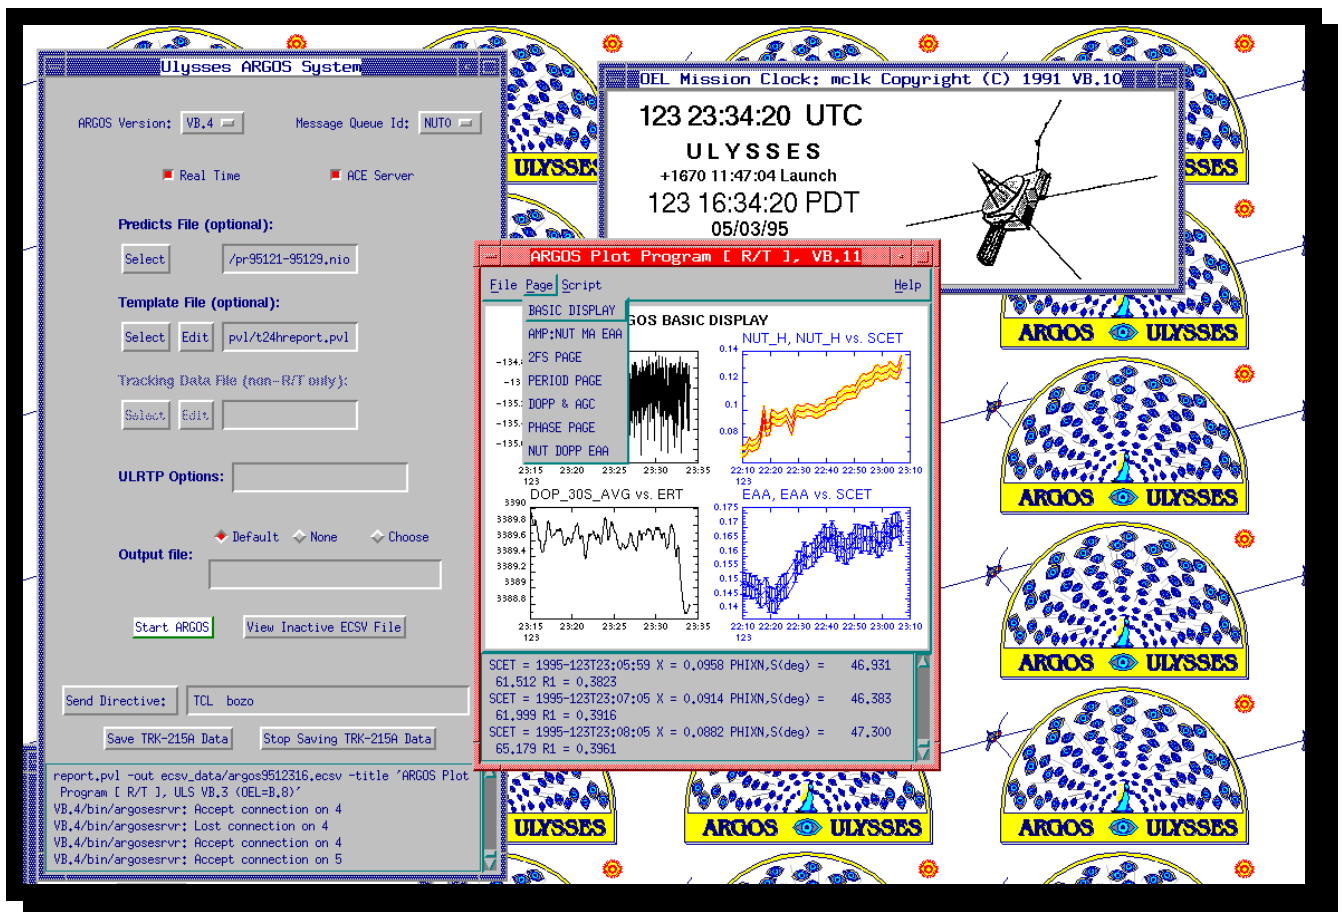




Attitude Reckoning from Ground Observable Signals
**A SYSTEM TO MONITOR
ULYSSES NUTATION AND THRUSTER FIRINGS
FROM VARIATIONS OF THE SPACECRAFT RADIO SIGNAL**



Raúl GARCÍA-PÉREZ ^{ESOC} & **Timothy P. McELRATH** ^{JPL},
Kevin J. MILLER ^{JPL}, **L. Alberto CANGAHUALA** ^{JPL}, and **L. Robert STAVERT** ^{JPL}

JPL: Jet Propulsion Laboratory, California Institute of Technology
ESOC: European Space Operations Center, ESA - Ulysses Flight Team
Office at the Jet Propulsion Laboratory

10 May 1995

(A paper for the Proceedings of the International Symposium on Space Dynamics, Toulouse, 19-23 June 1995)

**ARGOS: A SYSTEM TO MONITOR ULYSSES NUTATION AND THRUSTER
FIRINGS
FROM VARIATIONS OF THE SPACECRAFT RADIO SIGNAL**

Raúl GARCÍA-PÉREZ ^{ESOC} & **Timothy P. McELRATH** ^{JPL},
Kevin J. MILLER ^{JPL}, **L. Alberto CANGAHUALA** ^{JPL}, and **L. Robert STAVERT** ^{JPL}

JPL: Jet Propulsion Laboratory, California Institute of Technology

ESOC: European Space Operations Center, ESA - Ulysses Flight Team

Office at the Jet Propulsion Laboratory

ABSTRACT - *The predicted return of the solar-induced nutation of the spin-stabilized Ulysses spacecraft provided the impetus for the development of...*

ARGOS: (Attitude Reckoning from Ground Observable Signals) A **GROUND BASED** attitude measurement system to monitor spacecraft oscillations and attitude control manoeuvres.

Main features:

- *ARGOS does not use any onboard hardware or software. It only requires the spacecraft to transmit its carrier radio signal. This is vital because the telemetry is devoted to the science payload, and cannot provide high-rate attitude data.*
- *ARGOS reliably detects attitude dynamics in the 0.02° to 1.0° range with a 0.005° accuracy and a one-minute update rate. The attitude control delta-Vs are measured with an accuracy of up to 0.1 mm/sec.*

These features have attracted the interest of other projects. Galileo is already using ARGOS, and Mars Pathfinder and Cassini are planning to use it once they launch.

Apart from the realtime monitoring, ARGOS has improved the operational plans for the control of the Ulysses nutation anomaly. The ARGOS measurements have been used to refine the models of the nutation forcing and the spacecraft damping response, thus improving the predictions of the nutation timing and strength.

- Contents:**
1. Introduction
 2. Nutation Kinematic Model
 3. Processing Method
 4. Development
 5. Operational Results
 6. Conclusion

1. INTRODUCTION

The Ulysses mission is a cooperative venture between NASA and ESA to explore the heliosphere over the polar regions of the Sun for the first time in history. The Ulysses spacecraft is equipped to measure charged and neutral particles, dust, magnetic fields, electro-magnetic waves, and high energy radiation. Ulysses was launched in October 1990 on a trajectory designed to spend 234 days during 1994 and 1995 over the polar regions of the Sun. To achieve this unique out-of-the-ecliptic trajectory, Ulysses resorted to a gravity assist obtained from a high-latitude Jupiter encounter in February 1992.

The Nutation Anomaly: The spin-stabilized Ulysses spacecraft experienced a significant nutation level for a period of 46 days shortly after launch. Details are available in [2: Gienger *et al* 91], [5: García-Pérez 92], and [9: Crellin & Janssens 93]. A study of the phenomenon [1: Hoffman 90] revealed that it was due to solar heating effects on the spacecraft axial boom, and that it was a function of three geometrical factors: the *Sun-spacecraft-Earth angle*, the *Sun-spacecraft distance*, and the *spacecraft shadowing of the boom*.

These factors are combined into a “*nutation forcing function*,”¹ which is an accurate measure of the severity of the problem at a given time. Figure 1 shows the predicted nutation forcing function during the original occurrence of nutation in 1990 and for the periods during 1994-1995 and 2001 when nutation was predicted to return.




Figure 1: Nutation forcing function from 1990 to 2001

The nutation forcing function was the basis for the 1994-1995 nutation control operations campaign. Nutation control operations started in August 1994, the earliest time that a chance of nutation buildup existed.

The nutation control method: Ulysses’ passive dampers are not efficient enough to contain the nutation anomaly, and they must be supplemented with active damping from the “*conscan manoeuvre*.” This manoeuvre uses a continuous uplink from the tracking station as a beacon to determine the position of the Earth. The onboard conscan system commands attitude control pulses to reduce the spacecraft’s off-pointing from the Earth to specified levels, damping the nutation in the process.

Good nutation predictions are needed to reduce the cost of Ulysses operations and the impact on other projects: The conscan manoeuvre imposes a continuous tracking requirement on the ground stations supporting Ulysses. During the period when the spacecraft was at a large southerly declination this requirement was met with a combination of DSN (NASA/JPL Deep Space Network) stations and the ESA tracking station at Kourou. As these tracking stations are expensive resources which are in high demand by many projects, it was imperative to accurately predict the periods when their use would be necessary.

¹ The expression for the “nutation forcing function” is given in section 5.1

How ARGOS was conceived: The difficulty of obtaining nutation estimates from the spacecraft telemetry, as well as the increased risk due to the continuous manoeuvre configuration, drove the Ulysses Spacecraft Operations Team to explore *alternative sources of attitude and manoeuvre information*.

The directional gain pattern of the antenna used on the Ulysses spacecraft makes *radio signal strength measurements* an effective means of determining the attitude modes of the spacecraft [10: García-Perez 93]. These measurements, referred to as “AGC”, are provided once per second by the DSN tracking stations. In addition, the DSN stations simultaneously provide *Doppler* observations of the spacecraft carrier at a rate of ten times per second, which may be used to estimate in realtime the timing and delta-V magnitude of the manoeuvre pulses.

