Table 3. Species and charge states used during calibration.

HIS elements and charge states						
	1	2	3	4	5	6
H	Yes					
He	Yes					
C	Yes	Yes	Yes	Yes		
N	Yes	Yes	Yes			
O	Yes	Yes	Yes			
F	Yes	Yes				
Ne	Yes	Yes	Yes	Yes		
S	Yes	Yes	Yes	Yes		
Ar	Yes	Yes	Yes	Yes	Yes	Yes
Fe	Yes	Yes				
Kr	Yes	Yes	Yes	Yes		

Notes. "Yes" indicates the ion species was generated and measured, blank cells indicate the ion species was not able to be generated in the lab, and greyed out cells indicate the ion does not exist in nature.

coincidence logic; energy, elevation, and azimuth acceptance and resolution; TOF and energy resolution for various species; TOF and energy range; constant fraction discrimination thresholds on detectors; operation at high fluxes; and proton and alpha rejection capability. Table 3 shows the elements and the charge states collected during ground calibration.

Figure 13 shows the dependency of collected TOF bin and collected energy bin as functions of beam energy. Data collected using the Bern facility and the SwRI accelerator agree very well to each other and can be well approximated by a forward model that takes into account losses through the foil, and the dependency of detected energy to particle mass and energy.

HIS in-flight calibration. HIS will be calibrated in-flight to check the instrument performance parameters established during ground calibrations. Internal calibration of the SWA-HIS sensor will include validating instrument efficiencies and comparing the instrument response to various ions. Start, Stop, and SSD rates will independently measure instrument efficiencies, which will be monitored as a function of time, as detector responses may decline with age. In particular, the SSD may experience an increase in dead layer or reduced signal amplification over time, and MCP performance may degrade due to large cumulative radiation events. The MCP gain will be monitored and bias voltages increased to compensate for losses in gain. The measurements will be compared with our forward model for ion identification and adjustments made as necessary. Internal calibration capabilities include time and energy test pulse generators. Further calibration includes a test of the cross suite data link as well as cross-calibrations with SWA-PAS. Such internal test and suite-level cross calibrations will be performed routinely at UMich. Gain measurements for SSD and MCPs will be made every six months and after major SEP events, potential temperature increases near perihelion, or other spacecraft occurrences that could affect efficiencies. MCP and SSD voltage adjustments will be made as appropriate.

3.2.6. Summary of SWA-HIS specifications at delivery

The SWA-HIS calibration campaign had two major goals. Firstly, to establish all needed parameters to be able to invert the measurements to be collected in flight into physical units.

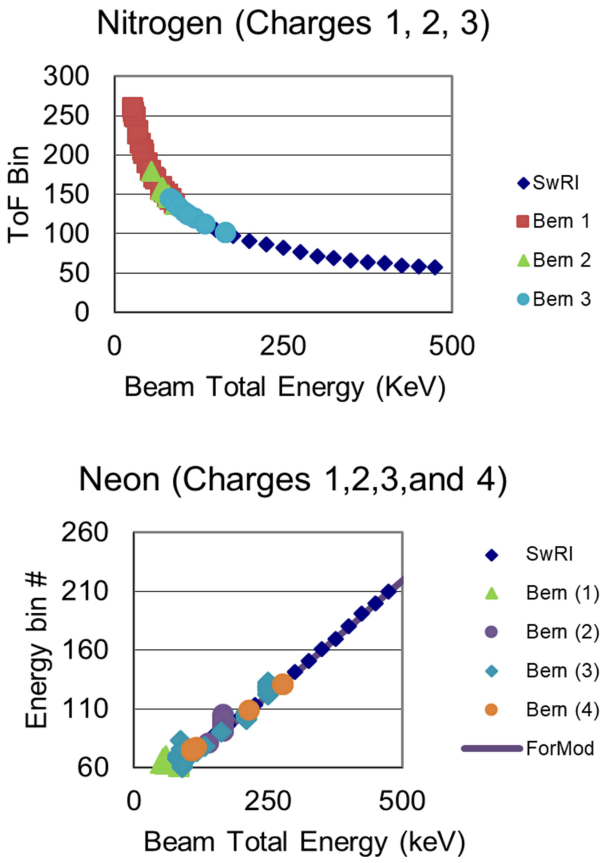

Fig. 13. Time-of-flight (*top*) and energy (*bottom*) response to beams of N and Ne as function of beam energy.

