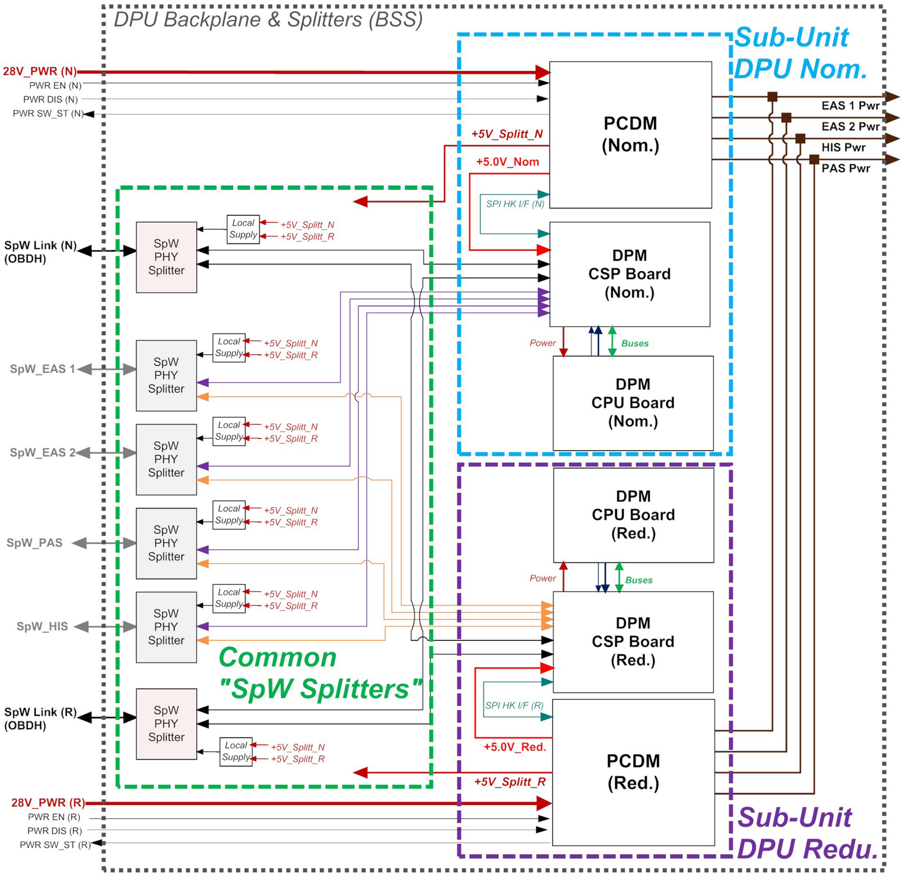

Fig. 28. Schematic of the SWA-DPU hardware architecture.

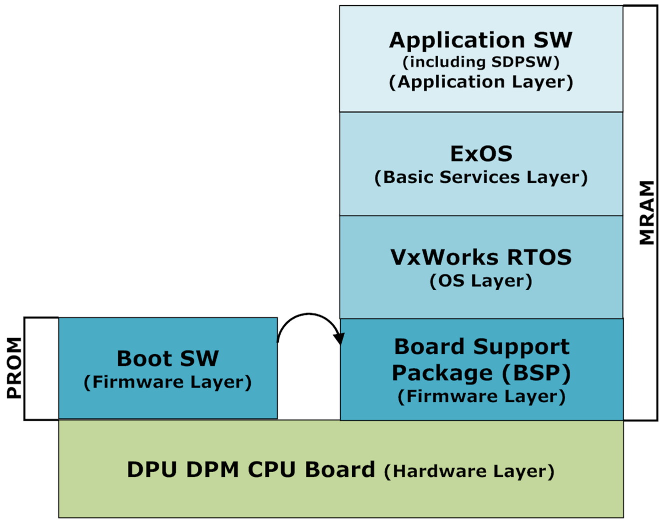

Fig. 29. SWA-DPU top-level software architecture.

the BSW initialises the communication through the main and redundant SpW interfaces with the spacecraft on board computer (OBC) in order to manage a reduced set of TM and TC services. These include Service 9 TC, to synchronise the SWA-DPU internal reference time with the spacecraft event time (SCET), Service 3 TM, to generate and transmit the SWA-DPU HK telemetry, and Service 6 TC, to perform check, dump, and load of memory blocks (for FSW patching).

The FSW is composed of three software layers: the application software layer (ASW), the application service software layer (ASSW), and the machine services software layer (MSSW). The ASW includes SWA-DPU TC and TM and process control management software (PCMSW), and the scientific data processing software (SDPSW). The PCMSW is a logical composition of sub-components which manage TC validation and execution, SWA-DPU and sensor state management, the SWA sensors (SWA-EAS1, SWA-EAS2, SWA-PAS, and SWA-HIS) themselves, and the science data acquisition. It also controls algorithm execution on acquired scientific data (i.e. moment computation, compression, etc.), SWA-DPU HK data handling, FDIR, and time management. The SDPSW provides post-processing functionalities on the SWA-EAS and SWA-PAS acquired science data (moment calculation, raw data compression, formatting of compressed stream for downlink). The ASSW includes the packet utilisation standard library software (PUS) and mission services for the ASW. It provides the set of functions for implementing the PUS services required for the SWA-DPU ASW. The MSSW includes the real time operating system (RTOS) kernel, the extended operating system service (EXOS) software, and the low level drivers in the board support package (BSP). The RTOS provides the SWA-DPU DPM runtime resource management, the basic mechanisms for task execution, and inter-process communication. The EXOS provides a library used by the ASW and ASSW layer components to gain access to the SWA-DPU hardware resources (memory, registers, time, SpW, sensors, etc.). The EXOS also efficiently executes some basic functions for the configuration and management of the system and its various subsystems. It provides appropriate functionalities to support the exchange of information among the subsystems and between the SWA-DPU and spacecraft. The BSP provides access services to all SWA-DPU internal hardware and physical interfaces of the system to the above software layers.