Raúl García-Pérez from the *ESA Spacecraft Operations Team* provided a *mathematical model* for determining the spacecraft attitude from the AGC data, as described below.

Both *AGC and Doppler* data were made available by *JPL* in realtime at the workstations in the Ulysses operations area. To this end JPL modified a part of their ground network which previously had not had channels for this type of data.

The *ARGOS processing system* was built by *Tim McElrath, Kevin Miller and the other JPL authors* to run on the local workstations of the Ulysses Project.

The name “ARGOS”: In Greek mythology, Argos is a hundred-eyed and ever-vigilant giant who served the goddess Hera (Zeus’s wife). Upon the giant’s demise, Hera decorated the peacock’s tail with 100 “eyes” in Argos’s honor. Conveniently for the authors, Argos is also the name of the faithful dog of Ulysses in “The Odyssey”. The system was named after these mythological and literary figures to convey the vision that “*ARGOS looks after Ulysses with 100 eyes*”. Officially, ARGOS stands for: “**Attitude Reckoning from Ground Observable Signals**” (One of the authors claims that ARGOS really means: “*A Raúl García-Pérez Original System*”; but he is obviously joking).

The way ARGOS works: Figure 2 provides an overview of the role ARGOS plays in the Ulysses operations environment. In summary, the spacecraft motion modifies the downlink radio signal, which is observed by the DSN and passed on to ARGOS, which performs separate processing on AGC and Doppler data to arrive at attitude estimates and manoeuvre information.




Figure 2: The way ARGOS works

2. NUTATION KINEMATIC MODEL

Raúl García-Pérez from the Ulysses Spacecraft Team created the following description of the nutation effect upon the AGC level.

Notation: While this description uses complex numbers to describe rotating vectors in compact formulae, their “complex” nature can be disregarded. The complex numbers here represent real vectors in a *polar projection* of a small region of the celestial sphere. In this “flattened” cap of the sphere, the distances (magnitudes) measure a few degrees of arc, while the phases go from 0 to 2π radians and also extend from $-\infty$ to $+\infty$ when they represent indefinite rotations. Under this convention the spacecraft motion can be expressed as the sum of a series of rotating vectors of the form:

$$\text{magnitude} \cdot e^{j(2\pi \cdot \text{freq} \cdot t + \text{phase})}$$

where: $r \cdot e^{ja} \equiv r \cdot \cos a + j \cdot r \cdot \sin a$ represents the real vector: $r \cdot \cos a \cdot \hat{x} + r \cdot \sin a \cdot \hat{y}$

The Model: The AGC level is a function “Xbeam(r)” of the earth location relative to the High Gain Antenna boresight. There is a bias “b” whose value comes from the rest of the link budget:

$$\text{AGC}(t) = \text{Xbeam}(\text{earth}(t)) + b \quad (2-1)$$

The expression “earth(t)” describes the Earth’s apparent motion as seen from the spacecraft-centered reference frame.

$$\text{earth}(t) = \frac{EAA \cdot e^{j\phi_E} - [nh \cdot \text{rose}(t) + ma \cdot \text{rosem}(t)]}{X1(t)} \quad (2-2)$$

“EAA” is the Earth Aspect Angle, and ϕ_E is the Earth phase angle. The nutation half-cone amplitude is “nh”, and “ma” is the half-cone amplitude of the “Meridian Anti-symmetric” oscillation mode produced by the 72.5m flexible wire booms. For brevity, this mode will be referred to as the “MA mode”. $X1(t)$ is the rotating unit vector of the spacecraft reference system,

$$X1(t) = e^{j(\omega_s \cdot t + \phi_s)}, \quad \omega_s = 2\pi f_s, \quad f_s = \text{spin frequency} \quad (2-3)$$

“rose(t)” is the description of the “*nutation rosette*” traced out by the spacecraft Z axis in inertial space taking as its origin the location of the momentum vector.

$$\text{rose}(t) = e^{j(\omega_s \cdot t + \phi_s)} \cdot \left[r1 \cdot e^{-j(\omega_n \cdot t + \phi_n)} + r2 \cdot e^{j(\omega_n \cdot t + \phi_n)} \right], \quad \omega_n = 2\pi f_n, \quad f_n = \text{nutation freq.} \quad (2-4)$$

and “rosem(t)”, is the rosette due to the MA mode.

$$\text{rosem}(t) = e^{j(\omega_s \cdot t + \phi_s)} \cdot \left[rm1 \cdot e^{-j(\omega_m \cdot t + \phi_m)} + rm2 \cdot e^{j(\omega_m \cdot t + \phi_m)} \right], \quad \omega_m = 2\pi f_m, \quad f_m = \text{MA freq.} \quad (2-5)$$

The parameters $r1$, $r2$, $rm1$, and $rm2$ are normalized ratios of certain moments of inertia. The normalization is chosen so that :

$$r1 + r2 \equiv 1, \quad rm1 + rm2 \equiv 1 \quad (2-6)$$

The phases associated with the spacecraft spin, Earth Aspect Angle, nutation, MA mode, and Xbeam offset direction are denoted by ϕ_s , ϕ_E , ϕ_n , ϕ_m , and ϕ_X , respectively. The angles ϕ_s and ϕ_E are not observable individually through AGC measurements. Instead, the angle defined as $\phi_c = \phi_s - \phi_E$ is used in the expressions that follow.

The “Xbeam” function used here is a symmetrical paraboloid. Based on measurements, the terms beyond 2nd order are not significant for off-pointing angles below 1° .

$$\text{Xbeam}(r \cdot e^{ja}) \approx -K \cdot \left(\left| r \cdot e^{ja} - X \cdot e^{j\phi_X} \right| \right)^2, \quad K \approx 5.0 \text{ dB}/(\text{°})^2, \quad X \approx 0.1^\circ, \quad \phi_X \approx 0.95 \text{ rad} \quad (2-7)$$

This approximation permits us to derive expressions for the amplitudes and phases of the main frequency components in the signal $\text{AGC}(t)$. Equations (2-1) to (2-7) can be used to derive the following:

$$\text{AGC}(t) \approx -K \cdot \left\{ \begin{array}{l} EAA + X^2 + nh^2 \cdot (r1^2 + r2^2) + ma^2 \cdot (rm1^2 + rm2^2) \\ + 2 \cdot r1 \cdot r2 \cdot nh^2 \cdot \cos(2\omega_n t + 2\phi_n) \\ - 2 \cdot EAA \cdot nh \cdot r1 \cdot \cos[(\omega_s - \omega_n)t + \phi_c - \phi_n] \\ - 2 \cdot EAA \cdot nh \cdot r2 \cdot \cos[(\omega_s + \omega_n)t + \phi_c + \phi_n] \\ - 2 \cdot EAA \cdot ma \cdot rm1 \cdot \cos[(\omega_s - \omega_m)t + \phi_c - \phi_m] \\ - 2 \cdot EAA \cdot ma \cdot rm2 \cdot \cos[(\omega_s + \omega_m)t + \phi_c + \phi_m] \\ - 2 \cdot EAA \cdot X \cdot \cos(\omega_s t + \phi_c + \phi_X) \\ + 2 \cdot nh \cdot X \cdot [r1 \cdot \cos(\omega_n t + \phi_n + \phi_X) + r2 \cdot \cos(\omega_n t + \phi_n - \phi_X)] \\ + A_{2s} \cdot \cos(2\omega_s t + \phi_{2s}) \end{array} \right\} + b \quad (2-8)$$

A small component at twice the spin frequency appears in some AGC data samples because the real Xbeam paraboloid is slightly asymmetrical. This component was added to (2-8), resulting in eight sinusoidal frequency components whose amplitudes and phases are listed in Table 1.

Table 1: AGC(t) Sinusoidal Components predicted by the Kinematic Model

Frequency	Symbol	Amplitude	Phase
Spin	f_s	$2 \cdot K \cdot EAA \cdot X$	$\phi_c + \phi_X$
Spin + Nutation	$f_s + f_n$	$2 \cdot K \cdot EAA \cdot nh \cdot r2$	$\phi_c + \phi_n$
Spin - Nutation	$f_s - f_n$	$2 \cdot K \cdot EAA \cdot nh \cdot r1$	$\phi_c - \phi_n$
Nutation	f_n	$2 \cdot K \cdot nh \cdot X \cdot (r1^2 + r2^2 + 2 \cdot r1 \cdot r2 \cdot \cos 2\phi_X)$	$\tan^{-1} \left[\frac{r2 \cdot \sin(\phi_n - \phi_X) + r1 \cdot \sin(\phi_n + \phi_X)}{r2 \cdot \cos(\phi_n - \phi_X) + r1 \cdot \cos(\phi_n + \phi_X)} \right]$
2 · Nutation	$2 \cdot f_n$	$2 \cdot K \cdot nh^2 \cdot r1 \cdot r2$	$2 \cdot \phi_n$
Spin + MA	$f_s + f_m$	$2 \cdot K \cdot EAA \cdot ma \cdot rm2$	$\phi_c + \phi_m$
Spin - MA	$f_s - f_m$	$2 \cdot K \cdot EAA \cdot ma \cdot rm1$	$\phi_c - \phi_m$
2 · Spin	$2 \cdot f_s$	A_{2s}	ϕ_{2s}

ARGOS measures the $\text{AGC}(t)$ components and derives: EAA , nh , ϕ_c , ϕ_n , ma , etc. (using the above expressions). If Xbeam requires higher-order terms, an explicit expression for $\text{AGC}(t)$ may not be obtainable, but it is possible to implement ARGOS as an iterative algorithm that fits the kinematic model to the observed $\text{AGC}(t)$, and estimates the kinematic parameters above.

As predicted by the kinematic model (Table 1), not all of the frequency components are significant at any given time. Figure 3 shows a real sample reflecting moderate oscillation amplitudes.