Secondly, to prove that the instrument has performance that will fulfil all Level-1 requirements.

Table 4 summarises the most important characteristics of the SWA-HIS sensor; rows colour-coded in blue show the project-defined Level-1 requirements. All properties of the SWA-HIS sensor are as expected from the design, and are compatible (and indeed most of the time better) with those needed to fulfil the Level-1 requirements. Some of the parameters reported in Table 4 are for an average configuration of the sensor: specifically, while the normal E/q resolution is ~8% and the geometric factor per pixel is $\sim 1 \times 10^{-5}$ ($\mathrm{cm^2\,sr\,eV\,eV^{-1}}$), by tuning the EA and IS voltages it is possible to increase the E/q resolution while lowering the GF, and vice versa. In-flight performance will be used to establish the most appropriate compromise between these two values.

It must be noted that for large azimuthal angles (ϕ) and for elevation angles (θ) out of the nominal orbit plane of the spacecraft ($\theta > 7°$, $\theta < -7°$), the centres of the elevation look-directions change slightly from the nominal. This property was expected for both SWA-HIS and SWA-PAS, and has been factored into the calibration files, such as to enable proper reconstruction of the VDFs.

The mass separation capabilities of SWA-HIS are illustrated in Fig. 14, which shows an energy versus TOF spectrogram obtained during laboratory testing. The track of Ar (doubly ionised in this series) is well separated from the track of Ne for all energies down to 100 keV. This energy corresponds to a charge state of 4, smaller than that expected for Ne in the solar wind. For higher charge state, the energy after PA will be larger (TOF smaller), and the mass resolution better.

Table 4. Summary of SWA-HIS performances at delivery.

	SWA-HIS L1 requirement or needed inversion parameter	SWA-HIS measured capability
Particle species	He^{2+}, C^{+4} to C^{+6}; O^{+5} to O^{+8}, Fe^{+6} to Fe^{+20}, Mg^{+6} to Mg^{+12}, Ne^{+6} to Ne^{+9}, Si^{+6} to Si^{+12}, He^{+}, C^{+}, and O^{+}	H, ^{4}He, ^{3}He, C, N, O, Mg, Ne, Si, Fe (see Table 3 for details on masses and charges measured)
Energy per charge range	0.5–60 keV/e	0.5–~78 keV/e
Energy per charge resolution ΔE/E at FWHM	Between 6 and 10%	~8–9%
Mass resolution M/ΔM at FWHM	4	>5
Mass/charge resolution (m/q)/Δ(m/q)	N/A	> 30
FoV		
In ecliptic	–33° to +66°	–33° to +66°
Out of ecliptic	±17°	±20°
Pixel resolution (Az × El)	6° × ~6° pixels	6° × ~4° pixels
G-factor/pixel (cm^2 sr eV/eV)	N/A	$\sim 1 \times 10^{-5}$
Time resolution	≤30 s	4 s (Burst mode); 30 s (Normal mode); 300 s (Normal Low Cadence mode)

Notes. Rows colour-coded in blue show the Project-defined Level-1 requirements.

One of the science goals of SWA-HIS is the characterisation of the full VDF for pickup ions. As their typical charge state is +1, they gain only a single 25 keV step from the post-acceleration, as opposed to multiples of this value for higher charge states. For energies <10 keV, they do not produce a detectable energy signal in the solid state detector and therefore must be identified by their TOF only. Figure 15 shows a histogram of counts collected for a beam energy of 0.3 keV (25.3 keV after post-acceleration) representative of the characteristics of pick-up ions in the inner heliosphere; the beam was composed of 20% each of H_2, He, N_2, Ne, and Ar, with H_2O and CO_2 as contaminants in the chamber. Even at this very low energy, TOF measurement enables clear identification of the single-charge species.

