Mars Atmosphere and Volatile Evolution (MAVEN) Mission

Solar Wind Ion Analyzer (SWIA)

PDS Archive

Software Interface Specification

Rev 1.0
SWIA Draft
3/24/2014

Prepared by

Jasper S Halekas
(jazzman@ssl.berkeley.edu)

MAVEN
Solar Wind Ion Analyzer (SWIA)

PDS Archive
Software Interface Specification

Rev. 1.0 SWIA Draft

March 24, 2014

Custodian:

3/25/2014

Jasper S Halekas
SWIA Principal Investigator
Date

Approved:

Digitally signed by David L. Mitchell
DN: cn=David L. Mitchell, o, ou=UCB SSL, email=mitchell@ssl.berkeley.edu, c=US
Date: 2014.03.25 09:21:31 -08'00'

David L. Mitchell
MAVEN SDWG Lead
Date

3/25/2014

Alexandria DeWolfe
MAVEN SDC Lead
Date

Ray Walker

Digitally signed by Ray Walker
DN: cn=Ray Walker, o, ou, email=rwalker@igpp.ucla.edu, c=US
Date: 2014.03.25 15:59:37 -07'00'

Raymond J. Walker
PDS PPI Node Manager
Date

Thomas H. Morgan Date
PDS Project Manager

Contents

List of Figures

List of Tables

1 Introduction

This software interface specification (SIS) describes the format and content of the Solar Wind Ion Analyzer (SWIA) Planetary Data System (PDS) data archive. It includes descriptions of the data products and associated metadata, and the archive format, content, and generation pipeline.

1.1 Distribution List

Table 1: Distribution list

Name	Organization	Email
Jasper S Halekas	UCB/SSL	jazzman@ssl.berkeley.edu
Dave Mitchell	UCB/SSL	mitchell@ssl.berkeley.edu
Alexandria DeWolfe	LASP/SDC	alex.dewolfe@lasp.colorado.edu
Steve Joy	UCLA/PDS/PPI	sjoy@igpp.ucla.edu
Ray Walker	UCLA/PDS/PPI	rwalker@igpp.ucla.edu
Joe Mafi	UCLA/PDS/PPI	jmafi@igpp.ucla.edu

1.2 Document Change Log

Table 2: Document change log

Version	Change	Date	Affected portion
0.0	Initial template	2012-Aug-24	All
0.1	Updated template	2013-Feb-13	All
0.2	Final template	2013-Feb-15	All
0.3	Revised for SWIA	2013-Mar-22	All
0.4	First partially completed version	2013-Aug-09	All
0.5	Fully completed version	2013-Sep-07	All
0.6	Updated format of data files, lids to lower case	2014-Jan-31	All
0.7	Updated file formats to make use of /novary support data, waiting for requirements to be agreed upon before generating new sample files	2014-Mar-06	All
0.8	Format edits by Joe Mafi, removed some context and XML schema from SWIA bundle	2014-Mar-14	All
1.0	Updated sections from original template to better reflect current archive plans	2014-Mar-24	All

1.3 TBD Items

Table 3 lists items that are not yet finalized.

Table 3: List of TBD items

Item	Section(s)	Page(s)
Full references for PDS4 Standards Reference, and Data Provider's Handbook documents (to be provided by PDS/PPI)	1.9	8
Sample labels (to be provided by PDS/PPI)	Appendices C, D, and E	51-53

1.4 Abbreviations

Table 4: Abbreviations and their meaning

Abbreviation	Meaning
ASCII	American Standard Code for Information Interchange
Atmos	PDS Atmospheres Node (NMSU, Las Cruces, NM)
CCSDS	Consultative Committee for Space Data Systems
CDR	Calibrated Data Record
CFDP	CCSDS File Delivery Protocol
CK	C-matrix Kernel (NAIF orientation data)
CODMAC	Committee on Data Management, Archiving, and Computing
CRC	Cyclic Redundancy Check
CU	University of Colorado (Boulder, CO)
DAP	Data Analysis Product
DDR	Derived Data Record
DMAS	Data Management and Storage
DPF	Data Processing Facility
E&PO	Education and Public Outreach
EDR	Experiment Data Record
EUV	Extreme Ultraviolet; also used for the EUV Monitor, part of LPW (SSL)
FEI	File Exchange Interface
FOV	Field of View
FTP	File Transfer Protocol
GB	Gigabyte(s)
GSFC	Goddard Space Flight Center (Greenbelt, MD)
HK	Housekeeping

Abbreviation	Meaning
HTML	Hypertext Markup Language
ICD	Interface Control Document
IM	Information Model
ISO	International Standards Organization
ITF	Instrument Team Facility
IUVS	Imaging Ultraviolet Spectrograph (LASP)
JPL	Jet Propulsion Laboratory (Pasadena, CA)
LASP	Laboratory for Atmosphere and Space Physics (CU)
LID	Logical Identifier
LIDVID	Versioned Logical Identifer
LPW	Langmuir Probe and Waves instrument (SSL)
MAG	Magnetometer instrument (GSFC)
MAVEN	Mars Atmosphere and Volatile EvolutioN
MB	Megabyte(s)
MD5	Message-Digest Algorithm 5
MOI	Mars Orbit Insertion
MOS	Mission Operations System
MSA	Mission Support Area
MSO	Mars Solar Orbital Coordinate System
NAIF	Navigation and Ancillary Information Facility (JPL)
NASA	National Aeronautics and Space Administration
NGIMS	Neutral Gas and Ion Mass Spectrometer (GSFC)
NMSU	New Mexico State University (Las Cruces, NM)
NSSDC	National Space Science Data Center (GSFC)
PCK	Planetary Constants Kernel (NAIF)
PDS	Planetary Data System
PDS4	Planetary Data System Version 4
PF	Particles and Fields (instruments)
PPI	PDS Planetary Plasma Interactions Node (UCLA)
RS	Remote Sensing (instruments)

Abbreviation	Meaning
SCET	Spacecraft Event Time
SDC	Science Data Center (LASP)
SCLK	Spacecraft Clock
SEP	Solar Energetic Particle instrument (SSL)
SIS	Software Interface Specification
SOC	Science Operations Center (LASP)
SPE	Solar Particle Event
SPICE	Spacecraft, Planet, Instrument, C-matrix, and Events (NAIF data format)
SPK	Spacecraft and Planetary ephemeris Kernel (NAIF)
SSL	Space Sciences Laboratory (UCB)
STATIC	Supra-Thermal And Thermal Ion Composition instrument (SSL)
SWEA	Solar Wind Electron Analyzer (SSL)
SWIA	Solar Wind Ion Analyzer (SSL)
TBC	To Be Confirmed
TBD	To Be Determined
UCB	University of California, Berkeley
UCLA	University of California, Los Angeles
URN	Uniform Resource Name
UV	Ultraviolet
XML	eXtensible Markup Language

1.5 Glossary

Archive – A place in which public records or historical documents are preserved; also the material preserved – often used in plural. The term may be capitalized when referring to all of PDS holdings – the PDS Archive.

Basic Product – The simplest product in PDS4; one or more data objects (and their description objects), which constitute (typically) a single observation, document, etc. The only PDS4 products that are *not* basic products are collection and bundle products.

Bundle Product – A list of related collections. For example, a bundle could list a collection of raw data obtained by an instrument during its mission lifetime, a collection of the calibration products associated with the instrument, and a collection of all documentation relevant to the first two collections.

Class – The set of attributes (including a name and identifier) which describes an item defined in the PDS Information Model. A class is generic – a template from which individual items may be constructed.

Collection Product – A list of closely related basic products of a single type (e.g. observational data, browse, documents, etc.). A collection is itself a product (because it is simply a list, with its label), but it is not a *basic* product.

Data Object – A generic term for an object that is described by a description object. Data objects include both digital and non-digital objects.

Description Object – An object that describes another object. As appropriate, it will have structural and descriptive components. In PDS4 a 'description object' is a digital object – a string of bits with a predefined structure.

Digital Object – An object which consists of real electronically stored (digital) data.

Identifier – A unique character string by which a product, object, or other entity may be identified and located. Identifiers can be global, in which case they are unique across all of PDS (and its federation partners). A local identifier must be unique within a label.

Label – The aggregation of one or more description objects such that the aggregation describes a single PDS product. In the PDS4 implementation, labels are constructed using XML.

Logical Identifier (**LID**) – An identifier which identifies the set of all versions of a product.

Versioned Logical Identifier (LIDVID) – The concatenation of a logical identifier with a version identifier, providing a unique identifier for each version of product.

Manifest - A list of contents.

Metadata – Data about data – for example, a 'description object' contains information (metadata) about an 'object.'

Non-Digital Object – An object which does not consist of digital data. Non-digital objects include both physical objects like instruments, spacecraft, and planets, and non-physical objects like missions, and institutions. Non-digital objects are labeled in PDS in order to define a unique identifier (LID) by which they may be referenced across the system.

Object – A single instance of a class defined in the PDS Information Model.

PDS Information Model – The set of rules governing the structure and content of PDS metadata. While the Information Model (IM) has been implemented in XML for PDS4, the model itself is implementation independent.

Product – One or more tagged objects (digital, non-digital, or both) grouped together and having a single PDS-unique identifier. In the PDS4 implementation, the descriptions are combined into a single XML label. Although it may be possible to locate individual objects within PDS (and to find specific bit strings within digital objects), PDS4 defines ‘products’ to be the smallest granular unit of addressable data within its complete holdings.

Tagged Object – An entity categorized by the PDS Information Model, and described by a PDS label.

Registry – A data base that provides services for sharing content and metadata.

Repository – A place, room, or container where something is deposited or stored (often for safety).

XML – eXtensible Markup Language.

XML schema – The definition of an XML document, specifying required and optional XML elements, their order, and parent-child relationships.

1.6 MAVEN Mission Overview

The MAVEN mission is scheduled to launch on an Atlas V between November 18 and December 7, 2013. After a ten-month ballistic cruise phase, Mars orbit insertion will occur on or after September 22, 2014. Following a 5-week transition phase, the spacecraft will orbit Mars at a 75° inclination, with a 4.5 hour period and periapsis altitude of 140-170 km (density corridor of 0.05-0.15 kg/km^3). Over a one-Earth-year period, periapsis will precess over a wide range of latitude and local time, while MAVEN obtains detailed measurements of the upper atmosphere, ionosphere, planetary corona, solar wind, interplanetary/Mars magnetic fields, solar EUV and solar energetic particles, thus defining the interactions between the Sun and Mars. MAVEN will explore down to the homopause during a series of five 5-day “deep dip” campaigns for which periapsis will be lowered to an atmospheric density of 2 kg/km^3 (~125 km altitude) in order to sample the transition from the collisional lower atmosphere to the collisionless upper atmosphere. These five campaigns will be interspersed though the mission to sample the subsolar region, the dawn and dusk terminators, the anti-solar region, and the north pole.

1.6.1 Mission Objectives

The primary science objectives of the MAVEN project will be to provide a comprehensive picture of the present state of the upper atmosphere and ionosphere of Mars and the processes controlling them and to determine how loss of volatiles to outer space in the present epoch varies with changing solar conditions. Knowing how these processes respond to the Sun’s energy inputs will enable scientists, for the first time, to reliably project processes backward in time to study atmosphere and volatile evolution. MAVEN will deliver definitive answers to high-priority science questions about atmospheric loss (including water) to space that will greatly enhance our understanding of the climate history of Mars. Measurements made by MAVEN will allow us to determine the role that escape to space has played in the evolution of the Mars atmosphere, an essential component of the quest to “follow the water” on Mars. MAVEN will accomplish this by achieving science objectives that answer three key science questions:

- What is the current state of the upper atmosphere and what processes control it?
- What is the escape rate at the present epoch and how does it relate to the controlling processes?
- What has the total loss to space been through time?

MAVEN will achieve these objectives by measuring the structure, composition, and variability of the Martian upper atmosphere, and it will separate the roles of different loss mechanisms for both neutrals and ions. MAVEN will sample all relevant regions of the Martian atmosphere/ionosphere system—from the termination of the well-mixed portion of the atmosphere (the "homopause"), through the diffusive region and main ionosphere layer, up into the collisionless exosphere, and through the magnetosphere and into the solar wind and downstream tail of the planet where loss of neutrals and ionization occurs to space—at all relevant latitudes and local solar times. To allow a meaningful projection of escape back in time, measurements of escaping species will be made simultaneously with measurements of the energy drivers and the controlling magnetic field over a range of solar conditions. Together with measurements of the isotope ratios of major species, which constrain the net loss to space over time, this approach will allow thorough identification of the role that atmospheric escape plays today and to extrapolate to earlier epochs.

1.6.2 Payload

MAVEN will use the following science instruments to measure the Martian upper atmospheric and ionospheric properties, the magnetic field environment, the solar wind, and solar radiation and particle inputs:

- NGIMS Package:
    - Neutral Gas and Ion Mass Spectrometer (NGIMS) measures the composition, isotope ratios, and scale heights of thermal ions and neutrals.
- RS Package:
    - Imaging Ultraviolet Spectrograph (IUVS) remotely measures UV spectra in four modes: limb scans, planetary mapping, coronal mapping and stellar occultations. These measurements provide the global composition, isotope ratios, and structure of the upper atmosphere, ionosphere, and corona.
- PF Package:
    - Supra-Thermal and Thermal Ion Composition (STATIC) instrument measures the velocity distributions and mass composition of thermal and suprathermal ions from below escape energy to pickup ion energies.
    - Solar Energetic Particle (SEP) instrument measures the energy spectrum and angular distribution of solar energetic electrons (30 keV – 1 MeV) and ions (30 keV – 12 MeV).
    - Solar Wind Ion Analyzer (SWIA) measures solar wind and magnetosheath ion density, temperature, and bulk flow velocity. These measurements are used to determine the charge exchange rate and the solar wind dynamic pressure.
    - Solar Wind Electron Analyzer (SWEA) measures energy and angular distributions of 5 eV to 5 keV solar wind, magnetosheath, and auroral electrons, as well as ionospheric photoelectrons. These measurements are used to constrain

the plasma environment, magnetic field topology and electron impact ionization rate.

- Langmuir Probe and Waves (LPW) instrument measures the electron density and temperature and electric field in the Mars environment. The instrument includes an EUV Monitor that measures the EUV input into Mars atmosphere in three broadband energy channels.
- Magnetometer (MAG) measures the vector magnetic field in all regions traversed by MAVEN in its orbit.