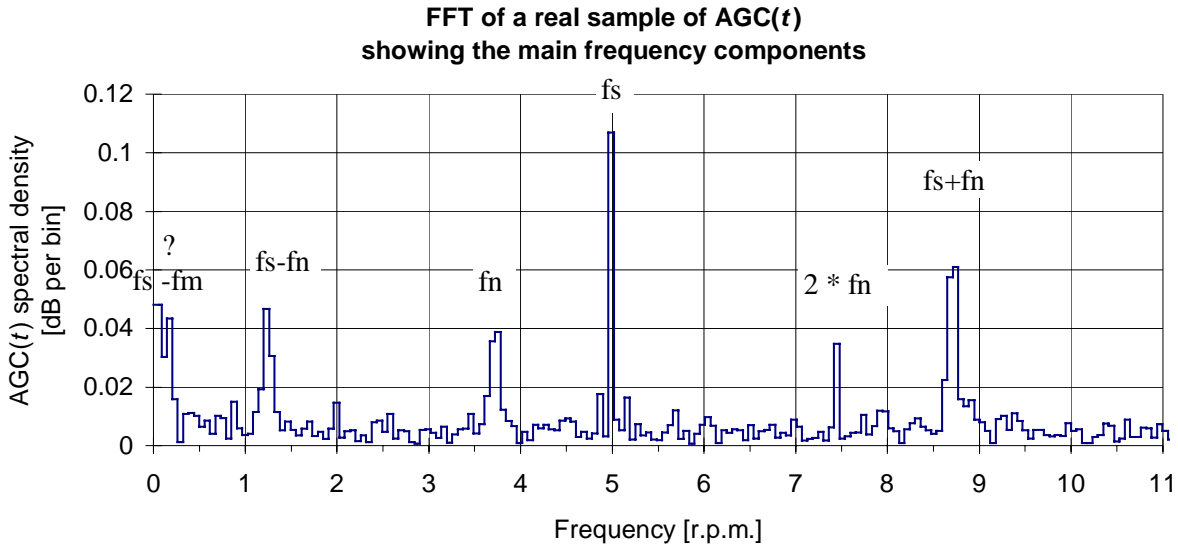


Figure 3: Sample FFT of a typical AGC time series

All five of the spin and nutation-related components are clearly visible above the noise level, and “ $fs-fn$ ” may be present as well, although there is enough low-frequency noise to make this questionable. Both the $2fs$ and $fs+fn$ components are not visible in this FFT.

Spacecraft parameters affecting frequency components: The values of $r1$ and ϕ_X exert a large effect on the relative magnitude of several frequency components, as described in Table 1. A typical value for $r1$ is 0.365, ($\Rightarrow r2 = .635$), causing the $fs+fn$ component to have almost twice the amplitude of $fs-fn$. This effect can be seen in Figure 3 when the total area under the peaks in the FFT at each frequency is considered. Even more importantly, since $\phi_X \approx 0.95 \text{ rad}$, then $\cos(2\phi_X) \approx -0.32$, and the amplitude of fn is 46 % of the maximum that would be possible (if $\phi_X = 0 \text{ rad}$). For this reason, the amplitude of $2*fn$ exceeds that of fn when $nh > 0.2^\circ$.

3. PROCESSING METHOD

3.1 Signal level (AGC) processing

Once an acceptable kinematic model of the of the spacecraft oscillation was developed, a processing method was required to produce estimates from the AGC data. Since the frequency of most of the oscillation components does not change very much, the problem of detecting a particular signal is not very difficult. However, the fact that there were up to eight usable signals in the data was a challenge, as was the large difference in the periods of the expected signals (from under 6 to over 300 seconds).

The current implementation is based on an FFT and a time domain least-squares fit, followed by the application of the kinematic model, but the authors have also explored some alternative methods that are worth mentioning at this point.

Historical evolution:

Kalman filter: The initial plan was to use a Fourier transform to initialize a Kalman filter, which would then follow even fairly rapid changes in the signal [14: Cangahuala 94]. The advantage of the Kalman filter is that all the information relating to any estimated parameter is correctly combined to provide the best estimate, even if several different signals depend on that parameter.

An example of this would be the nutation amplitude “ nh ”, which appears in five different signals.

The disadvantage of the Kalman filter approach is that it is dependent on the accuracy of the model, and that it may have difficulty when some of the frequency components cannot be observed (for instance due to very small values). Lastly, the difference in signal period makes it difficult to tune a Kalman filter that will correctly separate signal bias variations from the long-period terms.

Phase-lock loops: Another approach that was suggested later was to use eight bandpass filters, followed by eight digital phase-lock loops. The phase-lock loops would be able to follow strong signals quickly, and still provide an estimate of long period or weak signals by using long integration times. In addition, an out-of-lock condition for one signal can be reinitialized separately, and that signal can be excluded from the estimates until it is recovered. Such a process is also less sensitive to model errors, as the signal tracking does not depend on anything but the predicted frequency band. This approach may actually be very good, but unfortunately it was conceived after the end of the initial implementation, and resources have not been available to implement it so far.

Current implementation: It was recognized at an early stage that a sequential process including a Fourier transform, a least-squares fit, and the application of the kinematic model was an estimation process in itself. This process was the first to be implemented as an operational system, and it has worked so well that there has been no urgent need to replace it. As currently implemented, the attitude estimate process has several steps:

- *AGC Pre-processing:* Every 60 seconds, the available AGC points are searched for up to 1024 points. Gaps of up to 12 seconds are filled by interpolation of boundary values, and in the case of a long gap the process lowers the number of points down to a minimum (selectable by directive). As longer intervals become available, the number of points is gradually increased up to 1024. When there is a choice, the process selects the data from the station tracking in a “two-way” mode (for the definition of “two-way” see footnote 5).
- *A fast Fourier transform (FFT)* is performed on the data. Within the frequency range of each signal, the highest amplitude bin is found. The amplitude, frequency, and phase of the underlying signal is estimated from the highest bin and the two adjacent bins, according to a method summarized in an appendix at the end of this section (see Equations (3-1) and (3-2)).
- *Least-squares fit in the time domain:* The eight resulting signals are first checked against the expected frequency limits, then fit to the original time series using least-squares, and finally compared to the average noise level ². The least-squares fit is iterated until all the signals converge or the iteration limit is reached. The result of this detection process is a subset of “valid” signals.
- *Estimate kinematic parameters:* The validated set of signals is used to estimate the underlying kinematic parameters. The expressions used are chosen from Table 1 according to a complicated hierarchy which depends on the number of signals detected and also on which signal has the largest amplitude. Numerous limits are included to reduce the probability of a spurious result.
- *Estimate kinematic parameter variances:* The estimates are used to fit the original data once again using the full kinematic model, but only the variance of each parameter, scaled by the sum of the squares of the residuals, is used from this step. Consequently, noisy data or unexpected events can be detected by the large sigmas they produce. The sigmas and values are compared, and any estimates with a sigma-to-value ratio larger than a given threshold are regarded as invalid. The valid estimates are reported, along with their sigmas.

² An estimate of the background noise is made by summing the spectral power in regions devoid of any expected signals.

- *Accept operator directives and deliver system messages:* In addition to providing estimates, the processing of AGC data results in a number of warning and informational messages outside of the normal data flow. There are also several parameters of the estimation process, such as the maximum and minimum size of the initial FFT, that can be changed while the program is running. Both of these functions use mechanisms that will be described later.

Example of AGC processing: Let us illustrate the process with an example, which uses approximately the same AGC time-series whose FFT was shown in Figure 3.

- *AGC Pre-processing:* In this first step, 1024 points are selected, and no gaps are found in the data series.
- *The FFT* is then performed, with the same results as seen in Figure 3, and the background noise level is found to be 0.0114 dB. The initial estimate of each signal is obtained from the FFT, and is shown in Table 2.
- *Least-squares fit:* Next the program performs a least-squares fit of the time series, and iterates until convergence is reached. As can be seen from Table 2, some of the initial estimates are close to the final values and converge quickly, while others require several steps to converge. The estimated frequencies are tested at each step, and as a result the $2fs$ and $fs+fm$ signals are discarded. Once the least-squares fit is finished, the amplitudes are compared to the noise level, which does not invalidate any signals in this example.

Table 2: Example of amplitude convergence during iteration process (AGC units in dB)

	fs	fn	fn+fs	fn-fs	2fn	fs-fm
Initial	.183	.058	.089	.070	.104	.063
1st iter	.064	.059	.091	.057	.009	.053
2nd iter	.080	.059	.092	.059	.014	.043
3rd iter	.103	.059	.092	.059	.015	.045
4th iter	.106	.059	.092	.059	.025	.045
...
7th iter	.106	.059	.092	.059	.036	.045

- *The kinematic parameters* are estimated next, and due to the presence of all 5 spin and nutation parameters, a number of auxiliary parameters are estimated (including two different means of estimating ϕ_x , which result in values of 51.7° and 48.3°). The results are given in Table 3 along with the parameter variances. The fs signal has the largest amplitude, but the nutation nh is actually higher than EAA because of the effect of ϕ_x (explained at the end of section 2). Note that the uncertainty of the MA amplitude is almost half of its value, while all the other parameters shown here are much better determined. Due to the signals discarded earlier, no estimate is available for $rm1$ or any of the $2fs$ parameters. Lastly, all the valid estimates are provided to the display program, and the auxiliary estimates are sent as an informational message.

Table 3: Example of parameter estimates and sigmas

Parameter	Value	Sigma	Units
EAA	.106	.0031	degrees
nh	.143	.0050	degrees
ma	.051	.0218	degrees
r1	.3903	.0227	none
spin period	12.0473	.0020	sec
nut. period	16.1054	.0031	sec
MA period	11.6147	.0060	sec

Appendix: Method to estimate amplitude, frequency, and phase from the FFT. A complete derivation of the following method is given in [7: Stephens 92], but a short summary is given here.

