On-board data analysis techniques for space plasma particle instruments

D. W. Curtis, C. W. Carlson, and R. P. Lin

Space Sciences Laboratory, University of California, Berkeley, California 94720

G. Paschmann

Max Planck Institut für Extraterrestrische Physik, 8046 Garching b. München, Federal Republic of Germany

H. Rème and A. Cros

Centre d'Etude Spatiale des Rayonnements, Boite Postale 4346, 31029 Toulouse Cedex, France

(Received 4 April 1988; accepted for publication 1 November 1988)

This article describes the on-board data analysis techniques used on two recent space plasma particle instruments to compute meaningful parameters of three-dimensional plasma distributions in real time. These parameters may be transmitted with much higher time resolution than the full distribution within the limited telemetry. In addition, they greatly reduce the ground processing requirements.

INTRODUCTION

Space plasma particle instruments strive to measure three-dimensional velocity distributions with high energy, angle, and time resolution. A recent innovation in electrostatic analyzer design has greatly improved the ability to achieve this goal.[1] This analyzer can produce three-dimensional distribution functions with angular resolution down to a few degrees, covering an energy range from a few eV to 40 keV with $\delta E/E \sim 0.1$–0.2, in a one-half spin of the spacecraft (typically about 2 s). The amount of raw data generated by these instruments far exceeds the typical data rate allocation from the spacecraft telemetry system. For example, even a very coarse distribution, with 22.5° resolution and $\delta E/E = 0.3$, produces about 16 000 bits per second. However, since the plasma distribution functions can often be characterized by a smaller set of parameters, on-board microprocessors can be used to reduce the data volume by computing selected parameters from the measured distribution function in real time. The reduced data can be transmitted with high time resolution, together with snapshots or averages of the full distribution at lower time resolution. In addition to providing much better time resolution than would otherwise be available in a limited telemetry rate, production of physically meaningful parameters on-board greatly reduces the ground processing required to produce key parameters of the distribution function for synoptic plots and analysis.

This approach has been used successfully on two spacecraft missions in the last few years. On the Active Magnetospheric Particle Tracer Experiment/Ion Release Module (AMPTE/IRM) spacecraft the plasma instrument consists of two of the three-dimensional analyzers—one for electrons and one for ions.[2] On the Giotto Halley's Comet mission, the Rème Plasma Analyzer includes a single electron analyzer.[3] Both missions employ dedicated microprocessor systems to compute parameters of the measured distribution function which are telemetered with maximum time resolution. A different set of parameters was selected for each of these two missions, optimized for the types of plasmas expected to be observed and the available telemetry constraints. Future instruments might have a large repertoire of possible computations that could be selected by ground command or by an on-board algorithm as plasma conditions change.

Moments of the distribution function are computed in the AMPTE instrument: Density $= \int f(v) d^3v$; number flux vector $= \int f(v) v\, d^3v$; momentum flux tensor $= \int f(v) vv\, d^3v$; and energy flux $= \int f(v) vv\, d^3v$. These integrals are approximated by summing products of measured count rates with appropriate energy/angle factors over the sampled distribution. Moments provide a complete characterization of the distribution with a few numbers if the distribution is Maxwellian. For more complex distributions, moments still provide meaningful physical parameters such as density, temperature, bulk velocity, and heat flux which are directly related to plasma processes.

Often a fast electron plasma distribution can be characterized by a pitch angle distribution (PAD)—that is, the distribution has an axis of symmetry (the magnetic field direction). The measured distribution will have two-dimensional symmetry if the electron distribution is gyrotropic, and the bulk flow velocity of the plasma in the spacecraft frame is much less than the velocity of the detected particles. The three-dimensional distribution can then be reduced to a two-dimensional distribution (pitch angle and energy) without loss of information. The PAD requires far less telemetry space than the full distribution, so that it can be transmitted with high time resolution. The Giotto instrument computes the PAD of the electrons with an on-board processor, in real time. The symmetry direction of the distribution used to sort the data into a PAD is computed directly from the electron data, which gives an independent measurement of the magnetic field direction.

The hardware for both instruments was designed under the constraints of high reliability, low power, and ionizing radiation tolerance. In addition, these instruments were designed between 1979 and 1982, so that they do not represent the current state of the art in parts selection. Significant performance improvements could be achieved by substitution of the current generation of radiation hardened Complementary Metal-Oxide-Semiconductor (CMOS) microprocessor

 0034-6748/89/030372-09$01.30

technology (for example, the Harris 80C86RH and the soon-to-be-available Sandia/National 32C016). However, the speed increases are accompanied by increased power consumption, so that the speed-power product has not been greatly improved. Thus, these designs are still near the state of the art within the given power constraints.

I. AMPTE IRM 3D PLASMA INSTRUMENT MOMENT COMPUTATIONS

A. Hardware

The AMPTE instrument employs a set of four CMOS microprocessors (NSC800s running at 2 MHz) for its data system (Fig. 1). The main microprocessor controls the analyzers and three slave microprocessors, and interfaces with the spacecraft telemetry system. The analyzers are swept through their energy range 16 times per rotation of the spacecraft (4 s). The eight accumulators for each analyzer are read out 30 times per energy sweep, giving a distribution with 128 angular samples covering 4π steradians with a 22.5° resolution at each of the 30 energy steps during every rotation for each analyzer. The sweep rate and sampling time are automatically adjusted to the measured spin rate to maintain synchronization with the spin phase. The main processor transfers the eight accumulator outputs of each analyzer to the slave processors as they are collected (one slave gets the electron data and one slave gets the ion data). The transfer is made by holding the slave processor reset and copying the data directly into the slave processor memory. The slave processor is then released to do computations on the data. At the end of a spin, the main processor then retrieves the results from the slave processors. Note that the slave processor has no ROM (read-only memory). The slave processor program is loaded by the main processor at initial turn-on in the same way that data are transferred during operations. This allows the slave processor to be easily reprogrammed for different computations or calibration factors. All of the memory in the data system is maintained in a standby state by a contin-

