

The Double Star Plasma Electron and Current Experiment

A. N. Fazakerley¹, P. J. Carter^{1,*}, G. Watson¹, A. Spencer¹, Y. Q. Sun², J. Coker¹, P. Coker¹, D. O. Kataria¹, D. Fontaine³, Z. X. Liu², L. Gilbert¹, L. He^{1,**}, A. D. Lahiff¹, B. Mihaljčić¹, S. Szita¹, M. G. G. T. Taylor^{1,***}, R. J. Wilson¹, M. Dedieu³, and S. J. Schwartz⁴

¹Mullard Space Science Laboratory, University College London, Holmbury St. Mary, Dorking, Surrey, RH5 6NT, UK

²Center for Space Science and Applied Research, Chinese Academy of Sciences, P.O. Box 8701, 100080 Beijing, China

³Centre d'Etude des Environnements Terrestre et Planétaires, IPSL, 10–12 Av. de l'Europe, 78 140 Vélizy, France

⁴Space and Atmospheric Physics Group, The Blackett Laboratory, Imperial College London, Prince Consort Road, London, SW7 2BW, UK

* now at: UK Astronomy Technology Centre, Royal Observatory Edinburgh, Blackford Hill, Edinburgh, EH9 3HJ, UK

** now at: Department of Physics, University of Alberta, Edmonton, T6G 2J1, Canada

*** now at: Research and Scientific Support Department, ESA/ESTEC/SCI-SH, Keplerlaan 1, 2201 AZ Noordwijk ZH, The Netherlands

Received: 8 April 2005 – Revised: 13 July 2005 – Accepted: 23 August 2005 – Published: 8 November 2005

Part of Special Issue “Double Star – First Results”

Abstract. The Double Star Project is a collaboration between Chinese and European space agencies, in which two Chinese magnetospheric research spacecraft, carrying Chinese and European instruments, have been launched into equatorial (on 29 December 2003) and polar (on 25 July 2004) orbits designed to enable complementary studies with the Cluster spacecraft. The two Double Star spacecraft TC-1 and TC-2 each carry a Double Star Plasma Electron and Current Experiment (PEACE) instrument. These two instruments were based on Cluster Flight Spare equipment, but differ from Cluster instruments in two important respects. Firstly, a Double Star PEACE instrument has only a single sensor, which must be operated in a manner not originally envisaged in the Cluster context in order to sample the full range of energies. Secondly, the DPU hardware was modified and major changes of onboard software were implemented, most notably a completely different approach to data compression has been adopted for Double Star, which allows high resolution 3-dimensional distributions to be transmitted almost every spin, a significant improvement over Cluster. This paper describes these instruments, and includes examples of data collected in various magnetospheric regions encountered by the spacecraft which have been chosen to illustrate the power of combined Double Star and Cluster measurements.

Keywords. Space plasma physics (Instruments and techniques) – Magnetospheric physics (Solar-wind-magnetosphere interactions; Magnetospheric configuration and dynamics)

1 Introduction

The Double Star Project has placed two Chinese magnetospheric research spacecraft, Tan Ce 1 and 2 (TC-1, TC-2) carrying Chinese and European instruments, into a near-equatorial orbit and a polar orbit, respectively. TC-1 was launched on 29 December 2003 and has a 27.4-h period, 28° inclination and geocentric apogee/perigee of 13.37 R_E /1.09 R_E . TC-2 was launched on 25 July 2004 and has an 11.5-h period, 90° inclination and geocentric apogee/perigee of 7.01 R_E /1.09 R_E .

These orbits have been designed so that their MLTs (Magnetic Local Times) of apogee are aligned with each other and with that of the Cluster orbits, during the summer of 2004, when all spacecraft have their apogee in the magnetotail. The polar orbiting Cluster and TC-2 spacecraft will maintain this phasing, although the MLT of apogee of TC-1 drifts slowly apart from the other spacecraft. The NASA Polar spacecraft MLT of apogee is also fairly similar, and again phase locked with Cluster. If we treat Cluster as a single observation point (albeit with special small-scale multi-point measurement capabilities) we have a three or four spacecraft constellation well suited to simultaneously examining magnetotail processes both close to and far from the Earth, and similarly (on the dayside) to examining the dayside cusps at one or more sites while also making measurements at the low latitude magnetopause. This is a unique constellation, capable of making observations with which we can examine the drivers of global scale magnetospheric processes (in essence the final realisation of the original ISTP Geotail-Polar-Equator concept) and is orbiting at a time when upstream monitors, auroral and ring current imaging satellites, and sophisticated ground based ionospheric monitors are all active to provide

Table 1. Summary of resources used by a single Double Star PEACE instrument.

Mass	DPU	2.12 kg
	Sensor	1.92 kg
	DPU radiation shield	1.44 kg
	Sensor radiation shield	1.03 kg
	Total	6.52 kg
Power	Average	3.8 W
	Peak (during HV sweeps)	4.8 W
Telemetry	Science	4472 bps
	HK	152 bps

detailed supporting observations of the global context. A particular advantage of the Cluster-Double Star combination is that several Double Star instruments, including PEACE are identical or near identical copies of Cluster instrumentation, which helps in carrying out comparative data analysis studies. The project has been rather cost-effective as significant parts of the infrastructure (operational and for science data dissemination and analysis) developed for Cluster could be readily adapted for Double Star.

The PEACE instruments measure electrons from a few eV to 25 keV energies, covering the thermal plasma regime occupied by solar wind electrons, magnetosheath electrons and much of the magnetospheric plasma electron population. PEACE can measure the contribution of these electrons to magnetospheric currents, and measure convection of the plasma local to the spacecraft. PEACE can also detect field-aligned electrons which may have been accelerated elsewhere in the magnetosphere, providing an element of remote sensing capability (constrained to providing information on other regions linked by the magnetic field to the spacecraft location).

The two Double Star spacecraft TC-1 and TC-2 each carry a Double Star Plasma Electron and Current Experiment (PEACE) instrument. In this paper we describe how and why these differ from Cluster-PEACE instruments, and include examples of data collected in various magnetospheric regions encountered by the spacecraft, illustrating the power of combined Double Star and Cluster measurements.

2 Description of instruments

2.1 Introduction

The Double Star PEACE instruments were produced by reusing Cluster Flight Spare hardware, supplemented with some new components. Modifications to existing hardware were kept to a minimum. A complete flight-ready Cluster II PEACE instrument, consisting of two sensors (LEEA and HEEA) and a Data Processing Unit (DPU), was available.

In addition, the PEACE Flight Spare DPU from the ill-fated Cluster (I) mission of 1996 was available in partially disassembled form.

In order to provide two instruments, the Double Star TC-1 PEACE instrument was produced using the HEEA sensor and the Cluster II DPU, while the TC-2 PEACE instrument was produced using the LEEA sensor and the Cluster I DPU. Instrument resources are summarised in Table 1. The Cluster 1 DPU was upgraded to Cluster II specifications. Hardware modifications were made to both DPUs in order to ensure compliance with export control regulations, and extensive modifications to onboard software were required as a consequence, including removal of onboard moments determination capability.

The Cluster II PEACE instruments (Fazakerley et al., 2005¹, Johnstone et al., 1997) were built by a consortium consisting of Mullard Space Science Laboratory of University College London, Rutherford Appleton Laboratory (RAL) and the Norwegian Defence Research Establishment (NDRE, also known as FFI). The work to produce the Double Star PEACE instruments was carried out solely by MSSL, although the new science data compression software was developed at MSSL from a study led by colleagues at the Centre d'Etude des Environnements Terrestre et Planétaires (CETP). The instrument preparation phase proceeded on a very challenging schedule, as may be apparent from the interval between the agreement between CNSA and ESA to cooperate on the mission (9 July 2001) and the launch of the first spacecraft in December 2003.

2.2 Comparison of Cluster and Double Star design constraints

The most obvious change is the use of only one sensor unit for each Double Star instrument. Financial and time constraints prevented the production and calibration of two new sensor units, which would have been needed to produce two dual sensor instruments.

The telemetry resource available to PEACE on the Double Star spacecraft is significantly larger than the standard Cluster telemetry rate, and is still larger “per sensor” as data is generated from only one sensor on Double Star, rather than two on Cluster.

The power, telecommand and telemetry interface for the Cluster-derived payload was handled by the Payload Data Management System (PDMS) which was produced by the Centre for Space Science and Applied Research (CSSAR) in Beijing and designed to closely conform to the Cluster spacecraft-payload interface. This was done to facilitate the incorporation of the European payload onto the Chinese satellite bus.

The Double Star spacecraft both spend a significant fraction of their orbits within the outer electron radiation belt,

¹Fazakerley, A. N., Carter, P. J., Watson, G., et al.: The Cluster II Plasma Electron and Current Experiment, *Ann. Geophys.*, in preparation, 2005.




Fig. 1. Illustration of the principle of the Top Hat electrostatic analyser. The voltage applied to the analyser hemispheres diverts electrons (shown in blue) of a specific narrow band of energy and arriving within the acceptance angle $\Delta\Psi$ through the analyser to the detector, while electrons of other energies (e.g. in red) strike one of the analyser hemispheres and are not detected. The semi-annular microchannel plate amplifies the signal of an electron reaching it, and the resulting charge cloud is detected by one of 12 segments of the anode beneath, giving information about the electron arrival direction θ . The Top Hat design is able to provide a focussed spot in the detector plane for a parallel electron beam arriving from anywhere in the aperture plane. At times when the aperture looks sunward, the majority of photons (black wiggly arrow) pass through the aperture and do not find their way to the detector.

and also cross the proton belt. Cluster, with a higher perigee and a polar orbit, never encounters the proton belt and crosses quite rapidly through the outer electron belt. The original design specification for Cluster prohibited operations below 35 000 km altitude. The Cluster PEACE instrument was therefore designed to handle a 20–30 krad dose and without the specific anti-latchup protection that would be appropriate for a proton belt traversing instrument. Additional protection was therefore provided for the Double Star PEACE electronics subsystems, as described below.

Payload accommodation is handled differently on the Double Star spacecraft. The experiment platform has a smaller diameter and a cylindrical solar cell support structure both above and below the experiment platform, whereas on Cluster there was only one solar cell cylinder below the




Fig. 2. Illustration of instantaneous field of view of a Double Star PEACE sensor. Left: Side view of spacecraft, spin axis points close to Z_{GSE} , the “look” directions from which electrons arrive to be measured by the 12 anodes are indicated. Right: Top view of spacecraft, showing instantaneous field of view (elevation acceptance angle) also illustrating configuration of magnetometer booms (note that on TC-1 the STAFF boom is not deployed).

experiment platform. The main significance for PEACE was the need to adapt the PEACE thermal blanket design. The relationship between sun sensor look direction, magnetometer sensor orientations and PEACE look directions also differs from Cluster, requiring modification to software which uses this positioning information.

