<http://iopscience.iop.org/0264-9381/32/22/224019>)

[View the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 131.169.4.70

This content was downloaded on 08/01/2016 at 08:43

Please note that [terms and conditions apply](#).

Gravity Probe B data analysis: II. Science data and their handling prior to the final analysis

A S Silbergleit¹, J W Conklin^{1,2}, M I Heifetz¹, T Holmes¹, J Li¹, I Mandel^{1,4}, V G Solomonik¹, K Stahl¹, P W Worden Jr¹, C W F Everitt¹, M Adams¹, J E Berberian¹, W Bencze¹, B Clarke¹, A Al-Jadaan^{1,3}, G M Keiser¹, J A Kozaczuk¹, M Al-Meshari^{1,3}, B Muhlfelder¹, M Salomon¹, D I Santiago¹, B Al-Suwaidan^{1,3}, J P Turneaure¹ and J Wade¹

¹ HEPL, Stanford University, Stanford, CA 94305-4085, USA

² MAE Department, University of Florida, Gainesville, FL 32611, USA

³ King Abdulaziz City for Science and Technology, Riyadh, Saudi Arabia

⁴ School of Physics and Astronomy, University of Birmingham, Birmingham B15 2TT, UK

E-mail: francis1@stanford.edu.

Received 30 January 2015, revised 6 April 2015

Accepted for publication 7 May 2015

Published 17 November 2015



Abstract

The results of the Gravity Probe B relativity science mission published in Everitt *et al* (2011 *Phys. Rev. Lett.* **106** 221101) required a rather sophisticated analysis of experimental data due to several unexpected complications discovered on-orbit. We give a detailed description of the Gravity Probe B data reduction. In the first paper (Silbergleit *et al* *Class. Quantum Grav.* **22** 224018) we derived the measurement models, i.e., mathematical expressions for all the signals to analyze. In the third paper (Conklin *et al* *Class. Quantum Grav.* **22** 224020) we explain the estimation algorithms and their program implementation, and discuss the experiment results obtained through data reduction. This paper deals with the science data preparation for the main analysis yielding the relativistic drift estimates.

Keywords: general relativity, GP-B, data analysis



Content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

1. Introduction

Gravity Probe B (GP-B) was intended to measure the precession, or drift, of a free gyroscope orbiting the Earth with enough accuracy to check the predictions of general relativity (GR) found by L I Schiff in 1962 [4]. Its concept was very simple: put a star-tracking telescope and a high precision gyroscope on orbit, keep the telescope pointing at the guide star (GS) as an inertial reference, and measure the rate of change of the angle between the line of sight to the GS and the gyro spin (see figure 1 in [1] or [2]).

The implementation of this simple idea proved to be tremendously difficult: the effects to be measured were extremely small, and the number of disturbing factors, which had to be reduced by many orders of magnitude to ensure an accurate measurement, was huge. It is enough to note that the quality of the flight gyroscope had to be enhanced $\sim 10^7$ times as compared to the existing ones, and the quality of the tracking telescope $\sim 10^3$ times. Only the combination of cryogenics with space, supported by the development of numerous special break-through technologies, allowed us eventually to bring the experiment to a success. The whole experiment, including the science instrument and its development, the attitude and translation control (ATC) system of the spacecraft (S/C), and on-orbit operations, is described in [5–7].

Before launch we had anticipated the analysis of GP-B data to be just slightly more complicated than usual (see [8] and [2], section 4F). Several unpleasant on-orbit discoveries changed this expectation drastically. The three major ones were the changing gyro polhode period, significantly larger than expected patch effect torques on the gyroscopes, and data segmentation due to S/C anomalies. These unexpected circumstances required, on the one hand, a number of theoretical results to model the disturbances, and, on the other, the most careful data treatment combined with modern estimation techniques (signal filtering), to extract the relativity estimates from the available signals with satisfactory accuracy.

This paper provides a thorough description of the GP-B data analysis explained very briefly in [1]. The first paper [2] (abbreviated DA I below), contains detailed theoretical derivations of all the signal models needed for the analysis. The present second paper, DA II, describes the treatment of the science data preceding the estimation of relativistic drift rates. Estimation algorithms and their program realization (in particular, the main filtering tool, the ‘2-second filter’), as well as the resulting relativity estimates with their errors, are found in the third paper, DA III [3].

Some convenient terms and notations are used below to describe the structure of this paper. There were several steps in GP-B data handling on the ground prior to the main data analysis. The data after each step form what we call *levels*, which are numbered successively, from Level 0 raw telemetry data to Level 3 data ready for the relativistic drift estimation analysis. For brevity, we will usually write LN, ($N = 0, 1, 2, 3$), instead of Level N ; Level N data are kept in the Level N database denoted LN DB.

In section 2 we describe data segmentation due to S/C anomalies and the set of science signals. Ground data processing, starting with the raw data treatment (L0 to L1), and followed by the so called data preprocessing (L1 to L2), is presented in section 3. Section 4 deals with grading the science data based on various criteria. Final preparation of SQUID and telescope data (L2 to L3) for the main analysis, i.e., for the estimation of relativistic drift rates of GP-B gyroscopes, is explained in section 5. We formulate some conclusions in section 6.

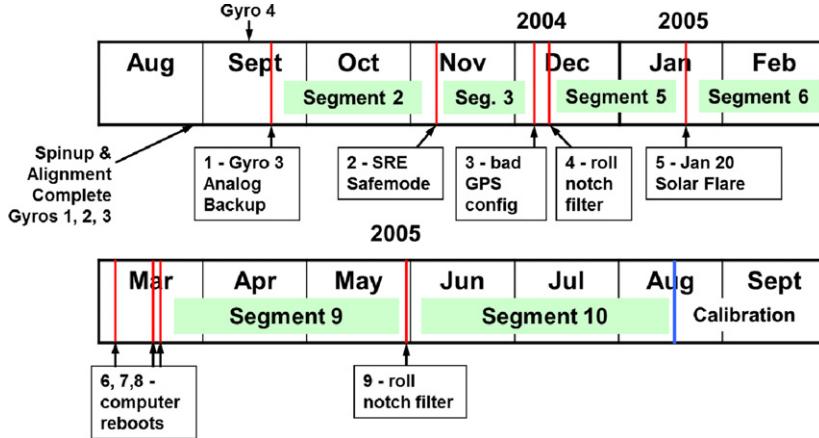


Figure 1. Nine spacecraft anomalies and ten science data segments, which comprise the science data collection phase of the mission.

2. Science data

2.1. Data segmentation

During the science data collection standard on-orbit operations were interrupted a total of nine times because of S/C anomalies. These had a variety of causes. Figure 1 shows the nine anomalies and ten science data segments.

The most frequent anomaly was a reboot of the spacecraft's flight computer. These reboots were probably caused by charged particle radiation interacting with the flight computer memory, causing 'double event upsets'. A different case is the anomaly which began on 20 January 2005, when the incidence of the solar proton flux with energies greater than 10 Mev went up from below 0.1 to over 100 particles $(\text{cm}^2 \text{ s sr})^{-1}$ for several hours, returning to within a factor of 10 of its quiescent level only after a day and a half. The flux of protons and secondary particles was such that the telescope photodiodes were receiving particle hits faster than their 10 Hz reset rate. So the signal from the science telescope was unusable, and the attitude control system had to rely only on the star trackers and rate gyroscopes. Over the course of one day the satellite roll axis drifted approximately 8 arc-min (12 mrad) from the GS before the particle radiation dropped to a level at which the quality of the telescope signal was satisfactory. Even then it took about 20 h after the flare started to reacquire the GS and re-establish the normal operating conditions.

The loss of data during the anomalies was less of a concern than the possibility of larger than expected torques acting on the gyroscopes during them. The data from each anomaly were carefully evaluated to determine which additional parameters should be included in the data analysis to account for any offset in the orientation of the gyroscope spin axis during these periods.

Of the ten available segments of science data, only six (segments 2, 3, 5, 6, 9, and 10) were deemed useful for the determination of the gyroscope relativistic drift rate. This was primarily due to the shortness of the four segments 1, 4, 7 and 8; each of them lasted eight days or less. The analysis of each new segment required additional parameters in the models describing the gyroscope dynamics and readout. These include new initial conditions for the gyro orientation, coefficients describing the gyro readout scale factor and misalignment

Table 1. Duration of each segment used in the science data analysis.

Segment number	Approximate start	Approximate end	Approximate duration
2	25 September 2004	10 November 2004	46 days
3	12 November 2004	4 December 2004	23 days
5	12 December 2004	19 January 2005	39 days
6	22 January 2005	4 March 2005	42 days
9	19 March 2005	30 May 2005	73 days
10	1 June 2005	23 July 2005	53 days

Table 2. Total number of days worth of science data included in the analysis of each gyroscope.

Gyroscope number	Number of days of data used
1	201
2	220
3	189
4	242

torque coefficient, and potentially new roll-polhode resonant torques. These new parameters are independent of the parameters from any other segment. Only the relativistic drift rate and telescope scale factor parameters are common for all the segments. As a result, the four short segments do not contribute in a statistically significant way to the estimation of the North-South and West-East relativistic drift rates.

The approximate beginning and ending dates, as well as the duration of the six long data segments used in the science data analysis are provided in table 1. They are approximate because they vary slightly from gyroscope to gyroscope; the exact choice of the start and end time of every segment for each particular gyroscope was based on the results of data grading (see section 4). The time spans shown in table 1 represent the maximum span used for at least one of the gyroscopes. For example, at the end of the science data collection phase of the mission, gyroscope 3 began its post-science calibration phase earlier than other gyroscopes, shortening the duration of segment 10 for this gyroscope.

The total duration of science data used for each of the four gyroscopes is shown in table 2. The smallest amount of data were 189 days for gyroscope 3, and the largest amount was 242 days for gyroscope 4. However, this does not directly translate into a larger statistical error in the estimate of the drift rate for gyroscope 3, because the total number of parameters to be estimated and their correlation to relativistic drift rate is also important. Indeed, gyroscope 2 contributed the least amount of information to the combined relativistic drift rate, because it had many more roll-polhode resonances leading to the largest number of estimated parameters.

2.2. Set of science signals

Here we describe the main set of GP-B signals used for the science data analysis, whose primary goal was to estimate the relativistic drift rates of the four science gyroscopes in the orthogonal North-South and West-East directions. This excludes signals used to determine the state of the GP-B spacecraft (e.g. control modes), and used to build confidence in the

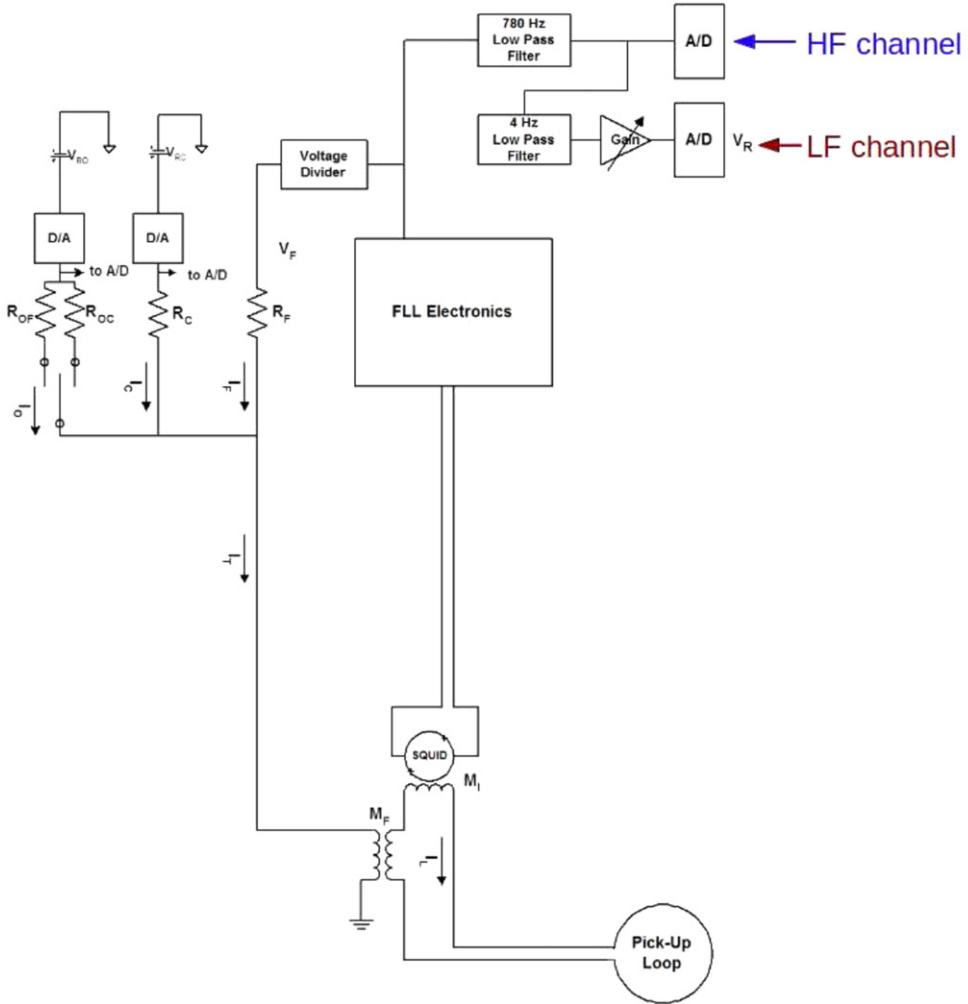


Figure 2. SQUID high frequency and low frequency readout electronics.

performance of the instrument (such as temperature monitors, gyroscope suspension signals, proton flux and so on), as well as the mathematical models describing the gyroscope dynamics and readout detailed in [2]. Only the signals that were directly involved in the parameter estimation are described here.