3.3. The SWA Proton-Alpha System (SWA-PAS)

3.3.1. SWA-PAS introduction

SWA-PAS is designed to measure the 3D VDFs of solar wind protons and alpha particles with high time cadence, energy and angular resolution. In the solar wind, the proton and alpha-particle thermal velocity is generally much smaller than the bulk velocity, which itself is usually strongly radially directed. Thus most of the individual ions within the distribution arrive at the spacecraft with at most a few 10's of degrees deviation from the radial direction. Thus SWA-PAS is mounted behind the spacecraft heat shield which has a dedicated cut-out in one corner. This cut-out allows a SWA-PAS FoV that covers the angular range of −22.5° to 22.5° in nine elevation (or polar) angle bins and has 11 azimuth angle bins covering a total width in this direction of 66°. In order to account for the expected aberration of the solar wind arrival direction due to the motion of the spacecraft, the FoV is offset from the solar direction by 9° such that the angular coverage is from −24° to +42°. In addition, the SWA-PAS sensor operates with a maximum resolution of 96 energy steps, from 200 eV/e to 20 keV/e.

The sensor will measure proton and alpha particle fluxes at a 4 s cadence, which will be returned to ground within the SWA telemetry stream. The SWA-DPU will also calculate moments of the associated distributions, which will be shared onboard with other instruments and returned to ground as a low-latency (<1 day following acquisition) data product for quick-look and mission planning purposes.

A complete 3D measurement of a VDF is a matrix of (96, 9, 11) elements (energy, elevation, azimuth), each element corresponds to the counts measured during a defined elementary time

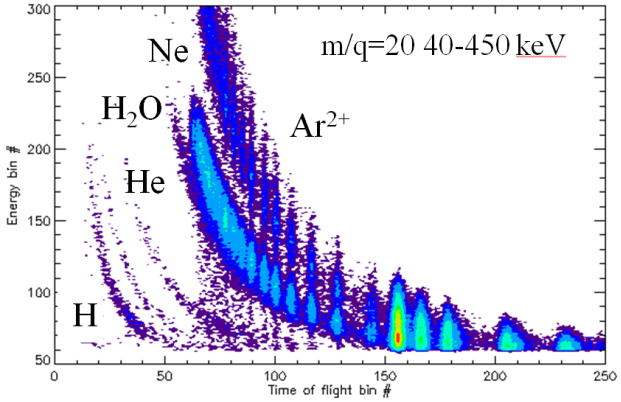

Fig. 14. E vs TOF spectrogram of m/q = 20 amu/e for various energy beams.

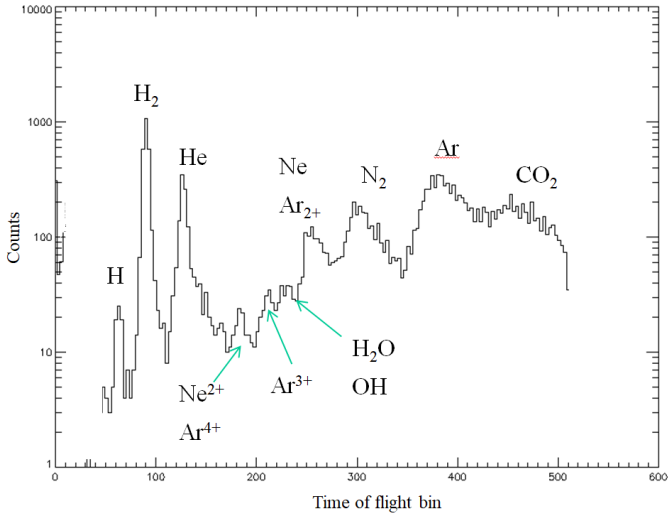

Fig. 15. Time-of-flight spectrum of a mixed beam of pickup ions of 0.3 keV energy.