2.3 Sensor description


Each sensor consists of a “Top Hat” electrostatic analyser with an annular microchannel plate chevron pair (MCP) and segmented anode to provide position sensitive detection of arriving electrons, together with a supporting Sensor Electronics Unit. Figure 1 illustrates the general principle of the PEACE electrostatic analyser, showing how photons typically travel right through the collimator (photons do not enter the collimator at all unless the sensor faces sunward), and only electrons with the selected energy are allowed to reach the detector. The “Top Hat” design ensures that all electrons on mutually parallel trajectories will be focused by the analyser to the same part of the MCP, whichever part of the aperture they enter from. The detector anode is divided into 12 equal parts, corresponding to 15° resolution in the “polar” angle. The PEACE analyser is mounted with its field of view fan lying perpendicular rather than tangential to the spacecraft surface, as shown in Fig. 2, in order to minimise the entry into the aperture of photoelectrons and secondary electrons from the local spacecraft surface. Consequently, only 180° of the 360° field of view available in principle with a Top Hat analyser are in use (hence the semi-annular MCP and anode). Figure 2 illustrates the direction from which electrons arrive to be counted on anodes 0 to 11, in terms of the spacecraft spin axis, which is intended to point close to the GSE (Geocentric Solar Ecliptic) Z direction. (Note that in practice some departure from this ideal spin axis orientation occurred as the mission proceeded).

Table 3. PEACE electrostatic analyser HV sweep modes.

Sweep mode	Energy steps per sweep	Step size	Sweep period		Sweeps per spin
			Measurement Interval/ T_{acc}	Flyback Interval/ T_{acc}	
LAR	60	1	60	4	16
MAR	60	2	30	2	32
HAR	30	2	15	1	64

The analyser HV Sweep Generator is capable of providing 88 distinct voltages, the first 16 of which are linearly spaced, and the remaining 72 of which are logarithmically spaced. For science data collection the analyser voltage is repeatedly swept through a range of values so that a range of electron energies is sampled during each “energy sweep”. Three energy sweep modes are available; Low, Medium or High Angular Resolution (LAR, MAR or HAR). Both a LAR and a MAR sweep jump down 60 consecutive voltage steps each sweep, but a LAR sweep duration is twice that of a MAR sweep. An HAR sweep covers 30 consecutive voltage steps per sweep in half the time of a MAR sweep. Each sweep begins with a “flyback” time interval during which the voltage returns to the highest voltage reached during the sweep; the flyback duration is defined differently for each sweep mode, as shown in Table 3.

Measurements are made by counting the number of electrons striking each anode during an accumulation period T_{acc} , defined as $T_{spin}/1024$ where T_{spin} is the satellite spin period. For a nominal 4.0 s spin, $T_{acc}=3.9$ ms. The accepted electron energy at any given moment is controlled by the electrostatic potential of the analyser inner hemisphere. There are 16 LAR sweeps, 32 MAR sweeps or 64 HAR sweeps per spin; hence the increase in spin phase (azimuthal) angle resolution implied in the sweep mode names. The LAR mode has the best energy resolution, as it changes analyser voltage by stepping down one level during each accumulation time. The MAR and HAR modes share the same (lower) energy resolution; in both cases the analyser voltage steps down two levels during an accumulation time. These “steps” are not near-instantaneous drops in voltage from one level to the next, after which the voltage sits at a level for the remainder of the accumulation time. Instead, the sweep voltage is allowed to decay exponentially between the top 72 levels, and does so over a period of about 1.8 ms, so that the “logarithmic” part of the sweep is very close to a smooth exponential decay in MAR and HAR modes, but somewhat more step-like in LAR mode. The calibration process involves determination of a “representative energy” for each “step” in each sweep mode. Table 3 summarises all these characteristics of the sweep modes. The MAR mode is used most often in orbit. In all cases, during a single spin, the instrument measures 11 520 values from the 12 anodes (excluding data collected during flybacks and fine zones, which is not useful). The energy range that is covered by a given sweep is controlled by

**Fig. 4.** Cross section through a PEACE Sensor electrostatic analyser.

the “sweep preset” value, which may be varied by telecommand.

The sensitivity of the MCP is controlled by adjusting the voltage applied to the microchannel plates. The capability is useful during operations, as the sensitivity is expected to decline with use in orbit, particularly if the sensor is exposed to high fluxes for a prolonged period.

The analyser head has been designed to minimise the number of photons, photoelectrons and secondary electrons which are able to reach the detector. Figure 4 shows a cross-section through an analyser head which illustrates some of the relevant design features. As discussed in Johnstone et al. (1997), computer simulations predicted that arriving electrons and UV photons which strike the lip where a hole is cut in the outer hemisphere would be a major source of secondary electrons and photoelectrons. Therefore a ring baffle was introduced in this region in order to reduce the surface area from which such electrons would be emitted on paths that reach the detector. Baffles in the roof of the input collimator are similarly intended to reduce the flux of primary electrons and UV photons reaching the inter-hemisphere gap region. The hemispheres and the entrance aperture are coated in black copper oxide, applied by the EBONOL-C process, which provides further absorption of UV photons and secondary electrons.




Fig. 5. A view of the TC-1 PEACE Sensor. Note the analyser head radiation shield material intended to protect the MCP, in particular the pie-slice plates and the enhanced thickness rim of the cylinder towards the aperture side of the sensor head.

Inspection of Fig. 4 demonstrates that all paths that penetrating radiation might take through the analyser head to reach the MCP will pass through a significant amount of material. The design is intended to put 5 to 6 mm of Aluminium or material of equivalent mass density in every direction. The resultant shielding is expected to be adequate to stop all electrons of energy up to 2 MeV (although not the Bremsstrahlung), and significant fractions of electron fluxes at higher energies up to perhaps 10 MeV, as well as protons up to 30 MeV. Figures 5 and 6 show photographs of a PEACE sensor including the enhanced radiation shielding on the analyser head.

The ground calibration work done on the sensors for Cluster II was repeated for Double Star in sufficient detail to confirm that there were no changes. In the process, the integrity of the sensors was confirmed, to ensure that no damage had occurred after their journey to and from the Cluster launch site (Baikonur, Kazakhstan).

Double Star TC-1 carries a DWP unit as part of the STAFF-DWP instrument, and DWP is host for the Particle Correlator software (see Cornilleau-Wehrin et al., 2005, this issue). In order to enable the utilisation of the Correlator capability on Double Star, it was decided to fly the HEEA sensor on the TC-1 spacecraft. Since HEEA has greater sensitivity than LEEA it would have been preferable to use LEEA on TC-1, at least from the perspective of magnetosheath measurements, since HEEA sensors often saturate in the high electron fluxes encountered in the magnetosheath. The Correlator software generates autocorrelation functions which may enable examination of the electron behaviour at higher time resolution than the ~ 3.9 ms PEACE measurements, and in particular to examine wave-particle interactions, e.g. enabling study of processes generating the diffuse aurora.




Fig. 6. TC-1 PEACE DPU and Sensor together with the radiation shields for the electronic sub-systems.

2.4 Data processing unit description

Figure 7 shows a block diagram of a PEACE Data Processing Unit (DPU), illustrating the main sub-systems and information flow

The TC-1 DPU could have been flown in its Cluster II Flight Spare configuration, except for the requirement that no components subject to export control regulations be used in Double Star payload provided by ESA. Unfortunately, the 64-kbyte SRAM components used for Cluster are export restricted. Within the constraints of predicted radiation dose and accommodation available within the existing DPU box, the only alternative was to replace the Cluster SRAMs with non-restricted 16-kbyte SRAMs. In addition, unused interfaces applicable only for Cluster (e.g. to the second sensor, the IEL link to EDI) were closed off.

For TC-2, more extensive changes were required, in order to bring the Cluster I heritage DPU up to Cluster II DPU build standard. In particular, a suitable power converter was purchased (the MCP specification changed from Cluster I to Cluster II, affecting power requirements) and new SI and PCU cards were built. However, the Cluster I T222 transputers were not replaced with Cluster II style T225 transputers.

A Double Star PEACE Data Processing Unit (DPU) thus consists of a power converter and switching card, a spacecraft interface card and two processor cards. Each processor card is loaded with a single Transputer, 16 Kbytes of SRAM, a 0.5 KByte Bipolar PROM and 64 Kbytes of EEPROM.

The 16 kbyte SRAM per processor constraint is very severe, as the RAM must:

- (i) accommodate the software used by the transputers during run time in order to achieve the required speed from the processor;
- (ii) house data arriving from the sensor (at an incoming data rate of 8 kbytes per second) – note that there is not sufficient RAM to store all the data delivered during one 4-s spin at one time;




Fig. 7. Block diagram of PEACE DPU.

(iii) contain the science telemetry buffer in which processed data is stored prior to transmission to the PDMS telemetry sub-system, (at an outgoing data rate of 2.2 kbytes per second).

The difficulties associated with this reduction in available RAM required that we develop new DPU processor software as discussed below.

On both spacecraft, the Inter-Experiment Link (IEL) between FGM (Carr et al., 2005, this issue) and the PEACE DPU was retained, to allow PEACE access onboard to raw FGM data, for use in production of onboard pitch-angles. Similarly, on TC-1 the IEL between PEACE and ASPOC is used to allow ASPOC to receive PEACE onboard estimated values of the spacecraft potential.

2.5 Radiation shields

Pre-launch analysis projected a total dose while shielded by 4 mm Aluminium equivalent as ~ 60 krad (Si) for TC-1 and ~ 67 krad (Si) for TC-2.

The PEACE sensor and DPU electronics were designed for the 2 year Cluster mission radiation environment (20–30 krad) and as such require additional protection for the Double Star environments. CNSA generously allowed us sufficient mass resource to provide this protection by building 4 mm thick aluminium radiation shields, which are placed around the existing DPU and Sensor. These are illustrated in Fig. 6. Together with the 0.6 mm thick DPU and Sensor housings, the shields should stop protons with energies up to 30 MeV and electrons with energies up to 1 MeV. According to the commonly used AP8 radiation belt mod-

els (Sawyer and Vette, 1976), the shield provides full protection against protons for $L > 3.5$. The outer electron belt is rather variable, and at times can have significant fluxes of > 1 MeV electrons. Further shielding is provided by spacecraft components including the solar cell arrangement, experiment platform and batteries, but is present for less than half of the 4π solid angle of possible arrival directions for energetic particles. The detector radiation shielding has been discussed in the sensor description (see Sect. 2.3 above).

2.6 Thermal blankets

The thermal blanket protection around the PEACE sensors on Cluster was designed by the PEACE team, and procured from EADS Astrium Ltd. Additional units were bought for Double Star. It was hoped that the blankets could be used in just the same way for Double Star as for Cluster, but accommodation constraints required some modifications to achieve closure with the spacecraft thermal protection system. A technique for adapting the Cluster-design PEACE blankets was developed. A newly purchased thermal blanket was used for TC-2 PEACE, however TC-1 PEACE used a Cluster Flight Spare blanket. The thermal blankets were effective, and in fact the Double Star spacecraft is a warmer environment for a PEACE instrument than a Cluster spacecraft.