3.4.4. Scientific data processing

The SDPSW is the component handling the processing of data collected by the SWA-EAS (1 and 2) and SWA-PAS sensors, while SWA-HIS processes scientific data on its own. Due to the limited bandwidth available to SWA for downlink, the full set of SWA raw data collected cannot be transmitted back to ground at all times. Indeed, even to return full resolution data products at relatively low time cadences requires compression rates (CR) in the range from 2 to 7.6, depending on the different kinds of data and their associated data product volumes. In order to return information on the nature of the solar wind at higher cadences, data processing within the SWA-DPU is used to evaluate derived scientific properties of the solar wind, in particular the moments of the particle VDF (Paschmann et al. 1998). These quantities are transmitted instead of the full raw data, with the transmission of full VDFs only at lower cadence or after appropriate compression. The processing of SWA-EAS1, SWA-EAS2, and SWA-PAS scientific data generated during their normal and burst modes is handled by the PCMSW that calls the services offered by the SDPSW, within its cyclic activities routine. Accordingly, the execution of scientific algorithm requests, such as "Compress Data Product" or "Calculate Moments", is distributed across several different run-time intervals. This is managed by means of dedicated activity tables defining the actions the SDPSW has to perform for each "activation" by the PCMSW. These tables define the software behaviour for both compression and moment calculation functions for each sensor and for each operating mode, according to the specific data collection rate and order. In this way, the SDPSW prepares data as they are ingested by the SWA-DPU from sensors, and, when the complete data sample to be processed has been reconstructed, the SDPSW's services start data compression (or moments calculation) and continue, with cyclic activation, until that operation is completed.

On board moments calculation. The SDPSW provides the functionality for the calculation of the number density (moment of order zero), number flux density (first order moment), pressure tensor (second order moment: computed in the satellite reference frame for SWA-EAS and SWA-PAS), and energy flux density vector (third order moment: computed only for SWA-EAS). Moments are computed starting from the counts accumulated in each elementary volume in phase space and with reference to the sensor' resolution in energy, azimuth, and elevation angles (Paschmann et al. 1998). From an operational point of view, all the equations defining the moments have been captured within a series of LUTs allowing the performance of the moments calculation by means of only sums and products, with the counts of each elementary volume in phase space "modulated" by a combination of these factors. This algorithm has an initialisation phase that is performed only at SWA-DPU start-up and must be repeated each time the configuration values are updated from ground. It is a deterministic set of computations, and can be completely verified on-ground. Algorithm configuration values are read from memory (MRAM) and used to evaluate a set of "constant" parameters used as LUTs at run-time. Some of these tables contain configuration parameters derived on the ground from the sensor physical and geometrical properties. Others parameters which are expected to show slow time variation are held in LUTs which can be readily updated from the ground. For some parameters, particularly those which are directly linked with the way sensors are commanded to sweep energies or elevations, the same inputs are shared for both sensor and SDPSW configurations. For example, both the SWA-EAS-Sequencer and the SWA-DPU-SDPSW rely on the same ground commandable parameters defining the hemisphere voltage ratio and maximum voltage. The output is the same hemisphere LUT composed of 64 values expressed in volts and eV. Other commandable LUT parameters include those fixing the range of variability and the precision at which a given set of moments will be telemetered to ground; this is done by defining the couplet of information (LSB, Offset) to be applied to the raw and scientific data in order to define the transmitted value within the data packet.

The calculation of accurate moments from the SWA-EAS sensor is dependent on identifying and removing data from those low energy bins which may be contaminated by the presence of spacecraft photo-electrons. Since such electrons are generally found at energies below the spacecraft potential, the baseline algorithm involves the use of a measurement of that potential passed to the SWA-DPU from the RPW experiment (Maksimovic et al. 2020), if it proves reliable, or otherwise using a fixed but ground-commandable level. This information is transmitted via the Service 20 IEL. The moment calculation then does not include SWA-EAS1 and SWA-EAS2 energy bins with ranges which are partially or fully below the value of the spacecraft potential. Moreover, the ranges of the energy bins above the spacecraft potential are reduced by the value of the spacecraft potential. In addition, moments are calculated as partial summations across three energy ranges (see Sect. 4.3 for details), so that even if the lower range is contaminated, some useful science data will remain in the upper ranges.