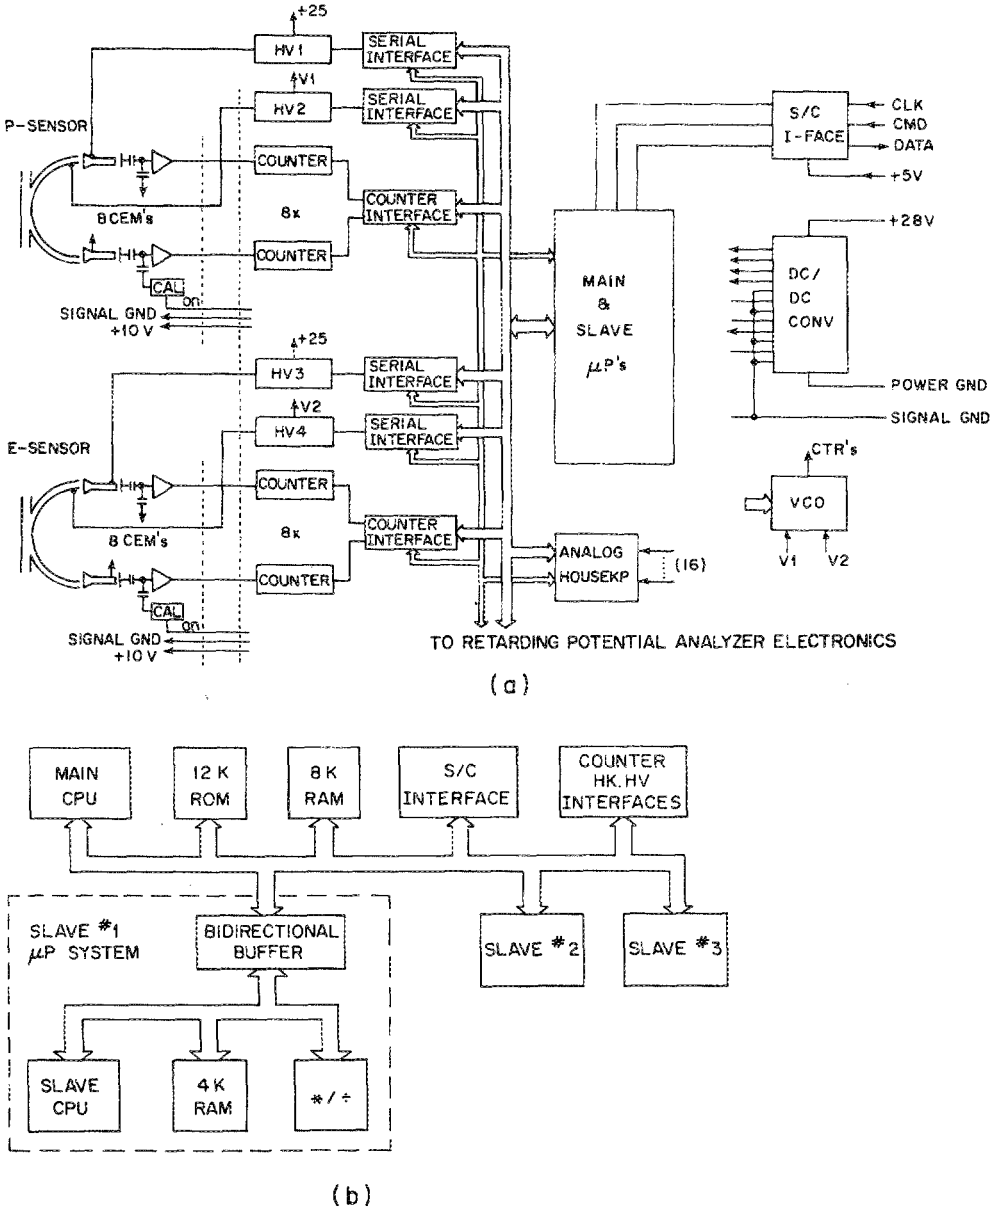

FIG. 1. AMPTE/IRM 3D plasma analyzer. (a) Block diagram of instrument showing sensors, sensor interfaces, processors, and spacecraft interface. (b) Block diagram showing details of the microcomputer system.

uous power source provided by the spacecraft when the rest of the instrument is turned off. This avoids the problem of having to reload program modification from the ground every time the power to the instrument is cycled off and on. All critical areas in the memories are check-summed periodically by the main processor to protect against memory upsets.

The slave processors also have dedicated CMOS multiply/divide units (two RCA CDP 1855s) incorporated to speed computations (the NSC800 has no multiply or divide instructions, and given the large number of multiplies required, a software multiply routine would be too slow). Some of the high-order address lines of the slave processor are used to control the multiply/divide unit so that data can be transferred between it and memory directly, rather than going through the microprocessor. This further improves the performance of the system for computations. The multiply units take about 17 μs to complete a 16×16 bit multiply once the coefficients are loaded. During this time the processor is free to do other tasks. Note that most of the next generation of processors (e.g., 80C86, 32C016) have multiply and divide instructions, eliminating the need for this hardware.

Three slave processors were included in the instrument. Two slaves are used to compute parameters for the two species of particles, and the third slave controls a peripheral instrument. The third slave is configured like the two computation slaves to provide a level of redundancy.

B. Software

The slave processors on AMPTE are used to compute moments of the distribution function—density, number flux vector, momentum flux tensor, and energy flux vector (Table I). During the data sampling interval, the slave processor takes the eight count rates transferred to it by the main processor and computes the corresponding moment increments. The increments are then summed into the moment registers. These operations must be completed by the start of the next sampling interval in order to allow the computations to keep up with real time.

The data and the computations have a wide dynamic range, but only 16 bits of precision are available in arithmetic operations. Doing all of the computations in floating point would be too time consuming. Instead, the eight counter values are normalized to a common exponent, and all computations up to summing the moment increments into the moment registers are done in 16-bit fixed point format. The moment increments are then shifted back into long integers and summed into the 40-bit-long moment registers. The counter interface hardware is designed to provide the counter data in both integer counts format and normalized (common exponent fixed point) format to simplify this processing.

In the following discussion, the coordinate system used is despun spacecraft coordinates, with the z axis parallel to the spacecraft spin vector and the x axis corresponding to the detector look direction at the beginning of each spin (at sun pulse time). The angle ϕ is the spin phase, and θ is the elevation (Fig. 2). All of the trigonometric functions are stored in look-up tables. The ϕ angle is a function of energy step and sweep number, and has 480 values. In order to minimize the size of the ϕ function tables, only the first quadrant of $\langle\cos(\phi)\rangle$, $\langle\cos(\phi)\sin(\phi)\rangle$, and $\langle\cos^2(\phi)\rangle$ functions are stored—the other quadrants and functions can be derived from these by symmetry considerations (the $\langle\ \rangle$ notation refers to the average over the $22.5^\circ\times22.5^\circ$ phase-space sample). The energy coefficients are also stored in look-up tables, in normalized format.