1.7 SIS Content Overview

Section 2 describes the Solar Wind Ion Analyzer (SWIA) sensor. Section 3 gives an overview of data organization and data flow. Section 4 describes data archive generation, delivery, and validation. Section 5 describes the archive structure and archive production responsibilities. Section 6 describes the file formats used in the archive, including the data product record structures. Individuals involved with generating the archive volumes are listed in Appendix A. Appendix B contains a description of the MAVEN science data file naming conventions. Appendix C, Appendix D, and Appendix E contain sample PDS product labels. Appendix F describes SWIA archive product PDS deliveries formats and conventions.

1.8 Scope of this document

The specifications in this SIS apply to all SWIA products submitted for archive to the Planetary Data System (PDS), for all phases of the MAVEN mission. This document includes descriptions of archive products that are produced by both the SWIA team and by PDS.

1.9 Applicable Documents

[1] Planetary Data System Data Provider's Handbook, TBD.

[2] Planetary Data System Standards Reference, TBD.

[3] PDS4 Data Dictionary – Abridged, Version 1.1.0.1, 21 November 2013.

[4] Planetary Data System (PDS) PDS4 Information Model Specification, Version 1.1.0.1.

[5] Mars Atmosphere and Volatile Evolution (MAVEN) Science Data Management Plan, Rev. C, doc. no.MAVEN-SOPS-PLAN-0068

[6] Archive of MAVEN CDF in PDS4, T. King and J. Mafi, 16 July 2013.

1.10 Audience

This document describes the interactions between the MAVEN Project, SWIA instrument team, and PDS, defining the roles and responsibilities of each in producing SWIA PDS archive products. It is also useful to those wishing to understand the format and content of the SWIA PDS data product archive collection. Typically, these individuals would include scientists, data analysts, and software engineers.

2 SWIA Instrument Description

The Solar Wind Ion Analyzer (SWIA) [See Fig. 1] is an electrostatic analyzer designed to measure solar wind and magnetospheric ions in the Martian system over an energy/charge range of ~5-25000 eV/q, and an angular range of 360x90 degrees [minus spacecraft obstructions]. The SWIA sensor is based on heritage from the Mars Global Surveyor Electron Reflectometer, Lunar Prospector Electron Reflectometer, Wind 3dp, FAST ESA, and THEMIS ESA instruments. The SWIA electronics are most directly based on those of THEMIS ESA, and the analyzer includes new deflection optics in order to provide a large field of view on the 3-axis stabilized MAVEN spacecraft. The SWIA sensor is mounted on the corner of the top deck of the spacecraft as shown in Fig. 1, positioned to ensure a clear field of view over both sides of the solar panel. For the nominal sun-pointed spacecraft orientation, the sensor is aligned such that the sun is centered in the sensor field of view, with an unobstructed field of view around the nominal solar wind flow direction.

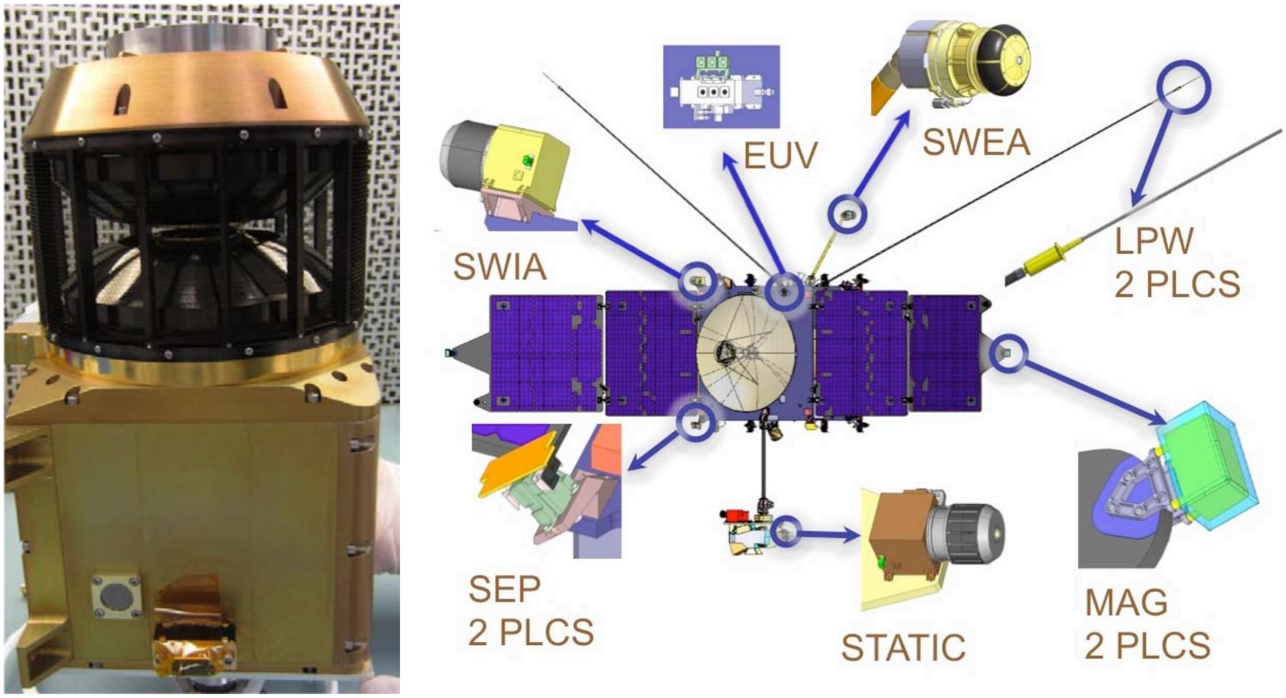

Figure 1: The SWIA instrument and its location on the spacecraft.

2.1 Science Objectives

SWIA provides measurements that satisfy the MAVEN level 1 requirement to determine density and velocity distributions of solar wind and magnetosheath protons from 1000 km/s to 50 km/s, with better than 15% energy resolution and better than 30 degrees angular resolution. MAVEN carries a suite of instruments that measure the significant energy inputs into the Martian system and the neutral and charged populations of escaping atmospheric gases, in order to determine how the former drives the latter, with the goal of characterizing the state of the upper atmosphere and its evolution over Mars' history. Within this framework, the main science objective for the SWIA sensor is to measure the properties of the energy input to the Martian system from the solar wind. Solar wind ion properties determine the solar wind and magnetosheath properties near Mars and constrain the nature of the solar wind interactions with the upper atmosphere, determine the ionization rates of neutrals from charge exchange, and determine the pickup acceleration of newly formed ions by the **v X B** electric field. SWIA measurements also contribute to the goal of determining the current state of the upper atmosphere and characterizing the non-thermal ion loss processes that occur, by providing high-cadence measurements of ion 3-d velocity distributions throughout the Martian magnetosphere.

In order to achieve these science goals, SWIA satisfies and in most cases significantly exceeds the following MAVEN Level 3 measurement requirements:

- SWIA shall measure energy fluxes from 10^7 to 10^{10} eV/[cm^2 s sr eV] w/ no worse than 25% precision
- SWIA shall measure ion flow velocities from 50-1000 km/s
- SWIA shall have energy resolution $\Delta E/E$ at least 15%
- SWIA shall have angular resolution of at least 30 degrees (10 degrees in Sun direction)
- SWIA shall have time resolution of at least 1 minute or better
- SWIA shall have a FOV of 180 x 40 degrees or better

2.2 Electrostatic Optics and Detectors

SWIA measures ions of a given energy by sweeping the negative voltage on the inner of two concentric toroidal hemispheres, ions of a given sensor phi angle (0-360 degrees) with a segmented charge collecting anode (24 anodes total, 10 with 4.5 degree resolution in the sun direction and 14 more with 22.5 degree resolution elsewhere) below a chevron pair of micro-channel plates, and ions of a given sensor theta angle (±45 degrees) by sweeping the positive voltage on the upper or lower deflectors [See Fig. 2]. A mechanical attenuator consisting of a "visor" with a slit centered in the field of view reduces the sensitivity in the sun direction when closed, in order to prevent saturation during periods of intense solar wind fluxes.

The SWIA electrostatic optics were simulated in detail, utilizing a Laplacian solver to derive the electrostatic potential produced by each charged surface in the analyzer, and a Runge-Kutta algorithm to trace charged particles through the resulting electrostatic fields. Our simulations indicate that (given an high voltage power supply that can produce an inner hemisphere voltage which can reach -4 kV) SWIA can cover ion energy per charges of up to ~31 keV/q. In normal operation, we sweep the inner hemisphere voltage in order to cover the range from 5 eV/q to 25 keV/q, utilizing a logarithmic energy sweep. SWIA can measure particles with sensor theta angles of up to ±45 degrees by placing a voltage of up to ~6.4 times the inner hemisphere voltage on one of the two toroidal deflector surfaces. In normal operation, each deflector voltage is scanned over its full range at each energy step, in order to cover as much of this angular range as possible. Given the maximum deflector high voltage supply output of +4 kV, full range deflection is only possible for ions with energy per charge of up to ~5 keV/q; above this energy, we sweep the deflectors so as to evenly cover the angular range accessible to the instrument.

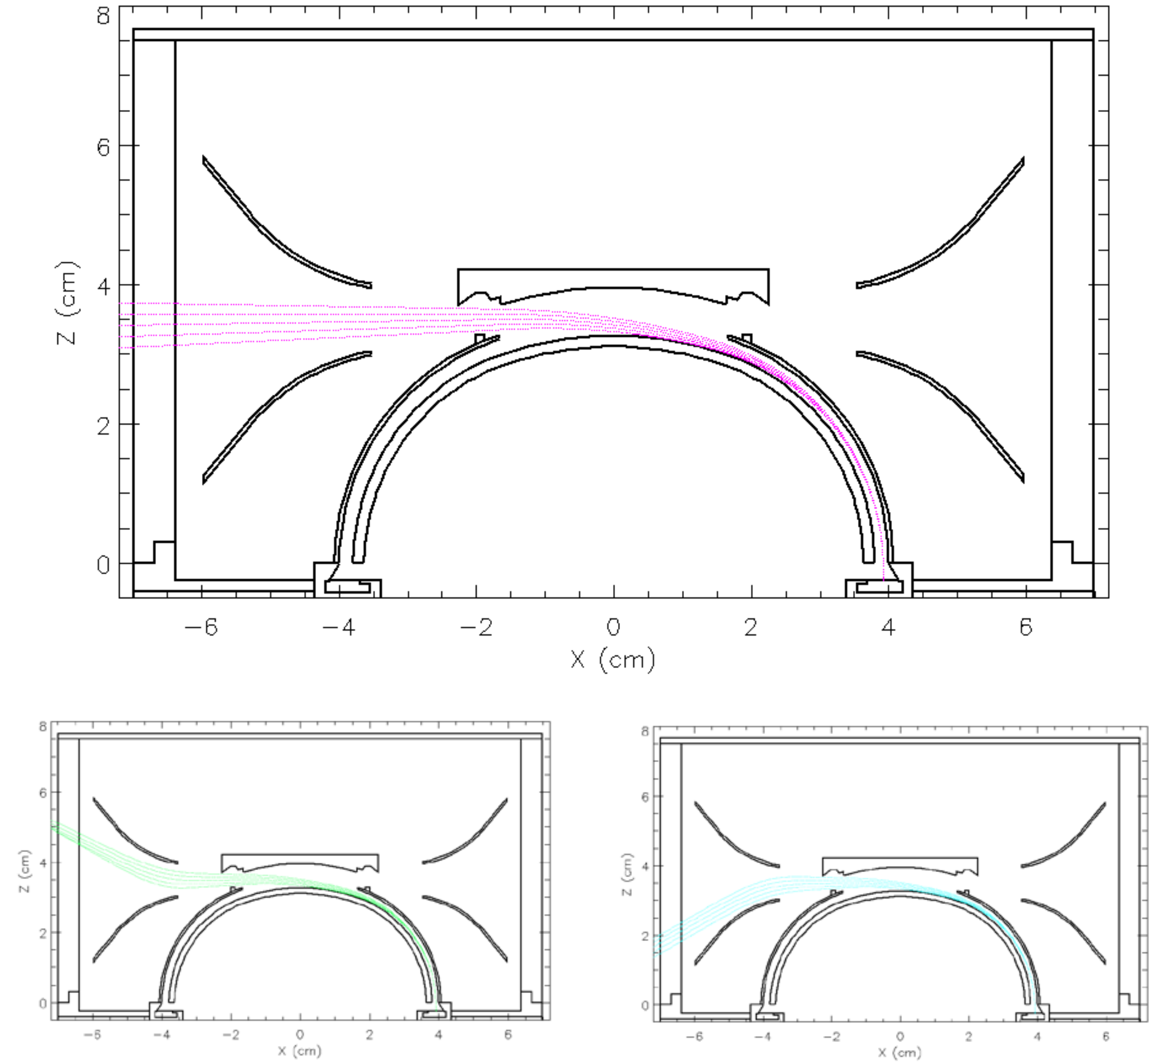

Figure 2: SWIA analyzer optics for no deflection (top) and selected up and down deflections, achieved by placing different combinations of voltages on the inner hemisphere and the two toroidal deflector surfaces.

When closed, the mechanical attenuator slightly improves the energy and theta angle resolution of the sensor, since it has the effect of collimating the response in the theta angle. As shown in Fig. 3, the simulated energy resolution of the sensor is ~14.5% with the attenuator open, and ~10% with it closed. The theta resolution is ~7 degrees with the attenuator open, and ~3 degrees with it closed, at zero deflection. The intrinsic phi resolution of the instrument is very narrow, on the order of a degree, thanks to the natural focusing properties of the top-hat electrostatic analyzer. However, we note that the actual phi resolution of the sensor is ultimately set by the convolution of this response with the anode width (4.5 degrees wide near sun, 22.5 degrees wide elsewhere) rather than solely by the intrinsically narrow analyzer phi resolution. The SWIA sensor angular response varies somewhat as a function of deflection angle, due to the focusing/de-focusing properties of the electrostatic optics [See calibration results in Section 2.7 for more details].

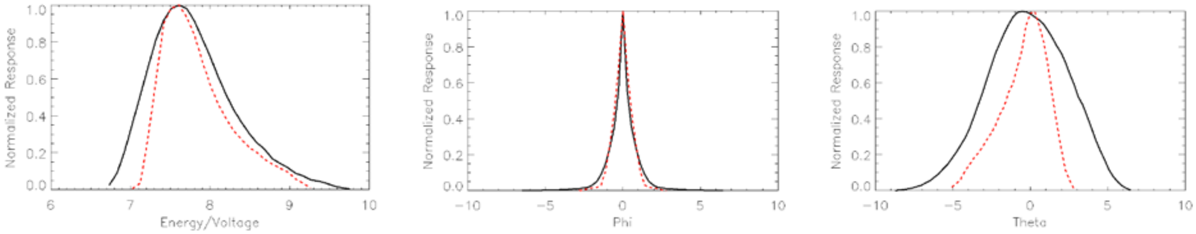

Figure 3: SWIA intrinsic energy, phi, and theta resolution at zero deflection, with attenuator open (black) and closed (red).