The science signals used in the data analysis include (a) timing data, (b) both low frequency (LF) and high frequency (HF) SQUID data, (c) telescope data, (d) the satellite roll phase, and (e) orbit data. All science signals are time-tagged to the satellite's 16.368 MHz quartz oscillator, using a time scale called *vehicle time*. The process by which the science signals were time-tagged on the satellite and then corrected on the ground via preprocessing is described in detail in section 3.2.

2.2.1. HF and LF SQUID signal. As the satellite rolls about the direction to the GS, magnetic flux through the pickup loop due to the London moment of the rotor varies at the satellite roll frequency. The amplitude and phase of this roll frequency signal are a direct measure of the

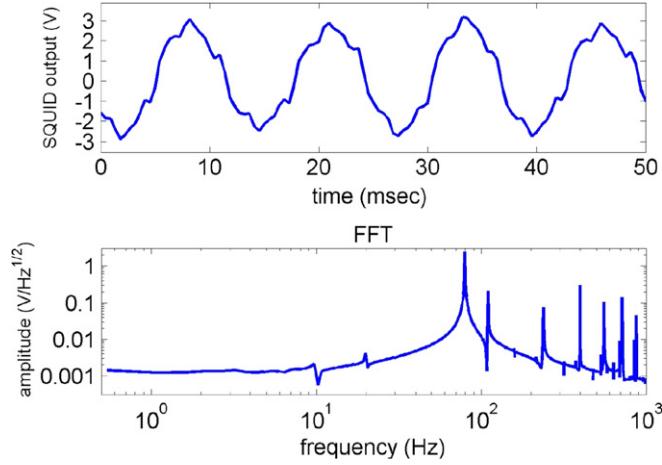


Figure 3. Time-series and FFT of a gyroscope 1 SQUID HF snapshot, 10 November 2004.

magnitude and direction of the misalignment between the gyroscope spin axis and the satellite roll axis, which is approximately aligned with the apparent direction to the GS. The magnetic London moment signal allows one to determine the angle between the gyro spin axis and the plane of the pickup loop. In addition to the London moment signal, there are also signals at harmonics of the gyroscope spin speed due to trapped magnetic flux [2, 9]. The SQUID readout electronics, shown schematically in figure 2, was designed to preserve the information from both of these sources. A HF channel is digitally sampled at 2200 Hz following a 780 Hz analog low pass filter, while the LF channel has an additional 4 Hz analog low pass filter whose output is also digitally sampled at 2200 Hz. A non-causal digital filter further attenuates the HF components of the over-sampled LF signal before it is included in the telemetry at a 5 *samples/s* rate.

The LF SQUID signal is considered the primary science signal, containing the relativistic drift rate information. There is one such signal for each of the four gyroscopes. The HF SQUID signal is packaged on board the S/C into 4096 sample, 2.2 kHz *snapshots* and sent to the ground along with the six lowest harmonics of the HF signal's FFT. The spectrum and part of the time-history of an example snapshot are shown in figure 3.

The FFTs of the HF channel are used throughout the Science Data Collection Phase to monitor the spin and polhode frequencies of the four gyroscopes. Precise determination of the these fundamental frequencies during the Science Data Collection Phase from the HF FFTs is described in [10]. We found the time histories of the rotor polhode motion and the readout scale factor using the HF snapshots. The snapshots were sent to the ground at roughly 40 s intervals during approximately half of each orbital period. Around 2×10^5 snapshots were acquired for each of the four gyroscopes over a roughly year long period starting in early September 2004 and ending mid-August 2005.

2.2.2. Telescope data. The cryogenic telescope [11–14] that was used to track the apparent position of the GS is a folded Cassegrain telescope with an aperture of 14 cm, a focal length of 3.8 m, and a field of view greater than 66 as, or 320 μ rad; here and elsewhere ‘as’ stands for ‘arc-second’, and ‘mas’ stands for ‘milli-arc-second’ (approximately 5×10^{-9} rad).

The light from the GS passes through four windows in the cryogenic probe. It is then reflected from the primary, secondary, and tertiary mirrors, entering an image divider assembly mounted on the front plate of the telescope. Within the image divider assembly, a beam splitter divides the light in half, and each half is focused on the vertex of a roof prism. If the telescope is pointed directly at a star, its light will be divided equally by the vertex of the prism. The light beams from both sides of the roof prism are sent to redundant pairs of silicon photodiodes, converting the light intensity to electric current.

Within the nominal ± 400 mas ($10 \mu\text{rad}$) linear range of the telescope, the ratio of the difference in the currents from the two sides of the roof prism to their sum is a direct measure of the angular deviation of the telescope axis from the apparent direction to the GS. (Note that to really linearize the telescope signal in this whole range, a cubic correction had to be introduced, see [2], section 6B.) Each pair of photodiodes measured the difference in light from both sides of the roof prism. Consequently, the total volume of telescope data used in the science analysis included the currents produced by each of the eight photodiodes, $\{i_{x,A}^+, i_{x,A}^-\}$, $\{i_{x,B}^+, i_{x,B}^-\}$, $\{i_{y,A}^+, i_{y,A}^-\}$, $\{i_{y,B}^+, i_{y,B}^-\}$. Here \pm represent the individual $+$ and $-$ currents for each detector pair, x and y denote the satellite body-fixed x and y axes, and A and B stand for the two redundant pairs of photodiodes corresponding to the two telescope readout electronics (TRE) boxes. These signals are combined using the telescope models described in [2]. The combined pointing noise from all four detector pairs was ≈ 0.48 mas per orbit.

Photodiodes are sensitive to particle radiation, which typically produces an offset in the snapshot data of the charge feedback loop. This in turn biases the photo current data output. For this reason it was necessary to eliminate any telemetry data points that were clearly charged particle hits. The number of energetic protons in the South Atlantic Anomaly made the telescope readout particularly vulnerable during those times when the satellite passed through this region, and the corresponding data were typically not used in the analysis.

2.2.3. Satellite roll phase. The entire S/C with the science instrument rolled about the line of sight to the GS with a period of 77.5 s. Therefore, both the telescope and the SQUID signals were modulated at the roll frequency. The satellite roll phase was thus needed to demodulate these signals and transform them into the inertial North-South, West-East reference frame [1, 2], and for many other purposes.

We determined the roll phase on the ground from the star tracker telemetry data. Two attitude reference platforms were attached to a graphite ring around the Dewar of the GP-B satellite, and one star tracker was mounted on each platform. The field-of-view of the star trackers was $8^\circ \times 8^\circ$, and the boresight axes were 50° and 60° away from the satellite roll axis, out of phase by 180° . A star tracker measures the position of the centroid of a star image on the plane of a charge coupled device (CCD). The telemetry data from the GP-B star trackers, particularly, the horizontal and vertical components of the centroid position, the star magnitude and the time of the centroid detection were preprocessed to determine the roll phase of the quartz block referenced to true north. Since the SQUID readout loops were fixed in the quartz block, the roll phase of each loop could be determined from the roll phase of the quartz block; this was done during ground preprocessing (see section 3.2).

2.2.4. Orbit data. Two redundant, custom GPS receivers on-board the satellite each provided both orbit and timing information. Together with the JPL Earth ephemerides, these data were used to calculate the total aberration of light from the guide star, which was necessary for calibration of the gyroscope and telescope readouts. While the GPS data was not directly used in the science data analysis, the aberration signal was a critical input; much smaller

contributions of the GS light bending and parallax were also computed from the orbit data and included in the measurement model [2].

3. Ground data processing: from L0 to L2 data

3.1. Conversion from raw (L0) to Level 1 data

GP-B telemetry data were downloaded from the satellite approximately once in 12 h. These data, called raw, or L0 data, were entangled binary files where various monitor values were mixed with the time stamps, numerous repetitions occurred, etc; L0 data were stored and kept in the L0 DB. Usually right after each download a special team at Stanford Mission Operation Center (MOC) started processing raw data using a special tool, a TCAD program package developed at the University of Colorado, which had been properly adjusted for GP-B needs. The result of this processing, a set of decoupled non-repetitive time ordered monitors (L1 data), was usually ready within a few hours from the data dump. This was stored as ASCII files in the L1 DB.

The L1 DB is basically a table whose first column contains vehicle time stamps (see [7] and more below) ordered throughout the whole mission; the row for every moment of time contains the values of all monitors marked by that particular time stamp. The monitors are ordered, all the values of a given monitor form one column in the table. This structure is convenient for reading stretches of L1 data with a monitor selection; we developed and used special programs to make reading from and writing to all our databases fast and efficient.

3.2. Preprocessing: L1 to L2 data conversion

3.2.1. Preprocessing goals. The next step in handling the GP-B data is their preprocessing. It takes a time sequence of L1 values for some specified monitor(s), processes them as described below, and outputs the corresponding set of L2 data to the L2 DB for permanent storage. This operation has a three-fold aim, namely: (1) preliminary data cleaning; (2) data compression (for signals with sampling rates *above* 0.5 Hz) or interpolation (for signals with sampling rates *below* 0.5 Hz); (3) time stamp correction and data synchronization (all L2 data points correspond to the same moments of vehicle time separated by 2 s intervals, namely, a data point at each twentieth 10 Hz strobe [7, 15]).

All the major science signals (timing and orbit data, LF SQUID signals, telescope signals, roll phase data, calibration signal, etc, as well as many auxiliary monitors giving temperature readings, magnetometer signals, etc) have been preprocessed. The block-diagrams of this and some other signal preprocessing, including FFT harmonics of HF SQUID signal, are in figures 4 and 5. The HF SQUID signal needed for trapped flux mapping (TFM) ([2, 9]), crucial for the GP-B data analysis, was processed in a special way described in section 5.4.