When using a discrete Fourier transform to detect a signal, the actual amplitude of a signal can be underestimated by up to 36% if the actual frequency falls near the edge of a frequency bin. Consequently, a much more effective estimator of the actual frequency can be obtained from using the Fourier coefficients of the largest bin and its nearest neighbors. If the complex coefficient of the largest bin is F_0 and its higher and lower frequency neighbors are F_h and F_l respectively, then it can be shown that the offset of the actual frequency from the center of the largest bin is:

$$\Delta k = \frac{1}{2} \cdot \left(\frac{F_l}{F_0 - F_l} + \frac{F_h}{F_h - F_0} \right) \quad (3-1)$$

If the actual amplitude is A and the actual phase is ϕ_0 , then:

$$A \cdot e^{j2\pi\phi_0} = \frac{1}{3} \cdot \frac{j2\pi}{e^{j2\pi\Delta k}} \left[\Delta k (F_l + F_0 + F_h) - (F_h - F_l) \right] \quad (3-2)$$

The above formulae are simplified from the expressions derived in [7: Stephens 92] by assuming that there is no zero padding of the Fourier transform, and the sampling of the Fourier coefficients is at the Nyquist frequency.

3.2 Doppler processing

The goal of ARGOS in terms of Doppler data processing is to provide a means of detecting attitude manoeuvre pulses as small as 0.25 mm/sec , either singly or in groups. The average spacing between pulses is 36 seconds, although values as small as 24 seconds or as large as 48 seconds are possible.

Doppler Predictions: In order to use Doppler data to the required accuracy, the motion of the tracking station and the spacecraft must be modeled with great care. The predicted Doppler observable is calculated in advance, because trying to calculate the predicted Doppler observable in realtime at ten points per second would be beyond the capabilities of the available equipment.

The predicted Doppler observable is calculated using the JPL Orbit Determination Program (ODP), which already contains all the trajectory³ and Earth platform models necessary to reach the specified accuracy. The points are calculated on ten minute centers and then smoothed with a fifth-order spline fit, because the value of this observable varies rather slowly. ARGOS then interpolates the spline coefficients in realtime to obtain the one-way light time and the Doppler observable for each point. Instead of frequency, the actual parameter used for the predicted Doppler observables is the normalized phase, which is a better representation of the actual observation process than the Doppler frequency.

Accuracy limiting factors: The Ulysses radio system consists of an S-band (2.1 GHz) uplink and an S- and X-band (8.4 GHz) downlink. During nutation control operations, only the X-band transmitter was used. The performance of the Doppler measurement systems over time scales of less than a minute is primarily limited by the effects of *charged particles on the S-band uplink*. Although *troposphere effects* on both the uplink and downlink signals are also a factor, most of the charged particle effects are probably due to solar plasma at large distances from the Sun⁴.

For navigation purposes, Doppler data from this system is assumed to have an accuracy of *1 mm/sec*, but the actual scatter of the residuals is often smaller than *0.1 mm/sec* after compressing the data to one point per ten minutes.

The *three-way*⁵ Doppler predicted observables have excessive noise levels, and cannot be used operationally. This is due to the limitations imposed by the existing (ODP) file formats for the input to this process. However, *one-way*⁵ and *two-way*⁵ Doppler data show good results.

Transmitter "ramps": The uplink frequency changes are not taken into account in computing the Doppler residual. This is not considered worthwhile given that the uplink frequency changes are few, very predictable, and easily distinguishable from the manoeuvres.

Doppler processing improvements: While navigation processing of Doppler data typically involves using differenced phase divided by the integration interval as a pseudo-velocity measurement, the high data rate available from DSN stations makes it possible for ARGOS to use different methods.

ARGOS removes the noise frequencies higher than one Hz by performing a slope fit to ten of the 0.1 second sample-rate phase residuals, and then reporting the slope as a velocity residual once per second. If the once-per-second data appears valid, ARGOS performs a slope fit over a 30.2 second interval once per second. This also removes the frequencies over 0.03 Hz, and suppresses the signal modulation due to the spacecraft spin. The resulting noise levels are small enough to permit the Doppler residuals *to achieve the goal* of detecting single pulses of *0.25 mm/sec*.

Figure 4 shows an example of the two Doppler processing methods for a data set that includes three manoeuvre events. The first part of Figure 4 shows the Doppler residual based on one second of measurements, and has a scatter of about 4 mm/sec. While the large manoeuvre at 01:00 can be detected, its details are not clear. The second part of Figure 4 shows the 30-second smoothed Doppler residual, which provides a much more accurate view of the events. The individual pulses (of which there are seven) of the first manoeuvre can almost be detected, and two single-pulse manoeuvres are clearly visible at about 01:12 and 01:21.

³ The trajectories used were obtained from the operational orbit solutions performed for trajectory reconstruction and prediction purposes. The methods used in these orbit solutions are described in References [4], [8], and [15].

⁴ The geocentric angle between the Sun and the spacecraft always exceeded 16° during nutation operations, so the uplink radio signal was never excessively affected by the solar corona.

⁵ *Two-way* is the mode where the spacecraft uses a coherent transponder and the same station is both receiving from and transmitting to it. *Three-way* is when the stations are different, and *one-way* is when the spacecraft is non-coherent or does not receive any uplink.

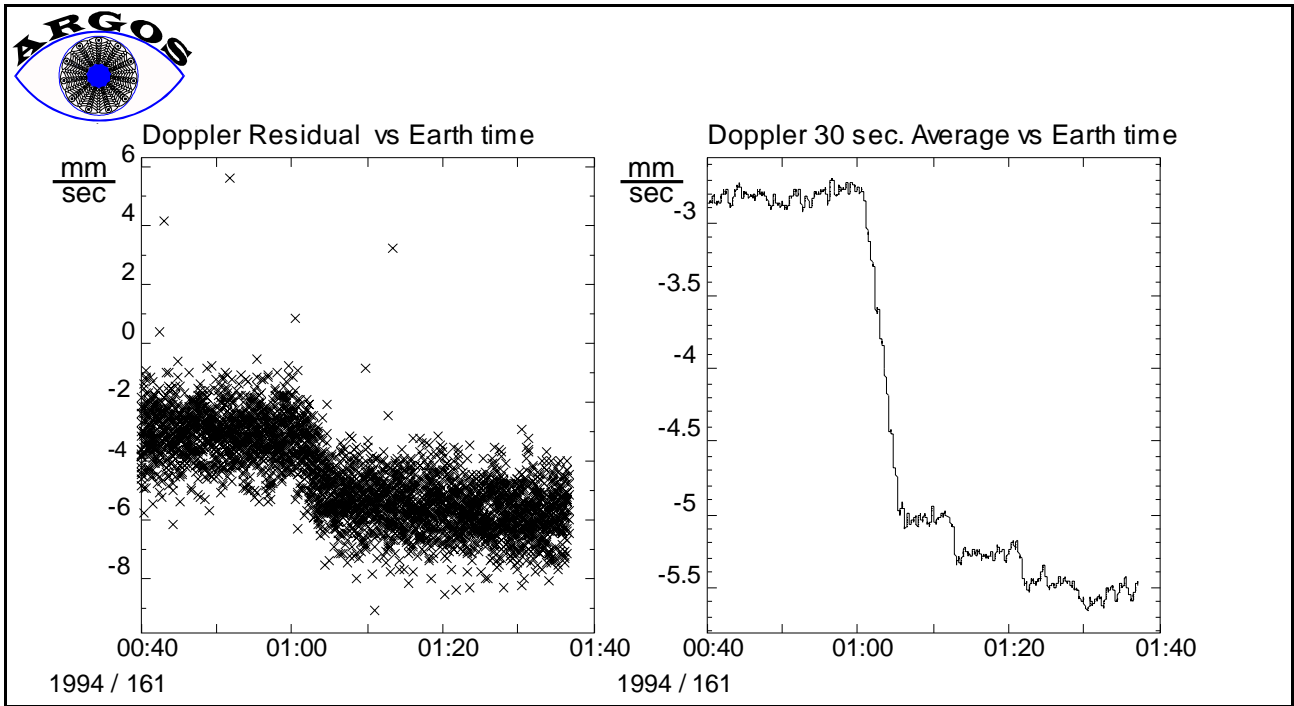


Figure 4: 1 second vs smoothed 30 sec Doppler including manoeuvre

4. ARGOS DEVELOPMENT

4.1 Data flow and interfaces

Tracking data (which contains both AGC and Doppler measurements) is generated by the DSN stations in the Trk 2-15a format, a variable-length bit-packed format with a long history. The tracking data samples are then packed into a 600-byte NASCOM block, and sent to JPL when the block is full or a time-out is reached. When 0.1 second sample rate Doppler is used, this results in a block containing 6 seconds of data being sent every 6 seconds. In addition to the normal routing of tracking data to navigation computers, changes in the JPL multi-mission ground system (AMMOS) were made to route Ulysses data to all workstations in the Ulysses operations area in a broadcast mode via the Ulysses LAN.

The ARGOS system was divided into four tasks:

- *Receive* the data from the network and pass it onto the processing task.
- *Process* the data to obtain attitude estimates and Doppler residuals.
- *Re-Transmit* the processed data to client workstations (optional).
- *Display* the processed data (described in 4.2).

One team was responsible for the “*Process*” task, while *K. Miller* was responsible for the remaining tasks. At the beginning of the development process the interfaces between the second task and the other tasks were clearly defined, allowing concurrent development to proceed for weeks at a time without requiring significant interaction across task boundaries.

Receive: Operationally, the same program is used to receive data from the network, read three different file formats, write the 600-byte blocks to a disk file, and pass text messages to downstream programs.

The interface between the “*Receive*” and the “*Process*” tasks consists of a stream of 600-byte NASCOM blocks interspersed with text data, all of which was sent through a message queue, which is an interprocess communication (IPC) method available in Unix operating systems.