The deflection angle as a function of deflector voltage (normalized by hemisphere voltage), and the geometric factor as a function of deflection angle are shown in Fig. 4. The predicted deflection angle as a function of deflector voltage is very linear over the entire angular range. The predicted sensitivity is uniform over +/-22.5 sensor theta angle (sufficient to cover the normal range of solar wind velocities, for nominal sun-pointed orientation of the spacecraft), with a graceful roll-off to ~50% at the edges of the range resulting from collimation by the deflectors. Note that the deflectors are serrated in order to eliminate any possible issues with scattered ions from the deflectors, and the internal surfaces are scalloped and blackened with Ebanol-C in order to eliminate both scattered charged particles and photons. The predicted sensor 360° geometric factor at zero deflection angle is 0.0236 cm^2 sr. With all grid transmissions and nominal MCP detector efficiencies folded in, the predicted sensor geometric factor is 0.0056 cm^2 sr, distributed proportionally over the 24 anodes (see below). Inflight calibrations will be utilized to determine the actual MCP efficiency, as described in section 2.8.

- *Simulated Instrument Geometric Factor, w/ grid transmissions and approximate MCP efficiency included (MCP efficiency to be confirmed on-orbit):*
  - Analyzer 360° geometric factor: 0.0236 cm^2 sr
  - Sensor 360° geometric factor with predicted efficiencies included: 0.0056 cm^2 sr
  - Large anode geometric factor: 0.00035 cm^2 sr
  - Small anode geometric factor: 0.000070 cm^2
  - Small anode geometric factor with attenuator in: 0.0000047 cm^2 sr
- *Measurable Flux Range:*
  - Measurable count rates per anode: few Hz to 2 MHz
  - Measurable diff. energy fluxes in small anodes: 5×10^4 to 7×10^{11} eV/(cm^2 s sr eV)
  - Measurable diff. energy fluxes in large anodes: 1×10^4 to 5×10^9 eV/(cm^2 s sr eV)
- *Required/Desirable Measurable Flux Range:*
  - Level 3 requirements: 1×10^7 to 1×10^{10} eV/(cm^2 s sr eV)
  - Lowest expected fluxes in magneto-sheath: 1×10^5 eV/(cm^2 s sr eV)
  - Highest expected fluxes for cold dense solar wind: 5×10^{11} eV/(cm^2 s sr eV)

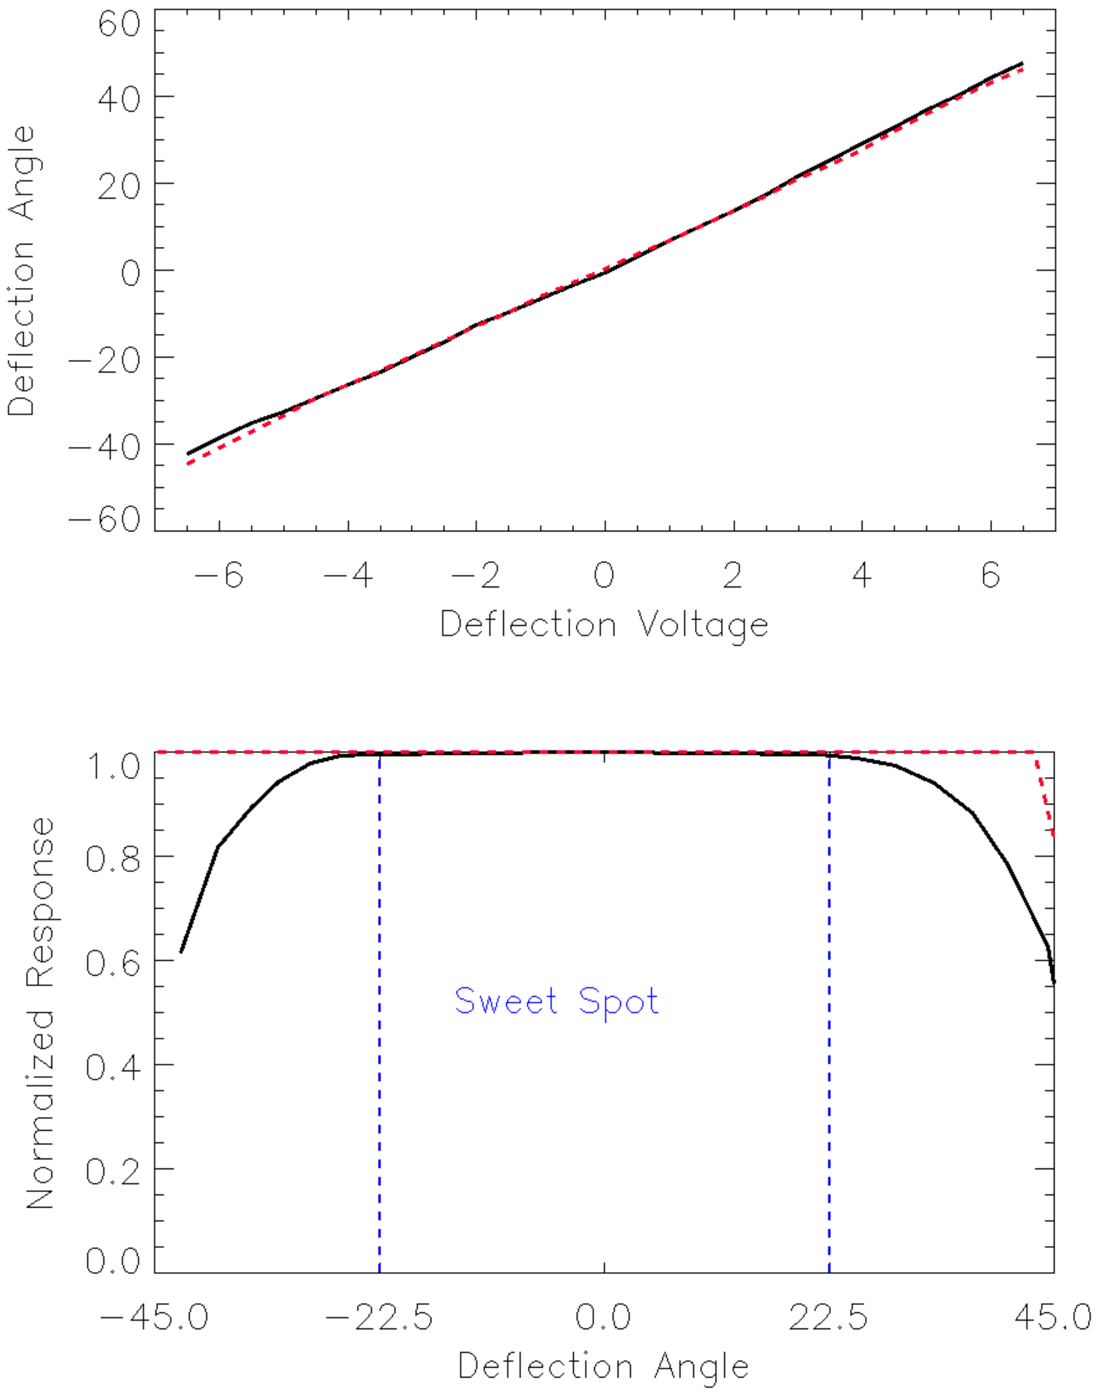

Figure 4: SWIA deflection angle vs. voltage (top), and geometric factor vs. deflection angle (bottom), for attenuator open (black) and closed (red). The +/– deflection voltages indicate the operation of the two deflectors; both deflector voltages are in fact positive.

The SWIA sensor utilizes a chevron pair of annular microchannel plate detectors to produce a secondary electron cascade (triggered by each incoming ion) to produce a pulse of ~10^6 electrons that can be registered by the electronics. The SWIA flight plates have a total stack resistance of 17.3 Mohms (at room temperature). This low resistance corresponds to a high strip recharge current that allows the sensor to count at a high rate on the order of 500 kHz, or higher for brief periods, before significant MCP droop occurs. Since the solar wind is only observed for a brief period of the energy/deflector sweep pattern, and the mechanical attenuator limits instrument sensitivity in the sun direction, MCP droop should not affect the sensor in normal operation.

2.3 Electronics

A block diagram of the SWIA sensor that details the key features of the electronics design is shown below. As described in section 2.2 above, each incoming ion triggers a secondary electron cascade in the microchannel plates, producing a charge pulse that is collected by one of the 14 22.5-degree anodes or 10 4.5-degree anodes that cover the full range of sensor phi angles. The microchannel plates are mechanically clamped (and thermally coupled) to the ANODE board, with a metal spacer seated on a metallized ring setting the gap between the output face of the channel plates and the anodes, and a 330 Kohm resistor from this ring to ground providing a small pre-acceleration voltage to focus the charge pulses from the channel plates to the anodes. High voltages for the channel plates and the inner hemisphere are carried on coaxial cables to the ANODE board, with the coax sheathes soldered to pads connected to the ground plane, and pass through the board on custom connections.

Charge pulses collected on the metallized anodes are carried on Hypertronics KA-17 connectors to the PREAMP/MCP board, which contains 24 Amptek A121 charge-sensitive preamplifiers. A 1 Mohm resistor from each anode to ground dissipates any DC charge buildup. The signals are capacitively coupled to the inputs of the preamplifiers, which have a digitally programmable threshold controlled by the FPGA on the DIGITAL board. The A121's are tuned with external resistors to produce a digital output pulse with a width of 50 ns, and to have a well-characterized fixed dead time of 100 ns (allowing count rates of up to 10 MHz periodic for each signal chain). The output signals are carried on MDM connectors to pigtails on the DIGITAL board, and ultimately to ripple counters in the FPGA. The preamplifiers can also be stimulated by a capacitively coupled test pulse signal generated by the FPGA on the DIGITAL board, allowing testing without an analyzer and/or without high voltage enabled. The digitally produced test pulse signal is divided down into four different frequencies, such that adjacent anodes do not share the same frequency (enabling testing for crosstalk). Each preamplifier is protected from high voltage discharges by clamp diodes.

The PREAMP/MCP board is shared with the high voltage supply for the microchannel plates, which produces a negative voltage of up to -2.5 kV that is applied to the input face of the channel plates. This voltage is controlled by a 0-4 V DAC output from the DIGITAL board. The MCP high voltage supply drives a significant resistive load (~2 kV across ~17.3 Mohms at room temperature), so the transistors are heat-sunk to the chassis to dissipate any heat buildup on the board. High voltage control lines and read-backs share the same MDM connectors with the preamplifier output signals. The high voltage return is connected to the board so that currents close appropriately, and connected to the board analog ground (which is connected to chassis).

The SWEEP/HVPS board has a single bi-polar raw supply that operates at a nominal voltage of 4.3 kV, with optically coupled high voltage outputs for the inner hemisphere (negative), and the two deflectors (positive), each with a range of 0-4 kV. These high voltages are routed from the SWEEP board to their respective destinations through custom connectors and coaxial cables, and the high voltage grounds are returned to the board so any currents can close, and connected to analog ground at the board (which is connected to chassis). The raw supply and each of the outputs are controlled by 0-4V DAC outputs from the digital board. The SWEEP board is connected to the DIGITAL board via stacking WTAX connectors.

The DIGITAL board contains an FPGA that controls the high voltage sweep, accumulates data, produces data products, and interfaces with the PFDPU. Digital counts from the preamplifiers for each signal chain are converted to a 3.3 volt level for the FPGA and then accumulated in ripple counters over 1.7 ms accumulation intervals, for each of the 96Ex24D steps in the sweep table (4-second full product cycle), and processed to produce the three data products described below in section 2.4. Short ~1 ms dead times at the end of each energy step, and a longer ~21.5 ms dead time at the end of the cycle, ensure that counts are not accumulated during high voltage retraces. The FPGA also produces the test pulse signal that goes to the PREAMP board, allowing testing of the electronics without the analyzer and/or without high voltage operating. The FPGA controls the DACs that set high voltage levels and preamplifier thresholds, and multiplexes and converts housekeeping values for all of the secondary and high voltages, as well as two thermistors. High voltage DACs are controlled utilizing a sweep table stored in SRAM, and product accumulations are stored in the SRAM (for those products that aren't transmitted to the PFDPU immediately). Digital and analog grounds are tied together on the board, and analog ground is connected to chassis.

Secondary digital voltages of 2.5, 3.3, and 5 volts, as well as analog voltages of -5, +5, 12, and 28 volts are produced in the LVPS (low voltage power converter) board, which is connected to the DIGITAL board through a stacking WTAX connector. All secondary voltages are isolated from the primary 28 input voltage from the PFDPU. Secondary returns are connected to chassis. The secondary 28 analog voltage, which powers the high voltage supplies, passes through an enable plug before going to the digital board (where it also passes through a FET switch controlled by the FPGA, which can only be opened after receipt of two software key codes), ensuring that high voltage can not be powered on when not desired.

Instrument temperature sensor readbacks, power for the one-time contamination cover (TiNi pinpuller), and heater power, are connected directly to the spacecraft through a high density D-sub connector. The sensor has three high-density D-sub connections to the PFDPU, which carry primary 28 to the LVPS, command, clock, and telemetry to the DIGITAL board, and power for the mechanical attenuator (which does not enter the SWIA electronics box).