Let us explain the goals (1)–(3). Preliminary data cleaning (1) was mostly outlier removal; typically, data points more than 3 standard deviations away from the average were considered outliers and removed from the data set. As for data compression (2), the majority of science signals had high sampling rates, like 10 Hz for the telescope signals, or 5 Hz for the four SQUID signals. Since the signals typically varied much more slowly, the rates and the corresponding large sets of data points were redundant for the analysis (in fact, even the chosen L2 time step of 2 s was initially thought to be too short, until the need to accurately model patch effect torques was established [8]). This was why the number of data points of frequently sampled signals was reduced between L1 and L2. Additionally, data analysis

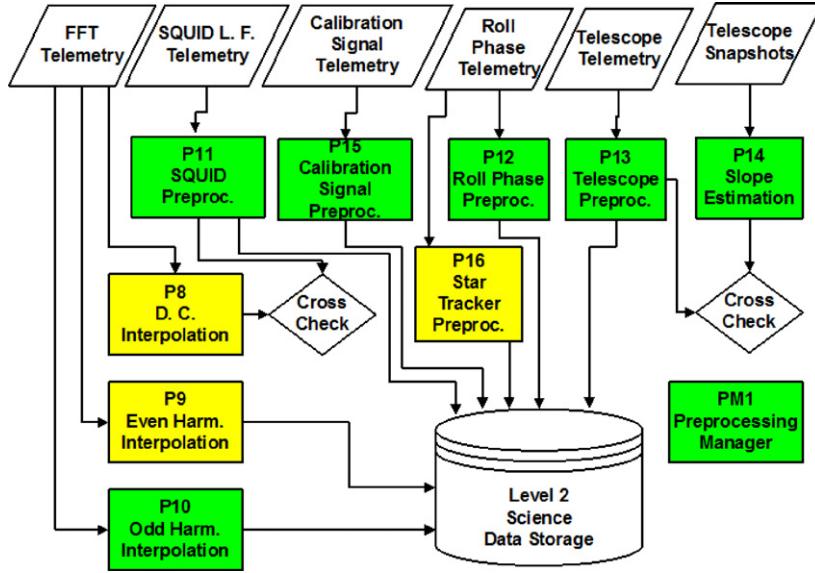


Figure 4. Preprocessing algorithms: SQUID, telescope and roll phase signals.

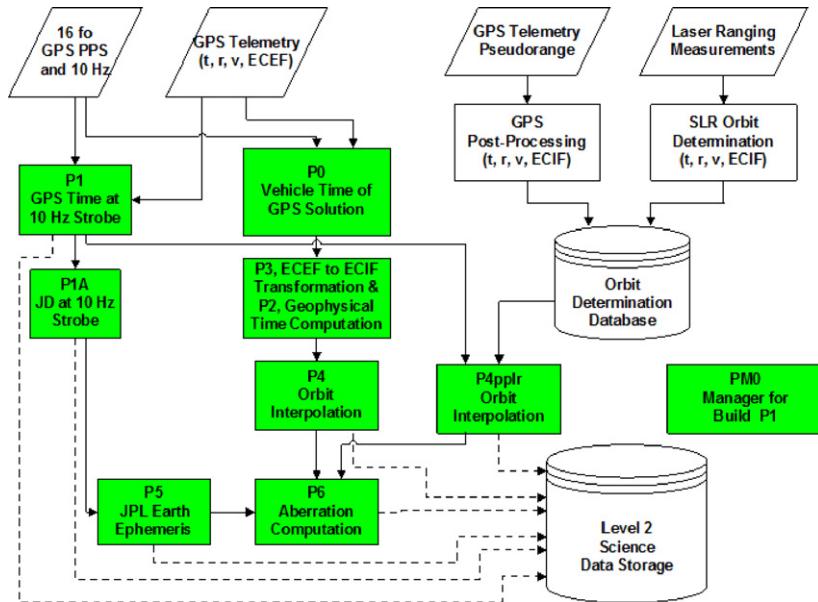


Figure 5. Preprocessing algorithms: timing and orbit data.

becomes far more difficult if different signals are not synchronized, i.e., their values do not correspond to the same instants. To specify the signal values at the same moments we had to interpolate the signals with lower sampling rates, as stated.

Finally, a time stamp correction (3) was needed because each monitor value was time stamped onboard the GP-B satellite at the moment it was placed in the telemetry format table

(TFT), *not* at the actual instant of the measurement. The latter usually differed significantly due to time delays in the onboard circuitry. So first the time step of each L1 data point was corrected, and then a single L2 data point at a predetermined moment of vehicle time was derived from a number of neighboring L1 data points (see more on preprocessing below in 3.2.2–3.2.5; GP-B timing is thoroughly described in a special paper [15]).

Additional issues with timing were fixed in the timing data preprocessing, see sections 3.2.3, 3.2.4. They stem from the fact that all the telemetry monitors are time stamped by the vehicle time, which is also used as the only time at Levels 2, 3 and in the whole analysis. But the GPS orbit data needed for calculating the orbital aberration, one of the two SQUID signal calibrators, were marked by the GPS time. This time therefore had to be accurately connected to the vehicle time, and to various geophysical times. These times participate in the conversion of GPS data to the proper inertial frame. Moreover, one of them, the so called Julian Date (JD), serves as an input to JPL ephemerides used to get Earth position and velocity, and thus the annual aberration, our second signal calibrator. The corresponding L2 outputs of timing and orbit data processing are shown in figure 6.

Preprocessing of a new stretch of L1 data normally started soon after the new data chunk was in place in L1 DB. Usually the L2 output was ready in L2 DB within 24 h of the data dump. First, the L1 data stretch, typically about 12 h duration, with carefully chosen start and end points (such that the values of *all* monitors to be preprocessed were present at both times), was given a *Run ID*. Apart from the proper header, the Run ID is a special data set recorded in the L2 DB containing information about the start and end time of the corresponding data stretch, and about all parameter values used in the preprocessing of all signals belonging to this Run ID. Thus the Run ID provided the possibility of fully tracking the data from L2 back to L1 (see section 3.2.2). After a Run ID was set up, any selection of monitors belonging to it could be quickly preprocessed at any time; the process was controlled by a special program called Preprocessing Manager also described in 3.2.2.

Now we move on to a more detailed description of the most important elements of signal preprocessing.

3.2.2. RunID concept and preprocessing manager programs. Here we describe a group of functions that are used to create and maintain sets of preprocessing configuration parameters.

The set of preprocessing configuration parameters is the RunID configuration set, or just RunID. A RunID set was created for each chunk of telemetry data (usually, 6–12 h stretch) after it was recorded in the L1 DB but before the start of preprocessing. Each RunID was stored in the L2 DB. The RunID parameter set includes the start and end vehicle times of the data time interval, the telemetry cycle number, and all configurable preprocessing parameters. For instance, it contains the widths of windows for the least square fit (LSQ) to outputting L2 data points, or the number of iterations in the outliers search, or the like. The values of these parameters can be adjusted as needed during preprocessing to produce the best possible results. Note that the non-modifiable (constant) parameters used by both preprocessing and analysis algorithms, were stored in a separate file. These are fundamental physical constants, some constant parameters of the GP-B instrument design, the fastest telemetry rate supported, etc.

When preprocessed data are recorded in the L2 DB, the RunID configuration set that was in effect at the time of preprocessing is stored along with the data. Each preprocessed data point always keeps the reference to RunID that was used to generate it. This approach allows tracing and re-preprocessing any portion of the L2 data with a different configuration of parameters. The time intervals defined in different RunID sets cannot overlap: each RunID set

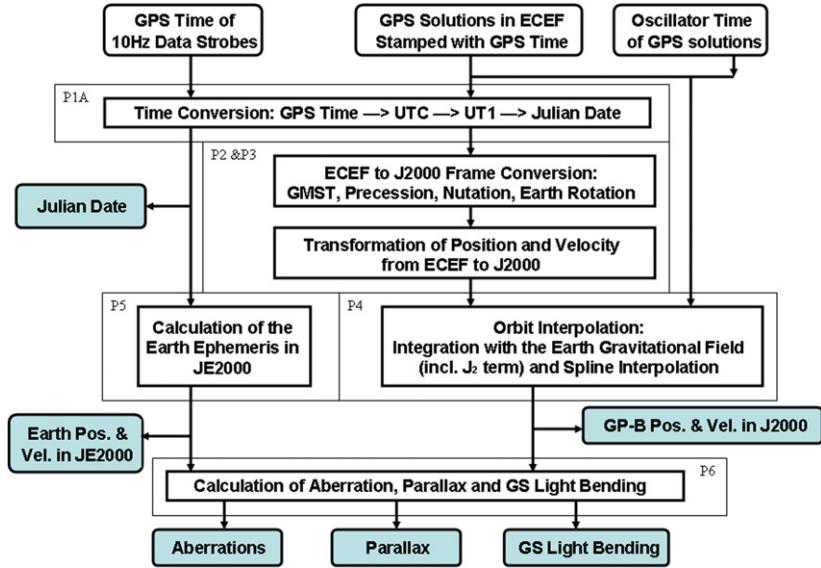


Figure 6. L2 outputs of timing and orbit data preprocessing.

covers its own time range. Usually, the time interval of the next RunID will be adjacent to the time interval defined for previously created RunID.

The preprocessing control group comprises four separate programs. The first one is used by any user to create a new RunID set. The other three implement the complete functionality of the configuration management, providing, in particular, tracking configuration parameters of any signal back from L2 data.

3.2.3. Preprocessing of timing data. The precision of GP-B measurements required accurate timing of the measured signals and was thus crucial for mission success. The GP-B timing system is described in detail in a separate paper [15] included in this issue. Here we give only a pertinent brief survey of timing data handling with emphasis on preprocessing of timing signals.

An onboard counter was incremented by 0.1 s at each 10 Hz data strobe [15] producing what is called *the vehicle time*. The science signals transmitted from the GP-B satellite to the ground station through the programmable telemetry were stamped with the vehicle time *when the data were written into the TFT, not when they were actually sampled*. Two redundant GPS receivers supplied an external reference for the time transfer between vehicle time and Coordinated universal Time (UTC).

On the ground, the effective sampling times of the GP-B science signals in UTC were determined in two steps: (1) the true vehicle time of the science signals sampling was determined; (2) conversion between the vehicle time and UTC was established in the timing data preprocessing. The GP-B timing requirements given in [15] contained the top level requirement, namely, that the effective sampling time of all science signals should be known relative to UTC with the accuracy of better than 0.1 ms, after ground preprocessing. It is important to note that the known constant timing offsets were compensated for in the ground processing of GP-B science data.

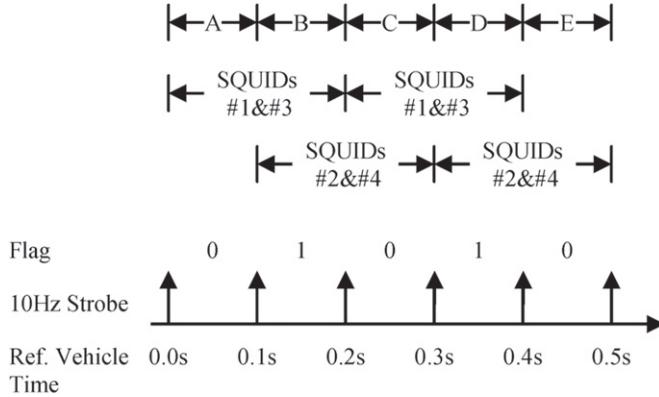


Figure 7. Timeline of SQUID signal lowpass filter.

As shown in figures 5 and 6, preprocessing the GP-B timing data was implemented in two programs: one determined the GPS time of 10 Hz data strobes, and the other determined the vehicle, or oscillator, time of GPS solutions. The oscillator time was derived from the 31 bit counter of the 16fo clock (see [15] for the definition), and it was equal to the vehicle time at the 10 Hz data strobes; therefore it was identical to the vehicle time to within 1/(16,368,000) part of a second.

The timing data preprocessing algorithms used two sets of telemetry data to establish the relation between the vehicle time and UTC. First, the timing data of the 16fo clock and the PPS (pulse-per-second, see [15] again) collected in the aft SRE box were used to determine the conversion between the vehicle time and the relative GPS time by means of LSQ fitting. Second, the timing data of the GPS data packet transmission and reception were used to resolve the ambiguity of the PPS, and the integer GPS second of the PPS was determined. The GPS time of 10 Hz data strobes and the vehicle time of GPS solutions were calculated, as one of the results. Since the difference between the GPS time and UTC is a known integer number of seconds plus a negligible fractional part (within a few hundred nanoseconds during the last several years), the conversion between the vehicle time and UTC was straightforward. The careful PPS data editing based on the residuals of the appropriate LSQ fit is described in detail in [15].

As mentioned, the vehicle time stamp attached to the GP-B telemetry data indicates the time when the data were written into the TFT, and not the actual time of the measurement (data sampling). The effective sampling vehicle time of GP-B science signals was determined by one of the following two methods during L1 to L2 data preprocessing. For less frequently sampled signals, the vehicle time of the signal sampling was recorded and included in the telemetry data; for more frequently sampled signals, the effective sample time was determined by subtracting the latency due to signal sampling, data processing, transfer and distribution, etc from the vehicle time stamp.