Data compression. Representative sample data have been produced by extrapolating real measurements by the Cluster mission to the conditions in which SWA is expected to operate. These have been evaluated for their information content, using both the data entropy measure and the actual CR. Results from this exercise demonstrate that SWA-EAS data CR can vary between ~1.9 and ~12.1, while for SWA-PAS data this ranges from ~3.4 to ~17.5. Data compression is performed by means of a customised implementation of the lossless Consultative Committee for Space Data Systems (CCSDS) 121.0 standard. This choice has been made as a trade-off between compression and computational efficiency. In fact, the CCSDS 121 technique allows full exploitation of the intrinsic structure exhibited by all data products, together with information theory methods, while remaining within the computational resource limitations. In particular, a specific customisation was designed for SWA-EAS data in order to reach the compression ratio needed, by detailed consideration of the specific data structure. SWA-EAS samples produce 3-dimensional matrices (cubes) whose dimensions cover the elevation and azimuth of incoming particle directions and their energy levels. These cubes are stored in a rolling buffer, within the SWA-DPU's memory, in a specifically defined one-dimensional array, mostly dependent on the acquisition sequence (simply meaning the first sample acquired is the first in the array). However, a preferential scanning sequence, which is different to that of the actual acquisition, has been identified from simulations using existing solar wind data. This sequence is identified by measuring similarities in each direction within the 3D data in order to identify which demonstrates the slowest variation rate. Thus, a further processing step is preliminarily applied to sampled SWA-EAS electron distributions in order to increase the algorithm's performance in terms of better prediction efficiencies, smaller prediction errors, and ultimately higher achievable compression ratio. This "complex reordering" mechanism has thus been designed in order to exploit the preferential direction in data similarity. It switches the data order after acquisition (this mechanism is identified as the "simple" reordering step), and then ensures that the highest degree of spatial continuity is established between contiguous samples (this mechanism is identified as the "complex" reordering step).

For SWA-EAS data, it has been found that a simple reordering of the data from the original order, that is elevation angle, energy level, azimuthal angle to energy level, azimuthal angle, elevation angle, brings clear improvements in the compression efficiency. In addition, the complex reordering step further improves compression performances by avoiding periodic jumps between acquisition directions. Re-ordered data are then passed to the unit-delay predictor and the standard pre-processing module. In the final compression step, the coder applies Rice's technique, in which several algorithms are concurrently applied to a block of consecutive pre-processed samples. The algorithmic option that yields the shortest encoded length for the current block of data are selected for transmission.

Book-keeping algorithm (BKA). The SWA-DPU ASW is responsible for controlling data collection, mode use, and imposing telemetry generation restrictions on each of the three SWA sensors separately, in order to keep each of them within the respective assigned allocations. In principle, during burst modes the sensors generate raw data at a rate which is significantly higher than their orbit-averaged allocation, while in normal mode they generate data products at a slightly slower rate than is consistent with the orbit allocation for each sensor. Thus, the book-keeping algorithm (BKA) is a software tool able to monitor and control the amount of burst mode used against the pro-rata expectation for any given point along the orbit. The principles underpinning the BKA are:

1. The BKA will be used by the ASW to assess the generation rate of science data by each of the three SWA sensors over an established time interval that starts at time T_0 and ends at time $T_0 + \Delta T$. ΔT is variable to allow for lessons learnt in flight, but the initial baseline should be the orbital period of the spacecraft or duration of a stable telemetry corridor (Sanchez et al., in prep.);

2. The ASW will hold record of two limits per sensor, set by the SWA team (changeable in flight to account for lessons learnt) representing: (a) the fractional level, O_S, against which the sensors may be allowed to become overdrawn against the pro-rata allocation, and; (b) the fractional level, U_S, against which an unacceptable "under-drawing" against the pro-rata allocation is deemed to have occurred;

3. At regular intervals the ASW will update the accumulated total volume, $V_{\mathrm{BM,S}}(t)$, of post-processed data which has been originated by sensor S since time T_0.

4. At regular intervals, the ASW will calculate $A_{\mathrm{BM,S}}(t)$, the expected pro-rata data accumulation for each sensor, S, since time T_0, based on the orbit-averaged allocation, $A_{\mathrm{BM0,S}}$, for that sensor:

$$A_{\mathrm{BM,S}}(t) = A_{\mathrm{BM0,S}} \times \frac{(t - T_0)}{\Delta T} \tag{1}$$

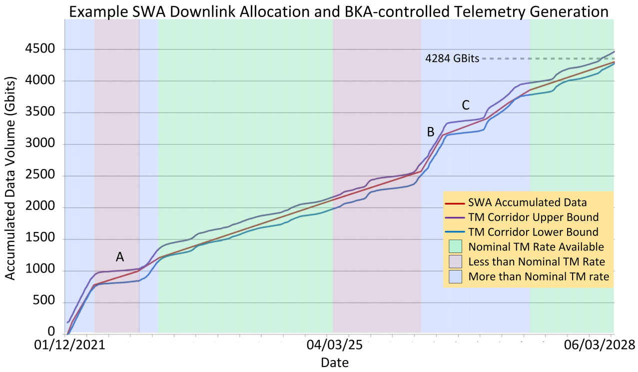

Fig. 30. Operation of the SWA BKA. Mission periods occur when less than (pink shading), more than (blue), and actual (green) nominal data rates will be available. The SWA BKA will steer data production (red line) through a project defined telemetry corridor with upper and lower limits (purple and blue lines respectively). Periods of low download rates (e.g. marked A and C) can be optimally negotiated by the BKA using higher rates on ether side (e.g. at point B).