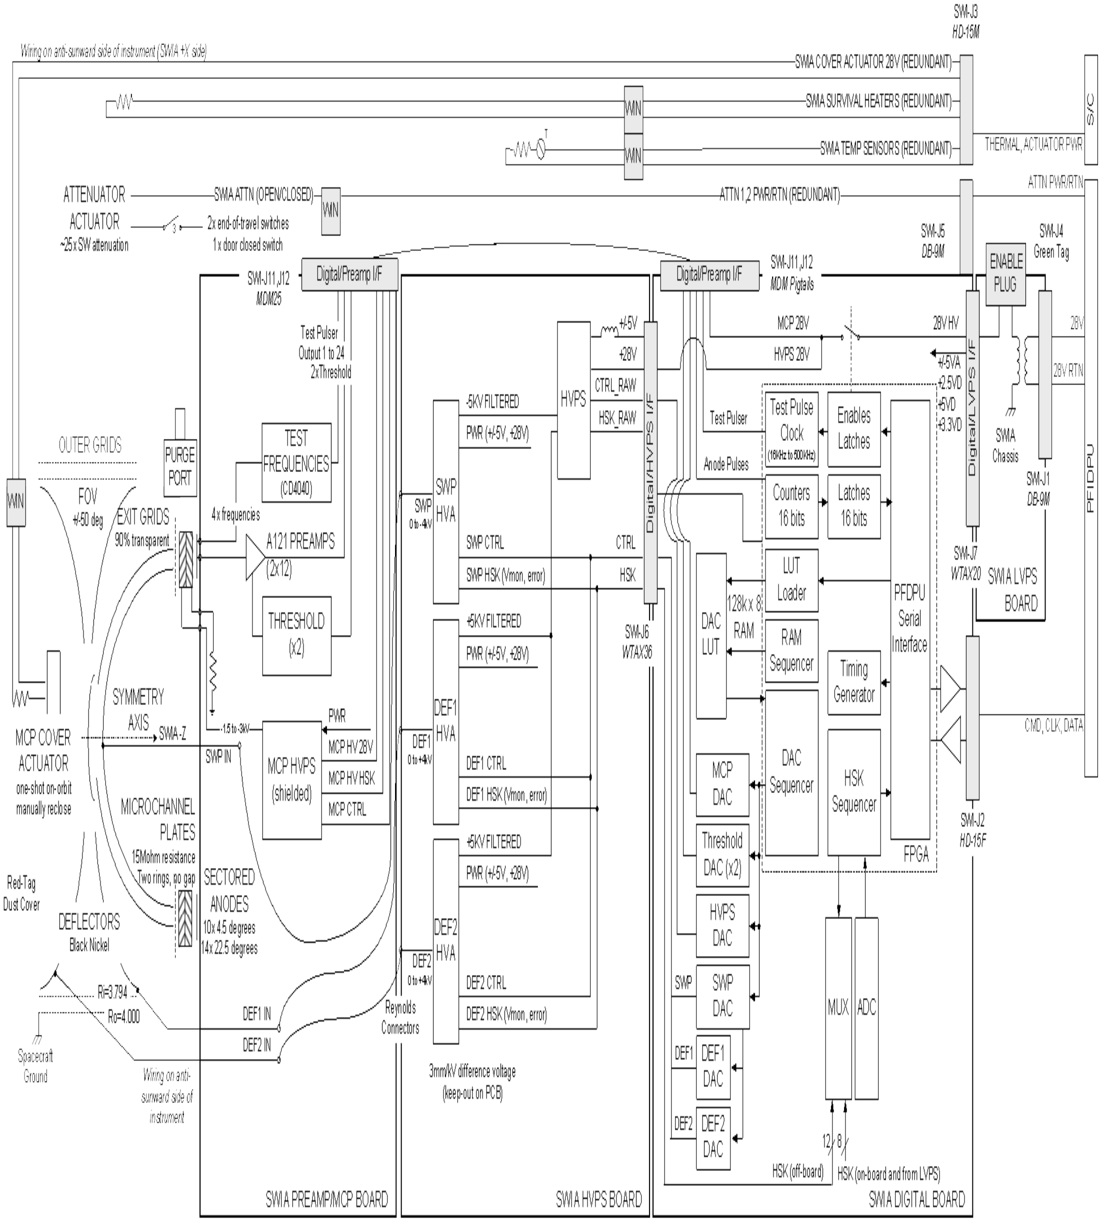

SWIA FUNCTIONAL BLOCK DIAGRAM

E. Taylor
SWIA Block Diagram, Revision 7
May 11, 2011

2.4 Measured Parameters

SWIA covers its entire energy range (96 energy steps), sweeping the deflectors over their full angular range (24 deflection steps) at each energy step, once every four seconds. At each sweep/deflector step, counts are accumulated for each anode for 1.7 milliseconds. From these accumulations, the SWIA FPGA produces three data products (as well as housekeeping and checksums). The P0 product, sent to the PFDPU immediately, contains 96E*24D values for each of the 24 anodes, resulting in a 96x24x24 array. The P1 product, sent to the PFDPU every other energy step, sums over the ten small anodes in two groups to produce a total of 16 22.5 degree bins, sums over groups of 6 deflection steps to produce 4 deflection bins, and sums over adjacent energy steps to produce 48 energy bins. The P2 product buffers all small anode counts for an entire cycle, finds the peak count rate, and sends back the small anode count rates for 48 energy steps and 12 deflection steps around the location of the peak to the PFDPU, with a delay of 4 seconds from data acquisition during which the peak-search and buffer selection occurs. P0, P1, and P2 are all sent from the SWIA FPGA to the PFDPU in the form of small messages of a few tens of words.

In normal operation on orbit, P1 and P2 provide the main data products, since P0 is formidably large and cannot be telemetered from Mars at a reasonable time resolution. The P2 product is designed to measure the very localized intense solar wind flux. The P1 product is designed to measure the more thermalized fluxes in the Martian magnetosphere, and also to search for pickup ions.

The P0, P1, and P2 messages can all be packaged by flight software into data products, packetized, compressed, and telemetered to the ground [P0 => Raw, P1 => Coarse, P2 => Fine]. In normal operation, only Coarse and/or Fine distributions are telemetered. Coarse distribution telemetry has the same angular resolution as the raw P1 product (16 phi bins x 4 deflection angles), and either the intrinsic P1 energy resolution (48 energy steps), or a binned resolution of 24 or 16 energy steps (produced by summing adjacent P1 energy steps in pairs or triplets). Fine distribution telemetry contains the same resolution and coverage as the raw P2 product (10 phi bins x 12 deflection angles x 48 energy steps), or the central subset of these quantities (6 phi bins x 8 deflection angles x 32 energy steps). Binning options are commandable and software-controlled, and can vary depending on telemetry mode [See Section 2.5 below for more detail].

Coarse products can be sampled or summed in powers of two in order to produce a lower time resolution than the intrinsic 4-second resolution, in order to fit within telemetry constraints. Since Fine distributions move in phase space as a function of time in response to changes in the plasma distributions, they can only be sampled. Coarse and Fine 3d products can each be sent either to Survey or to Archive data streams, with different binning options and time resolution possible. All Survey telemetry is returned to the ground, where it can then be used to select stored Archive data for intervals of interest for subsequent playback (subject to telemetry constraints – the amount of Archive data that can be telemetered depends in part on the efficiency of packet compression for the Survey data).

In addition to the full 3d products, an onboard moment computation produces the 13 quantities n, **v**, **p**, and **q** [density, velocity vector, pressure tensor, and heat flux vector] from either P1 or P2 products, taking into account the energy-angle dependence of the instrument response, and the effect of the attenuator when closed. Moment coefficients are stored in the flash memory in the PFDPU.

Finally, flight software produces energy spectra by summing the P1 counts over all angles to produce a 48-element array of accumulated counts. These energy spectra should only be used for quantitative purposes with some care, since there is no weighting for the reduced geometric factor at high deflection angles; instead the spectra represent a simple sum of counts as a function of energy. Their primary purpose is as a diagnostic that can be sent at high time resolution, though they can also be used for quantitative purposes if the distribution is narrow enough (e.g. in the solar wind).

2.5 Operational Modes

SWIA has only one hardware mode. The sensor, in normal operation, always operates from a single high voltage sweep table (loaded from the PFDPU EEPROM at instrument turn-on, stored redundantly in the SWIA FPGA, and check-summed to ensure fidelity, with a re-load to SWIA triggered by any change in the table checksum returned from SWIA), and covers the same range of energies and angles every four seconds. The SWIA sweep table can be changed or re-uploaded to the EEPROM if necessary, but only by ground command [Note: If the sweep table is changed, all moment coefficients also must be changed]. SWIA does have two "modes"; however, these modes simply consist of a different mix of Coarse/Fine telemetry appropriate for different regions of space, rather than an actual change in hardware operation.

SWIA always produces Coarse, Fine, Moment, and Spectra telemetry [as defined above in Section 2.4] at programmable cadences, sized such that the total telemetry fits within the available constraints during each mission phase. A number of options control how and when each of these products is packaged and telemetered; these options can change by telemetry mode and whether the product is bound for Survey or Archive data streams. First, for each telemetry mode and data stream, each product has an associated parameter *n* that determines how often the product is packetized and sent (for a given *n*, a product is sent every 2^n 4-second cycles). In the case of Coarse and Spectra telemetry, the product can either be a sum over the 2^n cycles, or a sample sent every 2^nth cycle. In the case of Fine and Moment telemetry, only sampling is possible, since these products cannot be gracefully summed. The Coarse and Fine 3-d products also have additional options that control the binning of energy steps (for Coarse distributions) and the sub-sampling of a selection of energy/angle bins (for Fine distributions), enabling higher time resolution at the expense of energy/angle resolution or coverage.

SWIA's two telemetry modes are designed to send the most useful mix of products for each plasma region through which MAVEN orbits. In "Solar Wind" mode, we include mostly Fine distributions appropriate for characterizing the narrow solar wind beam in both Survey and Archive data streams, with a smattering of Coarse distributions included to characterize the occurrence of pickup ions. In "Sheath" mode, on the other hand, only Coarse distributions appropriate for characterizing a more thermalized ion distribution are included in both Archive and Survey data streams. In both modes, spectra and moments are sent in the Survey data stream; however, in "Solar Wind" mode the moments are calculated from P2 products, while in "Sheath" mode they are calculated from P1 products. The switch from "Solar Wind" to "Sheath" modes is internally triggered, based on the ratio of the total counts contained in the P1 product to that in the P2 product. If most of the counts are contained in P2, the telemetry switches to "Solar Wind" mode. If the counts are more distributed, with significant counts not contained in the narrow energy/angle range covered by P2, then the telemetry switches to "Sheath" mode. The telemetry mode switches are based on two programmable count ratio thresholds, and are only allowed to occur at a programmable 2^n interval, ensuring that changes in the telemetry mode do not occur too often (for instance if MAVEN crosses an oscillating boundary multiple times in short succession, we intend that only one mode switch should occur).

In addition, SWIA has a mechanical attenuator to reduce the sunward instrument sensitivity that is also internally triggered, based on the peak count rate recorded by the sensor, with two programmable thresholds that trigger opening and closing of the attenuator. Similarly to the mode changes, attenuator actuations can only occur at a programmable 2^n interval, ensuring that the attenuator does not "shutter". In addition to this programmable 2^n limit, a hardware limit ensures that actuations cannot occur too frequently (~5 minute minimum timeout), and a time-based actuation limit prevents the attenuator from operating too often and causing excessive mechanical wear in any mission phase. The attenuation state, which is required to convert counts to physical units since the attenuator changes the geometric factor in the sunward field-of-view, is recorded in all telemetry

2.6 Operational Considerations

During normal operation, SWIA operates continuously in the same hardware mode, as described above in section 2.5.

The SWIA high voltage will be turned off if the spacecraft encounters densities higher than TBD, which may be encountered during deep dips. In such a situation, the high voltage will be disarmed automatically in response to a spacecraft-issued zone alert, and then autonomously ramped back up once the spacecraft returns to a normal density corridor. During these time periods, the SWIA sensor will return no counts since the micro-channel plates will not be operating.

There are very few other constraints on SWIA's operation, other than a restriction that not every type of message (P0, P1, P2, HSK, LUT) can be sent from the SWIA FPGA to the PFDPU simultaneously without exceeding the bandwidth and causing minor timing issues. Since the P0 messages are only for calibration purposes, this limitation has few consequences.

Some odd features in the data can exist around the time of a mode switch or attenuator actuation. If coarse distributions or spectra are in a summing mode, a non-2^n number of distributions may be contained in a packet near the time of a mode change; however, the number of accumulations will be reported correctly in the packet. Of more note is the fact that attenuator states and moments can change in the middle of a moment packet (containing 16 sets of moments), but only one value is provided in the packet header. As a result, if the time of a telemetry mode change or attenuator actuation is not known exactly, the conversion of the moments to physical units can have ambiguities that last for up to 15 sets of moments (typically 60 seconds). Similar issues can exist for the energy spectra (16 sets of spectra per packet).

2.7 Ground Calibration

The SWIA calibration utilized a vacuum chamber with a 3-axis manipulator to scan the sensor, and a controllable electron-impact ion source to produce a collimated beam of ionized residual gases in the chamber to stimulate the sensor. During calibration, we performed a large number of energy/angle scans of the instrument, at a variety of phi and theta angles, and at ion energies of 125 eV and 2 keV. As shown in Fig. 5 below, the energy response is constant to within a few percent around the analyzer, demonstrating that the hemispheres have very good concentricity. The analyzer constant varies from 7.6-7.8 around the analyzer, a variation of <2%. The measured SWIA energy resolution is ~14.5%, in good agreement with simulation results. We also performed a full scan over energies from ~5 eV up to 5 keV to ensure that the analyzer energy response is constant over energy.

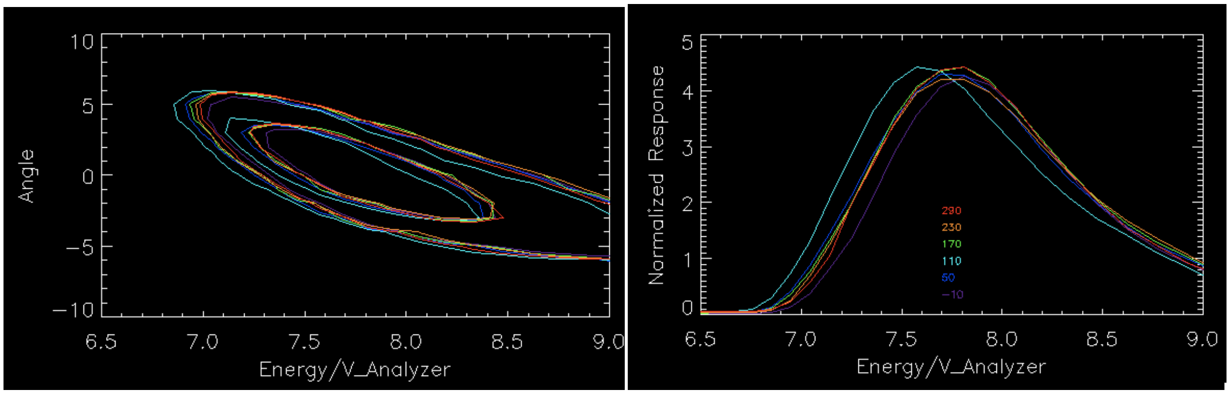

Figure 5: SWIA intrinsic energy/theta response for six different sensor phi angles, showing 0.1 and 0.5 response level contours with no deflection (left), and energy response integrated over all theta angles (right).

We also covered the entire deflection range in our calibration. The analyzer energy response is very constant over the deflection range, but the angular response does vary to some degree due to the focusing properties of the deflection optics, as shown in Fig. 6 below. The correspondence between measurement and simulation is quite good over the entire deflection range, with only a few minor exceptions. Our ion beam is a few degrees wide, so we cannot reproduce the narrow width of the theta angle response predicted by the simulation results at some deflection angles in calibration. Allowing for the slight broadening of the measurement due to beam width, the correspondence is very good.