The blankets were required in order to provide proper thermal insulation, and in addition to meet stringent electrical conductivity requirements to ensure that no differential charging of the spacecraft surface arises. Success in this respect ensures that the trajectories of electrons arriving at



Fig. 8. Photograph of TC-2 PEACE Sensor head showing arrangement of electrostatic conductive materials; spacecraft and PEACE thermal blankets and also solar cells above and below the experiment platform. Note the larger collimator of a LEEA sensor vs. the TC-1 HEEA sensor.

the sensor aperture will not be affected by localised electric fields from areas of spacecraft surface material which have become charged to different potentials than the surrounding spacecraft surface. The mounting system includes the use of a metal band to hold the blanket in place against the cylindrical Analyser Head, which is attached by a number of small screws. The screws are secured by small quantities of (non-conductive) epoxy, which are covered with a disk shaped patch of aluminium tape using conductive adhesive to ensure continuity of electrostatic conductivity in the vicinity of the aperture, as shown in Figs. 8 and 9. Extensive testing of the conductivity of the sensor head, thermal blanket and indeed the entire spacecraft surface was carried out by the spacecraft team, shortly before launch, to ensure that the required standard had been achieved. Some differences in conductivity were noted between newly purchased and older thermal blankets, with better performance from the older blankets.

2.7 Onboard software

2.7.1 Onboard software requirements

The software requirements for the PEACE DPU onboard the Double Star spacecraft can be summarised as:

1. To provide a software interface to service the data being streamed from the PEACE Sensor to the DPU.
2. To provide a software interface to the PDMS that is capable of receiving commanding and servicing signals, including telemetry requests.
3. To provide housekeeping telemetry data with which to monitor the health of the instrument.




Fig. 9. Close up of HEEA Sensor Head showing thermal blanket closeout; note electrically conductive material patches over the epoxy used to lock the screw heads.

4. To return as much as possible of the Sensor data via the DPU-PDMS telemetry interface using a scheme that can be interpreted on the ground.
5. To provide all data that is required to commission the instrument in real time, in the housekeeping telemetry.
6. To provide a command interpreter for all required DPU and Sensor functions.
7. To provide functionality to boot the DPU.
8. To provide software patching capabilities.

2.7.2 Double Star PEACE software architecture

One of the DPU processor cards is designated to service the PDMS interface (the Interface Processor, or IP), while the other handles the incoming data from the Sensor (the Science Processor, or SP). The IP code is designed to be entirely event driven, imposing significant timing constraints on the response time and latency that can be tolerated before an event is serviced. The SP code works in synchronisation with the spacecraft spin.

The software for the PEACE DPU is divided into three levels. Level 1 consists of boot code to initialise the Transputer on each of the processor cards, and provides basic functions to “peek” and “poke” bytes into memory. This code also allows execution of the Level 2 software. Level 2 provides a comprehensive set of software patching and memory analysis facilities. It also provides functionality to select a range of “Application Codes” that constitute the final level of software. The first two levels of the software are common to both the processor cards. The final level is specific to the Interface or Science Processor that it is running on.

There are several application codes available in the Double Star PEACE software. For two-processor operations, there are both a main application code and a backup application

code. Each contains full housekeeping data generation code as well as commanding, event handling, sensor data reception and spacecraft potential estimation code. In addition it is possible to swap the identity of the processor cards, providing redundancy should the link between the sensor and the SP processor fail – since the second processor is also linked to the sensor and could take on the SP functions instead.

There is also a single processor code that can generate a limited reduced resolution 3-D dataset (known as “3DR”) using quasi-log compression (as on Cluster II PEACE), for use in the event of failure of either processor.

The main two-processor application code is an implementation of JPEG compression designed to allow transmission of the full measured distribution at spin rate on a nearly continuous basis. In this scenario, the transmitted 3-D distribution contains a factor 8 more values than 3DR, and both moments and pitch angles are determined using ground data processing software. The method offers better data resolution and so is preferred. The energy resolution is twice that of 3DR, for the case of the most commonly used MAR sweep mode, allowing more accurate estimation of spacecraft potential in ground data analysis than from 3DR. We describe this implementation in more detail below.

The backup two-processor application code is a reduced capability version of the Cluster II code, which uses quasi-log data compression. The Cluster II capabilities of onboard moments determination and onboard magnetic field axis estimation (from symmetry of the measured distribution) were removed. The capability to store 3-D distributions was limited to the Cluster 3DR distribution applied to a single sensor. Onboard production of pitch angle distributions (adapted for the single sensor-only case), was retained, implemented essentially as for Cluster II, using FGM data from the Inter-Experiment Link. The reduced memory resources, compared to Cluster, made it impossible to retain the onboard moments determination software.

An onboard dead-time correction capability is not included, unlike the Cluster case where it was required since moments calculations are performed onboard. For Double Star, such calculations and prior dead time corrections are performed on the ground.

2.7.3 Commanding

In order to reduce (relative to Cluster) the number of commands transmitted to the spacecraft, the number of commands to be stored onboard the spacecraft and also the command checking workload during operations, we introduced the capability to store within PEACE up to 10 commanding macros with the possibility to use time delays between successive commands in each macro.

2.7.4 Double Star PEACE DPU data flow

Data arrives in the DPU from the sensor at a rate that would very quickly exceed the capacity of the RAM if no compression is carried out. Also, all data to be transmitted to Earth

must be buffered so as to be entered into the PDMS telemetry flow in an orderly manner. In addition the DPU must deal with signals from the PDMS.

The SP receives the data from the PEACE sensor in intervals of one sixteenth of a spin, termed “basic segments”. Each basic segment consists of 2 kbytes of data. After discarding data accumulated during sweep flybacks and data from the fine zones, 1440 bytes of science data remain to be processed. While part of the SP deals with incoming data from the Sensor, other routines are processing the basic segment that arrived immediately beforehand. The processing of one basic segment must be completed before the next basic segment is fully received from the sensor. An additional complication arises at a spin boundary (defined by the arrival of a new rephased sun pulse) since the act of processing the final basic segment of a given spin always occurs during the next spin. Thus the association of timing information with data generated on the SP during ground data processing is more complex than for the IP, which, being event driven, has no such constraint.

PEACE collects and produces science data at a rate synchronised to the ~ 4 s spin period of the satellite and not with the collection rate of the telemetry blocks from PEACE by the PDMS which occurs on a 5.15222 s cycle. In order to decouple these two activities, PEACE science data is telemetered in the form of “PEACE science telemetry packets” which can be identified anywhere in the science telemetry stream. The packet locations in the data stream are marked using a 4 byte synchronisation byte pattern, which is accompanied by information defining the length and structure of the PEACE science packet, and the data itself. Ground data processing software can therefore often recognize packets even in partially corrupted telemetry, however it can only properly reconstruct the PEACE data within the packets when the telemetry stream is complete.

2.7.5 Double Star PEACE house-keeping telemetry

Only house-keeping data was available in real time during instrument commissioning. The capability to transmit data from which crude energy-time spectrograms and also polar zone vs. time spectrograms could be constructed was put in place to support effective commissioning work. Specifically, this allowed us to characterise the prevailing plasma environment, and provided visualisation of the individual anode response in case of a problem emerging with a particular anode/MCP section. It was achieved simply by transmitting (i) the counts per spin data summed over adjacent pairs of energy bins and all polar and azimuth bins giving 15 distinct energies per spin, and (ii) summed over energy bins and azimuth bins giving 12 polars per spin. The HK data compression scheme limits the maximum reported value to 1×10^6 , safely in excess of usual measured values. This data set is also routinely available during normal science operations, allowing verification of the science data compression scheme.




Fig. 10. Example of onboard spacecraft potential estimation. Data is from TC-1 on 9/10 March 2004 in the solar wind.

2.7.6 The JPEG compression code

The available data rate for Cluster II PEACE in the standard telemetry mode used most of the time is 1257 bytes/spin (for a 4 s spin) or ~ 630 bytes/sensor/spin. By contrast the available data rate for Double Star PEACE is 2236 bytes/spin (again for a 4 s spin) or 2236 bytes/sensor/spin, a factor 3.6 greater than for Cluster. An uncompressed 3-D distribution amounts to 23 040 bytes/sensor/spin, so a compression factor of order 10 is required for Double Star.

A lossless compression scheme (typically achieving compression factors of order 2) would be inadequate to transmit the full measured distribution every spin, and does not satisfy the requirement for working within a fixed maximum size for all compressed buffers imposed by the telemetry packetisation scheme, because the compression rate in lossless schemes depends on the data complexity which varies with changes in the plasma environment. The publically available “norm JPEG” lossy compression algorithm we have adopted works within the afore-mentioned constraints of very limited storage size for software and data buffers, and limited processing time. It works with fast Discrete Cosine Transforms, thereby avoiding complex computations and is quick, due to a reduced number of operations.

The general approach for PEACE is to take data as it arrives in basic segments and to organize it into sets of two dimensional “images” of 8×8 pixels. The application of DCTs on such images had been extensively tested and is described in textbooks, for example (Nelson and Gailly, 1995; Pennebaker and Mitchell, 1993). In outline, the procedure works as follows:

1. Data in a basic segment is reorganized into a set of 8×8 “images”.
2. The values in the image are rescaled to a maximum of 512.
3. Each image is converted to a DCT coefficient matrix.
4. A weighting factor is applied to the DCT coefficients.

5. The DCT coefficients are bit-streamed into PEACE Science TM packets ready to be made available to the PDMS, for transmission to Earth and later conversion back to decompressed science data.

The process is discussed in greater detail in the Appendix.

2.7.7 Onboard estimation of spacecraft potential

An estimated value of the spacecraft potential can be computed every spin for each of up to 32 energy spectra. The energy spectra used for this consist of the lowest 16 energy bins of 8 consecutive energy sweeps for 4 consecutive polar anodes. The start polar anode and azimuth angle corresponding to the first energy sweep used can be commanded from the ground. The algorithm works by searching for an energy bin at which the gradient between neighbouring points in the energy spectra turns positive, starting from the highest energy bin. Next the search continues towards the lowest energy bin in the spectrum and identifies the bin with the absolute minimum count value as indicating the spacecraft potential. The rationale for the algorithm is explained in Johnstone et al. (1997). This algorithm was originally designed for use on Cluster LEEA data (sub-keV electron spectra) and is often quite effective in solar wind and magnetosheath plasma, although we note that not all energy spectra that a PEACE sensor may measure conform to the underlying assumed form. Also, measured energy spectra can be sufficiently spiky that this algorithm may fail on several individual spectra per spin. To provide some measure of reliability, the software is designed to return the mean value from only those spectra where the algorithm is successful, and also calculates and returns the variance to give some measure of the statistical significance of the mean value. An example of the onboard spacecraft potential estimation is shown in Fig. 10.

2.7.8 Achieving full instrument energy range coverage

The Double Star PEACE instruments have only a single sensor each. Each sensor was part of a pair of Cluster sensors that were designed to be used in combination to cover the full instrument energy range. For any one energy sweep a sensor can only cover about 2/3 of the full set of energies it can operate across.