5. The ASW will ensure that each sensor, S, does not produce so much BM data that the difference between the actual accumulated total volume, $V_{\mathrm{BM,S}}(t)$, of data from sensor S which has been sent to the spacecraft solid state mass memory (SSMM) since time T_0, and the pro-rata orbit allocation $A_{\mathrm{BM,S}}(t)$ does not exceed the fraction O_S of the remaining allocation. If the fraction is exceeded, then the SWA-DPU will disable trigger event capture and disable further optional scheduled burst modes until the excess is reduced to a factor of, at most, $M \times O_\mathrm{S}$. Specifically, the SWA-DPU will disable optional scheduled BM and trigger event capture if:

$$V_{\mathrm{BM,S}}(t) > A_{\mathrm{BM0,S}} \times \left[\frac{(1 - O_\mathrm{S})(t - T_0)}{\Delta T} + O_\mathrm{S} \right] \tag{2}$$

and will subsequently re-enable schedule BM and trigger event capture once:

$$V_{\mathrm{BM,S}}(t) < A_{\mathrm{BM0,S}} \times \left[\frac{(1 - MO_\mathrm{S})(t - T_0)}{\Delta T} + MO_\mathrm{S} \right] \tag{3}$$

6. In a similar way the SWA-DPU will ensure that $V_{\mathrm{BM,S}}(t)$ does not fall below the fraction U_S of the remaining allocation. If the fraction is not achieved, then the SWA-DPU will enable additional scheduled burst modes until the underspend is reduced to a factor of less than $M \times U_\mathrm{S}$ enabling additional scheduled BM;

7. In any case, the assessment period can be restarted once the period ΔT has elapsed. At this time, the accumulated data should have only a relatively small difference from the maximum allowed total for the orbit or period. Thus the assessment can be restarted by carrying over the small difference to the next assessment period;

8. The BKA is also able to handle ground commands for the ASW operation, which set the trigger-enable flag and control the amount of scheduled burst mode for limited specific periods, and automatically recover the required average telemetry rate in the following period;

9. All the parameters controlling the operation of the BKA are configurable in flight to allow the ASW to control the data production when the available telemetry rate is reduced below the nominal level. In practice this is likely to require (temporary) changes to the parameters O_S, U_S, and $A_{\mathrm{BM0,S}}$ by ground command for the duration of such periods, with possible compensatory changes to allow catch-up outside these periods;

10. The ASW is able to restart the BKA and resume correct operation in the event of a reboot or restart of the SWA-DPU.

The requirement to steer the SWA data accumulation through a telemetry corridor (defined for any given period at mission level) is equivalent to choosing a particular set of the BKA configuration parameters defined above. This will allow the SWA BKA to control SWA telemetry generation to remain within a defined telemetry corridor. This is illustrated in Fig. 30.

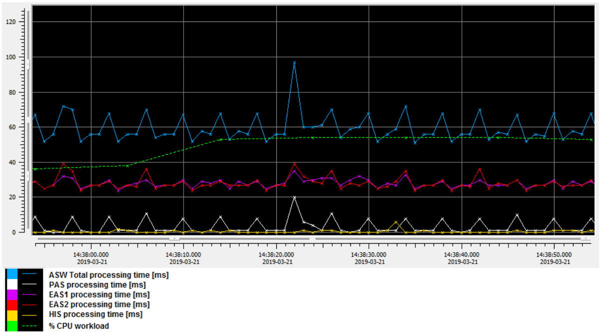

Fig. 31. Illustration of SWA-DPU performance. DPU activity enhancements associated with the 4 s cadences of SWA-EAS and SWA-PAS in normal mode are clearly visible.

3.4.5. SWA-DPU testing and characterisation

Figure 31 shows the ASW performance in the worst case condition occurring in periods in which all sensor scientific activities are activated, the compression algorithm is enabled, and a freeze of the trigger mode rolling buffer has been activated following the simulated receipt of a flag from the RPW instrument (Maksimovic et al. 2020) via PUS Service 20. Even under these conditions, the ASW performance remains under the required limit. In particular, we note that during each ASW processing slot (ASW activity cycles at 8 Hz) the ASW total processing time is lower than the maximum allowed time of 125 ms with the maximum time measured being 97 ms. In addition, the SWA-EAS and SWA-PAS processing times are largely uniform during science processing, and peaks are due to moment computation performed each 4 s and full 3D distribution extracted each 100 s. In contrast, SWA-HIS contribution is negligible as the ASW acts just as pass through for incoming SWA-HIS telemetry packets. Overall, typically the ASW CPU load is under 40% with a peak of 56% which is reached during RPW-trigger buffer freeze operations.

The data depicted in Fig. 31 are derived from housekeeping data dumped at its maximum frequency of 1 Hz (reporting the ASW processing time) and represents the maximum processing time measured in an observation window of 1 s for the management of each of the four SWA sensors (SWA-EAS1, SWA-EAS2, SWA-PAS, and SWA-HIS), and for the sum of all ASW activities (including the sensor management). In conclusion, the processor does not need to work under stressed conditions. It has more than 40 % margin in the worst case, and the time consumed by the ASW processing tasks is well under its maximum limit of 125 ms. So the risk of breaching operational limits, and the consequent possibility of losing a processing slot and the associated science data is very low (an overrun event report is generated in any such case).