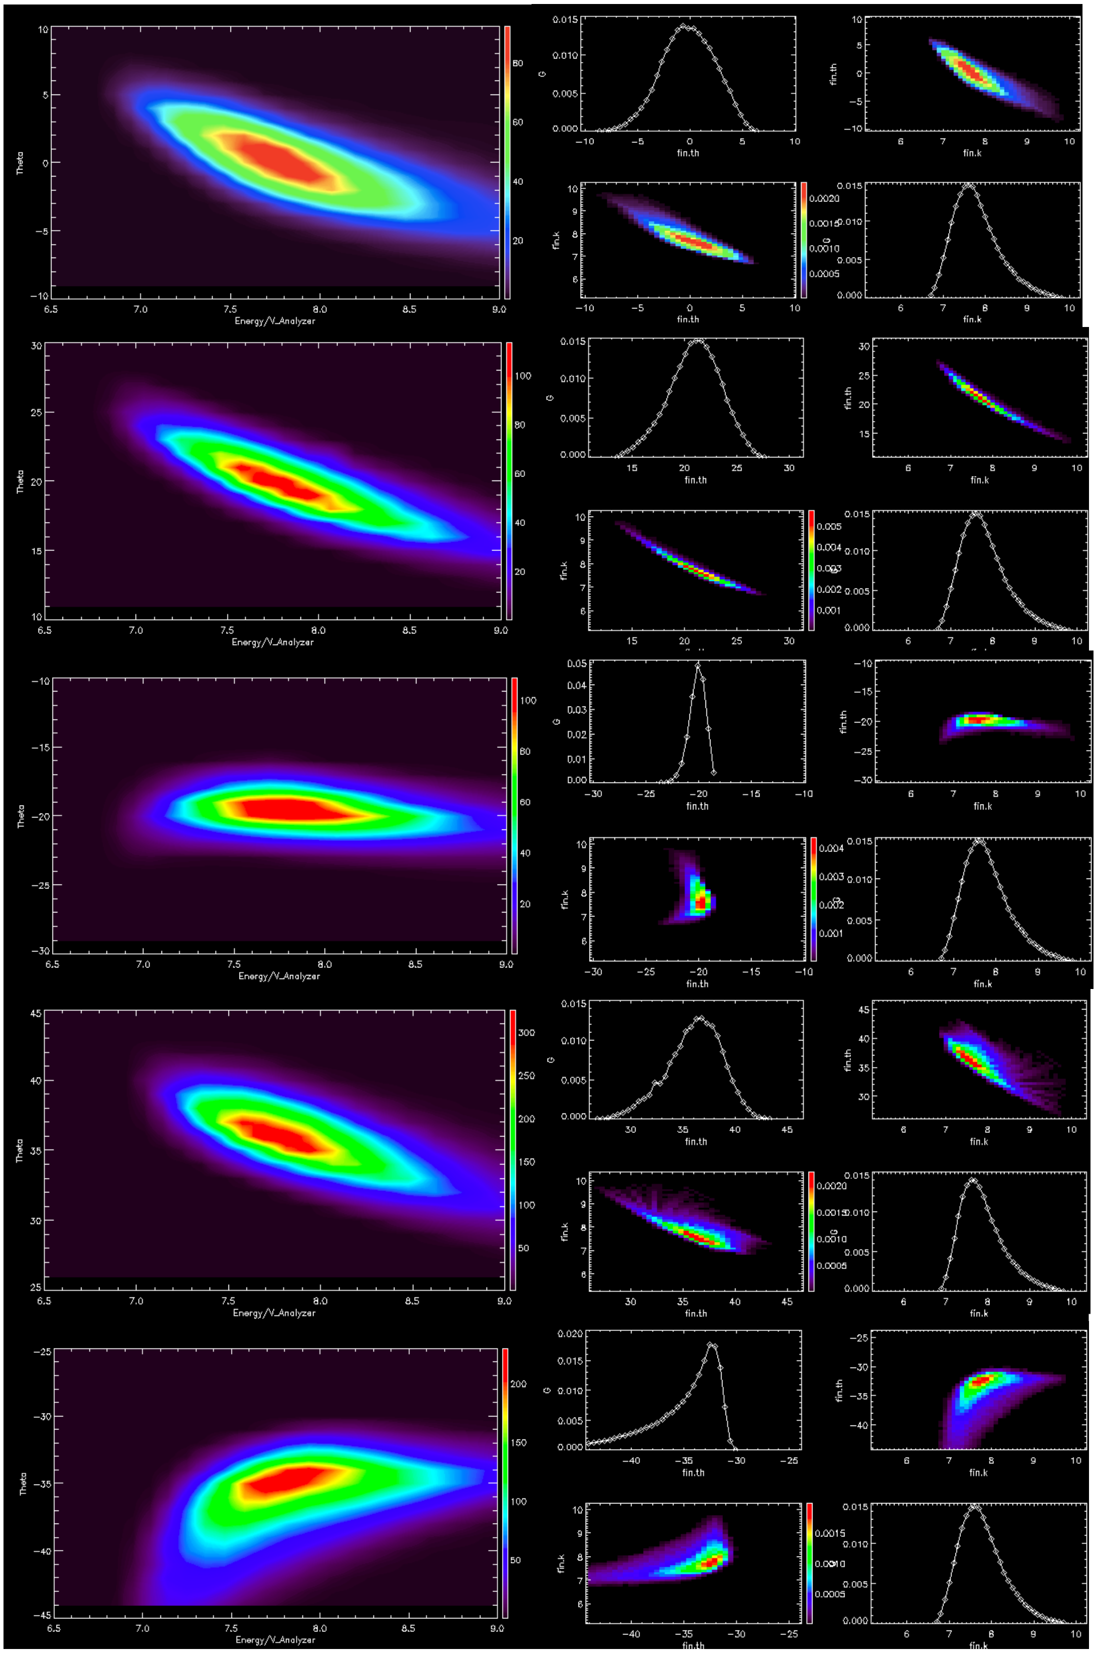
Theta
Energy/V_Analyzer
6.5
7.0
7.5
8.0
8.5
9.0
G
fin.th
fin.k
0.015
0.010
0.005
0.000
0.0020
0.0015
0.0010
0.0005
0.05
0.04
0.03
0.02
0.01
0.00
20
40
60
80
100
150
200
250
300

Figure 6: SWIA energy/theta response for five different deflection angles [0, +20, -20, +35, -35], comparing the measured response (left) and simulation results (right), for the same ratio of inner hemisphere to deflector voltage. Each simulation plot has four panels showing the integrated energy and angular response, and the energy-angle response with the axes in both senses. The measurement should be compared to the upper right panel in each case. Note the slightly different scales of measured and simulated plots.

Finally, we covered the entire sensor phi angle range in our sensor calibration, as shown in Fig. 7 below. The anode resolution is 4.5 degrees in the 45-degree wide region around the sun covered by the small anodes (and the P2 product), and 22.5 degrees outside of that. The dips in the overall response result from mechanical blockage by the ribs in the "spider plate" at the exit of the analyzer that holds the two hemispheres and keeps them concentric (there is no rib facing the sun; thus the increase rather than decrease at phi = 0). The large dip at phi = 180 results from the blockage by the harness cover, where cables carries deflector voltages and attenuator/actuator services to the upper part of the sensor. Some broadening in the phi response results from the spread of the electron cloud from the output face of the channel plates to the anodes, resulting in a moderate ~20-25% level of double-counting between adjacent small anodes.

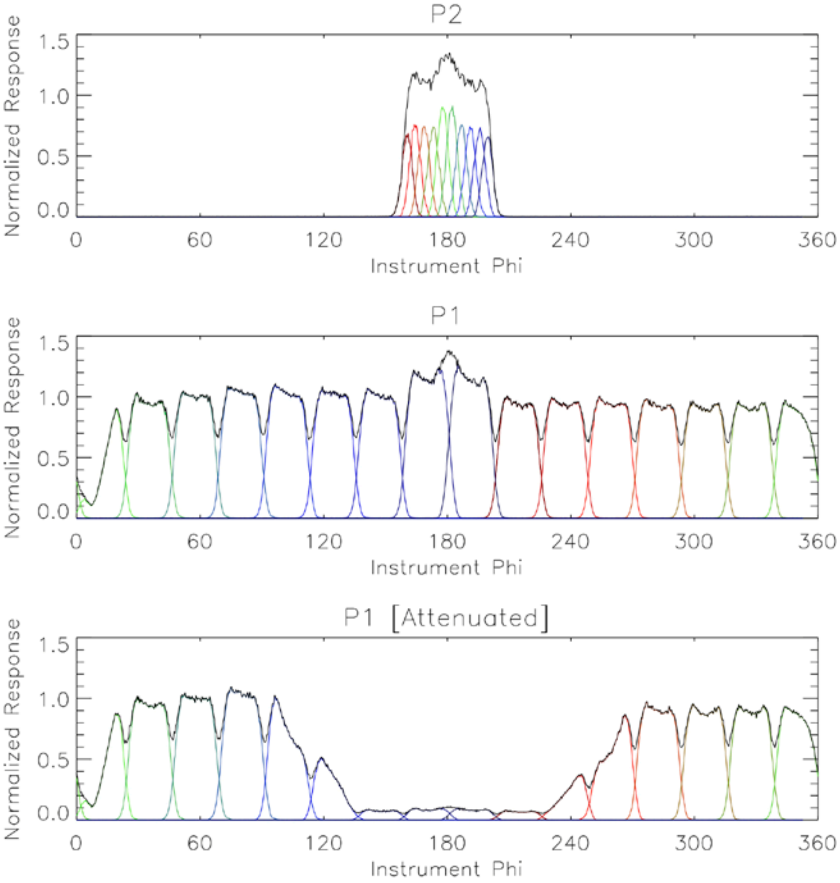

Figure 7: SWIA sensor phi response for coarse (bottom two panels) and fine (top panel) resolution products, measured during a full phi rotation of the instrument. Colored lines show response for each individual anode, white line shows the integrated response. Bottom panel shows response with attenuator closed, top two panels with attenuator open

2.8 Inflight Calibration

SWIA is required to have an in-flight calibration procedure to determine its absolute sensitivity to within 25%. The SWIA angular and energy responses and the geometric factor (minus detection efficiency) is determined on the ground (see above) to within ~10% by calibrations and electrostatic optics simulations. However, to obtain the absolute sensitivity, the detection efficiency must also be known. This efficiency depends on the microchannel plate (MCP) efficiency, which varies during the mission, especially in the first few months of operation. Thus, an in-flight calibration procedure is needed to measure and track this efficiency. The approach is to first determine the detection efficiency of STATIC and then to cross calibrate SWIA with STATIC. There are two independent methods of determining STATIC's detection efficiency, which provide redundancy and cross checks.

First, STATIC can be compared with LPW. LPW measures plasma density with two independent methods. I-V curves provide the total density, temperature, and spacecraft potential in the ionosphere with a 2-sec cadence. Additionally, plasma waves generally show a peak or cutoff at the plasma frequency, which provides an accurate measure of the total density (f_{pe} = 8.98 kHz × $n_e^{1/2}$). Measurements of this peak/cutoff, although not always present, provide an absolute calibration of density measured from the I-V Below 500 km, STATIC is oriented so that the ram direction is in the center plane of its field of view (FOV). The spacecraft velocity ranges from 3.9 to 4.2 km/s, so that the dominant ion (O_2^+) has a ram energy of 2.6-2.9 eV. The best measurements are obtained below ~200 km, where the ion temperature is ~0.03 eV, the spacecraft potential (measured by LPW) is of order –0.1 V, and collisional coupling with the neutral atmosphere limits the plasma bulk flow to less than ~0.3 km/s. Under these conditions, the ion distribution is beamed along the ram direction entirely within STATIC's FOV. Comparisons of STATIC and LPW measurements of the total density under these conditions provide an absolute calibration of STATIC to within 15%.

Secondly, STATIC can determine absolute START and STOP efficiencies from ratios START/(START+STOP) and STOP/(START+STOP) events, which can be combined with mechanical analyzer geometric factor knowledge from electrostatic optics simulations to get an absolute sensitivity, with error determined by mechanical tolerance and supported by the analyzer energy constant. This procedure will work whenever there is only one ion species present – for instance in the solar wind or outer magnetosheath. The results of this analysis should agree with those from the LPW comparison, providing a consistency check.

After STATIC is calibrated, SWIA can be calibrated to STATIC in the magnetosheath or in the solar wind at times with less intense fluxes. This calibration can be performed without any need to calculate a total density moment, since the two measurements overlap in both energy and angle coverage, and the sensors have the same analyzer electrostatic optics. We estimate that this calibration can be made with an accuracy of 5%, so that the absolute sensitivity of SWIA can be determined in flight to better than 20%.

3 Data Overview

This section provides a high level description of archive organization under the PDS4 Information Model (IM) as well as the flow of the data from the spacecraft through delivery to PDS. Unless specified elsewhere in this document, the MAVEN SWIA archive conforms with version 1.1.0.1 of the PDS4 IM [4] and version 1.0 of the MAVEN mission schema. A list of the XML Schema and Schematron documents associated with this archive are provided in Table 5 below.

Table 5: MAVEN SWIA Archive Schema and Schematron

XML Document	Steward	Product LID
PDS Master Schema, v. 1.1.0.1	PDS	urn:nasa:pds:system_bundle:xml_schema:pds-xml_schema
PDS Master Schematron, v. 1.1.0.1	PDS	urn:nasa:pds:system_bundle:xml_schema:pds-xml_schema
MAVEN Mission Schema, v. 1.0	MAVEN	
MAVEN Mission Schematron, v. 1.0	MAVEN	

3.1 Data Reduction Levels

A number of different systems may be used to describe data processing level. This document refers to data by their PDS4 reduction level. Table 5 provides a description of these levels along with the equivalent designations used in other systems.