While any one of several different IPC constructs could have been used, the message queue seemed most appropriate to this task. It allows multiple processes to send messages to one receiving process, which enables commands in text format to be inserted into the data flow so that the later tasks can be controlled in realtime. In addition, the message queue is insensitive to the termination of either the reading or the writing processes, which prevents failures of later tasks from affecting the network receiving task. During the course of ARGOS development, several different programs actually supplied the data flow through the message queue interface, which illustrates the flexibility of this IPC method.

Data was also supplied in a slightly more compact format in the form of files from the navigation computers, which was and still is useful as a backup data source.

Process: The input data is separated from the text messages, unpacked from either input format, and processed to obtain attitude estimates and Doppler residuals as described above.

The IPC method used for communicating between the “*Process*”, “*Re-transmit*”, and “*Display*” tasks is the “*simple pipe*”, which normally would mean that the termination of any program in the pipe causes all other programs to terminate. While this is not true of the display program, it is true of the processing program, although the reliability of the processes is such that this feature is acceptable. The data format⁶ is encapsulated comma-separated values (ECSV), which consists of a descriptive header followed by lines of comma-separated values interspersed with text directives. The directives are used for informational messages, comments, and display commands using Tool Command Language (TCL) [13: Wesley 94]. The ECSV data is reasonably compact and yet easily readable, and can also be imported into spreadsheets and other post-processors. Currently a total of 38 parameters (of which 15 are parameter uncertainties) are included in the ECSV data flow.

Re-transmit: This task may or may not be present, and consists of a server and a client program. The server makes the ECSV data available on the local network using a small extension of standard network services. Whenever a client program from another workstation requests a connection, the server sends the ECSV header and then starts sending the current ECSV data in realtime. The client program is usually piped to the display program, allowing two or more independent displays to be generated from a single processing program. The computation load of the display program is much lower than the processing program, allowing the display program to run on workstations that already carry a heavy load of other monitor, display, and command processes.

4.2 Graphical displays

The requirements for the graphical display program were provided primarily by the Ulysses Operations team, who were in turn heavily influenced by the excellent telemetry graphics provided by the Ulysses spacecraft Monitor & Control System (UMCS). The main features and capabilities of the graphical displays are the following:

- Clear full-color displays for a large number of parameters (~ 100).
- Easily configurable “pages” to show different combinations of parameters.
- Choice of time scales (i.e. ground or spacecraft time). They are both included in the ECSV file, and either one can be used as the independent coordinate.
- Scale modes: auto-scale, auto-scale enlarge only, and manual.
- Scale and location of each individual plot modifiable via a graphical user interface (GUI).
- Automatic scrolling to the left in quarter-width increments (whenever the edge of the plot is reached by the realtime plots).

⁶ Output from the processing task.

- Numerical value extraction of a data point selected with the mouse. The values are displayed in the message area, which also displays messages from the processing program.
- Logging of the ECSV data to a disk file, as well as printing of any displayed page on request.
- Realtime, and non-realtime modes.
- Room for additional functions. Functions are listed on pull-down menus configured using TCL scripts, allowing a simple means of changing printer settings, alarms, and other features.

The resulting program uses a text file in Parameter Value Language (PVL) format [3: CCSDS 91] to describe the appearance of numerous display pages, including the number and position of plots, the plot data type and appearance, and the text labels.

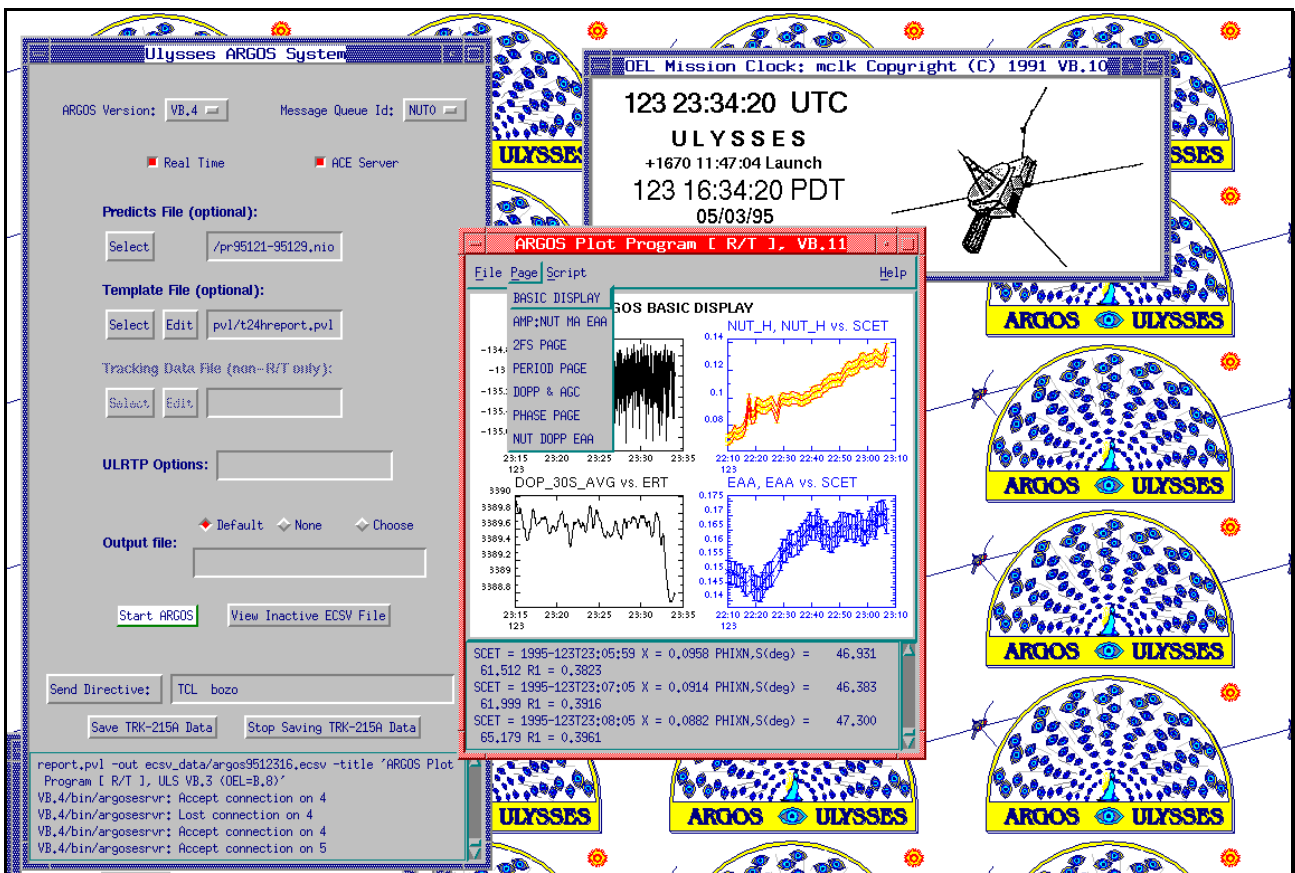


Figure 5: A sample of the ARGOS system screen

Control panel. A graphical user interface (GUI) provides a simple means of controlling the various options and modes of ARGOS. Plot formats, predict files, realtime vs. non-realtime modes, realtime directives, output files, and distribution options are all available in the GUI. Figure 5 shows a sample ARGOS screen, with the GUI on the left, the plot basic display and clock in the middle, and the colorful background with the “peacock logo” on the right. Note the message areas of both “argosplot” and “argosgui”. Most of the ARGOS spacecraft analysis work is performed using a larger *argosplot* window, but the basic display shown in Figure 5 conserves screen resources, and it is ideal for use as a client process on the busy workstations of the realtime mission operations team.

A sample of the “Basic Display” used frequently in operations is shown in Figure 6. The display shows raw AGC data and 30-second Doppler residuals with respect to Earth-receive time on the left column. The right column shows nutation magnitude and Earth aspect angle with respect to spacecraft time (denoted by a blue plot border). The 1 sigma error bars of the attitude estimates are linked by lines to form an error corridor⁷. A total of seven display pages have been typically employed in operational use.