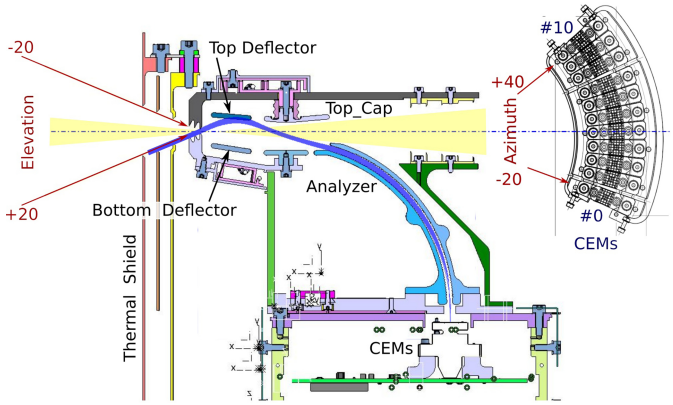

Fig. 16. *Left*: cut-away representation of the SWA-PAS analyser head. The cut-plane is that containing the Sun-spacecraft line and the direction of elevation deflection. The yellow shading represents the direct solar light and the blue curves show the extreme ion trajectories in this plane. The instrument has partial cylindrical symmetry extending from +40° to −20° out of the plane of the figure. *Right*: the insert shows the top view of the CEMs array detector arrangement and its angular range.

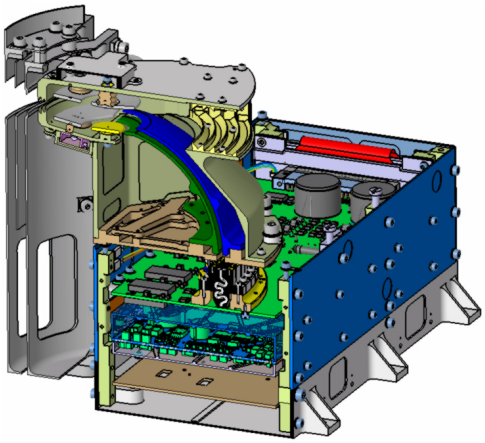

Fig. 17. Three-dimensional cut of the full SWA-PAS instrument illustrating the three main elements of SWA-PAS design: (1) the analyser (sensor head); (2) the electronics box containing the CEM detectors, HV sources, and other electronics boards; (3) the front thermal shield.

period by a specific bin. It is possible to reduce the number of energy and angular bins in one sample in order to significantly increase the sampling rate. SWA-PAS can automatically change the position of the reduced matrix to keep the maximum of the ion velocity distribution at the centre of the sampling window. Thus SWA-PAS is capable of taking a wide variety of "snapshot" and "burst" mode data products in which the time resolution of the sensor is increased. These high frequency modes consist of continuous acquisition of full or reduced VDFs, at second or sub-second cadences. The snapshot data products will be generated over periods of 8 s every 5 min in order to align with the operation of the RPW snapshot facility (Maksimovic et al. 2020; Walsh et al. 2020).

3.3.2. SWA-PAS design overview

Figure 16 shows the SWA-PAS sensor head which is comprised of the electrostatic analyser and an array of conventional channel electron multipliers (CEM) used to detect the solar wind ions. Since the instrument line-of-sight points to the Sun, the main driver of the sensor design was to provide reasonable defence against direct sunlight and heat radiation. In order to protect the sensor, it has a narrow aperture slit in front of the optical surfaces located outside of the solar light path through the sensor (yellow triangle in Fig. 16). This geometry avoids any interaction of the light beam with the internal elements of the sensor. The rear of the SWA-PAS analyser head is open to allow sunlight to pass straight through the head. In addition, the sensor has its own heat-shield system to locally protect the analyser.