In order to overcome this limitation, and achieve full energy range coverage from the single sensor, the DPU software has been adapted such that the sweep preset (which defines the energy range sampled) can be alternately switched back and forth between two ground-commanded states, upon reception of the rephased sunpulse (e.g. Figs. 11a and c). This is possible as the DPU triggers reception of the sunpulse on the rising edge of the pulse whereas the sensor triggers on the falling edge, and there is a window of 110 microseconds during which the DPU can sense the sunpulse has occurred and send a command to the sensor to switch the preset. The new preset is used for all sweeps in that spin until it is again switched back to its previous value on arrival of the next sunpulse. It is, however, imperative that the software does not




Fig. 11. (a) illustration of alternating preset mode; (b) illustration of option of returning only non-overlapping energy data from first sweep in a pair; (c) data plot showing coverage when operating as in mode outlined in (a).

exceed the 110 microsecond window or the sweep flyback will have started and the actual sweep energy coverage will not be as intended, at least for the first sweep of the new spin.

Using the sensor in this driven loop means that there is a considerable overlap region in the sweep energies measurement over two spins (i.e. $\sim 1/3$ of the full set of measurable energies). The software features the option of not telemetering the lower energy part of the data collected during the high energy band coverage spins. The data concerned is from the energy range that is sampled both on high energy band coverage spins and low energy band coverage spins (e.g. Fig. 11b), so that during a pair of spins, data from the full instrument energy range will be transmitted. However, the combined data volume sent from the pair of spins is reduced, allowing more efficient use of the telemetry resource and so resulting in fewer “missed spins”.

A further extension of the systematic variation of the sweep preset uses the spacecraft potential determined from the onboard algorithm to determine the lower sweep preset. In this way the lower sweep preset “floats” with the varying spacecraft potential. Both options can be used separately or combined.

2.7.9 Ground software

New EGSE software was developed to take advantage of the modified housekeeping data stream (including the coarse energy-time and polar-time spectrograms) and to enable its operation on a laptop PC. The new software was written in Java, providing the flexibility to run on both SUN workstations and PCs. New software (compared to Cluster) was required in any case to allow communication with the Double Star PDMS EGSE which used a multicast transmission protocol.

New telemetry data processing software was developed to enable the extraction and decompression of Double Star PEACE data from raw telemetry files, and to produce data files in the standard file format used in the PEACE database, with accurate time stamps attached. Thus the data files available to the science user appear essentially identical to Cluster PEACE data files and can be analysed with the same software tools, despite the different data compression method used for Double Star.

3 Instrument operations

3.1 Commissioning

TC-1 PEACE real-time contact commissioning was carried out on 21–23 January 2004, which coincided with Chinese New Year. The electronic systems were verified, including inter-experiment links, and then the Sweep HVs were cautiously exercised. During the following contact opportunity the MCP HVs were gradually brought up to operational levels. Finally, scientific operations were demonstrated and various modes exercised. Some unexpected plasma conditions were seen, which later proved to be the consequence of the arrival of an interplanetary coronal mass ejection, and in particular the “sheath” of compressed solar wind ahead of it. All tests were fully successful.

TC-2 PEACE commissioning achieved the same goals, but was carried out a little differently, partly due to the more limited contact opportunities associated with the polar orbit and partly in the light of TC-1 experience with resets in early operations (see Sect. 3.3 below). More time was spent in operations designed to cause sensor outgassing; initially by warming the sensor using the heat generated by the low voltage electronics, and later by warming the MCP itself using relatively low MCP voltages. These low risk operations, carried out using time tagged commands in between real-time contact opportunities were also used to maximise the time spent running the instrument, so as to expose any possible problems while real-time investigations were possible. Again the tests were fully successful.

3.2 Routine operations

In order to protect electronic systems and reduce MCP counting unconnected to plasma electrons, we avoid operations in the proton radiation belt altogether, and limit operations in the electron radiation belts by ensuring that PEACE is turned off inside the predicted $L=4$ L-shell. Penetrating radiation counts are still sometimes seen near $L=4$ as some of our data figures show, suggesting that if we did operate at lower altitudes, much of the data would be contaminated in the same way. Due to the fact that instrument turn-on takes several minutes, it is difficult to cover rapid low-altitude passes through the southern cusp and auroral regions on TC-2 while satisfying the constraints mentioned above.

Experience with Cluster has shown that PEACE MCP aging can be reduced by minimising exposure to photoelectrons. However on Double Star we have no independent measure of spacecraft potential and therefore need to measure at least the upper energy component of the photoelectron population to enable estimation of the spacecraft potential from PEACE energy spectra. As a compromise we usually raise the minimum measured energy at higher altitudes where the potential is expected to be larger, so as to keep visibility of the spacecraft potential while avoiding the higher fluxes of photoelectrons at the lowest energies. When ASPOC (Active Spacecraft Potential Control) (Torkar et al., 2005, this

issue) is operating, the potential is generally reduced and so PEACE is required to measure to lower energies in a given region than if ASPOC were off.

These constraints, and others, require fairly complex routine commanding activities which are carried out by the PEACE team.

3.3 Resets

Operations are from time to time disrupted by unplanned terminations of science data collection, and hence unwanted data gaps can occur. This applies to TC-1 and TC-2. These events are considered to be generally due to the space environment affecting the spacecraft subsystems and in particular affecting the spacecraft service signals to PEACE, such as the sun pulse, spin segment clock pulse or the power supply voltage. PEACE is designed to shut down its operations and power supply to the sensor under such circumstances. A trend for these events to occur during and following periods of enhanced radiation belt fluxes led to the suggestion that deep dielectric charging in the harness or possibly other subsystems could be generating false signals and hence causing PEACE to shutdown. An increase in such events was also noted on each spacecraft after the spacecraft ACS (Attitude Control Subsystem) computers were damaged by radiation. The operational workaround solution that we have adopted is to command regular off-on cycles. The large number of off-on cycles successfully handled by Cluster PEACE sensors gives us confidence that Double Star PEACE instruments can tolerate this without damage. If an instrument has experienced an unplanned turn off, the off-on cycle can restore operation and recover the planned data coverage for the rest of the orbit.

3.4 Illustration of data compression

Figure 12 illustrates how the number of correlation coefficients returned, varies from image to image in different plasma environments. During this example, PEACE was in MAR sweep mode. The top panel is a simple energy-time spectrogram showing magnetosheath plasma (high fluxes centered on 100 eV) interspersed with solar wind plasma (moderate fluxes at similar, but slightly lower energies) and finally magnetospheric plasma. The bottom panel shows the time history of the maximum count value per basic segment P_{\max} in red, which of course reflects the flux intensities shown in the spectrogram. The same panel also indicates the number of telemetered coefficients per basic segment C_{sum} in black, which tends to be smaller for more complex energy spectra, and which always falls short of the theoretical maximum of 768, i.e. if all 64 coefficients were sent for all 12 images in the basic segment. The remaining panels show the breakdown of the number of coefficients $C_{A,B,C,D}$ transmitted for each of 4 types of images, spanning the instrument energy range. These images correspond to four energy bands, with A the highest energy and D the lowest, as illustrated in Fig. A1, and are for data associated with the first group of

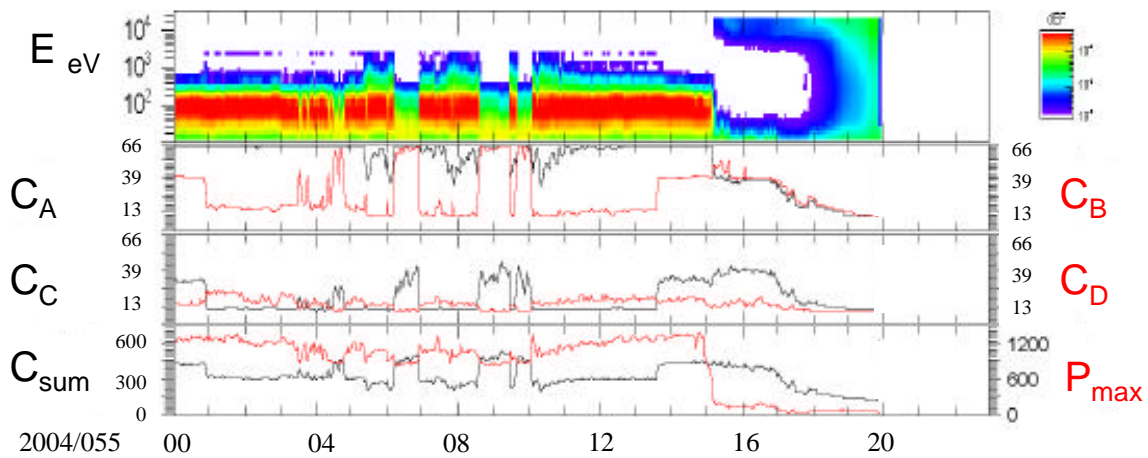


Fig. 12. Illustration of the variation with plasma environment of the number of transmitted JPEG coefficients per image.

four polars. All 64 coefficients can easily be returned in the available 10 bytes when there are very few measured counts, as we see for C_A , the black trace in the upper panel, which represents images of type A, at the highest energies. Only when the spacecraft enters the magnetosphere and PEACE sees energetic plasma sheet electrons, do significant fluxes of electrons appear at these energies, and many non-zero coefficients are then generated by the compression code so that the number of coefficients C_A , returned in the available 10 bytes per image falls, increasingly as the fluxes increase. At the lowest energies, the energy spectra include a photoelectron population below the spacecraft potential and a magnetosheath energy population above it. Many fairly large coefficients are required to describe the form of the resultant spectrum, so that C_D the number of coefficients for the type D images is particularly small. Generally speaking, the energy spectra covered by type B and C images have intermediate complexity and hence intermediate numbers of coefficients are returned. Note that the solar wind distributions are represented with smaller coefficients than the magnetosheath distributions in the images of types B and C, and so more can be transmitted. The jump in most values at 13:40 UT corresponds to a mode change and after 13:40 UT is an artefact of the plotting software which averages values from spins with higher and lower energy coverage (before 13:40 UT all spins cover lower energies).

3.5 Ground data processing

As noted in Sect. 2.9, the telemetry data is decompressed and converted to data files of the same form as Cluster PEACE data. In normal operations, the transmitted Double Star PEACE data corresponds to the full resolution 3-D distribution “3DF” data product. In order to support pitch angle determination, a support data product called “DFUNIT” is produced using magnetic field data supplied by the magnetometer team. Interpolation is used to generate magnetic field data at higher-than-measured time resolution, so as to produce a vector for each PEACE measurement. Also, the magnetic

field data are put into the coordinate frame which rotates with the PEACE sensor, after which standard PEACE team software can provide pitch angles straightforwardly from the 3DF data.

Since the Double Star PEACE data is provided to science users in the same form as Cluster PEACE data, analysis and plotting software designed for Cluster PEACE can be used quite straightforwardly with Double Star PEACE data. Most data plots in this paper were created with Southwest Data Display and Analysis System (SDDAS) software (e.g. see <http://www.sddas.org>) developed at the Southwest Research Institute (SwRI). The applicability of this software relies on the use of the underlying IDFS (Instrument Data File System, <http://www.idfs.org>) approach to data handling, which has been applied to Cluster PEACE by MSSL in co-operation with PEACE Co-Is at SwRI, and extended to Double Star PEACE by MSSL.