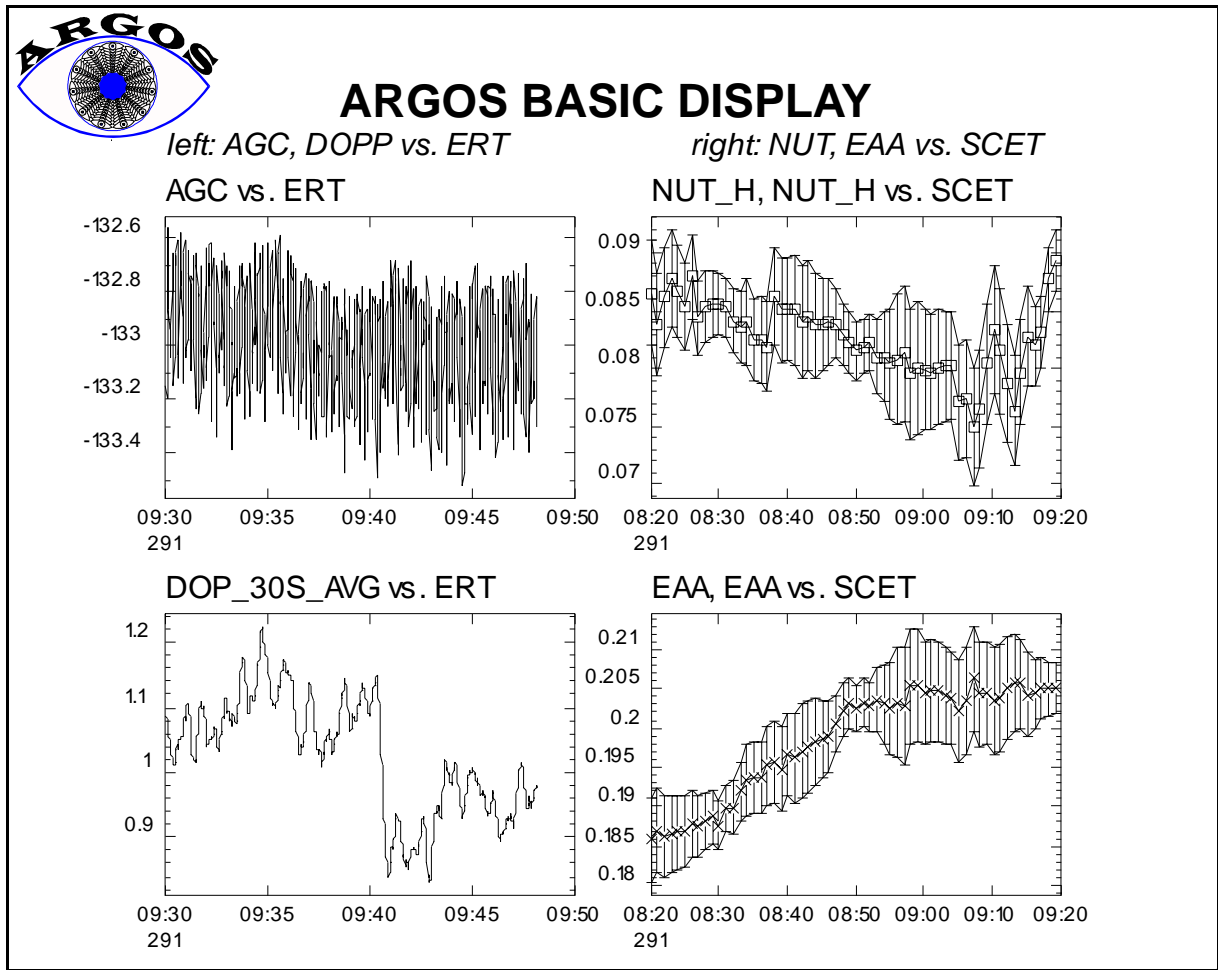


Figure 6: Example of the Basic Display

Figure 7 illustrates the flexibility of the display program, showing a different display layout. In this example the display shows three plots: the nutation magnitude, the 30-second smoothed Doppler, and the Earth aspect angle with respect to spacecraft time.

The attitude manoeuvre pulses are evident in the Doppler data in this example, and these pulses can be matched with changes in the attitude parameters. The attitude estimates take approximately seventeen minutes (1024 sec) to show the full effect of an attitude change due to the averaging nature of the attitude estimation process.

⁷ In the case of the nutation estimates, the plot displays the national colors of the Spanish instigator of the whole ARGOS effort.

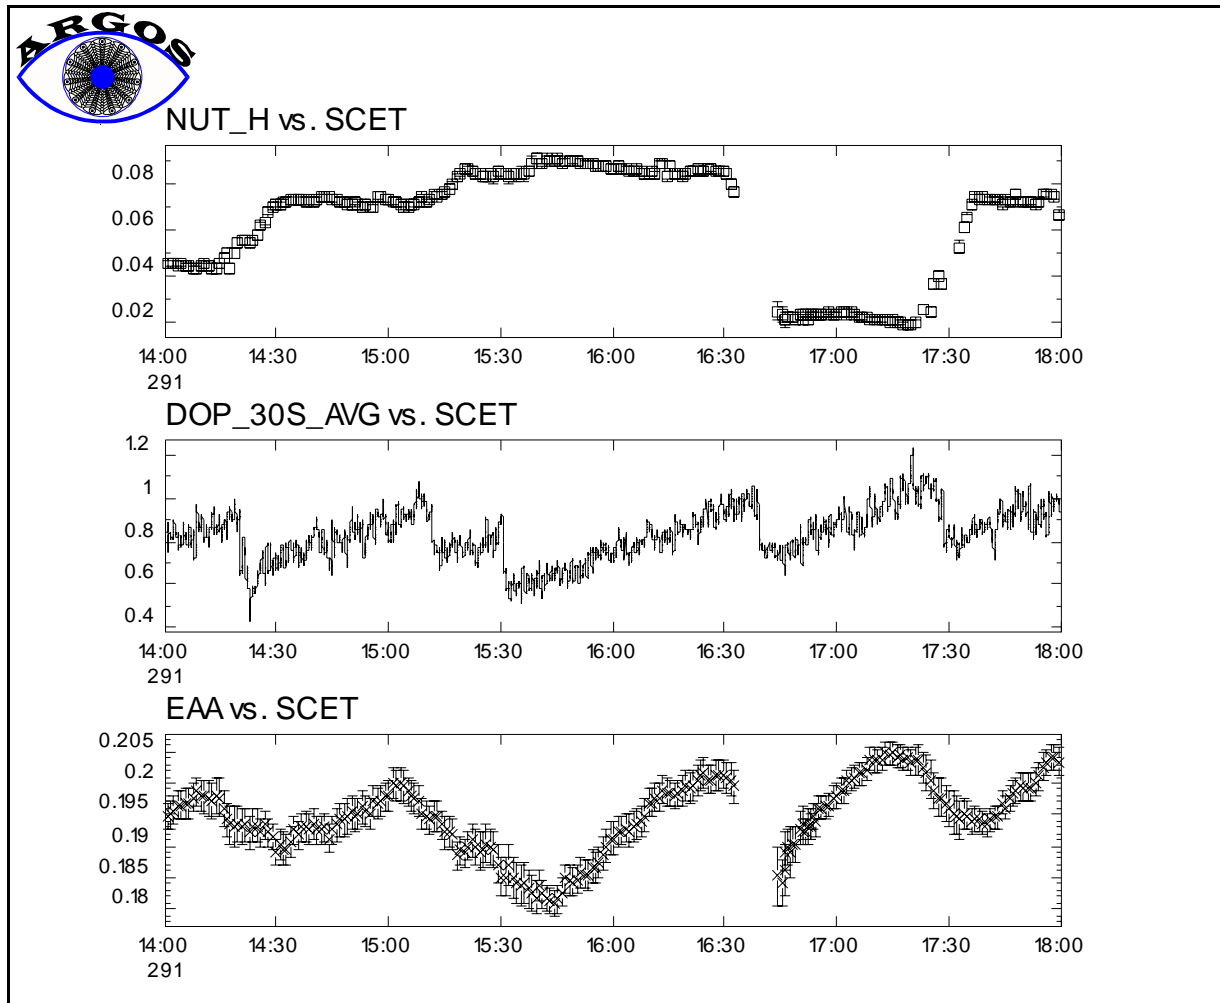


Figure 7: A display showing 4 hours of the parameters: Nutation, 30 second smoothed Doppler, and EAA

4.3 Development Schedule

The decision to develop the operational version of the ARGOS system was made in mid-February 1994. The programming effort was primarily performed by the JPL authors of this paper, with the assistance of several other JPL personnel. The ESA author coordinated this effort towards the Ulysses operational needs.

The severe time constraints resulted in the dedication of extra resources to the processing task. At one point there were two programmers working on the operational version of the processing task while another programmer refined attitude estimation algorithms. This approach was successful in producing a first operational version by late May 1994, which was used to support the first nutation control tests in early June 1994.

Further development continued throughout the summer, past the start of nutation operations on August 11, 1994. The last significant improvements were completed by early November, 1994, well before the first nutation forcing function peak. Other JPL projects, such as Galileo, Mars Pathfinder, and Cassini, are planning to make use of the ARGOS system, primarily in the manoeuvre monitoring and detection role. These projects will require further developments to ARGOS.

5. OPERATIONAL RESULTS

5.1 Nutation growth estimates

Exponential behaviour: The nutation anomaly produces an exponential behaviour⁸ in the nutation amplitude. The difference between the strength of the nutation forcing function and the efficiency of the passive nutation dampers determines whether the nutation rises or decays, and how fast it changes.

While the *nutation forcing function* provided a good model for the strength of the nutation anomaly, the ARGOS measurements of the nutation anomaly were necessary to validate and refine the predictions.

Relative slope in percent: ARGOS-derived estimates of the nutation amplitude from *manoeuvre-free arcs* are used to estimate the time constant of the underlying exponential function. The method used was to extract a number of estimates and transfer them to an Excel spreadsheet. The nutation estimates are first corrected for a nutation bias⁹ of 0.01°, and then fit with an exponential function. The results are given in terms of the *relative slope of the exponential*, because it is more meaningful operationally than the time constant. The relative slope is expressed as “percent of change per hour”.

$$\text{exponential relative slope} = \frac{\left[\frac{d(A \cdot e^{t/\tau})}{dt} \right]}{A \cdot e^{t/\tau}} = \frac{1}{\tau} \quad [\% / \text{hour}] \quad (5-1)$$

The nutation forcing function is an expression that evaluates the dynamic excitation caused by the thermal effects on the spacecraft axial boom, as derived in 1990 [1: Hoffman 90]. It is based in a simplified spacecraft geometry that represents the spacecraft’s body as a disk of radius r , and the axial boom as a straight cantilever (Figure 8).




Figure 8: Simplified spacecraft model used to derive the nutation forcing function

Two of the variables are well known: R is the “Sun-spacecraft distance”; and SAA is the “Solar Aspect Angle”. The other two (r and l) are not well determined because the spacecraft is

⁸ The exponential behaviour can be disrupted by external forces like the manoeuvre pulses.

⁹ 0.01° is a typical residual nutation level in the absence of solar-induced nutation or other dynamic disturbances.

an irregular box instead of a disk, and the boom is a 3-D elongated “S-shape” instead of a straight cantilever. The solar forcing function is dimensionless and can be normalized in several ways. The expression that is used in this paper is as follows:

$$\begin{aligned} \text{nutaton forcing function} &= \frac{1}{2 \cdot R^2} \left(1 - \frac{\cot(SAA)}{\frac{\ell}{r}} \right) \cdot \cos(SAA), & \text{for } \frac{\ell}{r} \geq \cot(SAA) \\ \text{nutaton forcing function} &= 0, & \text{for } \frac{\ell}{r} < \cot(SAA) \end{aligned} \tag{5-2}$$

It is clear that the parameter ℓ/r controls the shape and cut-off dates of the forcing function. We call this parameter the *effective shadowing ratio*, and its determination has been a prime objective of the ARGOS measurements, because it is the key to a better prediction and a more realistic operations plan.