The first step is to multiply each data sample by the efficiency of the detectors, and a cosine factor to normalize the phase-space sampling. In addition, the following dead time correction is applied:

$$C(\Theta)=\frac{CR_n(\Theta)\,[\text{efficiency}\,(\Theta)\langle|\cos(\theta)|\rangle]}{CR_{\max}-CR_i(\Theta)},$$

where Θ is the channeltron number (0–7); θ is the polar angle corresponding to channeltron number Θ; $CR_i(\Theta)$ is

TABLE I. Moment definitions. To within a multiplicative factor dependent on the analyzer geometric factor, the moments are given by the following sums.

Density:

$$N=\sum_E\frac{1}{V}(E)\sum_\phi\sum_\theta C(\theta,\phi,E)$$

Bulk velocity:

$$NV_x=\sum_E\sum_\phi\cos(\phi)\sum_\theta\cos(\theta)C(\theta,\phi,E)$$

$$NV_y=\sum_E\sum_\phi\sin(\phi)\sum_\theta\cos(\theta)C(\theta,\phi,E)$$

$$NV_z=\sum_E\sum_\phi\sum_\theta\sin(\theta)C(\theta,\phi,E)$$

Pressure tensor:

$$P_{xx}=\sum_E V(E)\sum_\phi\cos^2(\phi)\sum_\theta\cos^2(\theta)C(\theta,\phi,E)$$

$$P_{yy}=\sum_E V(E)\sum_\phi\sin^2(\phi)\sum_\theta\cos^2(\theta)C(\theta,\phi,E)$$

$$P_{zz}=\sum_E V(E)\sum_\phi\sum_\theta\sin^2(\theta)C(\theta,\phi,E)$$

$$P_{xy}=\sum_E V(E)\sum_\phi\cos(\phi)\sin(\phi)\sum_\theta\cos^2(\theta)C(\theta,\phi,E)$$

$$P_{xz}=\sum_E V(E)\sum_\phi\cos(\phi)\sum_\theta\cos(\theta)\sin(\theta)C(\theta,\phi,E)$$

$$P_{yz}=\sum_E V(E)\sum\sin(\phi)\sum_\theta\cos(\theta)\sin(\theta)C(\theta,\phi,E)$$

Heat flux vector:

$$H_x=\sum_E V^2(E)\sum_\phi\cos(\phi)\sum_\theta\cos(\theta)C(\theta,\phi,E)$$

$$H_y=\sum_E V^2(E)\sum\sin(\phi)\sum_\theta\cos(\theta)C(\theta,\phi,E)$$

$$H_z=\sum_E V^2(E)\sum_\phi\sum_\theta\sin(\theta)C(\theta,\phi,E)$$

where

E is the energy step number (1–30) covering an energy range from about 20 eV to 40 keV;

θ is the polar angle ($-90^\circ\pm90^\circ$) covered by the eight detectors of the analyzer (22.5° resolution);

ϕ is the spin phase angle (0°–360°), covered by sixteen samples at each energy during a spin (22.5° resolution);

$V(E)$ is the velocity corresponding to the energy step E (including an energy-dependent efficiency factor); and

$C(\theta,\phi,E)$ is the dead time- and detector efficiency-corrected count rate

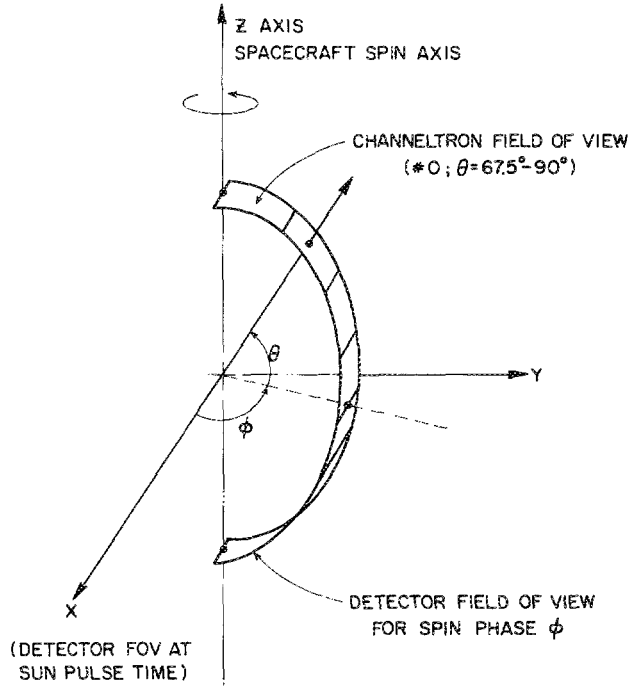

FIG. 2. Schematic of AMPTE/IRM 3D plasma instrument field of view and coordinates system.

the integer counter output; $CR_n(\Theta)$ is the normalized counter output, where the normalization is to a common exponent (power of 2) for all θ to improve the dynamic range of calculation; $CR_{\max}$, a constant, is the sample time interval divided by the detector dead time; efficiency (Θ) is the channeltron efficiency factor, normalized such that the maximum is 1.0; and the product [efficiency $(\Theta)\langle|\cos(\theta)|\rangle$] is stored in a look-up table. The $\langle|\cos(\theta)|\rangle$ term corrects for the oversampling of phase space by the analyzer near the poles. This operation requires eight multiplies and eight divides (one for each channeltron).