The QPEACE software suite (<http://www.space-plasma.qmul.ac.uk/QPEACE>) developed by one of the co-authors (SJS) is also able to work from IDFS data files and provides additional methods of examining PEACE data. Figure 13 shows an energy-time spectrogram, in which pitch angles are plotted within each energy band, generated by one of the QPEACE tools (QJAS), which used as input both the PEACE 3DF data and the DFUNIT magnetic field data set described above. Concentrating on 1 to 5 keV electrons, we see almost isotropic distributions initially, bi-directional field aligned beams around 02:35 UT, and strong fluxes at 90° pitch angle after 02:50 UT.

Moments are calculated from the transmitted velocity distributions, corrected using the estimated spacecraft potential, using “peacemoments” software developed at MSSL. These moments data are the basis for the Prime and Summary Parameter data supplied by MSSL to the Double Star Science Data System. Initial data of this type is produced with crudely estimated spacecraft potential and so may not be as accurate as later iterations; users are encouraged to select events with such data, but are advised to read the




Fig. 13. Example of energy-pitch angle-time plot to demonstrate the effectiveness of ground-based pitch angle production (3 September 2004, PEACE data from TC-1).

caveats provided with each file and to respect the “rules of the road” when using such data (<http://www.cluster.rl.ac.uk/ddms/rules.htm>). The same software can be used to calculate the electron contribution to electric current flow in the vicinity of the spacecraft.

3.6 Ground based spacecraft potential determination

The spacecraft potential is a critical parameter for producing accurate moments data from electron distributions. While an onboard potential estimate is often available, it is assumed that a more reliable time series could be generated using more complex algorithms during ground analysis of the transmitted data. Techniques are in the process of being evaluated, initially being developed using Cluster PEACE data where comparison with a reference measurement from EFW is possible, and thereafter being applied to Double Star PEACE data. Figure 14 is an example of one method being evaluated, which searches for an expected gradient change in phase space between the photoelectron spectrum and the natural plasma spectrum departure. The estimated potential is shown as black line overlaid on the spectrogram. This is work in progress at the time of writing.

4 Science highlights

Early science studies using PEACE data have focussed on dayside studies with TC-1 and Cluster, simply because that data was available first; some are reported in this journal issue, e.g. Dunlop et al. (2005), Marchaudon et al. (2005), and Pu et al. (2005).

Here we briefly describe four additional events to show the potential of the combined Double Star – Cluster dataset to examine large-scale magnetospheric processes. These events will be the subject of deeper study.




Fig. 14. Example of ground-based estimation of spacecraft potential.

4.1 Near-midnight apogees: magnetotail

On 3 September 2004, a close magnetic conjunction occurred involving both Double Star spacecraft and the Cluster flotilla, as illustrated in Fig. 15. TC-1 and Cluster lie down the tail near the current sheet illustrated in the model magnetosphere (upper plots) and show very similar trends in plasma sheet electron energies (lower plots) from 02:00 UT to about 02:38 UT, although after this time the plasma sheet seen at TC-1 is significantly more disturbed and energetic than that seen further down tail at Cluster. TC-2 is moving northbound, crossing progressively higher L-shells, and a few minutes before 03:00 UT crosses magnetic field lines that map close to the positions of TC-1 and soon after to Cluster, before entering the northern magnetotail lobe after 03:15 UT. Note that the actual magnetic mapping may differ in some details from the illustrative model shown here, but is expected to be reasonably representative.

At about 02:38 UT, a significant increase in both the plasma sheet electron energy and differential energy flux is seen at TC-1. At about the same time, the character of the plasma sheet seen at TC-2 also changes, and beams of energetic (~ 10 keV) electrons are seen by the Cluster quartet. Just prior to these events, at 02:36 UT, both Cluster and TC-1 see enhanced bi-directional electron fluxes at plasma sheet energies (\sim keV) (not shown). Preliminary magnetometer data from TC-1 show a reduction in B_x , together with an increase in B_z , suggesting that a dipolarisation of the magnetic field occurred near the location of TC-1. Similar signatures are seen at the four Cluster spacecraft, consistent with an Earthward moving dipolarisation front. Magnetic field deflections at TC-2 are suggestive of a field aligned current consistent with a substorm current wedge. The disturbance seems to propagate from Cluster to TC-1 to TC-2. Unfortunately, TC-2 electron pitch angle data are not available at the time of writing. This interval is being actively studied, to investigate how the plasmasheet electrons at TC-1 and TC-2 have become energised, and what is the exact sequence of events along the magnetotail from Cluster to the




Fig. 15. Above, 3 September 2004, orbital positions and OVT model magnetosphere in GSM, views from $-X$, Y and Z ; below, PEACE data from TC-1, TC-2 and C4. X-axis labels show UT time, L-shell, MLT and invariant latitude.

ionosphere. This example shows the potential of simultaneous multi-point data for investigating the sequence of events in substorms within the near- and mid-tail regions.

4.2 Near-noon apogees: cusp

During the afternoon (UT time) of 20 February 2005, both Cluster and TC-2 pass through the northern Cusp region while TC-1 is near perigee on the nightside, as shown in the upper part of Fig. 16. Cluster is travelling sunward and TC-2 anti-sunward, along trajectories whose model magnetic footprints come close together at about 15:45 UT. Throughout the period of interest, ACE data show IMF $B_z < 0$, so that reconnection equatorward of the cusps may be expected.

The energy time spectra show that from 14:00 to 16:30 UT, Cluster moves onto cusp field lines from a more mantle-like region. During the same interval, TC-2 moves from magnetospheric regions where it sees trapped \sim keV plasma sheet electrons onto cusp field lines by 15:00 UT. Between 14:25 and 14:55 UT, the density observed by Cluster is lower than in the cusp proper (judging by the spacecraft potential) suggesting that the injected electrons seen during this period may be arriving in advance of the cusp ions. Preliminary evidence from the TC-2 plasma ion instrument (LEID) supports this supposition. In addition, particularly strong fluxes of electrons seen at 15:00 and 15:10 UT on TC-2 appear to occur at about the same time as similar intensifications at Cluster. More detailed examination shows that the signatures appear first at the equatorward spacecraft, (initially TC-2, later Cluster) suggestive of poleward motion, consistent with these being signatures of Flux Transfer Events. The predicted orbit of the (NASA) Polar spacecraft appears to pass through or near the northern cusp at a still lower altitude. The combination of the Cluster and TC-2 transects of the cusp, perhaps supplemented with additional data from Polar, offers a potential opportunity for detailed examination of the ion and electron injection signatures, possibly leading to a study of the dayside reconnection site location(s) using particle energy dispersion signatures (e.g. see Lockwood, 1995; Trattner et al., 2002). The multi-spacecraft data set also offers the opportunity to investigate the larger scale structure of the cusp at this time.

4.3 Near-dusk apogees: flanks

A significant solar wind density enhancement was seen at ACE between roughly 06:00 to 18:00 UT on 27 December 2004, coinciding between 06:00 and 13:00 UT with solar wind speeds in excess of 500 km/s. An enhanced solar wind pressure and greater magnetospheric compression would be expected during the corresponding interval \sim 50 min later at Earth. A similar situation occurred on the previous day.

Figure 17 illustrates the spacecraft locations and PEACE data for the period 14:00 to 16:00 UT on 27 December 2004. According to the model magnetosphere presented in the figure (which was not designed to represent the effects of an enhanced solar wind pressure), TC-2 is expected to be ap-

proaching the mantle region tailward of the northern cusp, while TC-1 is inbound from an apogee near the flank magnetopause. Cluster is travelling north after passing through what is predicted to be a mantle region into the dawn flank magnetosphere.

The TC-1 PEACE data clearly show periodic variations in the flux of 1–10 keV electrons repeating with a 5 min cycle. Such behaviour had been seen continuously for several hours previously. These are tentatively identified as being related to Pc-5 pulsations, supported by corresponding perturbations in the magnetic field and plasma ion density and velocity. Given the flank location of the observations, and the consistently high solar wind speed, one possible driver for such oscillations may be magnetopause waves driven by the Kelvin-Helmholtz instability. TC-2 data show a similar periodicity between 13:30 and 13:55 UT as TC-2 moves towards what appears to be the mantle behind the northern cusp (note the relatively low fluxes). At about the same time Cluster appears to be in the magnetosphere rather than the magnetosheath (based on preliminary magnetic field data) and so presumably the high electron fluxes at \sim 100 eV indicate that Cluster is in the southern cusp from 14:00 to 15:00 UT. Cluster and TC-2 are in the southern/dawn and northern/dusk regions, respectively, so these cusp and mantle observations provide simultaneous measurements of the northern and southern cusp regions as defined by plasma entry at this time. However, at ACE the IMF orientation shows a negative B_y and positive B_z component, so we might have expected the cusp feet to be displaced dawnward in the northern hemisphere (Fuselier et al., 2003; Cowley et al., 1991) and further work is needed to reconcile the observations and expectations. The appearance and disappearance of \sim 2 keV electrons between 14:50 UT and 15:25 UT appear to coincide with sharp field rotations at ACE (allowing for solar wind travel time to Earth). The dataset clearly has the potential to be used as a detailed test case for models of the response of the cusp to interplanetary magnetic field orientation variations.

4.4 Magnetospheric response to CME

A more dramatic solar wind pressure enhancement (due to an ICME) occurred on 5 December 2004. In this case the IMF is enhanced to over 30 nT and the density is also significantly elevated, both for \sim 8 h, although the solar wind speed increases by only \sim 100 km/s to \sim 450 km/s. The spacecraft locations are similar to the 27 December example discussed above, except that Cluster was further south, nearer apogee and so expected to be outside the magnetopause even under more normal solar wind conditions. Figure 18 shows PEACE data for the entirety of 5 December 2004.