Figure 9 shows the observed and predicted nutation behaviour over a time span of 20 months in 1994-5. As a result of the values observed during 1994, a shadowing ratio $\ell/r \approx 5$ has been adopted because it provides the best fit to the observed nutation behaviour. This value is quite low compared to the value of 8 that could have been expected from the nominal spacecraft dimensions.

In any case, the nutation has been very accurately predicted since the beginning of 1995, thus permitting very realistic operational plans to be developed.

Predicted Ulysses solar induced nutation during 1994-95 & nutation increase rate measured by ARGOS




Figure 9: ARGOS measurements superimposed on the predicted nutation

5.2 Operational decisions

When a ground station fails during the period of high nutation increase rate, it may produce unacceptable levels of nutation in a short time. Above 0.5° the nutation starts to be a threat to the spacecraft integrity because it inflicts a bending fatigue to the axial boom root that the boom was not designed to sustain¹⁰. In such a case, there is a formal JPL “*spacecraft emergency*” declaration procedure to obtain another station. However, this procedure has to be thoroughly justified and explained to the other users, as the other station is normally supporting another project that will suffer a serious interruption of its scheduled activities.

ARGOS has provided evidence to demonstrate the severity of the nutation anomaly with the frequent measurements of the exponential increase rate. With the other monitor tools alone it would be impossible to see the nutation trend until a problem occurs, and the nutation grows rather large. Normally, the continuous nutation control manoeuvre prevents this kind of growth.

The emergency declaration decision was based on the current value of the MAGIC (Maximum Allowable Gap In Coverage), and this parameter was calculated from the ARGOS nutation increase rate measurements. The MAGIC happens to coincide with the “doubling time” of the nutation exponential because it is the time that the nutation takes to go from 0.25° (typical initial nutation induced by a ground station failure) to 0.5° (Maximum Allowable Nutation).

At the beginning of the nutation operations period, the MAGIC was 8 hours, but by the time of the first nutation peak it had dropped to 1 hour 20 minutes based on nutation growth estimates. For the onset of the second nutation peak, the predicted nutation behaviour had to be generated from predicted values on a daily basis, because the forcing function changed too fast for the measurements to be valid for more than one day (see April 95 in Figure 9). As the nutation rose to its peak on May 4th, the MAGIC dropped rapidly to values as low as 1 hour.

The gap in the nutation forcing function in February and March of 1995 provided the opportunity to conduct an important radio science Solar Corona Experiment (the first time in history that a high-inclination coronal sounding over a wide range of solar latitudes was done). This experiment would not have been possible in the middle of nutation control operations. However, an accurate prediction of the time when nutation growth would stop and re-start was required in order to develop a tracking schedule. Using nutation growth estimates from the first nutation peak, nutation operations were suspended from February 6 through March 27, rather than the shorter gaps predicted by the original estimates of the shadowing ratio.

5.3 Operational examples

Exponential nutation growth. Uplink variations occasionally caused periods of exponential nutation growth without any control pulses. One such event occurred on December 10, 1994, as shown in Figure 10.

The DSS 45 tracking station at Canberra experienced an uplink power fluctuation just before 01:00, causing about 5 spurious manoeuvre pulses a one-way light time later. These spurious pulses caused the Earth aspect angle to move away from the deadband that triggers the conscan manoeuvre, and allowed the nutation level to grow undisturbed for two hours. At the end of this time the combination of EAA growth and nutation level exceeded the conscan deadband and started generating conscan pulses. Over the next 4 hours, the conscan manoeuvre successfully reduced nutation to under 0.1° . Had nutation not been reduced, other measures such as changing the conscan deadband or the manoeuvre shown in the following example would have been considered.

¹⁰ There are no good means to estimate the amplitude of the flexible boom oscillation, but it is expected to be significantly larger than the main spacecraft body nutation amplitude.

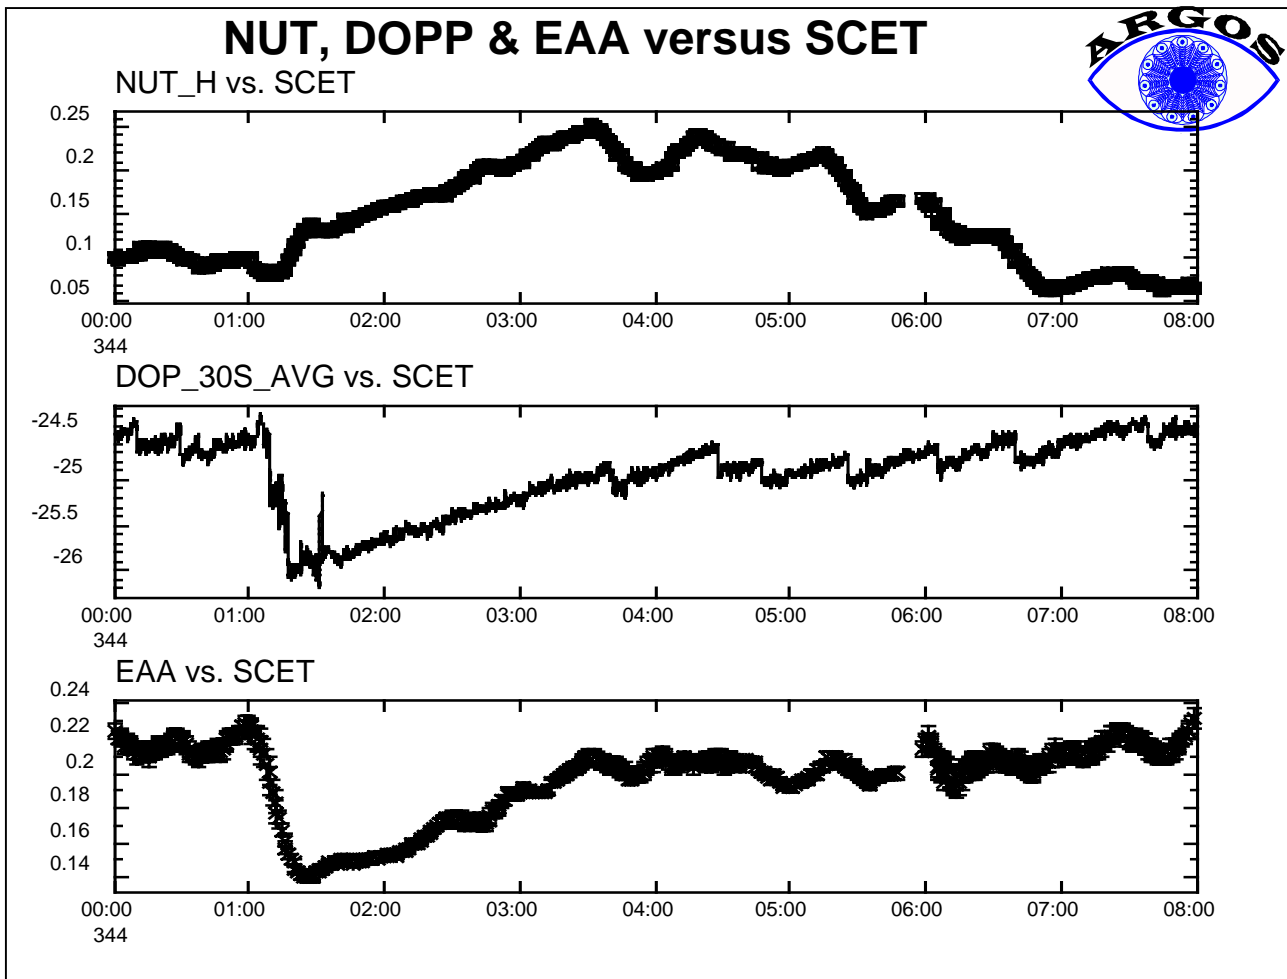


Figure 10: Exponential nutation growth

When nutation saturates conscan. During the peak periods, the nutation level showed a disturbing tendency to remain at levels of 0.3° to 0.5° despite the fact that the number of nutation control pulses increased as well. On the other hand, once the nutation was under 0.25° it was easy to bring it down further.

This situation occurs because the nutation growth is exponential, and the higher the value the faster it grows between pulses. In principle the manoeuvre pulses occur in response to the Earth drift rate, and do not become more frequent until the nutation is so large that it exceeds the size of the deadband, which is something that occurs just under 0.5° . At this point the nutation finds an equilibrium with the increased number of pulses, and it may stay trapped forever in a sawtooth pattern at high levels (0.3° to 0.5°). This also results in an increase in fuel consumption.

The solution to this problem is to artificially induce a series of consecutive conscan pulses to reduce the nutation below 0.25° . This can be done by changing from the wide to the narrow setting of the manoeuvre off-pointing deadband.

Once in the narrow deadband, the only known solution is to produce an artificial increment of the Earth drift rate by using an open-loop attitude manoeuvre to point the spacecraft away from the Earth, and then enable conscan to simultaneously reduce the Earth aspect angle and the nutation magnitude. This is called the "*McElrath*" manoeuvre after the author who first suggested it¹¹.

¹¹ Said author was amused to have finally had a wild idea that was actually useful (instead of dangerous) to spacecraft operations.

The author from the ESA Operations Team suggested that the open-loop part of the manoeuvre should contain precisely 12 pulses to avoid a nutation increase. According to a study conducted by ESA/ESTEC [12: Crellin 94], twelve pulses minimize the resonance between the open-loop manoeuvres and the main oscillation frequencies (nutation and MA).