4. SWA science operations

The suite of three SWA sensors plus its SWA-DPU will be operated in a variety of modes in order to address the overall Solar Orbiter mission science goals (Müller et al. 2020; Zouganelis et al. 2020). The SWA operations activities are distributed across three centres, primarily to allow key hardware institutes for the three sensors to address the health and calibration issues of their relevant sensor. The primary SWA operations group is based at UCL MSSL, and will take lead responsibility for SWA-level communications, coordinating activities across the suite, and for processing of the SWA-EAS data products. In this case, "communications" includes liaison and data exchange with the Mission Operations Centre (MOC) at the European Space Operations Centre (ESOC) in Darmstadt, Germany, the Solar Orbiter Science Operations Centre (SOC, Sanchez et al., in prep.) near Madrid, the Solar Orbiter Archive, as well as with partner SWA operations groups in Toulouse (for SWA-PAS related activities) and UMich (for SWA-HIS related activities). This section summarises the key operational issues, mode, and data product information likely to be of use to the scientific user of SWA data.

4.1. SWA telemetry

The baseline operations planning driving the original design of the SWA sensors and their operations was based on returning an orbit average data telemetry rate of ~14.5 kbps. Considering the required data products, the cadences at which they would be required, and the useful duty cycle for higher time resolution (burst mode) data products, a baseline target data production rate for the three sensors was established as ~5.5 kbps for SWA-HIS, ~4.5 kbps for SWA-PAS, ~4.3 kbps for SWA-EAS, and ~300 kbps for the SWA-DPU. Given the uncertainties associated with both the CRs achievable in flight and with the capture of data products associated with irregular triggers, the SWA team planned to maintain broad compliance with these rates by monitoring and internally controlling data taking by the SWA-DPU through the operation of the BKA (see Fig. 30 and associated text). Under this scheme, periods of poor compression ratio or numerous trigger responses would be compensated for by the cancellation of some planned burst mode intervals and switching off of the response to further trigger flags. In extreme cases, autonomous commanding of the sensors into low-cadence data taking would also occur. Conversely, if the sensors supply less data volume to the SSMM (due to a period of efficient data compression or lack of triggers) then the instruments can be commanded to take more burst mode data or to take higher volume data products at higher cadences.

More recent analyses of the mission orbits by the Solar Orbiter project have led to the realisation that there will be times when the instruments can produce significantly more data than the original baseline volumes, and conversely times when data taking cannot proceed at baseline rates without overfilling the SSMM data stores and overwriting previously taken data. For this reason the ESA SOC will, as part of the planning for the mission, define "telemetry corridors" for each instrument. These will define the upper and lower bounds for the data volume acquired by a given instrument at a given point in time. This thus sets the rate, as a function of time, and over periods that may be shorter than the orbital period, at which the instruments may accumulate data. For SWA, the prior existence of the BKA means that this development is readily accommodated in the instrument planning by simply changing the control parameters of the BKA, such that they match the requirements of the telemetry corridors. Although we describe later in this section the various modes and data products produced by SWA, the descriptions are largely based on the nominal behaviour, and we note that a key outcome of the wide range of telemetry rate availability will mean that there will be variances in SWA data cadences and resolutions to accommodate this.

4.2. SWA commanding

SWA commanding is the responsibility of the UCL MSSL SWA operations team, who lead the activity in close consultation with

SWA partners and the SOC. On the basis of mission-level planning by the Solar Orbiter Science Working Team (SWT) and Science Operational Working Group (SOWG), the SWA operations teams will have a baseline operational plan, defining instrument telemetry and power constraints against which to fix the detailed commanding of the three sensors for a given period. In general, the SWA operations team will attempt to plan for the sensors to operate to return the highest volume of best quality science data consistent with those constraints, in particular within the telemetry corridors defined by the SOC. However, coordinated data taking with other instruments (e.g. joint burst mode observations, Walsh et al. 2020) will also be factored in to maximise the potential science return. Compliance of the data taking by the SWA sensors will be monitored and controlled on board by the SWA BKA (see Fig. 30). Instrument operations request (IOR) files will be constructed and submitted to the SOC for compliance testing, before upload to the spacecraft a few weeks before execution.

4.3. SWA modes of operation

4.3.1. SWA-EAS modes of operations

The SWA-EAS instrument can operate in various modes that will return different subsets of the original 3D VDF data. Together with the HK data from each SWA-EAS, there are also various engineering modes that allow instrument health monitoring and fault diagnosis to be performed on a semi-regular basis (about once per week, for a limited duration) in order to ensure that the sensor is maintained in optimum configuration. Of more relevance here are the science modes. These are:

SWA-EAS normal mode. The two SWA-EAS sensors each send their respective sampling of the 3D VDF to the SWA-DPU every second. The SWA-DPU stores these data in a 5 min rolling buffer. Every 4 s the SWA-DPU selects the measurements from each sensor, performs a partial moment calculation (over three subsets of energy range and two angular ranges for each sensor), and adds the resulting 168 parameters to the SSMM for inclusion in the telemetry stream. Optionally, in the event of low counts, the SWA-DPU will add four consecutive measurements from SWA-EAS and then perform the moment calculation. Every 100 s of the full set of $64 \times 32 \times 16$ 3D measurements from each SWA-EAS is compressed and sent to the SSMM. In addition, every 100 s (offset by 50 s from the selection above) a single energy bin slice of the full 3D measurement of each SWA-EAS sensor is compressed and sent to the SSMM for telemetry as a low-latency data product (see Sect. 4.4.5). The data array dimensions for this product are thus $2 \times 1 \times 32 \times 16$;

SWA-EAS burst mode. On command, the SWA-DPU will place the SWA-EAS sensors into burst mode. The SWA-DPU will steer the SWA measurements with reference to the magnetic field unit vector provided by the MAG instrument (Horbury et al. 2020) over the Service 20 IEL feed at 0.125 s cadence. In response, the SWA-EAS sensor whose central plane of FoV passes closest to the magnetic field direction makes measurements at only two elevations (but at full energy and azimuth). These two elevations are chosen such that one set of observations includes the direction along the B-field direction, and the other along the anti-parallel direction. Given that only two elevations are sampled in this mode, the resulting $1 \times 64 \times 32 \times 2$ array of data can be captured every 0.125 s and transmitted to the SWA-DPU for addition to the SSMM and the telemetry stream. These data products will be reassembled on the ground to provide a measurement of the 2D PAD of electrons (with some limited gyrophase information) at 0.125 s cadence.