Table 6: Data reduction level designations

PDS4 reduction level	PDS4 reduction level description	MAVEN Processing Level	CODMAC Level	NASA Level
Raw	Original data from an instrument. If compression, reformatting, packetization, or other translation has been applied to facilitate data transmission or storage, those processes are reversed so that the archived data are in a PDS approved archive format.	0	2	1A
Reduced	Data that have been processed beyond the raw stage but which are not yet entirely independent of the instrument.	1	2	1A
Calibrated	Data converted to physical units entirely independent of the instrument.	2	3	1B

PDS4 reduction level	PDS4 reduction level description	MAVEN Processing Level	CODMAC Level	NASA Level
Derived	Results that have been distilled from one or more calibrated data products (for example, maps, gravity or magnetic fields, or ring particle size distributions). Supplementary data, such as calibration tables or tables of viewing geometry, used to interpret observational data should also be classified as 'derived' data if not easily matched to one of the other three categories.	3+	4+	2+

3.2 Products

A PDS product consists of one or more digital and/or non-digital objects, and an accompanying PDS label file. Labeled digital objects are data products (i.e. electronically stored files). Labeled non-digital objects are physical and conceptual entities which have been described by a PDS label. PDS labels provide identification and description information for labeled objects. The PDS label defines a Logical Identifier (LID) by which any PDS labeled product is referenced throughout the system. In PDS4 labels are XML formatted ASCII files. More information on the formatting of PDS labels is provided in Section 6.3. More information on the usage of LIDs and the formation of MAVEN LIDs is provided in Section 5.1.

3.3 Product Organization

The highest level of organization for PDS archive is the bundle. A bundle is a list of one or more related collections of products, which may be of different types. A collection is a list of one or more related basic products, which are all of the same type. Figure 8 below illustrates these relationships.

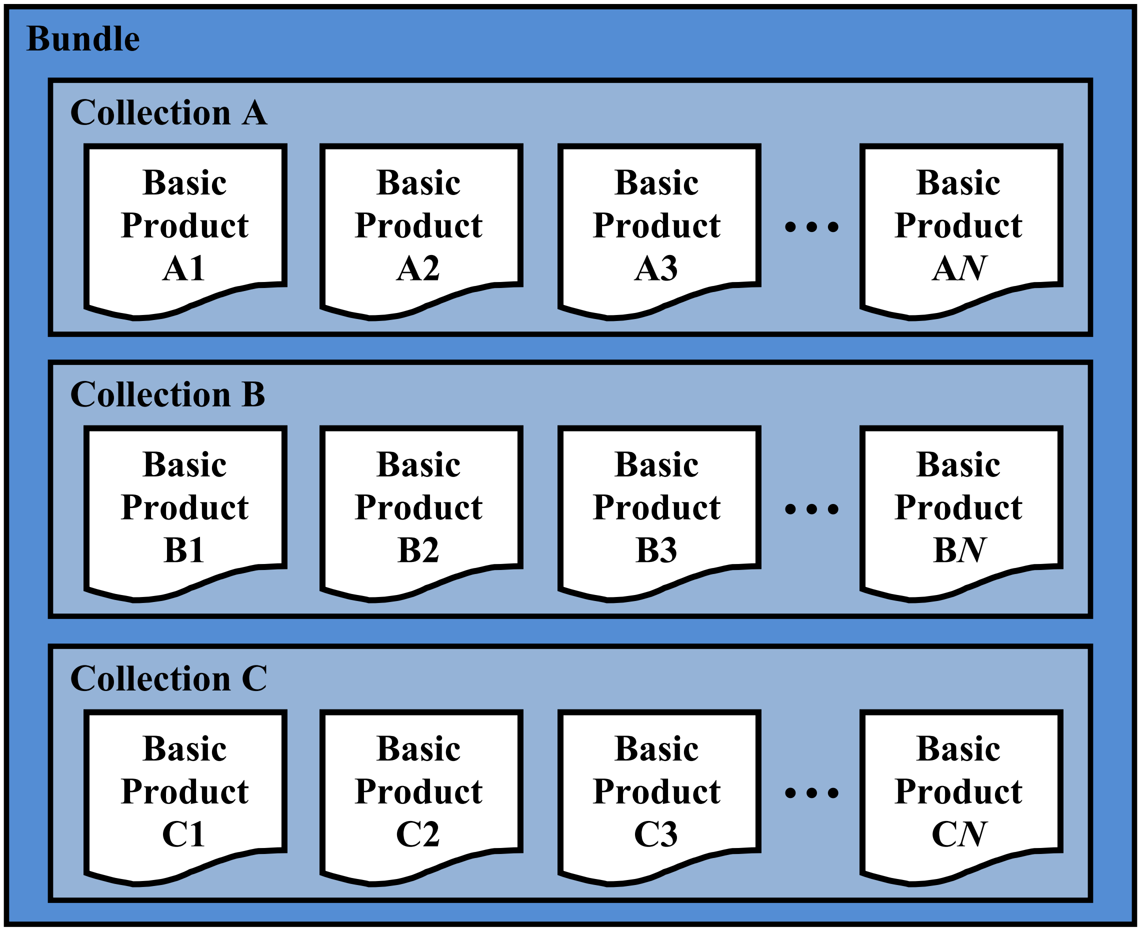

Figure 8: A graphical depiction of the relationship among bundles, collections, and basic products.

Bundles and collections are logical structures, not necessarily tied to any physical directory structure or organization. Bundle and collection membership is established by a member inventory list. Bundle member inventory lists are provided in the bundle product labels themselves. Collection member inventory lists are provided in separate collection inventory table files. Sample bundle and collection labels are provided in Appendix C and Appendix D, respectively.

3.3.1 Collection and Basic Product Types

Collections are limited to a single type of basic products. The types of archive collections that are defined in PDS4 are listed in Table 6.

Table 7: Collection product types

Collection Type	Description
Browse	Contains products intended for data characterization, search, and viewing, and not for scientific research or publication.

Context	Contains products which provide for the unique identification of objects which form the context for scientific observations (e.g. spacecraft, observatories, instruments, targets, etc.).
Document	Contains electronic document products which are part of the PDS Archive.
Data	Contains scientific data products intended for research and publication.
SPICE	Contains NAIF SPICE kernels.
XML_Schema	Contains XML schemas and related products which may be used for generating and validating PDS4 labels.

3.4 Bundle Products

The SWIA data archive is organized into 1 bundle. A description of the bundle is provided in Table 7, and a more detailed description of the contents and format is provided in Section 5.2.

Table 8: SWIA Bundles

3.5 Data Flow

This section describes only those portions of the MAVEN data flow that are directly connected to archiving. A full description of MAVEN data flow is provided in the MAVEN Science Data Management Plan [5]. A graphical representation of the full MAVEN data flow is provided in Figure 9 below.

All ITFs will produce calibrated products. Following an initial 2-month period at the beginning of the mapping phase, the ITFs will routinely deliver preliminary calibrated data products to the SDC for use by the entire MAVEN team within two weeks of ITF receipt of all data needed to generate those products. The SOC will maintain an active archive of all MAVEN science data, and will provide the MAVEN science team with direct access through the life of the MAVEN mission. After the end of the MAVEN project, PDS will be the sole long-term archive for all public MAVEN data.

Updates to calibrations, algorithms, and/or processing software are expected to occur regularly, resulting in appropriate production system updates followed by reprocessing of science data products by ITFs for delivery to SDC. Systems at the SOC, ITFs and PDS are designed to handle these periodic version changes.

Bundle Logical Identifier	PDS4 Reduction Level	Description	Data Provider
urn:nasa:pds:maven.swia.calibrated	Calibrated	Fully calibrated ion velocity distributions, energy spectra, and density, temperature, and velocity moments from onboard calculations. Tables of sensitivity and energy/angle maps included in files.	ITF

Data bundles intended for the archive are identified in Table 7.

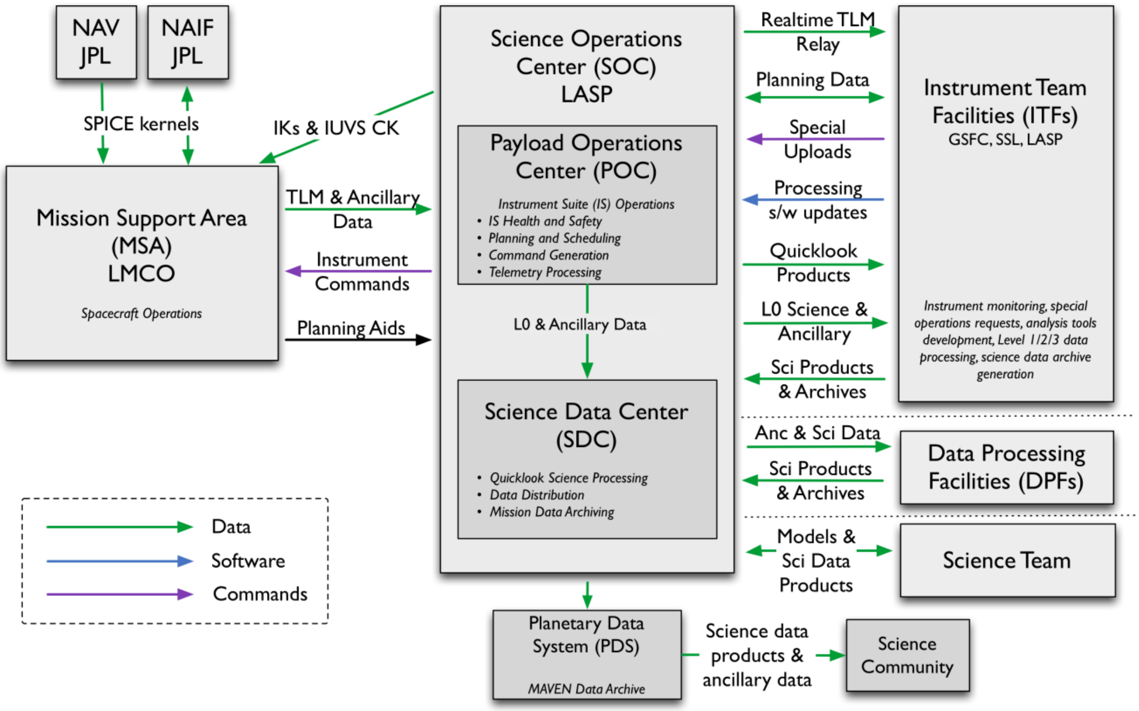

Figure 9: MAVEN Ground Data System responsibilities and data flow. Note that this figure includes portions of the MAVEN GDS which are not directly connected with archiving, and are therefore not described in Section 3.5 above.

4 Archive Generation

The SWIA archive products are produced by the SWIA team in cooperation with the SDC, and with the support of the PDS Planetary Plasma Interactions (PPI) Node at the University of California, Los Angeles (UCLA). The archive volume creation process described in this section sets out the roles and responsibilities of each of these groups. The assignment of tasks has been agreed upon by all parties. Archived data received by the PPI Node from the SWIA team are made available to PDS users electronically as soon as practicable but no later two weeks after the delivery and validation of the data.

4.1 Data Processing and Production Pipeline

The following sections describe the process by which data products in each of the SWIA bundles listed in Table 7 are produced.

4.1.1 Raw Data Production Pipeline

After receiving Level 0 data from the POC, the SDC will process the Level 0 into Quicklook science products using software provided by the SWIA ITF. The SDC will provide the SWIA ITF with Level 0 data files (consisting of compressed PF packets in their native format, one file per UT day for all PF Survey data, and one file per UT day for all PF Archive data), Quicklook science data and all ancillary data required for science processing. From this data, the SWIA ITF will generate Level 2 calibrated science data products. The science data products that the SWIA ITF delivers to the SDC will be stored by the SDC for the duration of the project, and will be made available to the MAVEN team. The SDC will deliver archival-quality science data products to the PDS for distribution to the public and long-term archiving in accordance with the SWIA-PDS SIS (this document) and the SOC-PDS SIS. The SDC will also be responsible for delivering Level 0 archives and non-SPICE ancillary data to the PDS for long-term archiving, in accordance with the SOC-PDS SIS and the Export Control Checklist.

4.1.2 Calibrated Data Production Pipeline

Calibrated SWIA Level 2 data will be produced from the raw level 0 PF data files by the SWIA ITF using IDL software, and provided for archiving in the PDS in CDF format. The data production pipeline will be run in an automated fashion to produce archival-ready files from the raw level 0 data.

Beginning as soon as possible but no later than 2 months after the start of science operations, the SWIA ITF will routinely generate Level 2 science data products and deliver them to the SOC. After the initial 2-month calibration period, the SWIA ITF will deliver preliminary Level 2 products to the SDC for distribution to the MAVEN team within two weeks of receiving all data required for science processing (including all SPICE kernels and other ancillary data required for processing) by the ITFs. Final Level 2 SWIA products will be delivered to the SDC as soon as they are complete, no later than needed to meet the PDS delivery schedule in Table 9.

The SWIA ITF does not plan to produce Level 3 products, instead using Level 2 as the final science products.

The SWIA ITF will deliver validated science data products and associated metadata for PDS archiving to the SOC two weeks prior to every PDS delivery deadline. The first PDS delivery will occur no later than 6 months after the start of science operations, and subsequent deliveries will take place every 3 months after the first delivery. The first delivery will include data collected during the cruise and transition phases in addition to the science data from the first 3 months of the mapping phase. Each subsequent delivery will contain data from the 3 months following the previous delivery. The final delivery may contain products involving data from the entire mission.

The SWIA ITF will also provide the SDC with data product descriptions, appropriate for use by the MAVEN science team in using MAVEN science data products and consistent with PDS metadata standards.

4.2 Data Validation

4.2.1 Instrument Team Validation

All SWIA data will be calibrated and converted to physical units by the SWIA ITF, then spot-checked by the instrument lead and his designees for accuracy and integrity.

4.2.2 MAVEN Science Team Validation

The MAVEN science team will work with the same SWIA products that will be archived in the PDS. If any calibration issues or other anomalies are noted, they will be addressed at the SWIA ITF by the instrument lead or his designees.

4.2.3 PDS Peer Review

The PPI node will conduct a full peer review of all of the data types that the SWIA team intends to archive. The review data will consist of fully formed bundles populated with candidate final versions of the data and other products and the associated metadata.

Table 9: MAVEN PDS review schedule

Date	Activity	Responsible Team
2014-Mar-24	Signed SIS deadline	ITF
2014-Apr-18	Sample data products due	ITF
2014-May to 2014-Aug	Preliminary PDS Peer Review (SIS, sample data files)	PDS
2015-Mar-02	Release #1: Data due to PDS	ITF/SDC
2014-Mar to 2015-Apr	Release #1: Data PDS Peer Review	PDS
2015-May-01	Release #1: Public release	PDS

Reviews will include a preliminary delivery of sample products for validation and comment by PDS PPI and Engineering node personnel. The data provider will then address the comments coming out of the preliminary review, and generate a full archive delivery to be used for the peer review.

Reviewers will include MAVEN Project and SWIA team representatives, researchers from outside of the MAVEN project, and PDS personnel from the Engineering and PPI nodes. Reviewers will examine the sample data products to determine whether the data meet the stated science objectives of the instrument and the needs of the scientific community and to verify that the accompanying metadata are accurate and complete. The peer review committee will identify any liens on the data that must be resolved before the data can be ‘certified’ by PDS, a process by which data are made public as minor errors are corrected.