To illustrate how it was done, we consider one very important example. The effective sampling time for the outputs of the onboard algorithm of the SQUID lowpass filter, i.e., the timing of SQUID science signal, was determined as follows. The inputs of the algorithm were digitized SQUID signals, sampled in two forward SRE boxes labeled A and B. The signals of SQUIDs #1 and #3 were sampled with a multiplexer and an A/D converter at 2200 Hz per channel in the forward SRE box A, and the signals of SQUIDs #2 and #4 were sampled using the same tools in the forward SRE box B. Since the sampling processes were run in parallel

and synchronized by the same clock signal, SQUID #1 signal was sampled in forward SRE A simultaneously with the SQUID #2 signal sampled in forward SRE B; the same was true for the SQUID #3 and #4 signals. Before sampling, the SQUID signals went through two lowpass analog filters with the cutoff frequencies of 780 Hz and 4 Hz, respectively, so they were delayed by 44.32 ms. The sampling of the signals led to an additional time delay of 0.19 ms for SQUIDs #1 and #2, and of 0.21 ms for SQUIDs #3 and #4.

The timeline of the SQUID science lowpass filter is shown in figure 7. In order to distribute the onboard processor load more evenly, during each interval of 0.1 s two SQUIDs just collect data, while the other two collect and process them. A ping-pong flag maintained in the algorithm synchronizes the data collection and processing. In phase A, the ping-pong flag value is 0, and the SQUID #2 and #4 data are collected and processed, so two filter outputs, i.e. weighted averages of 2200 points of SQUIDs #2 and #4 data, are computed. In phase B, the ping-pong flag value changes to 1, so now the SQUID #1 and #3 data are collected and the weighted averages of 2200 points of SQUIDs #1 and #3 data are computed. The reference vehicle time corresponding to the midpoint of the SQUID #1 and #3 2200 data points, which include 440 new points collected in phases A and B, is approximately -0.4 s. In phase C the filter outputs of SQUIDs #1 and #3 are transferred to the CCCA (Command/Control Computer Assembly), and the ping-pong flag is updated to 0 again. Since the execution of the algorithm in the aft SRE box is not synchronized with the data distribution into the TFT in the CCCA, the filter outputs of SQUIDs #1 and #3 may be written in the TFT in either phase D or phase E, and the attached vehicle time stamp would be 0.3 s or 0.4 s, respectively. The ping-pong flag, equal to 1 in phase D and to 0 in phase E, was written in a TFT together with the filter outputs; this was used to resolve the timing ambiguity during data preprocessing. Taking into account the time delays due to the signal sampling and the analog lowpass filters, the time latencies of the SQUID lowpass filter outputs were determined as follows: when the ping-pong flag attached to the outputs of SQUIDs #1 and #3 was 1 (or, equivalently, the ping-pong flag attached to the outputs of SQUIDs #2 and #4 was 0), the time latencies were 744.36 ms for SQUIDs #1 and #2, and 744.34 ms for SQUIDs #3 and #4; otherwise, the time latencies were 844.36 ms for SQUIDs #1 and #2, and 844.34 ms for SQUIDs #3 and #4.

The true sampling times of all other signals were established in similar, and often simpler, ways, using the appropriate information (flags, etc). As demonstrated in [15], the performance of the GP-B timing system was very satisfactory both in the ground tests and during the flight.

3.2.4. Preprocessing of orbit data (frame conversion). GP-B science data analysis required accurate knowledge of the S/C position and velocity in an inertial system centered on the solar barycenter. The velocity data were especially critical, since they were used to compute the aberration of light from the GS due to the orbital motion of the GP-B satellite, the main calibrator of the SQUID signal (see [2], section 4 E, equation (52), and section 6 A). Orbital data were also needed to calculate drift rates predicted by GR for a gyroscope moving on the real GP-B orbit.

However, the position and velocity of the S/C were initially measured relative to the Earth in an Earth-fixed coordinate system. Moreover, these measurements were discrete and usually not available at the time moments necessary for further processing (L2 time points). Several steps were thus required to transform discrete GPS or satellite-laser-ranging (SLR) orbital data into position and velocity information available at arbitrary times in the inertial coordinate system. These steps are:

(i) Transformation from the Earth-centered, Earth-fixed (ECEF) coordinates, in which satellite position and velocity were measured, to Earth-centered, inertially-fixed (ECIF) coordinates; this transformation accounts for the Earth's rotation, precession, and nutation.

(ii) Orbit interpolation within the ECIF coordinate system.

(iii) Transformation from ECIF coordinates to solar-center, inertially-fixed (SCIF) coordinates; this transformation accounts for the motion of the Earth relative to the solar barycenter.

We now describe calculations involved in these three steps.

(i) ECEF \rightarrow ECIF transformation.

The position and velocity are transformed from the ECEF frame to the ECIF frame by means of appropriate rotation matrices computed as functions of the Julian Date (JD; see [16]). The conversion matrices include Earth rotation, precession, and nutation factors; Earth polar motion (e.g., [17]) is not included, as being well beyond the required accuracy.

Denoting the precession, nutation, and rotation matrices by \mathbf{P} , \mathbf{N} , and \mathbf{R} , respectively, we can write the complete transformations as

$$\vec{r}_{\text{ECIF}} = \mathbf{PNR} \vec{r}_{\text{ECEF}} \quad (1)$$

and

$$\vec{v}_{\text{ECIF}} = \frac{d\vec{r}_{\text{ECIF}}}{dt} = \mathbf{PNR} \vec{v}_{\text{ECEF}} + \frac{d\mathbf{PNR}}{dt} \vec{r}_{\text{ECEF}}. \quad (2)$$

Since the rate of change of both the precession and nutation matrices is negligible in comparison to the rate of change of the rotation matrix, the second term can be approximated as

$$\frac{d\mathbf{PNR}}{dt} \vec{r}_{\text{ECEF}} \approx \mathbf{PNR} \mathbf{R}^{-1} \frac{d\mathbf{R}}{dt} \vec{r}_{\text{ECEF}} = \mathbf{PNR} (\vec{\omega}_E \times \vec{r}_{\text{ECEF}}), \quad (3)$$

where $\vec{\omega}_E = [0, 0, \omega_E]$ is the Earth angular velocity. The last equality comes about because the rotation matrix, with sufficient accuracy, is just

$$\mathbf{R} = \mathbf{R}_3(\omega_E t) = \begin{pmatrix} \cos \omega_E t & -\sin \omega_E t & 0 \\ \sin \omega_E t & \cos \omega_E t & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

so

$$\begin{aligned} \mathbf{R}^{-1} \frac{d\mathbf{R}}{dt} \begin{pmatrix} x \\ y \\ z \end{pmatrix} &= \begin{pmatrix} \cos \omega_E t & +\sin \omega_E t & 0 \\ -\sin \omega_E t & \cos \omega_E t & 0 \\ 0 & 0 & 1 \end{pmatrix} \omega_E \begin{pmatrix} -\sin \omega_E t & -\cos \omega_E t & 0 \\ \cos \omega_E t & -\sin \omega_E t & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} \\ &= \omega_E \begin{pmatrix} -y \\ x \\ 0 \end{pmatrix} = \vec{\omega}_E \times \vec{r}. \end{aligned}$$

By the equations (2) and (3), the overall velocity transformation can be written as

$$\vec{v}_{\text{ECIF}} = \mathbf{PNR} (\vec{v}_{\text{ECEF}} + \vec{\omega}_E \times \vec{r}_{\text{ECEF}}). \quad (4)$$

Note that the precise argument of the Earth rotation matrix, as used in our calculations, is the negative of the Greenwich Apparent Sidereal Time (GAST),

$$\mathbf{R} = \mathbf{R}_3(-\text{GAST}). \quad (5)$$

GAST is related in a known way [17] to the Greenwich Mean Sidereal Time, GMST, which, in turn, can be derived from JD by the formula [17–19]:

$$\begin{aligned} \text{GMST} = & (24110.54841 + 1.002737909350795\text{UT1}) \\ & + (8640184.812866 + 5.9006 \times 10^{-11}\text{UT1})T + (0.093104)T^2 - (6.2 \times 10^{-6})T^3. \end{aligned} \quad (6)$$

Here T is the Julian century, $T = d/36525$, and d is the JD relative to the J2000 frame, $d = \lfloor \text{JD} - 2451544.5 \rfloor - 0.5$ ($\text{JD} = 2451545.0$ corresponds to UT1 noon on 1 January 2000); UT1 is the universal Time. It differs from GPS time by an integer number of leap seconds (the difference between GPS and UTC) and an additional correction DUT1 (the difference between UT1 and UTC, see section 3.2.4); the latter does not exceed 0.9 seconds, and can be readily found online (e.g., [20]).

The nutation matrix, N , is found in [17, p 114, 115], and was used exactly as given there.

The precession matrix, P , was taken from [17, p 103, 104]. Its time dependence given in table 3.211.1 was simplified as allowed by the accuracy of the GP-B data analysis: several minor corrections were removed.

(ii) Orbit interpolation.

The position and velocity data for the GP-B S/C orbit, as measured by GPS or SLR, were semi-randomly spaced in time: for instance, the intervals between the GPS data packages were expected to be ~ 10 s, but in practice varied significantly, sometimes approaching 40 s and even more. The goal of this step was thus to produce position and velocity values at regular time intervals, e.g., every 2 seconds. The orbit was obtained by direct numerical integration of the motion equations in the inertial space using the initial data from GPS or SLR solutions.

The resulting algorithm performed numerical integration of the S/C equations of motion in the ECIF coordinate system over every interval between successive GPS or SLR measurements. A fourth-order Runge–Kutta integrator was used to integrate the equation

$$\frac{d^2\vec{r}}{dt^2} = -\nabla\Phi, \quad (7)$$

where \vec{r} is the ECIF position vector and Φ is the gravitational potential of the Earth. The latter includes the term with the Earth quadrupole moment $J_2 = 1.0826267 \times 10^{-3}$ (as in [2], section 2 B):

$$\Phi(r) = -\frac{GM_E}{r} \left[1 - J_2 \frac{3z^2 - 1}{2r^2} \right]; \quad (8)$$

here $r = |\vec{r}|$, G is the gravitational constant, M_E is the mass of the Earth. This integration procedure provided orbital data at scattered moments of time, set by accuracy requirements (the computational error was taken to be not greater than 1 m in position, 10^{-3} m s $^{-1}$ in velocity). To make the data available for any given time instant, a third-order spline approximation was made. The interpolation procedure had been successfully tested by comparing the results of the interpolation with some specially generated orbits.

(iii) ECIF to SCIF transformation.

The transformation from ECIF to SCIF coordinates in the J2000 reference frame was accomplished by adding the motion of the Earth in SCIF coordinates to the motion of the spacecraft in the ECIF coordinate system computed above, namely:

$$\vec{r}_{\text{SCIF}} = \vec{r}_{\text{ECIF}} + \vec{r}_E; \quad (9)$$

$$\vec{v}_{\text{SCIF}} = \vec{v}_{\text{ECIF}} + \vec{v}_E. \quad (10)$$

The position and velocity of the Earth, \vec{r}_E and \vec{v}_E , were obtained through the Earth ephemeris generated by the Jet Propulsion Laboratory [21]. Earth position and velocity vectors were calculated using interpolation coefficients of Chebyshev polynomials provided in the ephemerides files.

This algorithm was implemented in the S/W package (blocks P0, P2–P6 of the diagram in figure 5) that allowed prompt and reliable orbit data handling through the whole mission.

3.2.5. Other signal preprocessing.

Preprocessing of LF SQUID and telescope signals.

The LF SQUID signals from each of the four gyroscopes, and the eight photodiode currents (see section 2.2.2) representing the telescope signal were all preprocessed similarly. After reading the stretch of L1 data that was being preprocessed and the appropriate status flags, the time stamp of every L1 data point was corrected as described above. Next, all the data having at least one status flag ‘invalid’ were removed. The status flags, 3 for each SQUID signal and 2 for each photodiode current, characterized the state of the relevant onboard electronics at the moment of data sampling. This state, of course, had to be right for each any data point to be used in the analysis.