The entrance deflection system, comprising two deflector plates and the "top-cap" plate, steers the incident particles towards the entrance of the electrostatic analyser. The deflection of the particles is achieved by applying specific potential on the plates. Figure 16 shows the ion trajectories calculated for the case of the maximum deflector voltages. The spherical shape electrostatic analyser filters out ions with energies outside a narrow energy band and focuses the remaining particles on to the CEM array. The azimuthal position of the focal point corresponds to the azimuthal direction of the incident ions. Thus the instant count rate of each CEM is a function of the solar wind ions flux in the corresponding bin of the SWA-PAS 3D sampling matrix. The azimuth index of this bin is the CEM number, the energy index of the bin is defined by the electrostatic analyser instant voltage and the elevation index of the bin is defined by the deflector voltage to analyser voltage ratio. The averaged geometrical factor of one bin is $5{\times}10^{-6}$ cm^2 sr eV eV^{-1}. This geometrical factor restricts the maximum count rate for an individual CEM to 10^7 s^{-1} when Solar Orbiter is at its closest approach to the Sun. Moreover, SWA-PAS still can produce a statistically valid 3D ion VDF every 4 s when the spacecraft is at 1 AU distance from the Sun. Details of the SWA-PAS measurement scheme are described in Sect. 3.3.3.

Figure 17 shows the three mechanical sections of SWA-PAS: the titanium multi-layer heat shield, the sensor head, and the electronics box. The sensor head is mounted on the top of the electronics box and is insulated from it thermally and electrically. The CEM detectors are located inside the box on the dedicated CEM board. The fast-varying HV voltages for the optical surfaces of the analyser head are generated by the HVPS. Other boards provide the HVPS controls, data acquisition, the SWA-DPU interface and LV power. We provide the details of these in Sect. 3.3.4.

Figure 18 shows the SWA-PAS assembly and its position on Solar Orbiter. The upper right insert shows the SWA-PAS installed on the spacecraft. SWA-PAS looks towards the Sun through a rectangular cut-out in a corner of the spacecraft heat-shield which has a stepped profile which acts to reduce the stray light. The charged particles pass across this cut-out, enter through the slot in the SWA-PAS heat shield, and then enter the sensor aperture. Most of the SWA-PAS external surfaces are protected by a high temperature "black Kapton" MLI. Other surfaces, shadowed against the solar heat flux, as well as against the heat radiation from other Solar Orbiter units, are coated by PCBE white paint and serve as radiators.

3.3.3. SWA-PAS measurement principle

A complete 3D measurement of an ion VDF is a matrix of (96, 9, 11) elements (energy, elevation, azimuth). Each bin of this matrix corresponds to the counts measured during a defined elementary time period of ∼1 ms by a specific channeltron, for a given energy and elevation angle bin. SWA-PAS accumulates counts in all azimuthal bins simultaneously, but to cover all

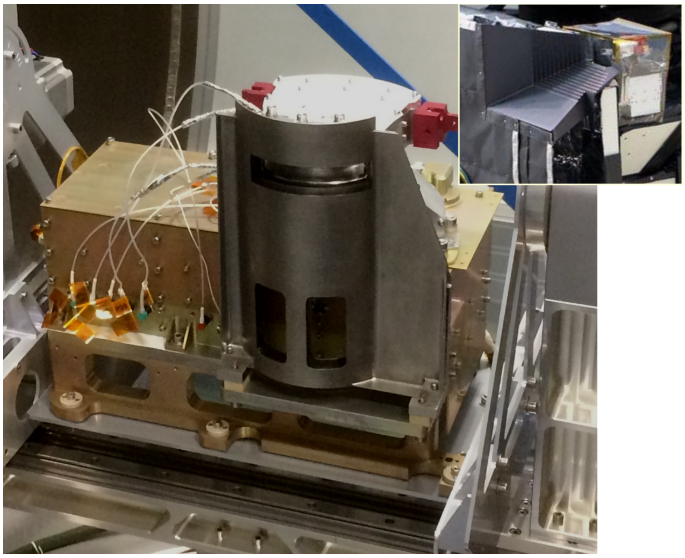

Fig. 18. Flight model of the SWA-PAS unit in the vacuum chamber during its calibration. The white analyser head is mounted on the top of the electronics box. Both are located behind the SWA-PAS-specific heat shield. Inset to the top right shows the sensor mounted behind its scalloped cut-out in the main spacecraft heatshield.