Next, the following set of sums is made across the eight samples:

$$S_0 = \sum_{\Theta} C(\Theta),$$

$$S_1 \sum_{\Theta} C(\Theta)\frac{\langle|\cos(\theta)|\cos(\theta)\rangle}{\langle|\cos(\theta)|\rangle},$$

$$S_2 = \sum_{\Theta} C(\Theta)\frac{\langle|\cos(\theta)|\sin(\theta)\rangle}{\langle|\cos(\theta)|\rangle},$$

$$S_3 = \sum_{\Theta} C(\Theta)\frac{\langle|\cos(\theta)|\cos^2(\theta)\rangle}{\langle|\cos(\theta)|\rangle},$$

$$S_4 = \sum_{\Theta} C(\Theta)\frac{\langle|\cos(\theta)|\sin^2(\theta)\rangle}{\langle|\cos(\theta)|\rangle},$$

$$S_5 = \sum_{\Theta} C(\Theta)\frac{\langle|\cos(\theta)|\cos(\theta)\sin(\theta)\rangle}{\langle|\cos(\theta)|\rangle}.$$

The θ function products are stored in look-up tables. The symmetries of the trigonometric functions are utilized to minimize the number of multiplies required. In addition, the functions are normalized such that their maximum is 1.0, again reducing the number of multiply operations required. (This normalization is removed as part of the ground-based processing.) Only 15 multiply operations are required for this step (each of the 5 sums is over 8 channeltrons; pairs of count rates are summed or differenced by symmetry across $\theta = 0$ before multiplying, leaving only 4 products per sum; the θ function coefficient normalization makes one of the coefficients in each sum 1.0, leaving three multiples for each of 5 sums, or 15 multiplies).

Next the ϕ angle functions and (normalized) velocity coefficients are applied to produce the partial moments for energy step E, spin sector ϕ, as follows:

$$dN = S_0[\text{efficiency}(E)/V(E)],$$
$$dNV_x = \langle\cos(\phi)\rangle[S_1\text{efficiency}(E)],$$
$$dNV_y = \langle\sin(\phi)\rangle[S_1\text{efficiency}(E)],$$
$$dNV_z = S_2\text{efficiency}(E),$$
$$dP_{xx} = \langle\cos^2(\phi)\rangle\{S_3[V(E)\text{efficiency}(E)]\},$$
$$dP_{yy} = \langle\sin^2(\phi)\rangle\{S_3[V(E)\text{efficiency}(E)]\},$$
$$dP_{zz} = S_4[V(E)\text{efficiency}(E)],$$
$$dP_{xy} = \langle\sin(\phi)\cos(\phi)\rangle\{S_3[V(E)\text{efficiency}(E)]\},$$
$$dP_{xz} = \langle\cos(\phi)\rangle\{S_5[V(E)\text{efficiency}(E)]\},$$
$$dP_{yz} = \langle\sin(\phi)\rangle\{S_5[V(E)\text{efficiency}(E)]\},$$
$$dH_x = \langle\cos(\phi)\rangle\{S_1[V^2(E)\text{efficiency}(E)]\},$$
$$dH_y = \langle\sin(\phi)\rangle\{S_1[V^2(E)\text{efficiency}(E)]\},$$
$$dH_z = S_2[V^2(E)\text{efficiency}(E)],$$

where E is the energy step number (0–29); efficiency (E) is the energy-dependent efficiency factor of the analyzer (nominally 1.0); and $V(E)$ is the velocity corresponding to energy step E. The energy function values—efficiency $(E)/V(E)$, efficiency(E), $V(E)$efficiency(E), and $V^2(E)$efficiency(E) are stored in look-up tables. Each energy function is normalized to a maximum value of 1.0, and then converted to an exponent (power of 2) and mantissa. The mantissa is used in this step, and the exponent is incorporated in the denormalization below. This requires 17 multiply operations. Finally, these moment increments are denormalized by the sum of the count rate normalization and the energy coefficient normalization exponents, and summed into the 40-bit moment registers.

This total of 40 multiplies, 8 divides, and 13 denormalization operations, together with many adds, table look-ups, etc., must be completed in one sector time interval of 1/512th of a spin. The nominal spin rate is 15 rpm, but significantly faster spin rates must be accommodated in case of a failure in the spin rate control of the spacecraft. The design goal was to be able to compute moments up to at least 20 rpm, or about 6 ms per sector. This requires an average of about 10 000 16-bit multiplies and a similar number of adds per second. This goal was met by coding entirely in assembly language, by minimizing overhead with the use of programming techniques such as heavy use of the processor's registers rather than memory, and by expanding some of the loops into straight line code.

Note that no approximations have been made in computing the moments as defined in Table I. Computations are typically done to 12-bit accuracy or better, which is more than the statistical significance of the data. The results are as

accurate as could be done on the ground, with far better time resolution than if only raw data were transmitted.

Figure 3 shows a plot made of selected moments computed on the AMPTE instrument in flight, and parameters derived from the moments.

II. GIOTTO RPA EXPERIMENT PITCH ANGLE SORTER

A. Hardware

The Rème Plasma Analyzer (RPA) on the Giotto spacecraft includes a mass analyzer (RPA2-PICCA) and a three-dimensional electron plasma analyzer (RPA1-EESA). The data system consists of three microprocessors: a main processor that interfaces with the spacecraft, the Electron Electrostatic Analyzer (EESA) instrument, and the other two processors; a slave processor to control the Positive Ion Cluster Composition Analyzer (PICCA) instrument; and a slave processor dedicated to doing on-board pitch angle distribution (PAD) sorting of the data (Fig. 4). The PICCA instrument and processor will not be discussed here.

The EESA analyzer is swept through its energy range 16 times per 4-s rotation of the spacecraft. Since the analyzer measures particles in a 360° fan, a full distribution is sampled every half rotation. The 16 accumulators are each read out 39 times per energy sweep, giving a measured distribution with 128 angular samples covering 4π steradians with a 22.5° resolution, at 39 energy steps during every half rotation. The main processor stores the data for 27 selected energies into a special 8192-byte memory during a half rotation. At the end of the half rotation, the main processor resets the PAD processor, and swaps the main processor's 8192–byte memory with the PAD's 8192-byte memory. This is accomplished by special hardware that multiplexes the address, data, and control signals for the two memories between the two processors such that while the main processor is addressing one of the memories, the PAD processor is addressing the other memory. This scheme places very little overhead on the main and PAD processors to transfer the data, and al-