TC-1 begins the day at perigee and PEACE is not turned on until outside the predicted $L=4$ L-shell. The arrival of the ICME driven shock shortly before 08:00 UT is clear as TC-1 enters the magnetosheath and remains in it for about 10 h before finally re-entering the magnetosphere after a series of multiple crossings of the magnetopause. These multiple

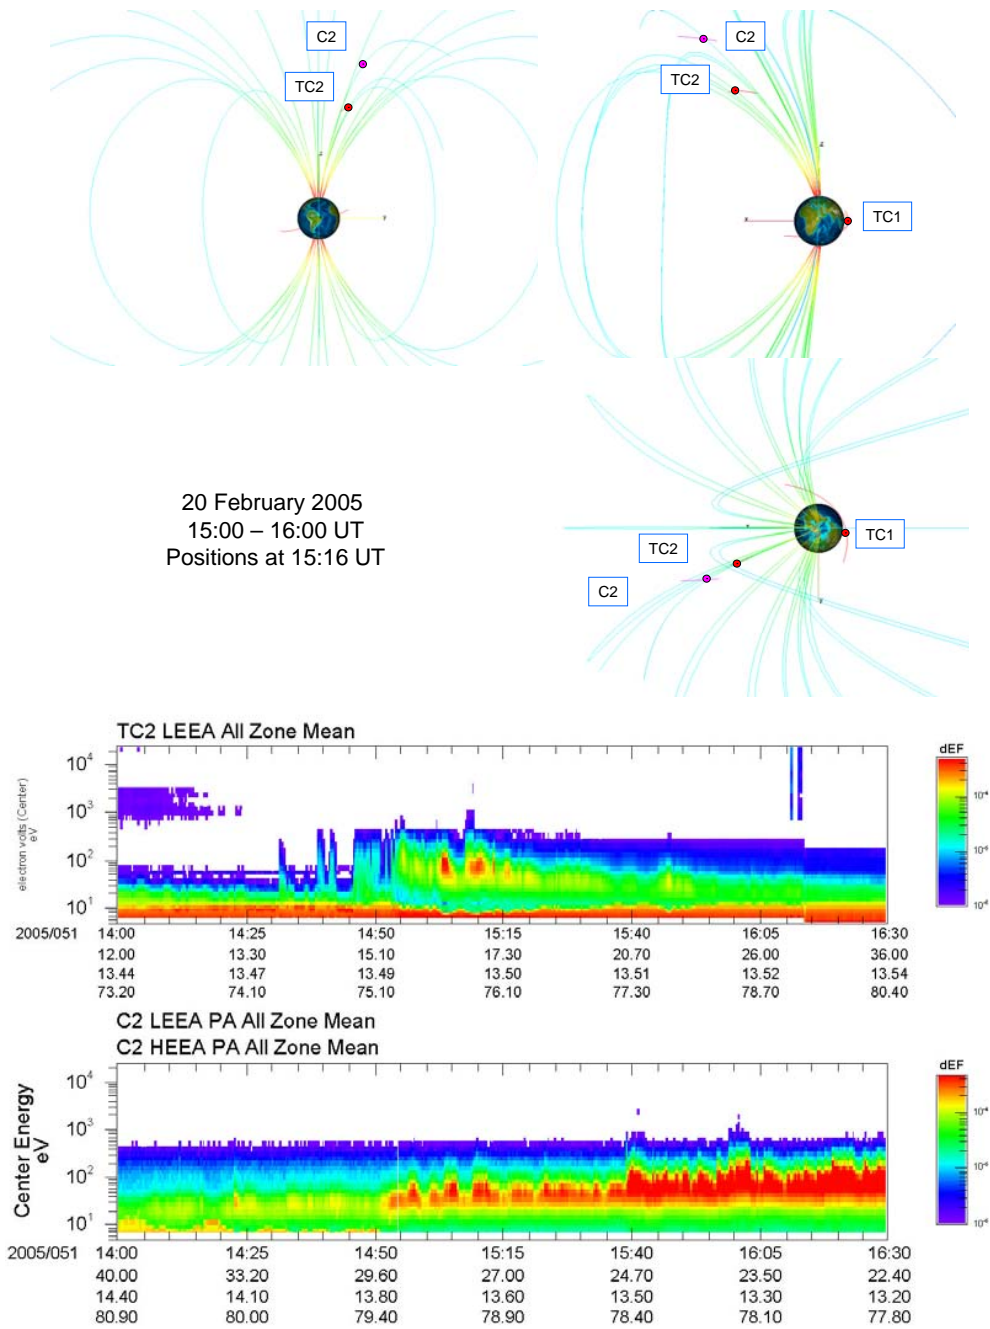
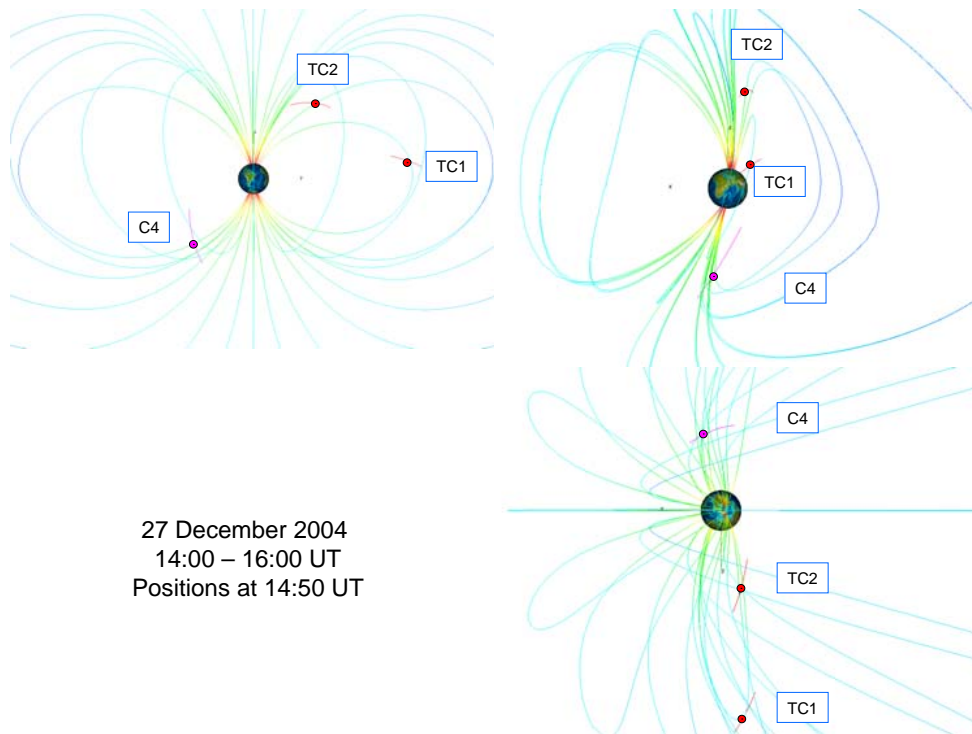


Fig. 16. Above, 20 February 2005, orbital positions and OVT model magnetosphere in GSM, views from X, Y and Z; below, PEACE data from TC-2 and C2 (TC-1 at perigee) X-axis labels show UT time, L-shell, MLT and invariant latitude.



27 December 2004
 14:00 – 16:00 UT
 Positions at 14:50 UT

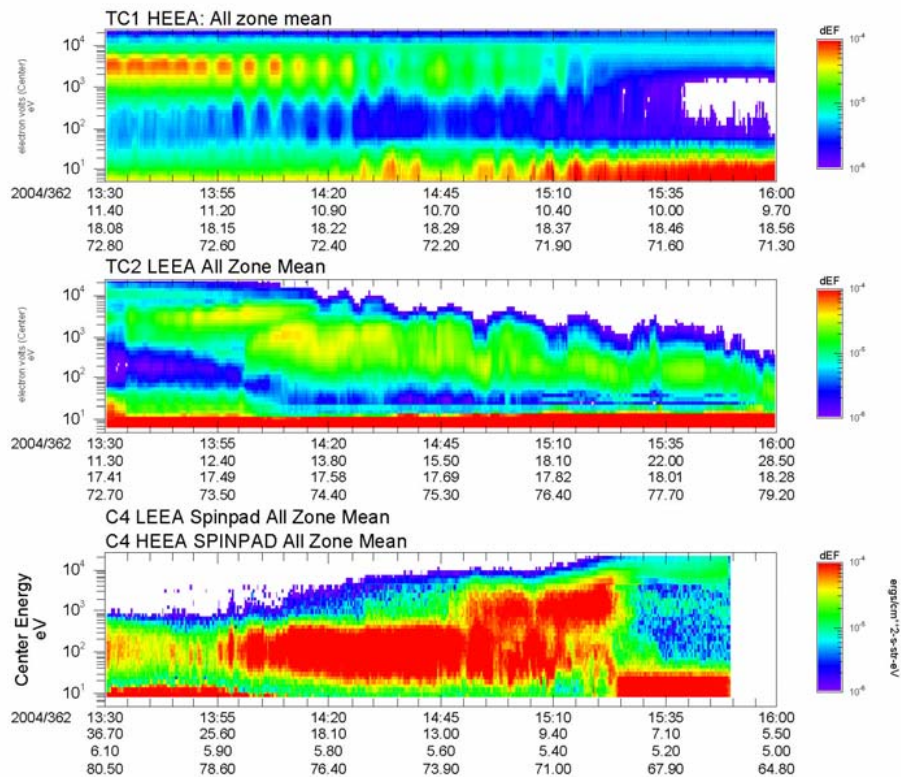


Fig. 17. Above, 27 December 2004, orbital positions and OVT model magnetosphere in GSM, views from X, Y and Z; below, PEACE data from TC-1, TC-2 and C4. X-axis labels show UT time, L-shell, MLT and invariant latitude.