The signal clean-up then removed all other data points that might be corrupt. There were several kinds of bad points: stuck points, statistical outliers, isolated points near the gaps in the signal and points that did not pass the ‘smoothness’ tests, etc. Stuck points are those that are numerically identical to the preceding values of the signal. This situation occurred whenever there was no update of the monitor at the next recording time, and, consequently, the ‘old’ value of the monitor was placed in the TFT. Statistical outliers are defined as those data points that are outside the n -sigma corridor around the signal (usually $n = 3$ was taken). The isolated data points or small clusters of points near significant gaps in the data might not properly represent the signal and are better removed. The same is true for the data points that violate some natural assumptions about the expected smoothness of the signal.

The result of these two preprocessing steps was a time-corrected ‘clean’ signal from which the L2 output data points were computed. A time window of chosen fixed width was moved in steps along the vehicle time axis so that at each step the center of the window coincided with the corresponding output time (an even integer second of the vehicle time). All points of the cleaned signal that occurred inside the window were weighted with appropriate coefficients taken from some digital filter, ‘kaiser’, ‘hamming’, ‘blackman’, etc. Then a weighted LSQ linear fit was performed within the current window, and the output L2 value was obtained from the resulting straight line at the window midpoint.

All these functions were carried out by specially developed MATLAB program packages; preprocessing a single Run ID stretch (~ 12 h of data) typically took just a few minutes of computing time.

The L2 SQUID and telescope signals preprocessed this way proved completely satisfactory during the flight for fast checking on the inertial gyro motion, quality of S/C pointing, etc. These L2 data were also successfully used at the initial stages of the relativistic drift estimation using the 2-floor estimator (section 5.4). However, the scale and complexity of the influence of patch effect torques and polhode damping on the science signal made evident the need to create the 2-second filter [3] and use it to get a satisfactory estimate of the relativistic effects. It also became clear that the quality of the SQUID and pointing signal preparation had to be better than that in the L2 DB. So the main science signals were again

reduced from L1 to what we called L3, in a more sophisticated and accurate way described below in sections 5.2 and 5.3.

Preprocessing of calibration signal.

A sinusoid of known stable frequency and amplitude was added intentionally to the SQUID signal to monitor readout scale factor changes. The commanded values of this sinusoid were telemetered to the ground providing the so called ‘calibration signal’ for the removal of the sinusoid from the SQUID signal prior to the analysis. The calibration signal was preprocessed much like the SQUID and telescope signals, so it included the time stamp correction, the signal clean-up, and the computation of the output L2 values by LSQ fitting in the moving window. The width of the window was an integer multiple of the calibration signal period, and it was fitted, naturally, to a sinusoid of known frequency, rather than a straight line.

There were no special status flags in this case, but there was one step preceding the time latency correction, namely, the frequency determination. This step allowed us to determine unambiguously all continuous time intervals where the calibration signal had the same commanded frequency. The point is that two different calibration signal frequencies, 1/62 Hz and 1/124 Hz were alternatively used. A special monitor showed which frequency was used at any given time (one such monitor was available for SQUIDs 1 and 3, the other for SQUIDs 2 and 4).

Preprocessing of roll phase signal.

The preprocessing algorithm for the roll phase determination included the following main steps: (1) star identification; (2) determination of continuous roll phase for each star tracker; and (3) determination of the roll phase of the quartz block.

Star identification builds a one to one correspondence between the star images on the CCD plane of the star tracker and stars in the star catalog providing thus precise information on the position of stars in the inertial North-South, West-East reference frame. The star identification algorithm was implemented with pairwise pattern matching between the star tracker telemetry data and the star catalog data using high occupancy bins of the histogram of roll angle separation, angular distance to the guide star, and star magnitude. Preprocessing results show that typically 10 to 25 stars per roll period were identified by each star tracker.

When a detected star image is identified, the roll phase of the star tracker at the time of detection can be derived from the position of the star in the inertial North-South, West-East reference frame, the alignment matrix of the star tracker, and the centroid position of the star image on the CCD plane of the star tracker. Since the rate of the telemetry data from the star tracker was highly variable, a set of roll phase data for each star tracker at the time of detection was fitted to a short term mode. The model included a second order polynomial and up to three harmonics of roll frequency, and the continuous roll phase of each star tracker was calculated using the fit model at the 2 s interval. The data for the fit were selected over several roll periods, depending on the frequency of detections. The RMS error of the fit was about 7 as, consistent with the instrument centroid accuracy. Mechanical distortions of the graphite ring led to a long term drift of the star trackers in roll referenced to the quartz block. The structure of the graphite ring constrained the differential and common modes of the drift. The roll phase difference between the two star trackers obtained from the differential mode of the drift and observed in preprocessing showed a high correlation with thermal gradients derived from the telemetry data of the graphite ring and Dewar temperatures. In order to estimate the long term drift of the star trackers, mechanical distortions of the graphite ring were modeled as linear functions of thermal gradients, and the parameters of the model were found by fitting this difference to the thermal gradients. The roll phase of the quartz block was determined from the roll phase of two star trackers whose long term drifts were estimated using the model

of mechanical distortions and corrected during preprocessing. Data analysis shows that the peak-to-peak value of the roll phase difference was about 40 as in the GP-B flight test, and the accuracy of the roll phase of the quartz block determined in the ground preprocessing was better than 10 as (50 μ rad).

4. Data grading

Selection of ‘good’, valid data is always done in physical experiments. Preliminary clean up of GP-B data (standard removal of statistical outliers and isolated data points, etc) was implemented during preprocessing described above. Moreover, many GP-B science signals were accompanied by the so called ‘flag’ monitors indicating the status of the device in use at the moment of data taking; so data points with inappropriate combinations of status flags were also eliminated.

Even after this, grading of GP-B science data was needed for marking those sections that were somehow corrupt or contained an unacceptably noisy signal. There were several causes for these disturbances that we discuss below; mostly, they were specific for GP-B. Once marked, the data were either excised completely or de-weighted, as appropriate.

De-weighting the data, widely used in signal processing, was implemented by assigning noise uncertainty to each data point. For example, if all ‘good’ data had an uncertainty of 1.0, then de-weighted data might have uncertainties of 2, 4 or 100 depending on how much they are to be suppressed. For weighted least squares fitting algorithms, the actual weights applied to the data were the square of the uncertainty, i.e. a de-weighted data point with a noise level assigned to 2.0, would have 1/4 of the impact that the same point has without de-weighting. A noise level of 10 virtually wiped out the effect of that data, as it was reduced by two orders of magnitude.

The rationale for using de-weighting instead of throwing out the data is somewhat empirical. Many signal processing algorithms work best with contiguous data. Brute force fit methods of throwing out data can impart spikes in the results when data gaps are included in the fits, e.g., frequency domain analyses. Using de-weighting instead provides contiguous data with more uniform time intervals between data points, thereby mitigating data-gap-induced spikes while still reducing the impact of noisy points that might otherwise have been discarded.

The first class of events that significantly affected the quality of the data were the so called S/C ‘anomalies’ that occurred due to either computer malfunctions disrupting data collection, or attitude control problems leading to poor S/C pointing or a total loss of it. During each of these anomalies hours or even days of data were lost, so, as described in section 2.1, the data were split into segments. This lack of continuity had serious consequences for the data analysis, which originally was planned to handle continuous data.

The second class of grades was due to telescope disturbances. Aside from occasional poor pointing caused by problems with the ATC, one recurring source of pointing errors was the South Atlantic Anomaly (SAA). The GP-B satellite was in a low Earth orbit almost exactly in the North/South plane, so periodically (3 or 4 consecutive orbits once in 12 h) it passed through the SAA. During these passages the telescope was bombarded with energetic protons that registered as large pointing offsets. An onboard algorithm was able to eliminate some of these outliers, nevertheless, the telescope was effectively blinded, and the S/C axis routinely went astray from the GS, with the telescope exceeding its 400 mas linear range approximately 2% of the time.

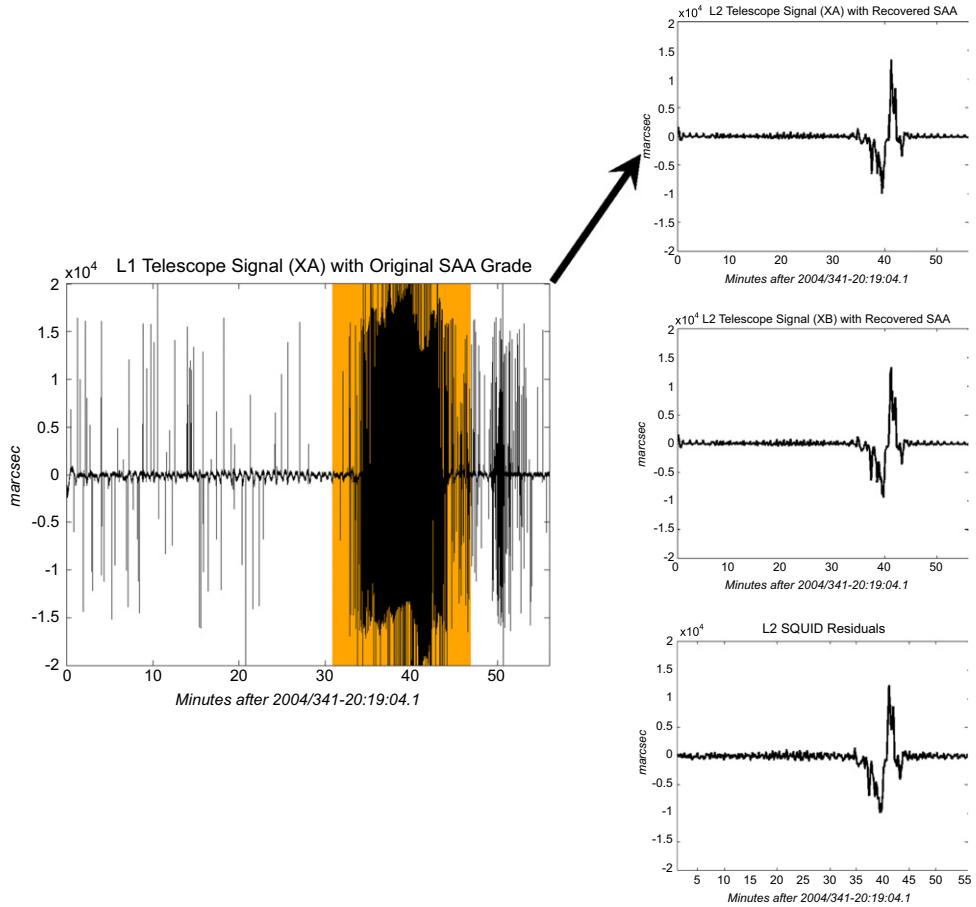


Figure 8. SAA data recovery: raw telescope signal with SAA region highlighted. On the right, the recovered telescope signals showing clear correspondence to SQUID residuals.

Recovery of the proton-induced outliers was successfully done in ground processing using an algorithm based on k-means clustering, a method often used in data analysis for grouping points that share some characteristic [22]. In this particular case, we took advantage of the fact that the outliers almost always registered as large, rapid jumps, which implied S/C motion well beyond its actual capability. This allowed us to restore the true telescope signal for use in the science data analysis, as illustrated by figure 8. Periods when the telescope exceeded its linear range were de-weighted in the analysis.