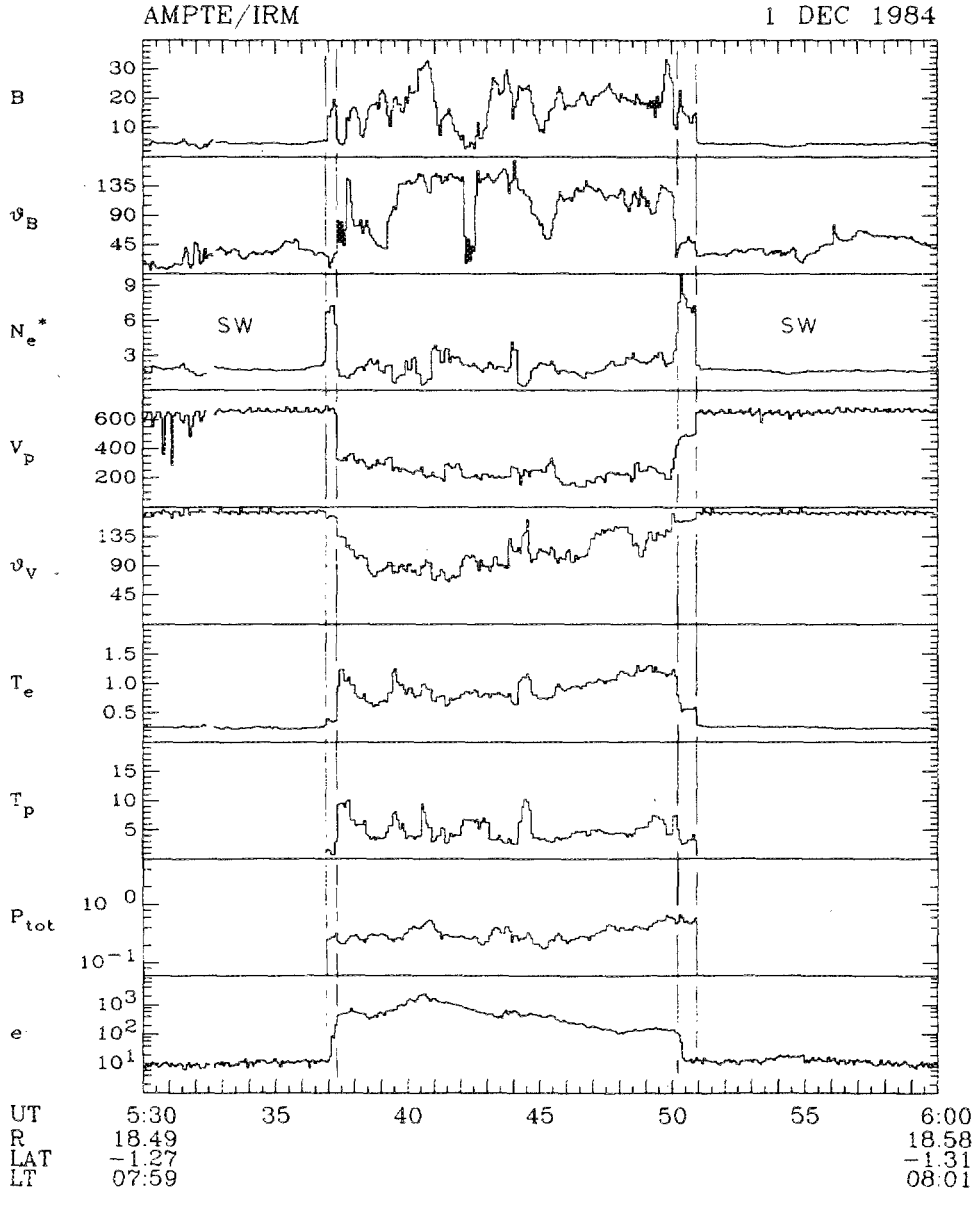

FIG. 3. 30 min of magnetic field and particle data from the AMPTE/IRM spacecraft. Particle data are derived from the moments computed on-board. The data show the occurrence of a region of hot slowly moving plasma, surrounded by cold, high speed solar wind (SW). Note the high time resolution in the plasma data.