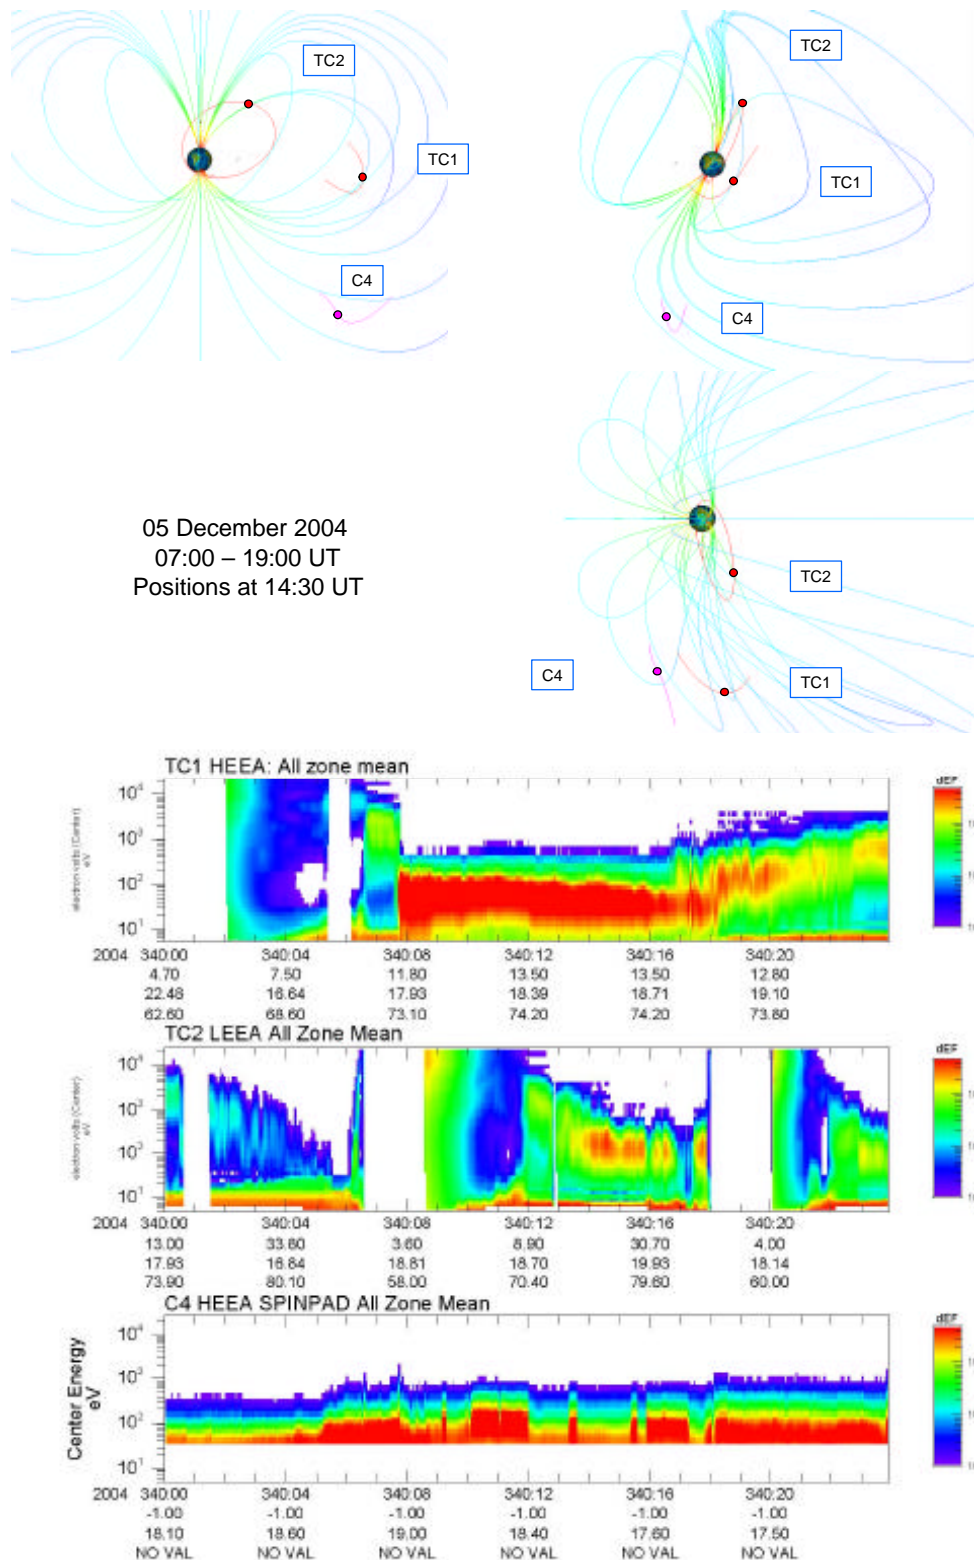


Fig. 18. Above, 5 December 2004, orbital positions and OVT model magnetosphere in GSM, views from X, Y and Z; below, PEACE data from TC-1, TC-2 and C4. X-axis labels show UT time, L-shell, MLT and invariant latitude.

crossings may again reflect magnetopause oscillations driven by fast-flowing solar wind. The magnetospheric electron population encountered thereafter by TC-1 is significantly colder than usual for that region (compare with 27 December) though its temperature rises steadily as TC-1 travels towards lower altitudes.

TC-2 is near perigee as the ICME shock arrives. The electron populations it encounters on three successive apogees are markedly different. At $\sim 01:00$ UT TC-2 sees a relatively low flux mantle plasma. After the ICME arrives it sees much denser plasma in the same region, with particularly strong enhancements after 14:00 UT, but one orbit later near the end of the day, the flux levels and temperatures are already falling. During this period the IMF is steadily northward, so cusp electron injections which may account for the aforementioned enhanced fluxes would be expected to originate at lobe reconnection sites. Such an explanation may account for their mean energy being significantly higher than the magnetosheath electron energy measured simultaneously by Cluster and TC-1. Alternatively the plasma may prove to belong to the same population as that which TC-1 sees when it returns to the magnetosphere. Sorting the electrons by pitch angle, which will be possible when the magnetic field data becomes available, may help clarify this point.

During this period Cluster is initially in the solar wind but spends significant amount of time in the magnetosheath despite the magnetospheric compression after $\sim 08:00$ UT. Further study is needed to understand what this tells us about the shape of the magnetosphere and the bow shock surface during this period. Cluster is near the bow shock as TC-1 re-enters the magnetosphere, providing a simultaneous reference for magnetopause and bowshock positions at that time.

Under prolonged periods of northward IMF, the Earth's plasma sheet becomes cold (Terasawa et al., 1997), but is rarely seen in the near Earth region except under very strong magnetospheric compression (Thomsen et al., 2003). There is an ongoing debate as to the source of this plasma (Borovsky et al., 1997). This event is currently under further study to address the question of the source of the cold dense plasma sheet.

5 Conclusions

In this paper we have described the Double Star PEACE instruments, concentrating on aspects which differ from Cluster PEACE instruments. Both Double Star PEACE instruments are returning good quality data. We have provided some examples of observations made by PEACE instruments on Cluster, and both Double Star spacecraft to illustrate the power of the combined dataset to investigate magnetospheric phenomena on the large scale.

Appendix A: Discussion of JPEG compression method

In Sect. 2.7.6 we introduced the JPEG Compression Code, and briefly outlined the method. In this appendix we discuss

the method in more detail.

(i) Decompose each basic segment into 8×8 "images"


The subdivision into 8×8 pixel images prior to data compression offers some useful advantages. For example, the contrast (difference between the highest and lowest values) will often be reduced in 8×8 images compared to larger images, which improves the compression rate and the quality of the reconstructed images. Also, by combining data of similar appearance (e.g. this is often the case for data collected at similar energy, polar and azimuth) we can design images with some symmetry, which also improves the compression rate. Our method involves some overlap of images, so that information sent twice in different images can be used to improve the quality of the reconstructed image after compression/decompression if necessary. The method is flexible enough to be used for all three sensor sweep modes (MAR, LAR or HAR). In all cases, 192 separate 8×8 images are compressed for each spin. As a basic segment contains data from 12 anodes, and in MAR mode 30 energies and 2 azimuths, the organisation of this data into 8×8 images is not quite straightforward. The mapping of data from a basic segment to 8×8 images is illustrated in Fig. A1 for the case of each sweep mode. In particular, for MAR we see that 12 images (labeled A to L) are made covering four energy bands for each of three groups of 3 anodes (polar zones).

(ii) Rescale the data in the images

In order to apply this compression technique within the PEACE DPU, where operations are performed on a 16-bit processor, it is necessary to rescale all counts data in images to values of 512 or less, before performing DCT operations. The inverse DCT operations can be computed (on the ground) as 32-bit words in case any result would slightly exceed the maximum value allowed by 16-bit computers (i.e. 32767). The peak count rate, P_{\max} , which is used to calculate the scaling factor, is determined using the full basic segment data including flyback data. Better compression may be possible if a different P_{\max} is used for each 8×8 image, but this would require that many more P_{\max} values are telemetered together with the data (to allow reversal of the scaling process in ground) and our resource constraints rule out this scenario. There is a possibility that no value in the images produced from the basic segment is set to 512 in the rescaling process, if the largest value in the basic segment is a flyback value. In that case, the effective range of values for image data would be smaller than 0 to 512, but since the flyback values are not expected to significantly exceed values in image data, the reconstruction should not be adversely affected. The value of P_{\max} is limited to a maximum of 8032 (since all values are capped at this level as a legacy of the Cluster quasi-logarithmic compression scheme) which is unlikely to be reached in a normal measurement. For P_{\max} less than 512 it is not in fact desirable to carry out the rescaling process, since the rescaled values will be larger, however the

Image


Each image consists of an 8 energy x 4 polar block from each of two azimuths, combined as shown with polars “reflected”, to optimise symmetry. Polars 1-4 from one azimuth are matched to polars 1-4 from the other azimuth, similarly for polars 5-8 and 9-12.



MAR Mode basic segment

30 energy x 12 polar x 2 azimuth


Four sets of 8 energies allows treatment of 32 energy bins, but there are only 30 to handle. Thus there is a small overlap in energy between images covering the lower energies.



HAR Mode basic segment

15 energy x 12 polar x 4 azimuth

Each azimuth only has 15 energies, so the images have an overlap in energy.



LAR Mode basic segment

60 energy x 12 polar x 1 azimuth

The 60 energies are separated into a set of 30 Even numbered energy bins and a set of 30 Odd numbered energy bins, which are then treated as for MAR mode




Fig. A1. Mapping basic segment data to 8x8 images for each sweep mode.

PEACE software constraints were so tight that a test to determine whether or not rescaling should be used could not be incorporated.

(iii) Convert each image to a DCT coefficient matrix

Two-dimensional DCTs are used to transform an image consisting of 64 pixels into the frequency domain where it is represented by an 8x8 matrix of 64 DCT coefficients:

$$\mathbf{DCT}(i, j) = \frac{C(i)}{2} \frac{C(j)}{2} \sum_{x=0}^7 \sum_{y=0}^7 \mathbf{pixel}(x, y) \times \cos\left(\frac{(2x+1)i\pi}{16}\right) \times \cos\left(\frac{(2y+1)j\pi}{16}\right)$$

$$C(i) = \frac{1}{\sqrt{2}} \text{ if } i = 0, \text{ else } 1 \text{ if } i > 0$$

Most images are composed of low-frequency information on to which higher frequency components are superimposed. The elements of the **DCT(i,j)** matrix with low *i, j* are coefficients representing the low frequency information, while elements with larger values of *i* and/or *j* represent higher frequency information. The original image can be recovered by applying the inverse operation:

$$\mathbf{pixel}(x, y) = \sum_{i=0}^7 \frac{C(i)}{2} \sum_{j=0}^7 \frac{C(j)}{2} \mathbf{DCT}(i, j) \times \cos\left(\frac{(2x+1)i\pi}{16}\right) \times \cos\left(\frac{(2y+1)j\pi}{16}\right)$$

$$C(i) = \frac{1}{\sqrt{2}} \text{ if } i = 0, \text{ else } 1 \text{ if } i > 0$$




Fig. A2. Bit stream treatment of an 8×8 image.

The large number of additions and multiplications involved in these formulas are considerably reduced by the implementation of fast DCTs (Pennebaker and Mitchell, 1993).

(iv) Apply a weighting factor to the DCT coefficients

A weighting process is introduced at this stage which reduces the significance of higher order coefficients. Fine details of the image may thus be less effectively handled, but the process ensures that the more important lower order coefficients can be well handled by the next stage in the telemetry scheme. The weighted coefficients $\mathbf{W}_{\text{DCT}}(\mathbf{i}, \mathbf{j})$ are given by

$$\mathbf{W}_{\text{DCT}}(\mathbf{i}, \mathbf{j}) = \frac{\text{DCT}(\mathbf{i}, \mathbf{j})}{\mathbf{WF}(\mathbf{i}, \mathbf{j})}; \text{ with } i = 0, 7 \text{ and } j = 0, 7$$

where the weighting factor, $\mathbf{WF}(\mathbf{i}, \mathbf{j})$ is usually chosen as:

$$\mathbf{WF}(\mathbf{i}, \mathbf{j}) = 16 \times ((1 + ((1 + i + j) \times Q)));$$

with $i = 0, 7$ and $j = 0, 7$

so that the process imposes a weight on the coefficients and this weight increases with the rank $(i + j)$ in the matrix, i.e. with the coefficient frequency. After the weighting has been applied, the low-frequency coefficients which carry the fundamental information are guaranteed to be larger than the higher frequency coefficients. In the bit streaming process to follow, this ensures that the lower frequency coefficients will be coded using a larger number of bits and so will be transmitted more faithfully. For the expansion phase, in ground data processing, the expression is similarly:

$$\text{DCT}(\mathbf{i}, \mathbf{j}) = \mathbf{W}_{\text{DCT}}(\mathbf{i}, \mathbf{j}) \times \mathbf{WF}(\mathbf{i}, \mathbf{j});$$

with $i = 0, 7$ and $j = 0, 7$

which restores the more important low-frequency coefficients much more accurately than the higher frequency ones.