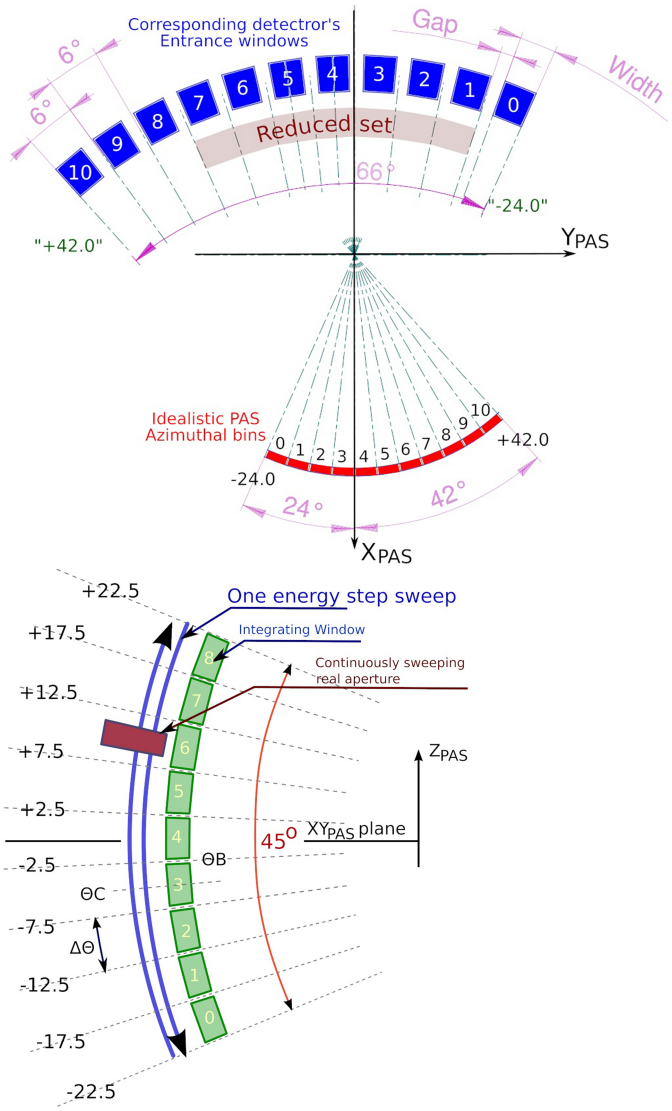

Fig. 19. SWA-PAS detector angular bins defined by the sensor geometry for the azimuth bins (*top panel*) and by deflector sweep for the elevation bins (*bottom panel*).

elevation and energy bins, SWA-PAS performs elevation sweeping and energy stepping (see Figs. 19 and 20). For a given energy, SWA-PAS makes a continuous sweep over the elevation range. The width of the instant elevation response is ~3°. To fill one bin, SWA-PAS opens the corresponding counter when the elevation scan enters the corresponding elevation bin, and stops the counter when the scan leaves the bin. As soon as one full elevation sweep is completed, SWA-PAS transfers to another energy step (Fig. 20). The transition time, ~2 ms, is that needed to stabilise the deflector HVPS. SWA-PAS transmits the data from the eleven CEMs to the SWA-DPU as soon as counts from one elevation bin are accumulated.

The duration of the full 3D sampling period consisting of (96, 9, 11) elements is exactly 1 s. The SWA-PAS design allows for the reduction of the number of energy steps and elevation angles, and also the use of seven rather than eleven channeltrons. Reduced VDFs are then obtained: for example, samples of (48, 9, 7) or (24, 3, 7), etc. Importantly, in these cases the reduced set of the energies is a subset of the energies of the full 3D sampling. The same is true for the elevation subsets. This allows placement of the 2D energy-elevation sub-window of the reduced sample in any position inside the full sampling window. If the duration of the reduced sample accumulation is less than the full sampling duration, SWA-PAS can perform several samplings per second. This fundamental advantage of the SWA-PAS operation allows the sensor to capture short-duration, high-cadence snapshots. The full set of the constants defining the SWA-PAS instant energy-elevation window is shown in Table 5.