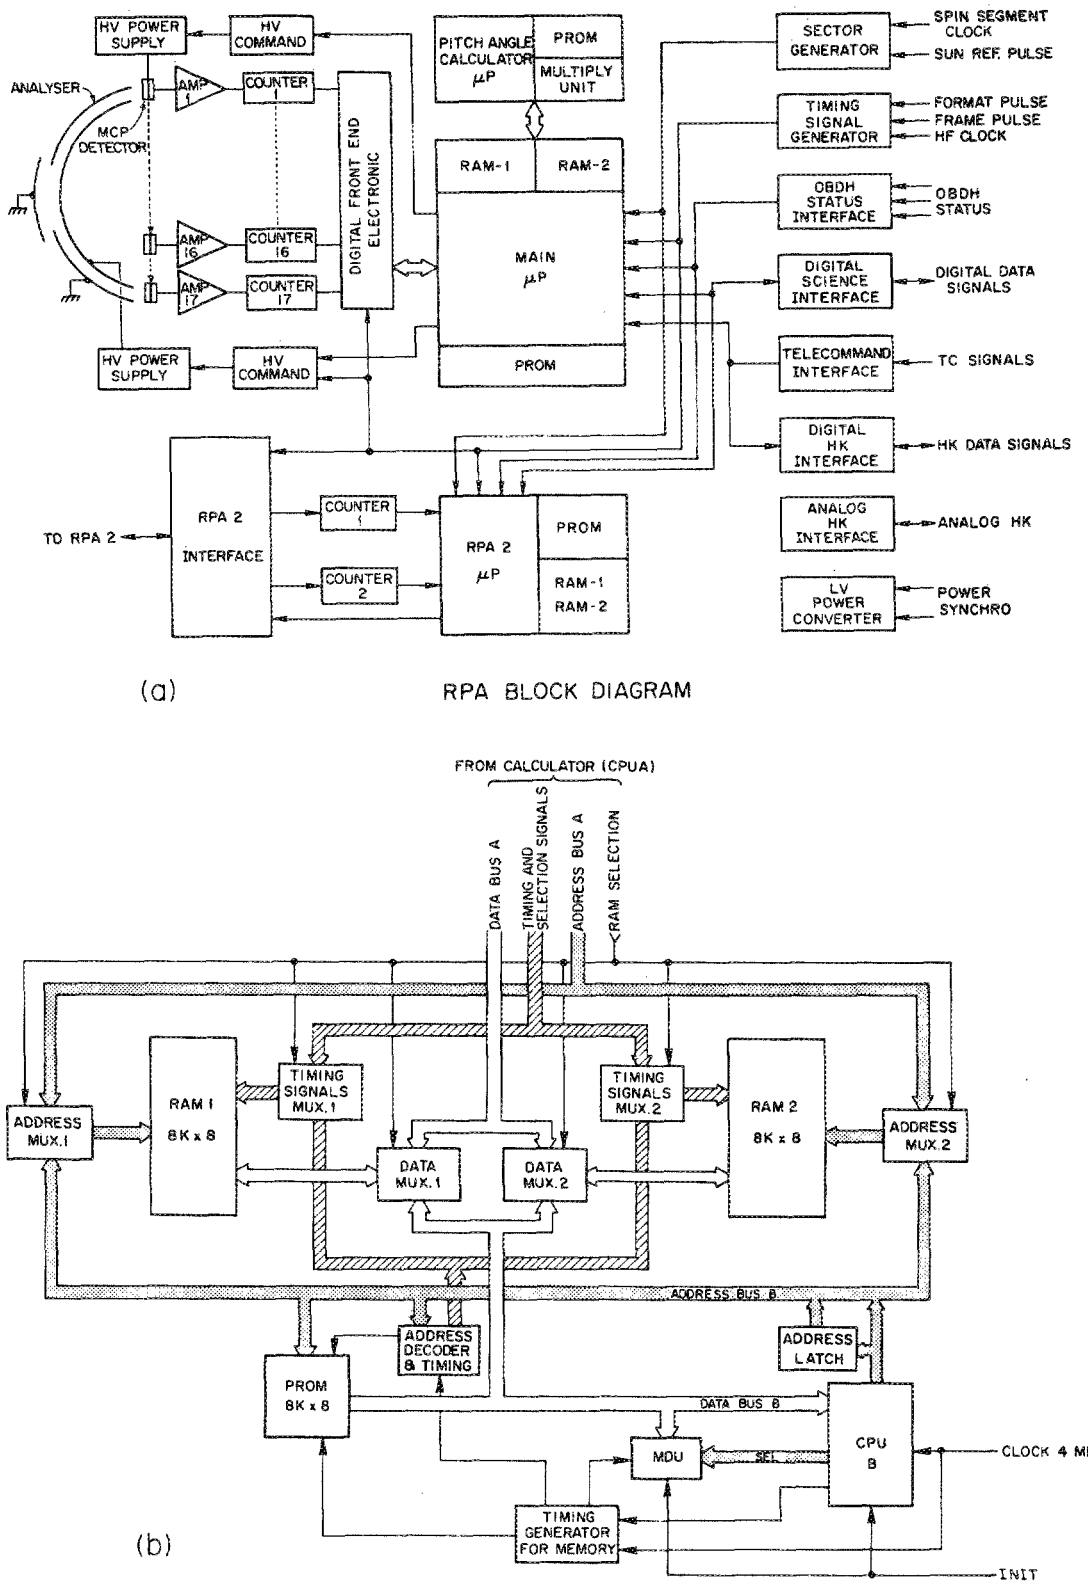

FIG. 4. Giotto RPA instrument. (a) Block diagram of instrument showing analyzer, analyzer interface, processors, and spacecraft interface. (b) Block diagram showing details of pitch angle calculator and "swap-RAM" system.

lows the PAD processor simultaneous access to a complete distribution for computations. The PAD processor places the results of its computations back into the special swap memory for the main processor.

The PAD processor consists of an NSC800 CMOS microprocessor running at 2.5 MHz, together with a hardware multiply/divide unit (two RCA CDP 1855s), 4096 bytes of ROM for program and constants, and the swap memory. The processor has no local RAM (random-access memory)—it stores all intermediate results in the swap memory. No provision was made for in-flight program modification because of the single pass nature of the mission, and to avoid the necessity of reloading the processor in case of a power failure. As in the AMPTE slave processors, the high order address lines were used to control the multiply/divide unit to improve the multiplier through-put.

B. Software

In order to compute the PAD, the symmetry direction of the distribution must be available. If the magnetic field direction were available on-board in real time, it could be used to sort the distribution into pitch angle bins. Unfortunately, the magnetometer on the Giotto spacecraft is not able to provide this information on-board because of large offsets and perturbations caused by the lack of a magnetometer boom and a magnetically noisy spacecraft. For this reason, the symmetry direction of the distribution is calculated directly from the electron data. This method has the advantage of providing an independent measurement of the magnetic field direction (but not its magnitude), which is useful in verifying that the magnetometer offsets had been removed correctly on the ground.

The symmetry direction is derived from the pressure tensor of the measured distribution. This tensor should have two degenerate eigenvalues if the distribution has an axis of symmetry. The nondegenerate eigenvector is then the symmetry direction of the distribution. In fact, the measured distribution will be nonsymmetrical at low energies due to the bulk flow of the plasma in the spacecraft frame and photoelectrons from the spacecraft, while at high energies there are problems due to low statistics and background. To avoid these problems, the pressure tensor is computed over a restricted energy range. In fact, two pressure tensors are computed: one covering the energy range from about 80 to 300 eV, and the second from 300 to 3600 eV. The distribution below 300 eV is sorted using the symmetry direction of the low-energy pressure tensor, while the distribution above 300 eV is sorted using that of the high-energy pressure tensor.

Once a symmetry direction is computed, the data are sorted into pitch angle bins by the dot product of the symmetry direction and the look direction of the analyzer for each sample. The inverse cosine of this product is the pitch angle. The data are sorted into eight pitch angle bins with 22.5° resolution, which is the same as the data sample resolution, so that a typical sample covers part of two adjacent pitch angle bins. The sample is split between the two bins proportionally to the fraction of the sample overlapping each bin. A second set of bins is used to accumulate the sample area summed into each pitch angle bin. These are used to normalize the final pitch angle sums.

The data transferred to the PAD processor by the main processor have already been dead-time corrected by the counter hardware using a look-up table and summing circuit. The first step of the PAD processing is to reduce the number of energy samples from 27 to 10 or 16 (depending on the telemetry format) by summing together adjacent energy samples. Next the data are normalized to a fixed point, 16-bit mantissa plus a common exponent form. The exponent is common to all samples of the same energy step, and is selected to maximize the dynamic range of the computations while avoiding overflows during the pressure tensor and the pitch angle sorting sums over angles. The exponent selection is based on the magnitude of the sum of all samples at each energy. The following detector efficiency and phase space sampling factor is then applied to the data:

$$C(E',\Theta,\Phi) = 2^{N(E')}[\text{efficiency }(\Theta)\langle|\cos(\theta)|\rangle] \times \sum_{E}^{E'} CR(E,\Theta,\Phi),$$

where E is the energy step number (0–26); E' is the reduced energy step number (0–9 or 1–15); Θ is the detector sector number (0–15); θ is the polar angle corresponding to detector sample Θ; Φ is the sweep number (0–7); $N(E')$ is the normalizing exponent for energy step E' (an integer); $CR(E,\Theta,\Phi)$ is the input count rate; and $C(E,'\Theta,\Phi)$ is the normalized count rate. The factor [efficiency (Θ) $\langle|\cos(\theta)|\rangle$] is contained in a look-up table, and is normalized to a maximum value of 1.0. These computations require 4096 multiplies and some shifting.