SWA-EAS triggered mode. Autonomously, and following the receipt of a trigger flag over the Service 20 IEL feed, the SWA-DPU will freeze the rolling buffer containing 5 min of 1 s-cadence samples of the full 3D velocity distribution from each SWA-EAS sensor. The SWA-DPU will transmit the resulting 300 samples of $2 \times 64 \times 32 \times 16$ data arrays to the SSMM for inclusion in the SWA telemetry stream. It is expected that the trigger flag will be set by the RPW instrument (Maksimovic et al. 2020) in response to an autonomous evaluation of whether combined in situ data suggests the passage of an event of scientific interest (e.g. an interplanetary shock) passed the spacecraft in the previous 5 min period.

4.3.2. SWA-PAS modes of operations

SWA-PAS is capable of producing one $96 \times 11 \times 9$ array of data every second. However, due to restrictions of telemetry and power, SWA-PAS will operate in various modes and states that will return different subsets of the original 3D data. These modes are:

SWA-PAS normal mode. SWA-PAS can be operated across a range of parameter space. However, typically in normal mode SWA-PAS will operate on a repeating 300 s cyclogram. The data taking activities during each 300 s cyclogram are illustrated in Fig. 32. The upper part of the figure shows SWA-PAS operation around a 300 s cyclogram boundary, while the lower part shows the operation at each of the two intervening 100 s boundaries. In each case the boundary is marked by the thick vertical dashed lines. The scheme is based on making a measurement of the proton-alpha particle VDF over 1 s every 4 s. This is achieved by taking a series of "normal samples", represented by the grey boxes marked "S" in Fig. 32 over 48 energies, seven azimuths, and five elevations. However, these samples are replaced near the 100 and 300 s boundaries by "full 3D" samples, represented by the orange boxes marked "FS", in which the measurement is extended to all nine elevations. This is to ensure (checked by the SWA-DPU) that the sensor scans to correctly capture the ion distribution in the "normal samples". One such distribution is taken near the 100 s boundary, while two are taken a few seconds before and after each 300 s boundary. Between the latter two full sample measurements, SWA-PAS will be commanded into snapshot mode (see below). Figure 32 also shows, for completeness, the commanding requirements for the scheme (green boxes) and an illustration of the power saving profile (pink lines) in which the sensor voltages are ramped down between measurements. Under this scheme the SWA-PAS sensor makes a sampling of the full 3D VDF of the protons and alpha particles every 4 s, which are passed by the SWA-DPU to the SSMM for telemetry to the ground.

SWA-PAS snapshot mode. For a period of 7 s every 300 s, the SWA-PAS sensor is commanded by the SWA-DPU into snapshot mode. During these periods the sensor will sample over 48 energies, seven azimuths but only three elevations. However, the measurement cadence will be raised to a rate of 4 per second over the 7 s snapshot period. The timing of this mode will be such that it covers the period that the RPW sensor (Maksimovic et al. 2020) will perform its nested snapshot mode. This nested timing of the observations will enable detailed analyses of wave-particle interactions through the combination of high-cadence field and particle measurements.

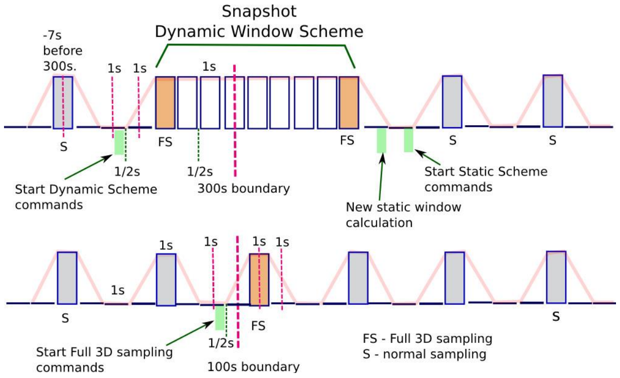

Fig. 32. Schematic illustration of the SWA-PAS cyclogram for data capture during SWA-PAS normal mode (for details see text).

SWA-PAS burst mode. On command, the SWA-DPU will place the SWA-PAS into burst mode. In this mode SWA-PAS returns samples of the full 3D proton and alpha particle distribution in a variety of formats, and for a limited period depending on telemetry constraints. Typically we expect to be able to support measurements of 48 energies, seven azimuths, and three elevations at 0.25 s cadence for periods of order 5 min. However, more complete samples over all nine elevations can also be made at 1 s cadence over similar periods, depending on scientific goals.