As the S/C passed behind the Earth during Guide Star Invalid (GSI), the GS position had to be estimated using onboard attitude control gyroscopes. It was vital to do this well, since the quality of the estimation determined the difficulty of acquiring the GS once it became visible again as the S/C emerged from the GSI zone. Unfortunately, the platform that supported the attitude control gyroscopes was not mechanically isolated from the effects of thermal expansion. Thus when the satellite was exposed to the Sun, the platform would bend, which led to inaccurate readings of the control gyroscopes. This resulted in the need for two

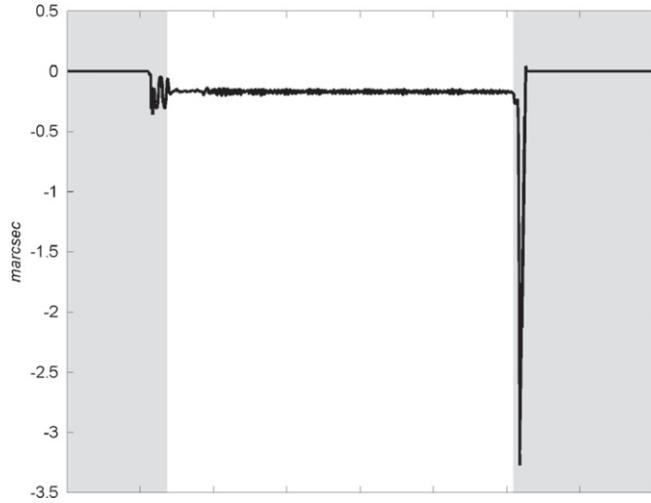


Figure 9. GSV (grey) and GSI (white) grade for one orbit of pointing. Strong pointing deviations are seen when the S/C transitions from one to the other (GS acquisition).

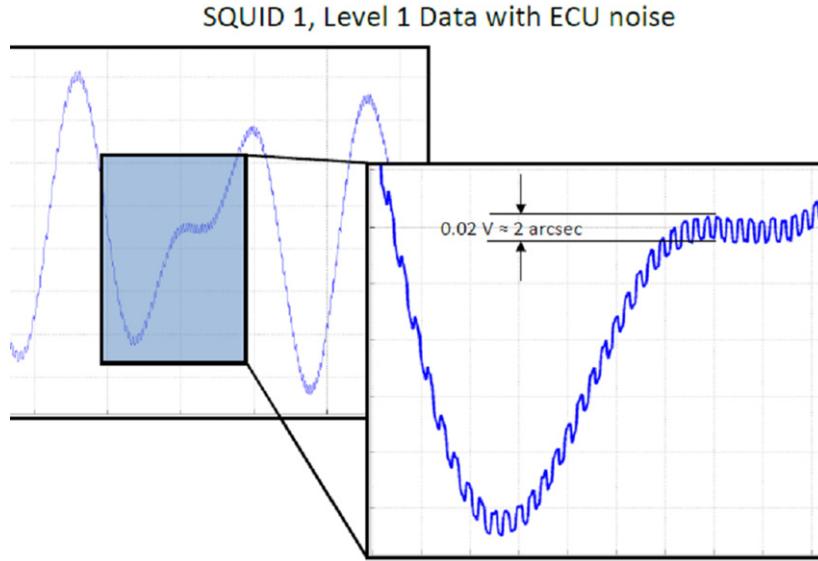


Figure 10. ECU noise in SQUID 1 signal.

more grades, one which flagged data when the S/C was not in eclipse, and another for the GS acquisition and loss periods (see figure 9).

The last class of grades covered disturbances in the SQUID signal itself. The first and foremost of them was the experiment control unit (ECU) noise. The ECU was an electronic instrument added to the spacecraft as its system health monitor. After about half of the science data collection period had elapsed, it was discovered that the ECU was unintentionally injecting a large (~ 2 as) signal into the SQUID data from science gyroscopes 1 and 3 (but *not* 2 and 4). At the smallest time scale, the ECU noise was observable as a square wave with an

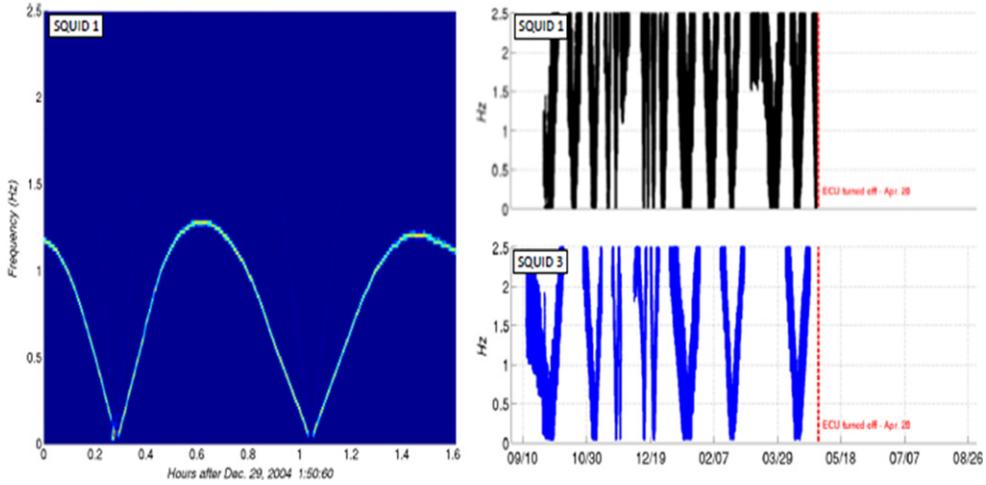


Figure 11. ECU noise spectrogram and the fundamental harmonics induced by this noise in SQUID 1, 3 signals.

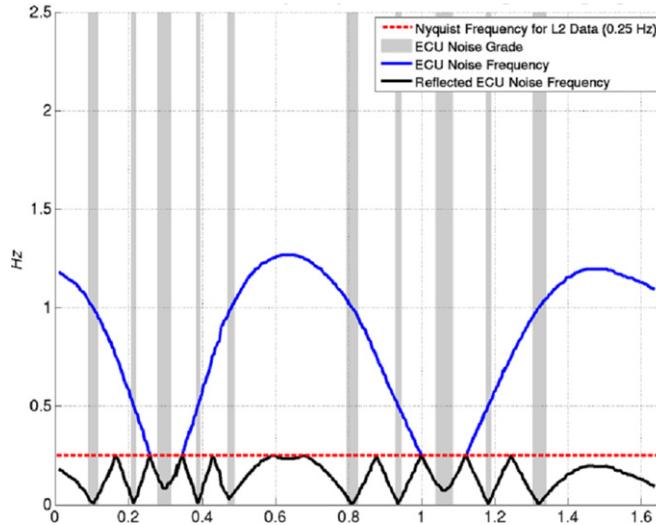


Figure 12. ECU noise frequency and graded regions.

approximately constant magnitude superimposed on the science signal (figure 10). By isolating the ECU noise and examining its spectrum, we found that the noise frequency varied sinusoidally at a short (~ 1 orbit) time scale, but rather irregularly at a long scale, on the order of weeks, as shown in figure 11.

The ECU noise was only of real concern when its frequency approached the roll frequency of ~ 12.9 mHz. Outside of these areas, frequency separation ensured that the ECU noise had a minimal effect. However, one side effect in converting raw L1 data ($10 \text{ samples s}^{-1}$ for SQUID data) to time-aligned L2 data ($0.5 \text{ samples s}^{-1}$) was that aliasing of ECU noise frequency made it effectively ‘reflected’ about the Nyquist frequency (see figure 12), resulting in additional ‘noisy’ periods in the L2 data. Attempts were made to filter

both the roll frequency and aliased ECU noise, but ultimately abandoned in favor of simpler de-weighting of the affected data. Only $\sim 18\%$ of SQUID data, out of $\sim 70\%$ total affected by the ECU noise, were de-weighted.

Other short periods when science data underwent de-weighting were: (a) the roll notch filter in the ATC system was turned off, (b) the drag-free mode was off, (c) roll rate disturbances were present, and (d) SQUID bias jumps occurred.

5. SQUID and telescope signal preparation: getting Level 3 data

5.1. The need for additional data preparation

We obtained the results of the GP-B experiment, i.e., the estimates of the relativistic drift rates, using the elaborate and sophisticated ‘2-sec filter’ described in detail in our paper [3]. The success of this approach depended strongly on careful and meticulous preparation of all major input science signals, i.e., LF SQUID signal, telescope signal, roll phase signal, etc. The goal was a serious improvement in the quality of the science data relative to the quality of the L2 data set delivered by preprocessing. The preprocessing algorithms were implemented in the code prepared before the flight; like all GP-B data analysis tools, they evolved dramatically after actual flight data were examined.

The original preprocessing algorithms were traditional and rather simple. Reduction of the science data from Level 1 (disentangled and time ordered telemetry data) to Level 2 (basic science data to be analyzed) implied synchronization of the SQUID, telescope, roll phase and other needed signals, as well as their compression to the equally spaced data points separated by 2-sec intervals. This was done by the application of the standard signal processing technology of digital filtering. Data grading, explained in the previous section, still consisted of removing or de-weighting data collected during SAA crossings, those affected by ECU noise, and the data close to the start and end of the Guide Star Valid part of each orbit. Only the GSV science data were planned to be used in the final analysis, which was to exploit a relatively simple parameter estimator, and to observe almost linear motion of the gyro spin axis in the NS—WE plane for all four gyroscopes. The slopes of these almost straight lines were supposed to present us with the values of the relativistic drift rates.

The reality, however, could hardly be more different. After it was discovered that all the gyroscopes were affected by the patch effect misalignment torques acting continuously, it became clear that data from the whole orbits, both GSV and GSI parts, had to be analyzed. It also became evident that many already prepared data analysis algorithms, including some preprocessing algorithms and the science estimator, had to be redesigned. As explained in detail below, there were two especially important reasons for redesigning preprocessing. Because of the way it was initially carried out, (a) the noise in SQUID and telescope L2 data proved to be strongly correlated, and (b) the time stamps of the L1 data points in every preprocessing window outputting an L2 data point were not spaced exactly evenly due to the time correction (section 3.2.3). Along with some other features, these made the quality of the L2 data insufficient for the high precision final analysis.

5.2. Preparation of LF SQUID and pointing signals

In the new approach, the relatively simple data reduction from L1 to L2 was replaced with a ‘ladder’ of interconnected intermediate data levels: from L1 to L1.5 to L2 to L3 (details follow). The process of going from one level to the other became adaptive and iterative. Indeed, quite often we had to go down to L2 with the data from some orbit, find

inconsistencies there, and then go back to L1, look, and eventually find, some ‘dirt’ in the SQUID and/or telescope signals that had not been appropriately handled in the preceding processing, and then repeat the data reduction, to remove those irregularities. This iterative process usually had to be performed a dozen times or more for both the LF SQUID and telescope signals.

5.2.1. L1 to L1.5 preprocessing. The L1 to L2 data reduction was split into two steps by creating an intermediate Level 1.5. The idea was to accomplish the following intermediate tasks: (1) adjust time stamps; (2) remove spikes (outliers) from the data; (3) interpolate the missing data; (4) check the telescope data feasibility; (5) estimate data noise levels.

Adjusting time stamps was a major reason for processing L1 to L1.5. The time correction technique was described in section 3.2.3 and in paper [15]. After time correction, the time tags of L1 data points were no longer separated by the same constant interval, so the data rate became variable. Most importantly, the LF SQUID signal presented a similar picture: its data rate was 5 Hz only approximately.

The next step in going to L1.5 was removing spikes from the data. We developed several spike removal procedures; some are described in section 3.2.1. Spikes (outliers) in both the LF SQUID and telescope signals were detected after removing principal spectral components—roll, twice roll, four times roll, roll \pm orbital, roll \pm twice orbital, dither, and calibration signal frequencies—from these signals. The residuals, representing the noise, were analyzed with the help of moving windows to determine the standard deviation. Then the data points outside the 3-sigma corridor, in most cases clearly visible, were declared ‘spikes’ and removed. The corresponding time tags of the removed outliers were declared ‘missing data point times’.

After the spike removal the data were interpolated using a small (shorter than a second) moving window, thereby filling in the missing data points. The data were interpolated assuming a consistent frequency content over the window. L1.5 data resulting from these operations retained the original L1 data rate, however, slightly modified by the time stamp adjustment; thus the L1.5 data rate was not constant.

The telescope channels required somewhat special processing because of the interpretation and use of these signals. Converting telescope data to a pointing signal required two telescope channels (the \pm channels, see section 2.2.2 or [2], section 6.2). Therefore we examined the validity of data pairs. If one of the data pair elements was missing, the pair was omitted and forwarded to the interpolation process. Further, the telescope signal was compared with the corresponding SQUID signal, which allowed us to find and remove more irregularities from the telescope data.