Table A1. Bit streaming: entropy encoding.

Bit count	Value amplitude
1	-1, 1
2	-3 to 2, 2 to 3
3	-7 to -4, 4 to 7
4	-15 to -8, 8 to 15
5	-31 to -16, 16 to 31
6	-63 to -32, 32 to 63
7	-127 to -64, 64 to 127
8	-255 to -128, 128 to 255
9	-511 to -256, 256 to 511
10	-1023 to -512, 512 to 1023

In the implementation used for PEACE, the “quality factor” Q in the weighting expression has been set to 1, the smallest practical value within system constraints, but detailed verification that this is the optimum value has not been carried out at the time of writing.

(v) Bit-stream the DCT coefficients into PEACE Science TM packets

The final step of the compression procedure is the encoding of the quantized 8×8 DCT coefficient matrix as a bit stream. The matrix is first read by a zig-zag sequence from the low to high frequency coefficients, as illustrated in Fig. A2. Thus the matrix is read in priority order, starting with the most significant, lowest rank coefficients.

Entropy encoding is then used to compress the DCT coefficients into a bit stream. The technique involves using a number (“bit count”) to identify the range of numbers that a DCT coefficient belongs within, and then more bits to identify where in that range it lies, as illustrated in Table A1. The coefficient can often be represented by fewer bits using this method than if it had been represented by a standard 16-bit number, especially when the coefficient has a small value. When many coefficients are small or zero, the resulting compression is particularly effective.

As a result, a bit stream is produced for each image, which is a series of priority ordered, entropy encoded coefficients. The process is readily reversed during decompression on the ground. The bit stream must be incorporated into PEACE Science Telemetry packets to be collected by the PDMS telemetry handler.

The PEACE science packets each hold data from two consecutive basic segments, i.e. a total of 24 8×8 images, as well as a packet overhead (which includes the synchronisation pattern) and the packet checksum. The PEACE science packets are designed to incorporate the JPEG bit stream data, together with the number of coefficients that have been telemetered for each image (in order to verify the decoding of the bit stream on the ground) and the maximum count value for each basic segment (P_{max} , used to rescale the data in the decompression stage).

The telemetry rate for PEACE Science Data is fixed, thus there is a maximum size of PEACE science packet which is consistent with sending a spin of data every spin, which corresponds to a constraint that only 10 bytes/image are available to store the encoded DCT coefficients. To maximise the usage of the available space, the JPEG bit stream data for a given image is loaded into the packet until the available 10 bytes is filled. If the full set of coefficients for a particular image cannot be stored within 10 bytes, then a truncated set is transmitted for that image. If all 64 coefficients of the image can be stored and the space is not completely filled, the extra room is available for the next set of coefficients, and so on until all 12 images in the basic segment are processed. In practice, quite complex images are usually found at low energies (e.g. representing the photoelectron/plasma boundary) which will be represented by many quite large coefficients and so any unused bytes from images recorded at higher energies in the sweep will usually be mopped up by the low energy data.

The constraint of 10 bytes/image can be increased by telecommand to as much as 15 bytes/image (this is the upper limit for Double Star PEACE, due to a buffer constraint) in order to allow more coefficients to be transmitted and hence, for cases where truncation is found at 10 bytes/image, a more faithful decompression becomes possible. The penalty in this case is that the PEACE science packets are larger, and so fewer complete spins can be transmitted each minute. An example of the consequent missing spins is seen in Fig. 11c, centered at ~06:50:47 UT. Note that even with 10 bytes/image, some missing spins occur in practice.

A consequence of our method is that if the complete JPEG bit stream is not properly received, it is not possible to know which bits represent the bit count and which represent the values, making data reconstruction impossible. It is similarly not possible to know where an 8×8 block starts and ends, if the data has been corrupted. Checksums on PEACE data packets, are used to give high confidence (but not a 100% guarantee) that we know when the bitstream is uncorrupted, for each packet.

Acknowledgements. The PEACE team thanks the Chinese National Space Agency and the European Space Agency for making the Double Star project a reality. We thank the engineers and scientists of the Centre for Space Science and Applied Research (CSSAR), under the leadership of J. Wu, who created the PDMS system which provided Cluster-style interfaces to the European Cluster-heritage instruments, amongst other things, and whose professional approach was essential to the collaboration. We also thank the engineers of the Chinese Academy of Space Technology (CAST), in particular S. Yuan, for striving to produce spacecraft meeting the stringent mission requirements. We particularly thank the ESA team and their Astrium colleagues for their unstinting efforts in support of the European payload, and the CSSAR engineer assigned to the PEACE team (our co-author Y. Q. San) whose help was invaluable. We thank all those involved in the Commissioning and now in Operations Phase of the mission, in China and Europe. We thank the OVT Team (<http://ovt.irfu.se>) for their useful tool, used to produce orbit plots in this paper.

The ESA funded the work to provide, deliver and test flight-standard Double Star PEACE instruments. The UK Particle Physics and Astronomy Research Council supports Double Star operations.

The PEACE PI would especially like to thank the members of the team which produced the DS PEACE instruments, typically contributing far beyond their formal working hours.

Topical Editor T. Pulkkinen thanks G. Parks and G. Paschmann for their help in evaluating this paper.

References

- Borovsky, J. E., Thomsen, M. F., and McComas, D. J.: The Superdense plasma sheet: Plasmaspheric origin, solar wind origin, or ionospheric origin?, *J. Geophys. Res.*, 102, A10, 22 089–22 097, 1997.
- Carr, C. M., Brown, P., Zhang, T. L., Gloag, J., Horbury, T., Lucek, E., Magnes, W., O'Brien, H., Oddy, T., Auster, U., Austin, P., Aydogar, O., Balogh, A., Baumjohann, W., Beek, T., Eichelberger, H., Fornaçon, K.-H., Georgescu, E., Glassmeier, K.-H., Ludlam, M., Nakamura, R., and Richter, I.: The Double Star magnetic field investigation: instrument design, performance and highlights of the first years observations, *Ann. Geophys.*, 23, 2713–2732, 2005.
- Cornilleau-Wehrin, N., Alleyne, H. S. C., Yearby K. H., de la Porte de Vaux, B., Santolik, O., Parrot, M., Belmont, G., Rezeau, L., Le Contel, O., Roux, A., Attié, D., Robert, P., Bouzid, V., Herment, D., and Cao, J.: The STAFF-DWP wave instrument on the DSP equatorial spacecraft: description and first results, *Ann. Geophys.*, 23, 2785–2801, 2005.
- Cowley, S. W. H., Morelli, J. P., and Lockwood, M.: Dependence of convective flows and particle participation in the high latitude dayside ionosphere on the X and Y components of the interplanetary magnetic field, *J. Geophys. Res.*, 96, 5557–5564, 1991.
- Dunlop, M. W., Taylor, M. G. G. T., Davies, J. A., Owen, C. J., Fazakerley, A. N., Pitout, F., Pu, Z., Laakso, H., Zong, Q.-G., Bogdanova, Y., Shen, C., Nykyri, K., Cargill, P., Carr, C. M., Escoubet, C. P., Lavraud, B., Lockwood, M., Milan, S. E., Phan, T. D., Rème, H., and Sonnerup, B.: Coordinated Cluster/Double Star observations of dayside reconnection signatures, *Ann. Geophys.*, 23, 2867–2875, 2005.
- Fuselier, S. A., Mende, S. B., Moore, T. E., Frey, H. U., Petrinec, S. M., Clafin, E. S., and Collier, M. R.: Cusp dynamics and ionospheric outflow, *Space Sci. Revs.*, 109, 285–312, 2003.
- Johnstone, A. D., Alsop, C., Burge, S., Carter, P. J., Coates, A. J., Coker, A. J., Fazakerley, A. N., Grande, M., Gowen, R. A., Gurgiolo, C., Hancock, B. K., Narheim, B., Preece, A., Sheather, P. H., Winningham, J. D., and Woodliffe, R. D.: Peace: A Plasma Electron and Current Experiment, *Space Sci. Rev.*, 79, 351–398, 1997.
- Lockwood, M.: Location and characteristics of the reconnection X line deduced from low-altitude satellite and ground-based observations, 1. theory, *J. Geophys. Res.*, 100, A11, 21 791–21 802, 1995.
- Marchaudon, A., Owen, C. J., Bosqued, J.-M., Fear, R. C., Fazakerley, A. N., Dunlop, M. W., Lahiff, A. D., Carr, C., Balogh, A., Lindqvist, P.-A., and Rème, H.: Simultaneous Double Star and Cluster FTEs observations on the dawnside flank of the magnetosphere, *Ann. Geophys.*, 23, 2877–2887, 2005.
- Nelson, M. and Gailly, J.-L.: *The Data Compression Book*, M&T Books, New York, 1995.

- Pennebaker, W. B. and Mitchell, J. L.: JPEG still image data compression standard, Van Nostrand Reinhold, New York, 1993.
- Pu, Z. Y., Xiao, C. J., Huang, Z. Y., Fu, S. Y., Liu, Z. X., Gang, Z. Q., Dunlop, M. W., Carr, C. M., Reme, H., Dandouras, I., Fazakerley, A., Phan, T., Zhang, T. L., Zhang, X. G., Zhang, H., and Wang, X. G.: Double Star TC-1 observation of component reconnection at the dayside magnetopause: a preliminary study, *Ann. Geophys.*, 23, 2889–2895, 2005.
- Sawyer, D. W and Vette, J. I.: AP-8 Trapped Proton Environment for Solar Maximum and Solar Minimum, National Space Science Data Center, NASA Goddard Space Flight Center, NSSDC/WDC-A-R&S, 76–06, 1976.
- Terasawa, T., Fujimoto, M., Mukai, T., Shinohara, I., Saito, Y., Yamamoto, T., Machida, S., Kokubun, S., Lazarus, A. J., Steinberg, J. T., and Lepping, R. P.: Solar Wind control of density and temperature in the near-Earth plasma sheet: WIND/GEOTAIL collaboration, *Geophys. Res. Lett.*, 24, 935, 1997.
- Thomsen, M. F., Borovsky, J. E., Skoug, R. M., and Smith, C. W.: Delivery of cold, dense plasma sheet material into the near-Earth region, *J. Geophys. Res.*, 108, A4, 1151, doi:10.1029/2002JA009544, 2003.
- Torkar, K., Arends, H., Baumjohann, W., Escoubet, C. P., Fazakerley, A., Fehringer, M., Fremuth, G., Jeszenszky, H., Laky, G., Narheim, B. T., Riedler, W., Rüdener, F., Steiger, W., Svenes, K., and Zhao, H.: Spacecraft potential control for Double Star, *Ann. Geophys.*, 23, 2813–2823, 2005.
- Trattner, K. J., Fuselier, S. A., Peterson, W. K., Boehm, M., Klumpar, D., Carlson, C. W., and Yeoman, T. K.: Temporal versus spatial interpretation of cusp ion structures observed by two spacecraft, *J. Geophys. Res.*, 107, A10, 1287, doi:10.1029/2001JA000181, 2002.