4.3.3. SWA-HIS modes of operations

The SWA-HIS sensor measures heavy ions arriving at the sensor aperture over a range of azimuths and elevations and distributed over a spectrum of energies, together with a measure of the TOF of these particles across a known distance in order to determine species. Each individual heavy ion entering SWA-HIS represents an event. The rates of events over different ranges (energy, elevation, TOF) are recorded, which can form the basis of a determination of a phase space distribution for individual elements. SWA-HIS will provide the full PHA, and the rates of PHA, at different cadences and resolutions since events can be selected to fill the available telemetry rates. The selection algorithm works as follows: All ion event words from an E/q scan are divided into five priority ranges, based on their average abundance in the solar wind. Priority ranges are defined by large Energy-TOF boxes, defined separately for each E/q. Events are chosen at random from each of these ranges, with more events selected from ranges containing less abundant ions. In cases where there are insufficient events present in a given range, that number is added to those to be taken from the next range (and so on). On this basis, we anticipate that in normal mode SWA-HIS will provide packets to the SSMM for telemetry which correspond to a cadence of 30 s for Helium ions and 300 s for heavier ions. Limited periods of burst mode data are possible within the telemetry constraints during which (if there are sufficient counts) Helium data will be returned every 4 s and heavier ions at 30 s. However this mode can only be run on average 1% of the time due to telemetry constraints. Conversely a low-cadence mode will also be employed at larger distances from the Sun or during periods of telemetry restriction, in which the data return would be significantly slower than in normal mode.

SWA-HIS also has a number of engineering modes which, together with the housekeeping data, allows instrument health monitoring and fault diagnosis to ensure that the sensor is maintained in optimum configuration.

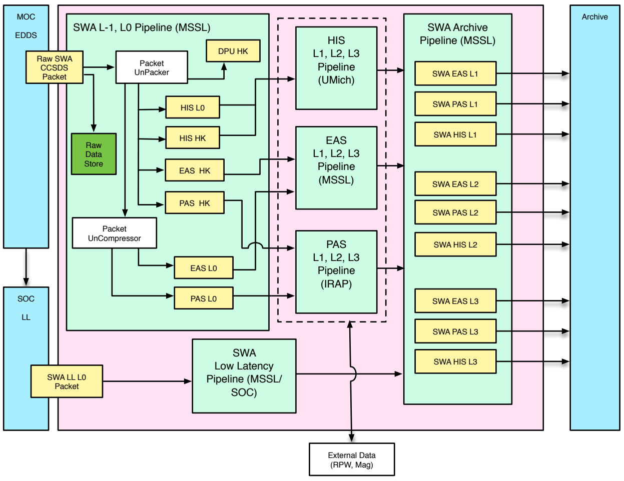

Fig. 33. Schematic of the structure of the SWA data processing pipelines from receipt of telemetry by the MOC to deposit within the ESA science archive.

4.4. SWA data products

In this subsection we detail the SWA data products that we expect to be made available through the Solar Orbiter Archive for the purposes of supporting the science goals of the mission and the community. It is expected, in accordance with the mission Science Management Plan, that best-effort science-quality level-2 data will be made generally available through the archive at 90 days after the receipt of the relevant telemetry by the ESOC. However, some SWA data products may need additional effort to produce or understanding of calibration parameters may evolve, so it is likely that further releases of the data may occur to update the available data to "current best quality" at various points after the initial release.

The SWA data processing flows from the MOC data delivery service through a series of processing pipelines and back out to the SOC archive and other archives as illustrated in Fig. 33. There are six distinct pipeline elements (shaded light green in the figure) dealing with the processing and calibration of the various data products collected by the three SWA sensors and distributed among the SWA operations teams at UCL MSSL, IRAP, and UMich. The remainder of this section provides a top level description of the data types and products produced from these pipelines.

4.4.1. SWA house-keeping and engineering data

The sensor teams for all three of the science sensors within SWA, and the SWA-DPU team have defined a set of house-keeping and engineering parameters for priority download from the spacecraft. It is anticipated that these are available immediately after the ground station pass following their acquisition. An archive of these data products will be maintained by the instrument teams and the project, although it is not foreseen that these will routinely be released to the wider community. However, these data can be made available to researchers should a specific need be identified.

4.4.2. SWA Level 0 and Level 1 data

Data collected by the SWA sensors and SWA-DPU are returned in the telemetry stream as individual CCSDS packets. Once on the ground, these packets are decommutated into relevant files containing the separate SWA-EAS, SWA-PAS, and SWA-HIS data, and uncompressed to form the L0 raw data packets. These files are still in CCSDS format and are saved and archived in binary format. The process used to create these data files is a simple C code that searches on the data type, subtype, and SID. Data which are compressed are passed through a decompressor. The decommutated and uncompressed CCSDS packet files are grouped into appropriate files for each sensor and each 24 h period. They are stored at UCL MSSL and made available to the wider SWA team.

The SWA L1 data are the uncalibrated, uncompressed L0 data converted into CDF format. The individual sensor teams are responsible for generating SWA L1 data from the L0 packets. These CDF files will have the CCSDS header data and the CCSDS science data combined. All data products are stored as CDF files according to "SOL-SGS-TN-0009 Metadata Definition for Solar Orbiter Science Data". As well as being stored in the SWA Master Repository at UCL MSSL with the L0 data, these files will be be converted to NetCDF format to fit the AMDA tool specification and stored in the Centre de Données de la Physique des Plasmas (CDPP) data archive.