We used one more criterion for the validity of the telescope/pointing data, namely, the sign criterion. The normalized pointing signal was computed by the formula (79) [2] as the ratio of the (weighted) difference of the \pm photodiode currents over the sum of the same currents, namely

$$z_T = \frac{i^+ - w i^-}{i^+ + w i^-}. \quad (11)$$

(the subscripts indicating the TRE box and the pointing direction used in section 2.2.2 are dropped here). It was important to constantly check the sign of the above denominator: according to the TRE design, it had to be the same throughout the mission. But in the real data this denominator sometimes changed its sign, indicating some disturbance had made the pointing signal meaningless. So if the denominator sign changed, we first rejected the data pair, and then tried to fill it in by interpolation.

Another important task was to estimate the noise levels of the L1.5 SQUID signals. This estimation was done in the frequency domain by removing the roll and roll \pm orbital frequencies from the signal in small windows of 1.5 – 2 s width. Instead of using the conventional FFT, which requires evenly separated data points, we employed specially designed frequency domain algorithms that did not assume a constant data rate. They used a harmonic representation of all frequencies of interest, such as, e.g., the roll and roll \pm orbital frequencies, in order to accommodate varying data rates and data gaps. After subtracting the main harmonics from the signal, the statistics of the resulting time domain residuals were recorded for each channel. These recorded noise levels were then smoothed further. The resulting SQUID noise level data was a slowly varying signal, not noisy itself.

5.2.2. L1.5 to L2 processing: the Z signals. Processing from L1 to L2 involved: (1) converting all channels of various data rates to the synchronized signals with data points at 2-second intervals (0.5 Hz rate, see also section 3.2.1), (2) assigning data grades to all the channels, and (3) converting L1.5 noise estimates into new L2 estimates by means of the state estimation technique.

Converting the SQUID signals to data points 2 s apart was done in the frequency domain with full information retention, so that the roll component magnitude in L2 remained consistent with that at L1.5. This was critically important for the relativity estimates; it was discovered that the original pre-flight preprocessing algorithms did not preserve the magnitude of the roll component.

The new algorithm for converting L1.5 to L2 was specifically designed to ensure that the roll frequency amplitude was precisely the same in L2 as it was initially in L1.5, and to obtain uncorrelated L2 noise. In this algorithm a set of 2 s width windows was formed with their centers (midpoints) on the centers of the chosen 2 s intervals. A simple estimator was designed to find the amplitude and phase of the roll frequency component in each window using all the signal points inside it. From these estimates, the value at the window midpoint was entered into L2 database.

One more important part of the L1.5 to L2 data reduction was determining the L2 data noise level. In this procedure the L1.5 noise was adjusted by using noise uncertainties obtained from the state estimator. For instance, there were states (estimated parameters) known to be constant over an orbit, so these states were estimated using orbit-wide windows. The statistics of the variations in these state estimates was compared to the Cramer–Rao Lower Bounds (CRLB; see e.g. [23]) from the L1.5 noise levels. The new noise level could be computed from the ratio of the CRLB uncertainty to the collected statistics. Using the resulting L2 noise estimates, the mathematical uncertainty for the states matched the CRLB. These per-orbit ratios were appropriately smoothed over the entire mission, and the resulting smoothed ratio values were applied to the L1.5 recorded noise data to generate the final L2 SQUID signal noise estimate (sigma-noise) for every data point.

There were also numerous data consistency checks for the L2 SQUID signal reduction; other checks compared the telescope and SQUID data. These two signals were supposed to conform consistently. Thus a telescope data point was rejected if it was identified as an outlier relative to the SQUID data (the method of removal and the 3-sigma criterion for a data point to be removed are explained in the previous section). Basically, the difference between the telescope pointing and the SQUID pointing was filtered, and spikes were thus detected. More outliers in the telescope data were removed using the analysis of this difference. An additional difficulty, of course, was that the filter innovations required the SQUID and telescope signal values be synchronous, therefore the deleted telescope data were interpolated to fill in the

gaps in place of the deleted spikes. All in all, in the resulting signals stored in L2 DB, it was guaranteed that any triplet of telescope/SQUID data was really valid.

The output of this rather complicated processing was the four SQUID signals, one for each gyroscope, and two pointing signals, from A and B TRE boxes, that we called the Z-signals, for brevity. All the data points were uncorrelated with each other, and each data point had its own recorded uncertainty (sigma-noise); those uncertainties were used also as the data grading quality parameters.

The structure of the Z-signals as the sets of uncorrelated data points with the known noise level allowed us to effectively use them as the main input to the 2-sec filter [3].

5.3. L2 to L3 processing: 2-floor estimator and gyro orientation profiles

The 2-floor estimator [24] was originally designed as the main analysis tool aimed at obtaining the final relativity estimates. It was built, of course, around the general model of the LF SQUID signal. The model, rather complicated by patch effect torques and polhode damping, was derived in DA I, namely, in sections 4–6 and appendices A–C. According to this derivation, the measurement equation (52) from [2] is

$$z(t) = C_g(t) \times \left[(\tau_{ns} - s_{ns}) \cos(\phi_r + \delta_r) + (\tau_{we} - s_{we}) \sin(\phi_r + \delta_r) \right] + \text{bias} + \text{noise.} \quad (12)$$

(The term $C_g(t)\alpha$ is now included in the bias.) Here $z(t)$ is the LF SQUID signal, $s_{ns}(t)$ and $s_{we}(t)$ are the North–South and West–East projections of the gyro spin vector (gyro orientation), $\tau_{ns}(t)$ and $\tau_{we}(t)$ are the corresponding components of the pointing vector (S/C orientation; precise definitions of these quantities are found in [2], sections 1.2 and 1.4); $C_g(t)$ is the time-varying scale factor, $\phi_r(t)$ and $\delta_r(t)$ are the roll phase and the roll phase offset, respectively.

The models for the scale factor ([2], section 4 and appendix A), the gyro motion ([2], section 5 and appendix C), and the roll phase offset include a set of unknown parameters, such as torque coefficients and the like. The S/C pointing $\tau_{ns}(t)$, $\tau_{we}(t)$ was calculated from GP-B orbit data, JPL ephemeris, telescope data (for GSV only), and SQUID data (for GSI only) by the formulas of section 6 and appendix B from [2].

The multiplicative structure of the measurement equation (12) makes it essentially nonlinear. The 2-sec filter [3] that provided the final estimates of relativistic drift rates [1] used the same measurement model (12), so it was intrinsically nonlinear as well (this was the core reason why L1.5 to L2 preprocessing was needed). Whatever the estimation method, this nonlinearity has to be specially dealt with.

An iterative multidimensional LSQ estimator, the 2-sec filter, was chosen to complete the analysis of GP-B data, particularly for the relativity estimates. The main problem posed by nonlinearity is recognized in the estimation theory and practice as the possibility for the iterations to converge to a wrong, local, minimum of the quadratic cost function, instead of the global minimum sought, giving a wrong set of parameter estimates as the answer. The known remedy for this is to start the estimation process with an initial state vector value not too far from the one giving the global minimum. The 2-floor estimator, a specifically designed Kalman filter, provided proper initial conditions (estimates and errors of all the needed parameters, particularly, torque coefficients) for the nonlinear 2-sec filter.

The 2-floor estimator got its name from its algorithmic structure, where the analysis of data was carried out in two sequential steps we called ‘floors’. Its 1st floor was a modification of the pre-launch scheme [24] that used batch processing. At the 2nd floor, the more accurate

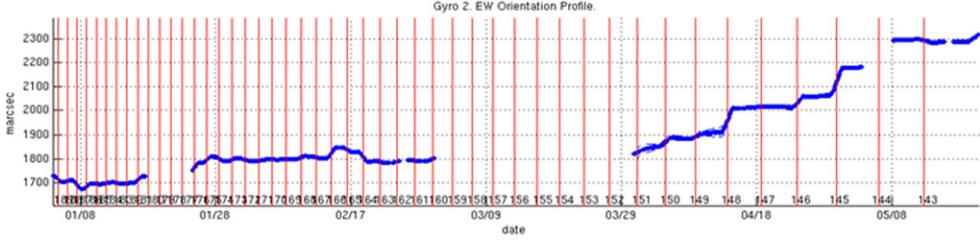


Figure 13. L3 orientation profile ('slope') in WE direction for gyroscope 2 (data segments 5, 6, and 9).

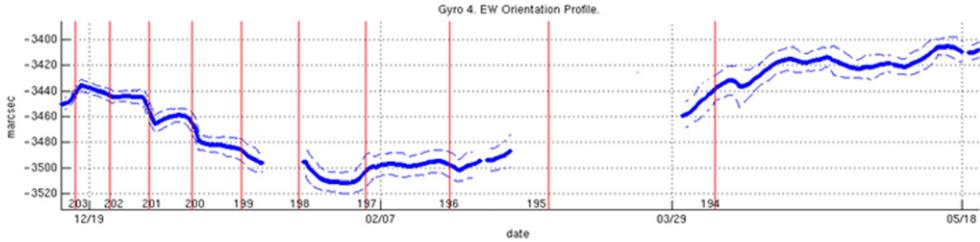


Figure 14. L3 orientation profile ('slope') in WE direction for gyroscope 4 (data segments 5, 6, and 9).

modeling of the scale factor and misalignment torque coefficient took place, integrating information from all the batches. It was natural and convenient to choose the batch length equal to one orbit (~ 97 min).

At the 1st floor all Level 2 SQUID data points (2 s rate) belonging to the GSV part of each orbit were collected in one batch. Then, to address the $1/f$ nature of the SQUID noise (to 'flatten' the noise), a band-pass filter was applied to the data. The frequency band of this filter was from a half to one and a half times the roll frequency.

In the signal model we assumed that within one GSV: (a) the scale factor, the roll phase offset, and the bias remained constant, and (b) gyro orientation described by $s_{ns}(t)$ and $s_{we}(t)$ changed according to equation (71) from [2] with a constant misalignment torque coefficient $k = \text{const}$. These assumptions are quite reasonable on the grounds of low change rates of the pertinent dynamic variables. No resonance torques were included.

The S/C pointing, $\tau_{ns}(t)$ and $\tau_{we}(t)$, during GSV part of each orbit was computed by the formula (72) of [2] as the sum of the orbital and annual aberrations (equation (74)), the guide star parallax (equation (75)) and its light bending by the Sun (equation (76)), and the inertial pointing error. The last was calculated from the telescope signal using formula (80) with the known telescope scale factor and nonlinearity coefficient (all the referenced equations are found in section 6 of [2]). For GSI periods, pointing was determined from the four SQUID signals as explained in [2], section 6.1; particularly, formula (77) was used.

The state vector of this estimator was observable: estimates of C_g and δ_r were computed based on the orbital aberration signal at roll \pm orbital frequency, and then, at the 2nd floor, the parameters of the gyro motion in the NS and WE directions were estimated. Finally, the gyro position (orientation), i.e., s_{ns} and s_{we} , was computed from those estimates at the midpoint of the GSV time interval. These orientations were treated as 'one-point-per-orbit' output for each

orbit. The set of all these one-per-orbit estimates represented the gyro orientation profile, which was stored as the L3 signal for all 4 gyroscopes.

The WE orientation profiles (with one-sigma errors) for gyroscopes 2 and 4 are plotted in figures 13 and 14, respectively. Red vertical lines show where patch effect resonance torques, and therefore ‘jumps’ in the orientation, occur; gyroscope 2 has many more roll-polhode resonances than gyroscope 4 because of the much larger variation in its polhode period (see [1, 25], and [2], appendix C, sections 3.2 and 5). On the other hand, the error corridor is hardly discernible in the gyro 2 plot.

All gyro orientation profiles were carefully examined for outliers, both visually and with the proper S/W tools. Whenever an outlier was observed, the corresponding orbit was flagged out, and we re-examined the L2 data for this particular orbit. The orbit was analyzed carefully once again, and in most cases the reason for the outlier was determined. The LF SQUID and/or telescope data were then re-preprocessed from Level 1.5 to Level 2. Thus the process of generating the gyro orientation profiles became iterative, which allowed us to significantly improve the overall quality of the L3 data.