The pressure tensor is computed in three steps, similar to the AMPTE moment computation:

$$S_1(E',\Phi) = \sum_{\Theta} \frac{\langle\cos^2(\theta)|\cos(\theta)|\rangle}{\langle|\cos(\theta)|\rangle} C(E',\Theta,\Phi),$$

$$S_2(E',\Phi) = \sum_{\Theta} \frac{\langle\sin^2(\theta)|\cos(\theta)|\rangle}{\langle|\cos(\theta)|\rangle} C(E',\Theta,\Phi),$$

$$S_3(E',\Phi) = \frac{\Sigma_\Theta\langle\sin(\theta)\cos(\theta)|\cos(\theta)|\rangle}{\langle|\cos(\theta)|\rangle} C(E',\Theta,\Phi).$$

The symmetry of the trigonometric functions is used to reduce the number of multiplies required. For each E' and Φ, 12 multiplies are required, for a total of 1440 per distribution. Next,

$$dP_{xx} \sum_{\Phi} \langle\cos^2(\phi)\rangle S_1(E',\Phi),$$

$$dP_{yy} = \sum_{\Phi} \langle\sin^2(\Phi)\rangle S_1(E',\Phi),$$

$$dP_{zz} = \sum_{\Phi} S_2,$$

$$dP_{xy} = \sum_{\Phi} \langle\sin(\phi)\cos(\phi)\rangle S_1(E',\Phi),$$

$$dP_{xz} = \sum_{\Phi} \langle\cos(\phi)\rangle S_3(E',\Phi),$$

$$dP_{yz} = \sum_{\Phi} \langle\sin(\phi)\rangle S_3(E',\Phi),$$

where dP_{ij} is the i,jth partial (delta) pressure tensor element for energy step E', and ϕ is the azimuthal angle (spin phase) corresponding to energy step number E' and energy sweep number Φ. All of the trigonometric functions are stored in look-up tables. The ϕ angle can take on 128 values, so interpolation is used between a limited number of samples of the ϕ functions to reduce the size of the look-up tables. In addition, only the first quadrant of the functions is stored, the rest being obtained by symmetry relations. These computations require a total of 600 multiplies.

Next,

$$P_{ij} = \sum_{E'} V_n(E') dP_{ij} 2^{V_e(E') - N(E')},$$

where P_{ij} is the i,jth pressure tensor element; $V(E')$ is the velocity corresponding to the median energy of step E'; $V_n(E')$ is the fractional part of $V(E')$; and $V_e(E')$ is the

exponent of $V(E')$ (an integer). The P_{ij} terms are summed into 32-bit fixed point registers. These computations require a total of 90 multiplies and 90 shifts.

Under the assumption that the distribution is two-dimensional, the nondegenerate eigenvector of the pressure tensor is computed:

$$P_{nn} = P_{nn} - (P_{xx} + P_{yy} + P_{zz})/3 \quad (n = x,y,z), \tag{1}$$

$$a = [P_{xx}(P_{xx} - P_{yy}) + P_{yy}(P_{yy} - P_{zz}) + P_{zz}(P_{zz} - P_{xx})]/9 + (P_{xy}^2 + P_{yz}^2 + P_{xz}^2)/3, \tag{2}$$

$$b = P_{xx}P_{yz}^2 + P_{yy}P_{xz}^2 + P_{zz}P_{xy}^2 - P_{xx}P_{yy}P_{zz} - 2P_{xy}P_{yz}P_{xz}, \tag{3}$$

$$\alpha = 2\,\mathrm{sign}(b)\sqrt{(a)}, \tag{4}$$

$$P_{nn} = P_{nn} + \alpha, \tag{5}$$

$$\text{Find } i \text{ such that } |P_{ii}| = \text{minimum}(|P_{nn}|), \tag{6}$$

$$B_i = P_{jj}P_{kk} - P_{jk} \quad (i,j,k \text{ cyclic}),$$

$$B_j = P_{ik}P_{jk} - P_{ij}P_{kk},$$

$$B_k = P_{ij}P_{jk} - P_{ik}P_{jj}, \tag{7}$$

$$B_n = B_n/\sqrt{(B_x^2 + B_y^2 + B_z^2)}. \tag{8}$$

B is the symmetry axis of the distribution function.

These calulations cannot be easily done in fixed point, but they need be done only twice per distribution (once for the high-energy part of the distribution, and once for the low energy part). A floating point library was implemented using the hardware multiplier. The numbers are stored as a 16-bit mantissa, a 7-bit signed exponent, and a sign bit. The square-root function is computed with a look-up table. Because ROM program space is limited, and the calculation time for this step is an insignificant fraction of the total calculation cycle, the computations are made by a custom designed software interpreter. Each operation is coded in a single byte containing a function code and a parameter address. Reverse Polish Notation (RPN) is used, with a 16-parameter memory and a stack of intermediate values. Functions included "sum parameter to top of stack," etc. These computations require a total of 116 floating point operations.

The final step of the computation is to sort the distribution function into pitch angle bins about the symmetry direction B. First, the dot product of B with the look direction of each sample of the distribution is taken:

$$B\cdot V = B_z\langle\cos(\theta)\rangle + [B_x\langle\cos(\phi)\rangle + B_y\langle\sin(\phi)\rangle]\langle\sin(\theta)\rangle,$$

where V is the look direction of the sample. The terms $B_z\langle\cos(\theta)\rangle$, $B_x\langle\sin(\theta)\rangle$, and $B_y\langle\sin(\theta)\rangle$ need be computed only once for four values of the θ functions (the others are related by symmetry). The two products with ϕ functions are done for each of the ϕ samples. This requires a total of 264 multiplies per distribution.