SWA-EAS L1 data. UCL MSSL is responsible for generating the SWA-EAS L1 data from the L0 source. The SWA-EAS L1 data products are, for each SWA-EAS sensor, as follows:

– Normal Mode Spectra in counts, one set for each SWA-EAS sensor. The angular bin directions are in the SWA-EAS sensor reference frames;

– Burst Mode spectra in counts, from one sensor viewing the magnetic field direction. The angular bin directions are in the relevant SWA-EAS sensor frame;

– Triggered Mode Spectra in counts, one set for each SWA-EAS sensor. The angular bin directions are in the SWA-EAS sensor reference frames. These data are at the highest cadence of one 3D sweep per second for a period of 5 min;

– Partial moments calculated on board (6 sets per SWA-EAS sensor) in physical units. The frame references are the SWA-EAS sensor reference frames;

– Engineering mode and calibration data.

SWA-PAS L1 data. IRAP is responsible for generating the SWA-PAS L1 data from the L0 source. The SWA-PAS L1 data products are as follows:

– Normal Mode Spectra in counts. The angular bin directions are in the SWA-PAS frame;

– Normal Mode Snapshot Spectra. The angular bin directions are in the SWA-PAS frame;

– Burst Mode spectra. The angular bin directions are in the SWA-PAS frame;

– Onboard moments in physical units and in the SWA-PAS frame of reference;

– Engineering and Calibration data;

An example of the expected output from the SWA-PAS sensor is shown in Fig. 34. We used the SWA-PAS Proto-Flight Model (PFM) calibration data to simulate the SWA-PAS response, together with very high sampling rate solar wind data from the Faraday cup instrument BMSW onboard of RadioAstron mission (Šafránková et al. 2013) as an input for the simulation. The resulting E-T colour-coded spectrogram in the lower panel of the figure illustrates to the expected output in the SWA-PAS normal mode when the instrument takes a full 3D ion VDF every 4 s.

To illustrate the SWA-PAS ability to resolve the proton and He^{2+} peaks, we have made also a realistic simulation of SWA-PAS energy spectra for two different solar wind conditions (Fig. 35). In both cases the He^{2+} temperature is factor 4 higher than proton temperature. We can see that the He^{2+} peak should be easily resolved by the sensor.

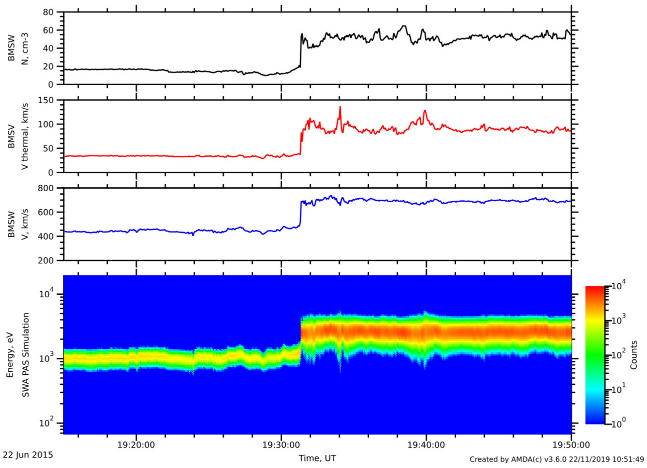

Fig. 34. Simulation of a SWA-PAS Energy-Time Spectrogram. *Three top panels*: input parameters to the simulation as measured by the BMSW instrument onboard the RadioAstron mission. *Bottom panel*: expected sensor response integrated over all angular bins. The time of the measurements is shifted relative to UT by about 1.5 h.

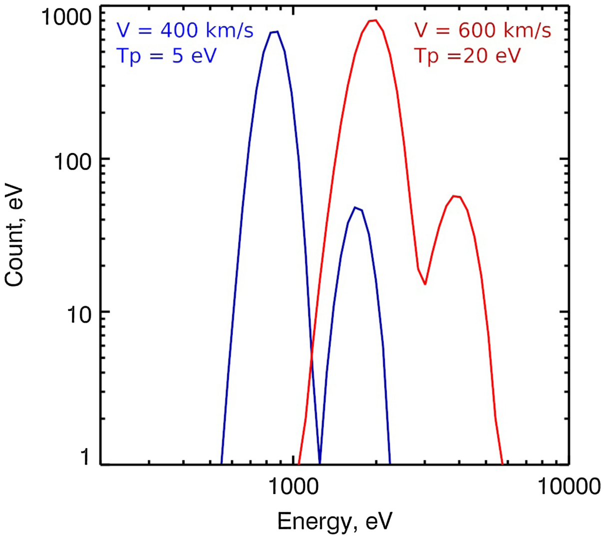

Fig. 35. Simulated examples of SWA-PAS energy spectra for different solar wind conditions. The temperature of He^{2+} peak is assumed as Tp × 4.

SWA-HIS L1 data. UMich is responsible for generating the L1 SWA-HIS data from the L0 source. These data will be in CDF format. SWA-HIS L1 data products are as follows:

– Ion Event (PHA) words: Individual ion event data, containing full information on incident angles (elevation and azimuth), E/q, TOF, and SSD energy in digital units. The number of PHA words telemetered for each energy scan is fixed (the same for each scan) but configurable, based on available resources. A sampling algorithm is employed to select the sample sent to the ground (see Sect. 4.3.3).

– Priority Rates: Counts of PHA events within a priority range, as a function of E/q and elevation. See below for description of typical use;

– Sensor Rates: Counts of unclassified ion event words on the SWA-HIS detectors (start MCP, stop MCP, SSD) as a