In nominal operation, one VDF (one 1 s 3D VDF sample) will be measured each 4 s, with 3 s idle time between samples. This provides the basis to determine the density, velocity, and pressure of the solar wind at the required 4 s cadence.

3.3.4. SWA-PAS design details

From the functional point of view, SWA-PAS can be divided into five subsystems (see Fig. 21) as follows: (1) the electrostatic analyser, (2) the CEM board with the detectors, (3) the HVPS board providing HVs for the sensor head optical surfaces, (4) the FPGA board hosting the necessary digital data processing, and (5) the DC-DC converter providing the corresponding LV power. All electronics boards are located the SWA-PAS electronics box. It is not expected that the electronics box will need to be heated when the spacecraft is far from the Sun. However, for safety reasons, an operational heater is located on the side of the electronics box. The heater is powered via the same primary power line as SWA-PAS itself. The FPGA provides very simple open-loop heater control.

The SWA-PAS electrostatic analyser has no direct heritage in past space plasma instruments. Since corresponding instruments flown on HELIOS and on Bepi-Colombo are, or were, on spinning spacecraft, their plasma instrument designs did not need to include elevation steering. The Parker Solar Probe ion instrument is completely protected by the spacecraft heat shield and does not need special design to accomodate direct exposure to sunlight. However, the SWA-PAS sensor head must look directly at the Sun, which has driven the need to design the SWA-PAS analyser from the scratch. No existing coating can survive in such conditions, thus the entrance collimator and the

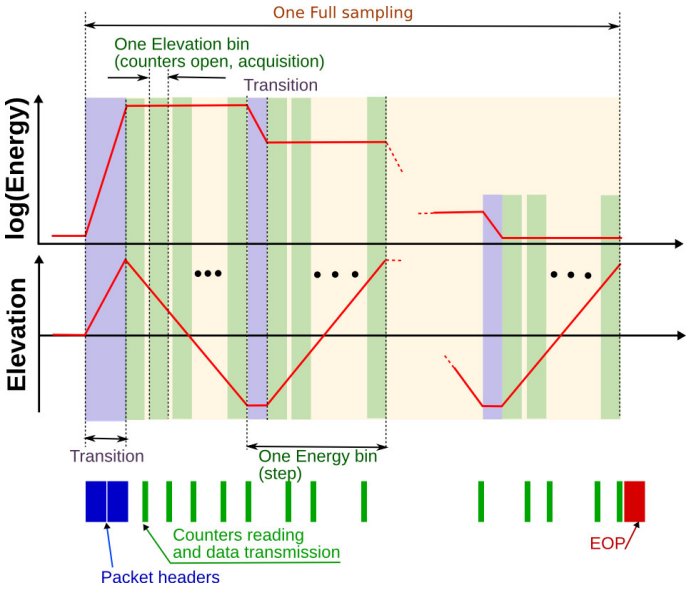

Fig. 20. SWA-PAS Energy – Elevation sampling waveform. The transition time, 2 ms, is needed to stabilise the deflector HVPS. Without this delay, the slope on the HV is perturbed at higher elevations.

Table 5. Constants controlling the SWA-PAS reduced energy-elevation window.

Parameter	Description	Possible values
Ne	Number of energy bins	2, 4, 6, ... 92
Se	Start position of the energy window	Any if the window does not exceed 0 or 95
Nel	Number of elevation bins	1, 3, ... 9
Sel	Start position of the elevation window	Any if the window does not exceed 0 or 8

deflector plates exposed to sunlight are made of polished aluminium. The electrostatic analyser spheres, less impacted by thermal flux, are scalloped and coated by Ebonol "C" UV absorbing film. The insulation of the deflector plates must be maintained up to 10 keV, since the steering of 20 keV charged particles requires at least 5.5 keV sweeping voltage. We describe the analyser properties in more detail in Sect. 3.3.5.