The computed dot product is used to index a look-up table. This table gives the bin number to sum the sample into, and the fraction of the sample to sum into that bin. The remainder of the sample is summed into the adjacent sample:

$$\mathrm{PAD}(\alpha,E') = \frac{\Sigma_{\Theta,\Phi}\mathrm{OF}[\alpha,B\cdot V(\Theta,\Phi)]C(E',\Theta,\Phi)}{16\Sigma_{\Theta,\Phi}\mathrm{OF}[\alpha,B\cdot V(\Theta,\Phi)]\langle|\cos(\theta)|\rangle},$$

where α is the pitch angle bin number; OF is the overlap function, which gives the fraction of sample Θ, Φ that should be summed into bin α; and PAD is the two-dimensional pitch angle distribution. Note the denominator, which normalizes the PAD by the sample area summed into each bin. This normalization gives the PAD in counts per $22.5\times22.5°$ sector. These computations require a total of 3840 multiplies.

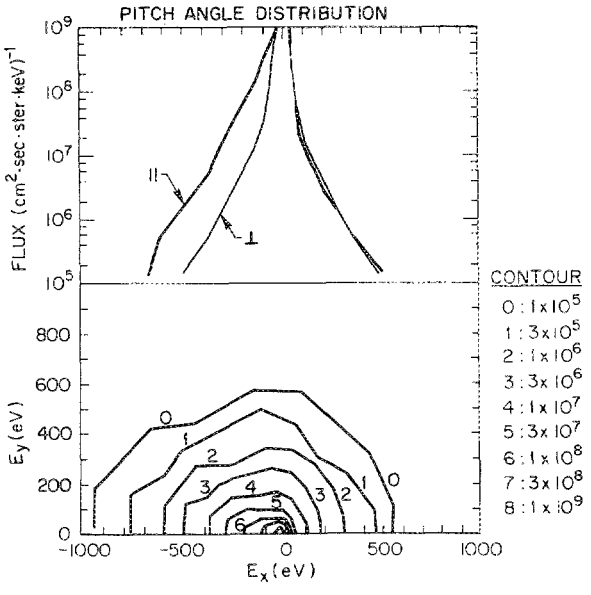

FIG. 5. Sample plot showing a pitch angle distribution computed on-board Giotto RPA instrument. Lower panel shows contours of constant flux as a function of energy. Upper panel shows cuts through the distribution along the symmetry axis (∥) and perpendicular to the symmetry axis (⊥).

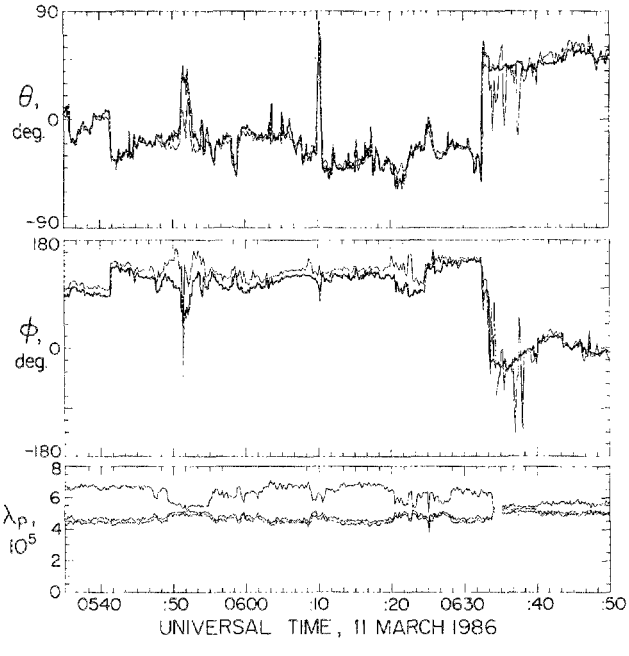

FIG. 6. Example of a comparison between magnetic field direction, as measured by the Giotto magnetometer (heavy) and the computed symmetry direction of the electron pressure tensor. The third panel shows the eigenvalues of the pressure tensor matrix (arbitrary units).

The whole computation requires over 10 000 fixed point multiplies, plus a number of normalizations, floating point computations, etc. These computations must be completed in a half rotation period, which is 2 s. Running the processor at 2.5 MHz, this was accomplished with a 10% margin. In order to achieve this goal, the code was written entirely in assembly language, and carefully optimized for speed and compactness.

C. In-flight results

Figure 5 shows a typical solar wind electron pitch angle distribution. The lower panel shows a contour map of the distribution function, while the upper panel shows three cuts through the distribution: two along the symmetry direction (marked with $\|$), and one perpendicular to the symmetry direction (marked with a $\perp$—repeated over the other two for reference).

Figure 6 shows a comparison of the computed symmetry direction and the magnetic field direction measured by the magnetometer experiment. Note that when all three eigenvalues are close to the same value, indicating a near isotropic distribution, the symmetry direction is poorly defined. The small systematic offset of the symmetry direction from the magnetic field direction indicated on this figure is probably caused by variations in the efficiency across the detector which were not included in the on-board detector efficiency correction factors. Magnetic field data are supplied courtesy of Dr. F. M. Neubauer of the Institut für Geophysik und Meterologie der Universität Cologne, Germany.

ACKNOWLEDGMENTS

This work was supported in the U. S. by NASA under Contract No. NASW-3575 and Grants Nos. NAGW-1298 and NGL-05-003-017, and in Europe by CNES under Grant No.1212.

[1]C. W. Carlson, D. W. Curtis, and W. Michael, Adv. Space Res. **2**, 67 (1983).
[2]G. Paschmann, H. Loidl, P. Obermeyer, M. Ertl, R. Laborenz, N. Sckopke, W. Baumjohnn, C. W. Carlson, and D. W. Curtis, IEEE Trans. Geosci. Remote Sensing **GE-23**, 262 (1985).
[3]H. Reme, F. Cotin, A. Cros, J. L. Medale, J. A. Sauvaud, C. d'Uston, A. Korth, A. K. Richter, A. Loidl, K. A. Anderson, C. W. Carlson, D. W. Curtis, R. P. Lin, and D. A. Mendis, in *The Giotto Mission—Its Scientific Investigations*, edited by R. Reinhard and B. Battrick (ESA Noordwijk, The Netherlands, 1986), p. 33.