In addition to verifying the validity of the review data, this review will be used to verify that the data production pipeline by which the archive products are generated is robust. Additional deliveries made using this same pipeline will be validated at the PPI node, but will not require additional external review.

As expertise with the instrument and data develops the SWIA team may decide that changes to the structure or content of its archive products are warranted. Any changes to the archive products or to the data production pipeline will require an additional round of review to verify that the revised products still meet the original scientific and archival requirements or whether those criteria have been appropriately modified. Whether subsequent reviews require external reviewers will be decided on a case-by-case basis and will depend upon the nature of the changes. A comprehensive record of modifications to the archive structure and content is kept in the Modification_History element of the collection and bundle products.

The instrument team and other researchers are encouraged to archive additional SWIA products that cover specific observations or data-taking activities. The schedule and structure of any additional archives are not covered by this document and should be worked out with the PPI node.

4.3 Data Transfer Methods and Delivery Schedule

The SOC is responsible for delivering data products to the PDS for long-term archiving. While ITFs are primarily responsible for the design and generation of calibrated and derived data archives, the archival process is managed by the SOC. The SOC (in coordination with the ITFs) will also be primarily responsible for the design and generation of the raw data archive. The first PDS delivery will take place within 6 months of the start of science operations. Additional deliveries will occur every following 3 months and one final delivery will be made after the end of the mission. Science data are delivered to the PDS within 6 months of its collection. If it becomes necessary to reprocess data which have already been delivered to the archive, the ITFs will reprocess the data and deliver them to the SDC for inclusion in the next archive delivery. A summary of this schedule is provided in Table 9 below.

Table 10: Archive bundle delivery schedule

Bundle Logical Identifier	First Delivery to PDS	Delivery Schedule	Estimated Delivery Size
urn:nasa:pds:maven.swia.calibrated	No later than 6 months after the start of science operations	Every 3 months	TBR

Each delivery will comprise both data and ancillary data files organized into directory structures consistent with the archive design described in Section 5, and combined into a deliverable file(s) using file archive and compression software. When these files are unpacked at the PPI Node in the appropriate location, the constituent files will be organized into the archive structure.

Archive deliveries are made in the form of a “delivery package”. Delivery packages include all of the data being transferred along with a transfer manifest, which helps to identify all of the products included in the delivery, and a checksum manifest which helps to insure that integrity of the data is maintained through the delivery. The format of these files is described in Section 6.4.

Data are transferred electronically (using the *ssh* protocol) from the SOC to an agreed upon location within the PPI file system. PPI will provide the SOC a user account for this purpose. Each delivery package is made in the form of a compressed *tar* or *zip* archive. Only those files that have changed since the last delivery are included. The PPI operator will decompress the data, and verify that the archive is complete using the transfer and MD5 checksum manifests that were included in the delivery package. Archive delivery status will be tracked using a system defined by the PPI node.

Following receipt of a data delivery, PPI will reorganize the data into its PDS archive structure within its online data system. PPI will also update any of the required files associated with a PDS archive as necessitated by the data reorganization. Newly delivered data are made available publicly through the PPI online system once accompanying labels and other documentation have been validated. It is anticipated that this validation process will require no more than fourteen working days from receipt of the data by PPI. However, the first few data deliveries may require more time for the PPI Node to process before the data are made publicly available.

The MAVEN prime mission begins approximately 5 weeks following MOI and lasts for 1 Earth-year. Table 9 shows the data delivery schedule for the entire mission.

4.4 Data Product and Archive Volume Size Estimates

SWIA data products consist of files that span one UT day, breaking at 0h UTC SCET. Files vary in size depending on the telemetry rate and allocation.

4.5 Data Validation

Routine data deliveries to the PDS are validated at the PPI node to insure that the delivery meets PDS standards, and that they the data conform to the standards defined in the SIS, and set in the peer review. As long as there are no changes to the data product formats, or data production pipeline no additional external review will be conducted.

4.6 Backups and duplicates

The PPI Node keeps three copies of each archive product. One copy is the primary online archive copy, another is an onsite backup copy, and the final copy is an off-site backup copy. Once the archive products are fully validated and approved for inclusion in the archive, copies of the products are sent to the National Space Science Data Center (NSSDC) for long-term archive in a NASA-approved deep-storage facility. The PPI Node may maintain additional copies of the archive products, either on or off-site as deemed necessary. The process for the dissemination and preservation of SWIA data is illustrated in Figure 10.

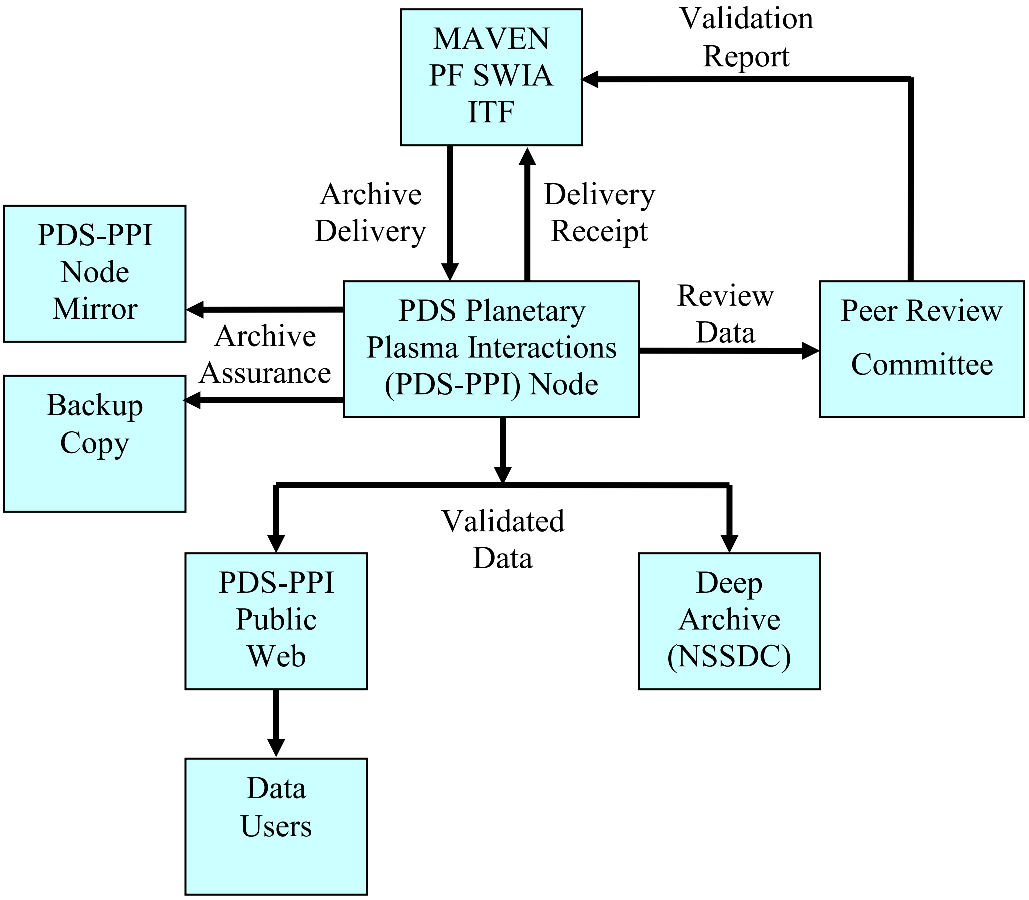

Figure 10: Duplication and dissemination of SWIA archive products at PDS/PPI.

5 Archive organization and naming

This section describes the basic organization of an SWIA bundle, and the naming conventions used for the product logical identifiers, and bundle, collection, and basic product filenames.

5.1 Logical Identifiers

Every product in PDS is assigned an identifier which allows it to be uniquely identified across the system. This identifier is referred to as a Logical Identifier or LID. A LIDVID (Versioned Logical Identifier) includes product version information, and allows different versions of a specific product to be referenced uniquely. A product's LID and VID are defined as separate attributes in the product label. LIDs and VIDs are assigned by the entity generating the labels and are formed according to the conventions described in sections 5.1.1 and 5.1.2 below. The uniqueness of a product's LIDVID may be verified using the PDS Registry and Harvest tools.

5.1.1 LID Formation

LIDs take the form of a Uniform Resource Name (URN). LIDs are restricted to ASCII lower case letters, digits, dash, underscore, and period. Colons are also used, but only to separate prescribed components of the LID. Within one of these prescribed components dash, underscore, or period are used as separators. LIDs are limited in length to 255 characters.

MAVEN SWIA LIDs are formed according to the following conventions:

- Bundle LIDs are formed by appending a bundle specific ID to the MAVEN SWIA base ID:

  urn:nasa:pds:maven.swia.<bundle ID>

  Since all PDS bundle LIDs are constructed this way, the combination of maven.swia.bundle must be unique across all products archived with the PDS.

- Collection LIDs are formed by appending a collection specific ID to the collection's parent bundle LID:

  urn:nasa:pds:maven.swia.<bundle ID>:<collection ID>

  Since the collection LID is based on the bundle LID, which is unique across PDS, the only additional condition is that the collection ID must be unique across the bundle. Collection IDs correspond to the collection type (e.g. "browse", "data", "document", etc.). Additional descriptive information may be appended to the collection type (e.g. "data-raw", "data-calibrated", etc.) to insure that multiple collections of the same type within a single bundle have unique LIDs.

- Basic product LIDs are formed by appending a product specific ID to the product's parent collection LID:

  urn:nasa:pds:maven.swia.<bundle ID>:<collection ID>:<product ID>

  Since the product LID is based on the collection LID, which is unique across PDS, the only additional condition is that the product ID must be unique across the collection.

A list of SWIA bundle LIDs is provided in Table 7. Collection LIDs are listed in Table 10.

5.1.2 VID Formation

Product version ID's consist of major and minor components separated by a "." (M.n). Both components of the VID are integer values. The major component is initialized to a value of "1", and the minor component is initialized to a value of "0". The minor component resets to "0" when the major component is incremented.

5.2 SWIA Archive Contents

The SWIA archive includes the calibrated (MAVEN level 2) bundle listed in Table 7. The following section describes the contents of this bundle in greater detail.

5.2.1 SWIA Calibrated (MAVEN Level 2) Science Data Bundle

The swia.calibrated Level 2 Science Data Bundle contains fully calibrated data in physical units, consisting of Coarse and Fine resolution 3d distributions and energy spectra and moments from onboard computations.

Table 11: swia.calibrated Level 2 Science Data Collections

Collection LID	Description
urn:nasa:pds:maven.swia.calibrated:data.coarse_svy_3d	Full 3d ion distributions in units of differential energy flux from SWIA coarse survey data
urn:nasa:pds:maven.swia.calibrated:data.coarse_arc_3d	Full 3d ion distributions in units of differential energy flux from SWIA coarse archive data
urn:nasa:pds:maven.swia.calibrated:data.fine_svy_3d	Full 3d ion distributions in units of differential energy flux from SWIA fine survey data
urn:nasa:pds:maven.swia.calibrated:data.fine_arc_3d	Full 3d ion distributions in units of differential energy flux from SWIA fine archive data
urn:nasa:pds:maven.swia.calibrated:data.onboard_svy_mom	Ion density, temperature, and velocity moments in physical units and coordinates calculated onboard from SWIA coarse or fine data (depending on mode)
urn:nasa:pds:maven.swia.calibrated:data.onboard_svy_spec	Ion energy spectra in units of differential energy flux calculated onboard from SWIA coarse data
urn:nasa:pds:maven.swia.calibrated:document	Documents related to the swia.calibrated bundle.
urn:nasa:pds:maven.swia.calibrated:browse	

5.2.1.1 swia.calibrated.coarse_svy_3d Data Collection

SWIA coarse survey 3d collections contain files with time-ordered fully calibrated ion distributions in units of differential energy flux derived from the SWIA Coarse distribution Survey telemetry, as well as a header of ancillary information needed to interpret the data.

The data files contain a time-ordered array with time in Epoch time, Mission-Elapsed-Time (MET) and Unix time (Seconds since 1970-01-01/00:00), a 48 energies X 4 deflection (theta) angles X 16 anode (phi) angles array of data, the attenuator state (1 = open, 2 = closed), the data binning format (0 = 48 energies, 1 = 24 binned energies, 2 = 16 binned energies), and the number of accumulations per data product, at each time step.

For data with 24 or 16 binned energies, a full 48-energy data structure is provided for commonality and ease of use, but the 48 energy steps contain the binned counts divided by the binning factor and duplicated in pairs or triplets. For example, for a case with 16 binned energy steps, the first three steps of the 48 energies contain the number of counts in the first binned energy step divided by three. This re-binning scheme ensures that moments or other sums computed from binned Coarse data still come out right, and use of Coarse data with different binning schemes is transparent to the end user.

The data files contain a 48-element list of energies, two 48x4-element arrays of theta angles for attenuator open and closed, two 48x4-element arrays of relative sensitivities for attenuator open and closed, a 16-element array of phi angles, two 16-element arrays of anode geometric factors for attenuator open and closed, and integration time, for use with the coarse data products. The files also contain the energy resolution of the sweep tables and full sensor geometric factor. All of these support data are stored as /novary records in the CDF files.

The SWIA ITF will produce these products, with one file per UT day, with the naming convention mvn_swi_l2_coarsesvy3d_<yyyy><mm><dd>T<hh><mm><ss>_v<xx>_r<yy>.cdf

5.2.1.2 swia.calibrated.coarse_arc_3d Data Collection

SWIA coarse archive 3d collections contain files with time-ordered fully calibrated ion distributions in units of differential energy flux derived from the SWIA Coarse distribution Archive telemetry, as well as a header of ancillary information needed to interpret the data.

The data files contain a time-ordered array with time in Epoch time, Mission-Elapsed-Time (MET) and Unix time (Seconds since 1970-01-01/00:00), a 48 energies X 4 deflection (theta) angles X 16 anode (phi) angles array of data, the attenuator state (1 = open, 2 = closed), the data binning format (0 = 48 energies, 1 = 24 binned energies, 2 = 16 binned energies), and the number of accumulations per data product, at each time step.

For data with 24 or 16 binned energies, a full 48-energy data structure is provided for commonality and ease of use, but the 48 energy steps contain the binned counts divided by the binning factor and duplicated in pairs or triplets. For example, for a case with 16 binned energy steps, the first three steps of the 48 energies contain the number of counts in the first binned energy step divided by three. This re-binning scheme ensures that moments or other sums computed from binned Coarse data still come out right, and use of Coarse data with different binning schemes is transparent to the end user.