It is worthwhile to note in conclusion that, in addition to providing the proper initial conditions for the final analysis with the help of the 2-sec filter, the gyro orientation profiles played an important informative role. They gave a clear picture of gyro motion in both NS and WE directions, alerting us to all serious, and often meaningful, irregularities in it. The NS orientation profiles allowed us to see the geodetic drift of the GP-B gyroscopes throughout the mission, and even to make a preliminary estimate of its rate.

5.4. Preparation of HF SQUID signals for TFM

Before the estimation process for the rotor polhode motion and time-varying scale factor, known as TFM, could begin, we had to extract harmonics of the fundamental frequency, $\omega_f/2\pi$ (see below), from the snapshot data. We analyzed each snapshot in a multistage process.

The analysis started with a fast Fourier transform of the snapshot data. This was used to make an initial estimate of the fundamental frequency (essentially, gyro spin just slightly corrected, see below) from the position of the largest peak in the spectrum. With this initial determination of the spin frequency we fit the following model to the amplitudes A_n , B_n of each spin harmonics n :

$$\begin{aligned} z^i(t) = & z_0^i + C_1^i \cos \omega_{\text{cal}}(t - t_i) + D_1^i \sin \omega_{\text{cal}}(t - t_i) \\ & + \sum_{n=1}^N A_n^i \cos n\omega_f^i(t - t_i) + B_n^i \sin n\omega_f^i(t - t_i). \end{aligned} \quad (13)$$

In this equation $\omega_{\text{cal}}/2\pi$ is the known calibration frequency that was intentionally injected into the SQUID signal to monitor readout scale factor changes, and t_i is the known reference time at the midpoint of snapshot i . Parameters z_0^i , C_1^i , D_1^i , A_n^i , B_n^i , and ω_f^i were to be estimated. Notice that all of them enter the expression linearly, except for ω_f^i , which makes the overall fit nonlinear.

For each snapshot, parameters z_0 , C_1 , D_1 , A_n , and B_n , with n up to $N = 51$, were estimated using a standard LSQ fit; ω_f was treated as known with its value taken from the FFT. Harmonics with statistical error larger than 50 % of their amplitude did not contribute significantly to the result and were dropped, and the estimate repeated for the remainder. Accurately determining ω_f^i required the special fitting strategy described next.

After this preliminary fit, about 10 mV rms of so-called ‘switching’ noise remained, from the original ~ 3 V signals. This irregular waveform, related to the operation of the data system, repeated in each snapshot with only slow variations over periods of weeks. Neither spin nor calibration signals could be precisely measured with this noise present, because of its high harmonic content: the switching noise had harmonics near harmonics of all the spin frequencies, interfering with both amplitude and frequency determination. Nonlinear fits including the spin frequency could not be expected to converge. We averaged the switching noise over each week and subtracted the result from every snapshot before the next and final stage of the snapshot analysis.

The first stages of analysis ensured that iterations at the last stage converged reliably to a meaningful result. The preliminary fits for z_0 , C_1 , D_1 , A_n , B_n were used as initial values to the full model, now including ω_f as one of the parameters to be estimated. We used the NonlinearRegress function in Mathematica 5.0 to implement this nonlinear fit, which resulted in a much better estimate of ω_f and other parameters. Spin frequencies could be determined to $\sim 1 - 5$ μ Hz, while the size of the amplitude residuals dropped to 2 – 10 parts in 10000, depending on the gyroscope and time. The correlation coefficient between the residuals and spin signal was typically less than 0.001; any remaining signals were small, noisy and too poorly characterized to remove. About 99.99 % of the roughly one million snapshots were automatically processed in this manner.

The model (13) was fit to every snapshot for all four gyroscopes for n up to $N = 51$. These fits provided two pieces of information needed to study the rotor dynamics and trapped magnetic flux. The first piece was the complex harmonics of the fundamental spin plus or minus roll frequency (depending on the direction of the gyro spin) at time t_i , $\tilde{H}_n(t_i)$,

$$\tilde{H}_n(t_i) = \frac{1}{2}A_n^i - i \frac{1}{2}B_n^i. \quad (14)$$

We use the tilde to denote that these harmonics are not exactly the harmonics $H_n(t)$ that are needed to perform TFM. The latter must evolve from a single reference time origin t_{ref} , while the measured harmonics, $\tilde{H}_n(t_i)$, each have their own independent time origin t_i based on the time when the snapshot was taken. The difference between the two time references results in a phase shift that is taken into account by writing

$$H_n(t_i) = \tilde{H}_n(t_i) e^{-in(\phi_s(t_i-t_{\text{ref}}) + \phi_r(t_i-t_{\text{ref}}))}. \quad (15)$$

Equation (15) simply connects the harmonics extracted from snapshots $i = 1, 2, \dots, N$ by accounting for the elapsed phase at times t_1, t_2, \dots, t_N relative to a specified time origin, t_{ref} . The elapsed phase is the sum of the rotor’s spin phase $\phi_s = \int_{t_{\text{ref}}}^{t_i} \omega_s dt$ and the satellite’s roll phase $\phi_r = \int_{t_{\text{ref}}}^{t_i} \omega_r dt$. For gyroscopes 1 and 3, $\phi_s > 0$, and for gyroscopes 2 and 4, $\phi_s < 0$.

The second useful piece of information is the estimate of ω_f . This angular rate is actually $\omega_f = \omega_s + \omega_r - \Omega_p$, where Ω_p is the polhode rate. This is due to the fact that the dipole component of the magnetic potential shifts the measured rate by Ω_p . The integral of $\omega_f + \Omega_p$ is equal to $\phi_{s0} + \phi_r$, where ϕ_{s0} is the spin phase of the corresponding symmetric gyroscope. Recall that the observed spin phase of an asymmetric rotor ($I_1 \leq I_2 \leq I_3$) consists of a secular component corresponding to the spin rate ω_s and a modulation at twice the polhode rate Ω_p . Since ϕ_r is known to high accuracy from the spacecraft’s star sensors, ω_f provides a good starting point for estimation of the symmetric spin phase ϕ_{s0} , needed to initialize the TFM estimation routine [1, 2, 9].

6. Discussion and conclusions

Successful, i.e., reliable and prompt, science data handling during the GP-B flight came from tedious pre-flight work on the overall data processing scheme, proper algorithms, and their program implementation.

This work included, in particular, a number of end-to-end tests of the preprocessing and analysis S/W, although the software was rather simple as compared to the tools that had to be developed later (see sections 4, 5.1–5.3 for details). During these tests one team of researchers generated a stretch of L1 data, putting in specific values of some parameters ('realistic' or somewhat 'enhanced'), and preprocessing the generated data down to L2. Noise of various kinds which had been anticipated, and in many cases detected experimentally, was always added to these L1 data, which were thus far from ideal.

Subsequently another team, not knowing the parameter values, analyzed the data to estimate the parameters and compare them with those used in Level 1. These tests were carried out with the most important signals, such as the SQUID and telescope signals, timing and orbit data, as well as with signals needed to control the state and health of the instrument, such as various temperature signals, magnetometer signals, and more. The tests not only allowed us to find and remove various S/W bugs, and make other improvements in the code; they also seriously built up our confidence in the process of data reduction and analysis, demonstrating the possibility to really estimate relativistic drift rates.

A very important role in the effective data processing was played by our RunID and Program Management system (section 3.2.2). Much attention was also paid to enhancing program performance speed, predominantly when writing to databases and reading from them. The standard reading/writing tools provided with the databases were much too slow for our purposes, so a number of appropriate special programs were developed, and then modified to achieve an even higher speed. The final programs were tens of times faster than the standard ones.

It is true that, as described in section 5, several effects were discovered that required us to re-preprocess the SQUID and telescope signals, and create new tools for the analysis of the re-preprocessed L3 data (still using, by the way, much originally preprocessed L2 data, such as the orbit-related quantities, or the roll phase data). Moreover, a very complex HF SQUID signal analysis (TFM) had to be developed and carried out, as explained in [2, 9], here in section 5.4, and in [3], section 2. Nevertheless, the knowledge and skills gained during pre-flight preparations for data treatment were very valuable for this post-flight work.

Finally, it is worthwhile to note that the success of new filters for both LF and HF SQUID signal analysis was not only due to the advanced estimation technologies, but first of all due to the physics-based modeling [2] that laid the ground for them.

All in all, we hope that the description of GP-B data reduction and analysis given in this paper and in [3] may be helpful for other space and on-ground experiments that require sophisticated handling and analyzing of large arrays of data which vary in their nature and structure.

Acknowledgments

This work was partly supported by the King Abdulaziz City for Science and Technology gift to Stanford University. The authors are grateful to all people who contributed to the GP-B data reduction, directly or indirectly.

References

- [1] Everitt C W F *et al* 2011 Gravity Probe B: final results of a space experiment to test general relativity *Phys. Rev. Lett.* **106** 221101
- [2] Silbgerleit A S *et al* 2015 Gravity Probe B data analysis: I. Coordinate frames and analysis models *Class. Quantum Grav.* **32** 224018
- [3] Conklin J W *et al* 2015 Gravity Probe B data analysis: III. Estimation tools and analysis results *Class. Quantum Grav.* **32** 224020
- [4] Schiff L I 1960 Possible new experimental test of general relativity theory *Phys. Rev. Lett.* **4** 215–7
- [5] Everitt C W F *et al* 2009 The Gravity Probe B data analysis: status and potential for improved accuracy of scientific results *Space Sci. Rev.* **148** 53–70
- [6] Keiser G M *et al* 2009 Gravity Probe B *La Riv. Nuovo Cimento* **32** 555–89
- [7] Everitt C W F *et al* 2015 The Gravity Probe B test of general relativity *Class. Quantum Grav.* **32** 224001
- [8] Heifetz M *et al* 2009 Gravity Probe B data analysis *Space Sci. Rev.* **148** 411–28
- [9] Silbgerleit A S *et al* 2009 Polhode motion, trapped flux, and the GP-B science data analysis *Space Sci. Rev.* **148** 397–409
- [10] Salomon M *et al* 2011 Nanohertz spin frequency determination for the Gravity Probe B gyroscopes *Rev. Sci. Instrum.* **82** 125110
- [11] Wang S *et al* 2003 *Proc. SPIE Cryog. Opt. Syst. Instrum.* **5172** 108
- [12] Wang S *et al* 2006 *Proc. 24th Int. Conf. on Low Temperature Physics* p 1621
- [13] Keiser G M *et al* 2009 Gravity Probe B, atom optics and space physics *Proc. Int. School of Physics 'Enrico Fermi'* **168** p 482 IOS Press
- [14] Wang S *et al* 2015 The design and performance of the Gravity Probe B telescope *Class. Quantum Grav.* **32** 224008
- [15] Li J, Keiser G M, Lockhart J M, Ohshima Y and Shestople P 2015 Timing system design and tests for the Gravity Probe B relativity mission *Class. Quantum Grav.* **32** 224014
- [16] *Resolution B1 on the use of Julian Dates by The XXIIIrd Int. Astronomical Union General Assembly*, http://www.iers.org/nm_10910/IERS/EN/Science/Recommendations/resolutionB1.html
- [17] Seidelmann P K (ed) 1992 *Explanatory Supplement to the Astronomical Almanac* (Mill Valley, CA: University Science Books)
- [18] McCarthy A and Dennis D 1996 *IERS Conventions, IERS Technical Note* **21**
- [19] Aoki S *et al* 1982 The new definition of universal time *Astron. Astrophys.* **105** 359–61
- [20] <http://tf.nist.gov/pubs/bulletin/leapsecond.htm>
- [21] <http://iau-comm4.jpl.nasa.gov/>
- [22] http://en.wikipedia.org/wiki/K-means_clustering
- [23] Kay S M 1993 *Fundamentals of Statistical Signal Processing: Estimation Theory* (Englewood Cliffs, NJ: Prentice Hall)
- [24] Heifetz M I, Bencze W, Holmes T, Silbgerleit A S and Solomonik V G 2009 The Gravity Probe B data analysis filtering approach *Space Sci. Rev.* **148** 411–28
- [25] Keiser G M, Kolodziejczak J and Silbgerleit A S 2009 Misalignment and resonance torques and their treatment in the GP-B data analysis *Space Sci. Rev.* **148** 383–96