The CEM board is the uppermost board in the SWA-PAS electronics box. It serves as a base-plate for eleven ceramic CEMs mounted via two semi-metallised Ultem holders (Fig. 22). A corresponding set of eleven anodes lies under the CEM assembly, surrounded with a grounded guard fencing, made as a part of the outer golden metallisation layer. The anodes are internally routed and capacitively coupled to the hybrid charge-sensitive pre-amplifier-discriminators. The pre-amplifier digital output signals exit the board via logical buffers and level shifters and are passed to the FPGA board where the detected particles are counted. Each pre-amplifier-discriminator can register 6×10^6 Poisson-distributed pulses per second and their thresholds are set by fixed resistors to $\sim1.5\times10^5$ e^-. In order to save electrical power, the pre-amplifiers are switched off during the idle phases of the measurement (independently for the central and lateral sections).

The channeltron exits are kept at a potential ~-100 V, while the CEM entrance grids are biased to the CEM bias (-3900 V maximum). Two CEM variants with different resistances are used in the SWA-PAS detector to decrease the total power consumption. CEMs with nominal resistance of 80 MΩ, allowing maximum count rate 10^7 s^{-1}, are used for the azimuthal sectors 2−6. Sectors 0, 1, 7−10 utilise CEMs with nominal resistance of 200 MΩ (maximum count rate 10^6 s^{-1}) (see Fig. 19 and Table 6). Two independent HV converters (HVC) provide CEM bias for the central CEMs (sectors 1−7) and the lateral CEMs (sectors 0, 8−10). The count versus HV bias profile of both types of central CEMs are very similar, so it is unnecessary to provide separate voltages on them. Both HV converters work at 25 kHz, synchronised by a phase-locked loop (PLL) circuit to the 5 kHz clock signal provided by the FPGA board. When working with the reduced set of azimuthal sectors, the lateral pre-amplifiers and the lateral CEM HVC are always off. To minimise electromagnetic emissions and channeltron noise, the CEM board provides a single point connection between the SWA-PAS internal electrical and mechanical (chassis) grounds.

Table 6. SWA-PAS CEM properties.

CEM #	0	1	2−6	7	8−10
Resistivity, MOhm	200	200	80	200	200
Central or lateral	L	C	C	C	L

Two temperature sensors are glued near the two outermost channeltrons to monitor their operational temperatures. The CEM operational temperatures are a critical factor for detector lifetime and the SWA-DPU continuously monitors the corresponding HK values. If the temperature becomes too high, the CEMs are switched off immediately.

The SWA-PAS HVPS board provides four very fast varying HV voltages for the electrostatic analyser optical surfaces. Figure 23 shows the wave-forms of HVPS outputs. The deflector voltage linearly sweeps from -2000 V up to 5500 V over 10 ms. To provide the necessary waveform, the HVPS switches the output voltage sign and changes the internal gain of the source to transfer from the highest voltages to the lower voltage range needed for the low energy measurements.

The SWA-PAS HVPS provides a static voltage source, ±6500 V, which powers four opto-coupler-based fast HV drivers. The design includes dedicated opto-couplers, maintaining 12 keV bias, for the HVPS. The HPVS instant output is controlled by FPGA via a very fast LVDS data line.

As shown in Fig. 21, the FPGA code functionality is divided into several subsystems: (1) SpW interface; (2) SWA-DPU command decoder; (3) operational heater control; (4) HVPS control; (5) CEM HV control; (6) 11 fast counters; (7) HK acquisition unit; and (8) sequencer. The HV control system is a fast multi-channel LVDS control of the DACs integrated with the dedicated boards. The HK control is a multiplexing multi-channel analogue-to-digital converter control. Several channels needed for critical control, such as the operational heater, can be measured and transmitted to SWA-DPU at a factor of 20 times faster than other HK parameters. The counters can count up to 10^7 events per second.

The main FPGA subsystem is the sequencer, which gives significant flexibility to the system. The sequencer is an embedded programmable controller which allows simple mathematical operations, performs simple execution control loops, and "if-else" switching. The sequencer calculates HV values for the next substep using appropriate constants held in a small look-up table (LUT), and sends the commands to the HVPS board. It opens and closes the counters at the required times and forms