Figure 11 is the ARGOS display for one of the “McElrath” manoeuvres. Before the manoeuvre the AGC data shows the characteristic nutation modulation pattern (in this scale the $fs-fn$ modulation is particularly visible). The nutation level varies around 0.35° , and the Earth aspect angle is well under the 0.125° deadband. During the open-loop manoeuvre the Earth Aspect Angle increases, while the nutation level decreases slightly. Then there is a short break between manoeuvres indicated by the 1-minute plateau in the Doppler residuals around 08:35. Finally, the conscan manoeuvre reduces the nutation below 0.15° and the Earth aspect angle to 0.1° . The attitude estimates smooth the actually sharper EAA and nh profiles of this manoeuvre, but its overall effect is very clear.

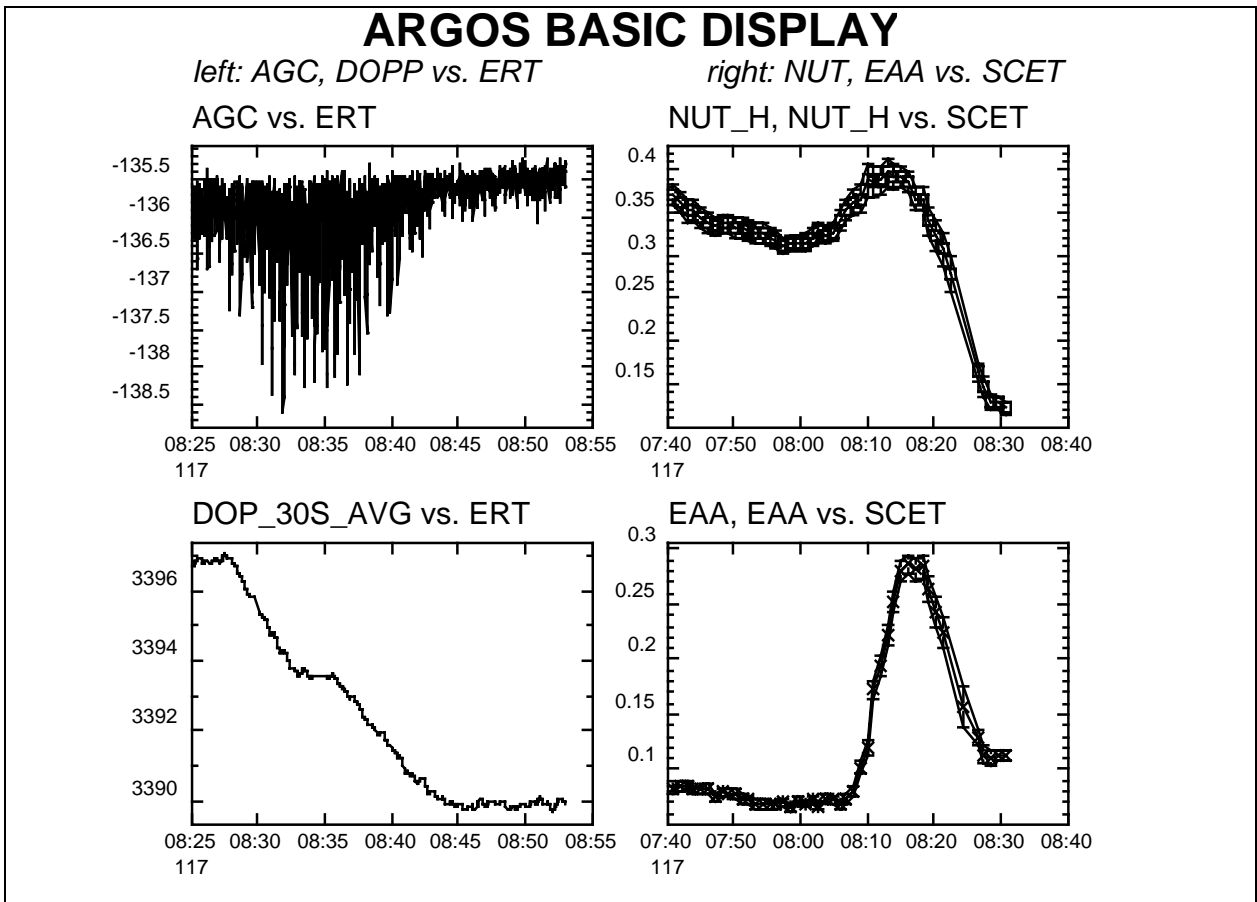


Figure 11: The “McElrath” manoeuvre as seen from the basic display.

6. CONCLUSION

The ARGOS system has developed into a capable realtime monitor and non-realtime analysis tool, producing attitude estimates and manoeuvre information without the need of spacecraft telemetry or any dedicated onboard subsystem.

ARGOS has been invaluable in observing the Ulysses nutation anomaly and clarifying the associated operational decisions. Operational modes that would have been conducted with uncertainty have been made routine, and the use of scarce tracking resources has been optimized. ARGOS has also provided a wealth of data for the analysis of the spacecraft dynamics conducted by the experts at ESA/ESTEC [11: Crellin & Janssens 94].

During contingencies that interrupt the telemetry (like the Ulysses spacecraft anomaly "DNEI", and some ground segment problems) ARGOS has been the only available spacecraft monitoring tool.

ARGOS is expected to continue to contribute after the end of the Ulysses nutation operations in 1995, due to its current or planned use by several other JPL projects. Ulysses will continue to use it as manoeuvre monitor and contingency tool. ARGOS will be especially vital when the Ulysses nutation returns even more strongly in the year 2001.

7. ACKNOWLEDGEMENTS

This work was conducted under the terms of the memorandum of understanding between NASA and ESA for the Ulysses Mission.

The work of the JPL authors was carried out at the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, under contract with the National Aeronautics and Space Administration (NASA).

The work of the ESA author was carried out as part of his duties within the ESA/Ulysses Flight Operations Team Office located at JPL, and as a member of the ESA/European Space Operations Center / Mission Operations Department (ESA/ESOC/MOD).

The authors would like to acknowledge the work of P. Mullen, D. Richardson, and others in organizing and implementing the ground system changes necessary for ARGOS. We would also like to take the opportunity to especially thank those who contributed to the understanding of the nutation anomaly, including: H. Hoffman, E. Crellin, F. Janssens, the ESOC Flight Dynamics Team and so many others.

The authors are appreciative of the many useful software libraries made available by members of JPL's Navigation Systems Section, including Gene Goltz, K.J. Lee and George Rinker, who also did some ARGOS-specific programming. The support and encouragement that the authors received from their respective management is gratefully acknowledged.

Mrs. Dorothy McElrath and Mrs. Eugenia García-Lizana are to be commended for tolerating the long hours that their respective husbands spent in developing ARGOS. In addition, Mrs. McElrath waded through the streams of obscure technical grammar in this paper, and offered many useful proofreading suggestions. Finally we thank M. Greutert for making her PC available to us.

REFERENCES:

- [1: Hoffman 90] H. Hoffman "The Ulysses nutation forcing function", Personal communication from NASA Goddard Space Flight Center to JPL.
- [2: Gienger *et al* 91] G. Gienger and all the members of Ulysses Flight Dynamics Team "Ulysses Flight Dynamics Report. Vol. 2: Post axial boom deployment nutation anomaly. Part I: Investigations and results" Issue I, ESOC/ECD/OAD, 24 June 1991.
- [3: CCSDS 91] "Parameter Value Language Specification," *Red Book*, June 1991.
- [4: Gordon *et al* 91] Gordon, H. J., Luthey, J. L., McElrath, T. P., Menon, P. R. "Ulysses Orbit Determination", *AAS/AIAA Astrodynamics Conference, AAS 91-472*, 21-24 August 1991.
- [5: García-Pérez 92] Raúl García-Pérez "Ulysses prime Mission Phase. Requirements to prevent the 1994-1995 Nutation Anomaly", internal report. ESOC/MOD/SMD Ulysses flight operations office at JPL, 1 August 1992.
- [6: García-Pérez 92] Raúl García-Pérez "Ulysses Log 1992" *Proceedings of the*

- second international symposium on Ground Data Systems for Spacecraft Operations. SPACEOPS 92. 16-20 November 1992. JPL publication 93-5, 1 March 1993.*
- [7: Stephens 92] S. A. Stephens, "An Analysis of FFT Tone Acquisition," JPL Interoffice Memorandum 335.1-92-14, May 14, 1992 (internal document).
- [8: McElrath *et al* 92] McElrath, T. P, Tucker, B., Criddle, K. E., Menon, P. R., Higa, E. S. "Ulysses Navigation at Jupiter Encounter," *AIAA/AAS Astrodynamics Conference, AIAA 92-4524*, 10-12 August 1992.
- [9: Crellin & Janssens 93] E. Crellin, F. Janssens "On the oscillatory motion of Ulysses following Axial antenna deployment." Internal communication, ESA/ESTEC/EWP 1720, June 1993.
- [10: García-Perez 93] Raúl García-Pérez "How to read nutation on the modulation of the Ulysses signal" Ulysses ESA at JPL internal communication. 6 November 1993.
- [11: Crellin & Janssens 94] E. Crellin, F. Janssens "Treatment of the AGC signal for Ulysses" Interoffice memorandum, ESA/ESTEC/WMM, 4 January 1994.
- [12: Crellin 94] E. Crellin "Response to thrusting for Ulysses", Interoffice memorandum, ESA/ESTEC/WMM, 20 January 1994.
- [13: Wesley 94] A. Wesley "TCL and the Tk Toolkit" ISBN 0-201-63337-X, April 1994.
- [14: Cangahuala 94] L. A. Cangahuala "ARGOS Adaptive Kalman Filter", JPL interoffice communication. 20 July 1994 (internal document).
- [15: McElrath *et al* 95] McElrath, T. P and Lewis, G. D. "Ulysses Orbit Determination at High Declinations," *Proceedings of the NASA/GSFC Flight Mechanics/Estimation Theory Symposium*, 16-18 May, 1995.