The data files contain a 48-element list of energies, two 48x4-element arrays of theta angles for attenuator open and closed, two 48x4-element arrays of relative sensitivities for attenuator open and closed, a 16-element array of phi angles, two 16-element arrays of anode geometric factors for attenuator open and closed, and integration time, for use with the coarse data products. The files also contain the energy resolution of the sweep tables and full sensor geometric factor. All of these support data are stored as /novary records in the CDF files.

The SWIA ITF will produce these products, with one file per UT day, with the naming convention mvn_swi_l2_coarsearc3d_<yyyy><mm><dd>T<hh><mm><ss>_v<xx>_r<yy>.cdf

5.2.1.3 swia.calibrated.fine_svy_3d Data Collection

SWIA fine survey 3d collections contains files with time-ordered fully calibrated ion distributions in units of differential energy flux derived from the SWIA Fine distribution Survey telemetry, as well as a header of ancillary information needed to interpret the data.

The data files contain a time-ordered array with time in Epoch time, Mission-Elapsed-Time (MET) and Unix time (Seconds since 1970-01-01/00:00), a 48 energies X 12 deflection (theta) angles X 10 anode (phi) angles array of data, the starting energy step, the starting deflection step, the attenuator state (1 = open, 2 = closed), and the data format (0 = 48 energies X 12 deflections X 10 anodes, 1 = 32 energies X 8 deflections X 6 anodes), at each time step.

For the case with data format = 1, a full 48x12x10 array is still provided, but with only the central values containing non-zero counts. This allows Fine data products with different formats to be mixed in a transparent way.

In order to determine which elements of the energy and angle tables in the supporting data to use, it is necessary to use the starting energy step and deflection step. In other words, the correct energies for each distribution will consist of the 48 elements of the 96 energies ranging from the starting energy step (0 to 48) to the starting energy step + 47. Similarly, the correct deflection angles for each distribution will consist of the 12 elements of the 24 deflection steps ranging from the starting deflection step (0 to 12) to the starting deflection step + 11.

The data files contain a 96-element list of energies, two 96x24-element arrays of deflection angles for attenuator open and closed, two 96x24-element arrays of relative sensitivities for attenuator open and closed, a 10-element array of phi angles, two 10-element arrays of anode geometric factors for attenuator open and closed, and integration time, for fine data products. The files also contain the energy resolution of the sweep tables and full sensor geometric factor. All of these support data are stored as /novary records in the CDF files.

The SWIA ITF will produce these products, with one file per UT day, with the naming convention mvn_swi_l2_finesvy3d_<yyyy><mm><dd>T<hh><mm><ss>_v<xx>_r<yy>.cdf

5.2.1.4 swia.calibrated.fine_arc_3d Data Collection

SWIA fine archive 3d collections contain files with time-ordered fully calibrated ion distributions in units of differential energy flux derived from the SWIA Fine distribution Archive telemetry, as well as a header with ancillary information needed to interpret the data.

The data files contain a time-ordered array with time in Epoch time, Mission-Elapsed-Time (MET) and Unix time (Seconds since 1970-01-01/00:00), a 48 energies X 12 deflection (theta) angles X 10 anode (phi) angles array of data, the starting energy step, the starting deflection step, the attenuator state (1 = open, 2 = closed), and the data format (0 = 48 energies X 12 deflections X 10 anodes, 1 = 32 energies X 8 deflections X 6 anodes), at each time step.

For the case with data format = 1, a full 48x12x10 array is still provided, but with only the central values containing non-zero counts. This allows Fine data products with different formats to be mixed in a transparent way.

In order to determine which elements of the energy and angle tables in the supporting data to use, it is necessary to use the starting energy step and deflection step. In other words, the correct energies for each distribution will consist of the 48 elements of the 96 energies ranging from the starting energy step (0 to 48) to the starting energy step + 47. Similarly, the correct deflection angles for each distribution will consist of the 12 elements of the 24 deflection steps ranging from the starting deflection step (0 to 12) to the starting deflection step + 11.

The data files contain a 96-element list of energies, two 96x24-element arrays of deflection angles for attenuator open and closed, two 96x24-element arrays of relative sensitivities for attenuator open and closed, a 10-element array of phi angles, two 10-element arrays of anode geometric factors for attenuator open and closed, and integration time, for fine data products. The files also contain the energy resolution of the sweep tables and full sensor geometric factor. All of these support data are stored as /novary records in the CDF files.

The SWIA ITF will produce these products, with one file per UT day, with the naming convention mvn_swi_l2_finearc3d_<yyyy><mm><dd>T<hh><mm><ss>_v<xx>_r<yy>.cdf

5.2.1.5 swia.calibrated.onboard_svy_mom Data Collection

SWIA onboard moment collections contain files with time-ordered ion moments converted to physical units and coordinates, as computed onboard from Coarse and Fine ion distributions, as well as a header with ancillary information needed to interpret the moments.

The data files contain a time-ordered array with time in Epoch time, Mission-Elapsed-Time (MET) and Unix time (Seconds since 1970-01-01/00:00), the calculated ion density, three components of temperature in instrument coordinates and Mars Solar Orbital coordinates, three components of velocity in instrument and Mars Solar Orbital coordinates, the attenuator state (1 = open, 2 = closed), and the telemetry mode (1 = ‘Sheath’, 0 = “Solar Wind”) that defines which distribution the moments are calculated from, at each time step.

The SWIA ITF will produce these products, with one file per UT day, with the naming convention
mvn_swi_l2_onboardsvymom_<yyyy><mm><dd>T<hh><mm><ss>_v<xx>_r<yy>.cdf

5.2.1.6 swia.calibrated.onboard_svy_spec Data Collection

SWIA onboard moment collections contain files with time-ordered angle-averaged ion energy spectra in units of differential energy flux, as computed onboard from Coarse ion distributions, as well as a header with ancillary information needed to interpret the spectra.

The data files contain a time-ordered array with time in Epoch time, Mission-Elapsed-Time (MET) and Unix time (Seconds since 1970-01-01/00:00), a 48-element array of angle-averaged differential energy fluxes, the attenuator state (1 = open, 2 = closed), and the number of accumulations used to compute the spectra, at each time step.

The data files also contain a 48-element list of energies, and the intrinsic energy resolution of the sweep table, integration time, and full sensor geometric factor. All of these support data are stored as /novary records in the CDF files.

The SWIA ITF will produce these products, with one file per UT day, with the naming convention
mvn_swi_l2_onboardsvyspec_<yyyy><mm><dd>T<hh><mm><ss>_v<xx>_r<yy>.cdf

5.2.1.7 swia.calibrated Document Collection

The SWIA calibrated data document collection contains documents which are useful for understanding and using the SWIA Calibrated (MAVEN Level 2) Science Data bundle. Table 11 contains a list of the documents included in this collection, along with the LID, and responsible group. Following this a brief description of each document is also provided.

Table 12: SWIA Calibrated Science Data Documents

Document Name	LID	Responsiblility
MAVEN Science Data Management Plan	urn:nasa:pds:maven:document:sdmp	MAVEN Project
MAVEN SWIA Archive SIS	urn:nasa:pds:maven.swia:document:sis	SWIA Team
MAVEN SWIA Instrument Paper	urn:nasa:pds:maven.swia:document:instpaper	SWIA Team

MAVEN Science Data Management Plan – describes the data requirements for the MAVEN mission and the plan by which the MAVEN data system will meet those requirements

MAVEN SWIA Archive SIS – describes the format and content of the SWIA PDS data archive, including descriptions of the data products and associated metadata, and the archive format, content, and generation pipeline (this document)

MAVEN SWIA Instrument Paper – describes the instrument operation and data products.

While responsibility for the individual documents varies, the document collection itself is managed by the PDS/PPI node.

6 Archive products formats

Data that comprise the SWIA archives are formatted in accordance with PDS specifications [see *Planetary Science Data Dictionary* [4], *PDS Data Provider's Handbook* [2], and *PDS Standards Reference* [3]. This section provides details on the formats used for each of the products included in the archive.

6.1 Data File Formats

This section describes the format and record structure of each of the data file types.

6.1.1 Calibrated data file structure

SWIA calibrated data files will be archived with PDS as Common Data Format (CDF). In order to allow the archival CDF files to be described by PDS metadata a number of requirements have been agreed to between the SWIA ITF and the PDS-PPI node. An early version of these requirements are detailed in the document *Archive of MAVEN CDF in PDS4* (T. King and J. Mafi, July 16, 2013). All parties will agree upon the final requirements before sample files are produced. These CDF files will be the same ones used and distributed by the SWIA ITF internally. The contents of the SWIA CDF files are described in the tables below.

Table 13: Contents for swia.calibrated.coarse_svy_3d and swia.calibrated.coarse_arc_3d calibrated data files

Field Name	Data Type	Description
EPOCH	EPOCH	Spacecraft event time for this data record (UTC Epoch time from 01-Jan-0000 00:00:00.000 without leap seconds), one element per ion distribution (NUM_DISTS elements)
TIME_TT2000	TT2000	UTC time from 01-Jan-2000 12:00:00.000 including leap seconds), one element per ion distribution (NUM_DISTS elements)
TIME_MET	DOUBLE	Mission elapsed time for this data record, one element per ion distribution (NUM_DISTS elements)
TIME_UNIX	DOUBLE	Unix time (elapsed seconds since 1970-01-01/00:00 without leap seconds) for this data record, one element per ion distribution (NUM_DISTS elements)
NUM_ACCUM	INTEGER	Number of four second accumulations per distribution, one element per ion distribution (NUM_DISTS elements)
ATTEN_STATE	INTEGER	Attenuator state (1 = open, 2 = closed, 3 = cover shut), one element per ion distribution (NUM_DISTS elements)
GROUPING	INTEGER	Resolution of coarse 3d ion distribution (0 = 48x12x10, 1 = 24 energies, 2 = 16 energies), one element per ion distribution (NUM_DISTS elements)

COUNTS	FLOAT	48x4x16-element array of Coarse product counts, one array per coarse 3d distribution (NUM_DISTSx48x4x16 elements). See notes below, which also apply to this table.
DIFF_EN_FLUXES	FLOAT	48x4x16-element array of differential energy fluxes [eV/(cm^2 s sr eV)] computed from the COUNTS array, the full sensor geometric factor GEOM_FACTOR, the relative phi and theta sensitivities G_THETA and G_PHI (or G_THETA_ATTEN and G_PHI_ATTEN), the accumulation time ACCUM_TIME, and the number of accumulations per distribution NUM_ACCUM, one array per distribution (NUM_DISTSx48x4x16 elements). For convenience, in cases with binned data with 24 or 16 energies (GROUPING = 1 or 2) a full 48-energy data structure is still provided for commonality and ease of use, but the 48 energy steps contain the binned counts divided by the binning factor and duplicated in pairs or triplets. For example, for a case with 16 binned energy steps, the first three steps of the 48 energies contain the number of counts in the first binned energy step divided by three. This re-binning scheme ensures that moments or other sums computed from binned Coarse data still come out right, and use of Coarse data with different binning schemes is transparent to the end user.
GEOM_FACTOR	FLOAT (NOVARY)	Full sensor geometric factor [cm^2 s sr eV/eV]
DE_OVER_E_COARSE	FLOAT (NOVARY)	Energy resolution of coarse distributions
ACCUM_TIME_COARSE	FLOAT (NOVARY)	Accumulation time for each sample, nominally 12*1.7 ms to account for the summation over 2 energy and 6 deflection steps for each element of the P1 product that is used to make the coarse distributions [An additional summation over the 16x4 angle elements of the coarse distribution is required to use this for energy spectra]
ENERGY_COARSE	FLOAT (NOVARY)	48-element array of energies (eV) covered by the distribution
THETA_COARSE	FLOAT (NOVARY)	48x4-element array of theta angles (instrument coordinates) covered by the distribution with the attenuator open
THETA_ATTEN_COARSE	FLOAT (NOVARY)	48x4-element array of theta angles (instrument coordinates) covered by the distribution with the attenuator closed

G_THETA_COARSE	FLOAT (NOVARY)	48x4-element array of relative sensitivity as a function of theta with the attenuator open
G_THETA_ATTEN_COARSE	FLOAT (NOVARY)	48x4-element array of relative sensitivity as a function of theta with the attenuator closed
PHI_COARSE	FLOAT (NOVARY)	16-element array of phi angles (instrument coordinates)
G_PHI_COARSE	FLOAT (NOVARY)	16-element array of relative sensitivity as a function of phi with the attenuator open
G_PHI_ATTEN_COARSE	FLOAT (NOVARY)	16-element array of relative sensitivity as a function of phi with the attenuator closed
NUM_DISTS	INTEGER (NOVARY)	Number of coarse 3d distributions in the file

Table 14: Contents for swia.calibrated.fine_svy_3d and swia.calibrated.fine_arc_3d calibrated data files

Field Name	Data Type	Description
EPOCH	EPOCH	Spacecraft event time for this data record (UTC Epoch time from 01-Jan-0000 00:00:00.000 without leap seconds), one element per ion distribution (NUM_DISTS elements)
TIME_TT2000	TT2000	UTC time from 01-Jan-2000 12:00:00.000 including leap seconds), one element per ion distribution (NUM_DISTS elements)
TIME_MET	DOUBLE	Mission elapsed time for this data record, one element per ion distribution (NUM_DISTS elements)
TIME_UNIX	DOUBLE	Unix time (elapsed seconds since 1970-01-01/00:00 without leap seconds) for this data record, one element per ion distribution (NUM_DISTS elements)
ATTEN_STATE	INTEGER	Attenuator state (1 = open, 2 = closed, 3 = cover shut), one element per ion distribution (NUM_DISTS elements)
GROUPING	INTEGER	Coverage of fine 3d ion distribution (0 = 48x12x10 elements, 1 = 32x8x6 elements), one element per ion distribution (NUM_DISTS elements)
ESTEP_FIRST	INTEGER	Address of starting energy step for P2 distribution, one element per ion distribution (NUM_DISTS elements)
DSTEP_FIRST	INTEGER	Address of starting deflection step for P2 distribution, one element per ion distribution (NUM_DISTS elements)

COUNTS	FLOAT	48x12x10-element array of Fine product counts, one array per fine distribution (NUM_DISTSx48x12x10 elements). See notes below, which also apply to this table.
DIFF_EN_FLUXES	FLOAT	48x12x10-element array of differential energy fluxes [eV/(cm^2 s sr eV)] computed from the COUNTS array, the full sensor geometric factor GEOM_FACTOR, the relative phi and theta sensitivities G_THETA and G_PHI (or G_THETA_ATTEN and G_PHI_ATTEN), the accumulation time ACCUM_TIME, and the number of accumulations per distribution NUM_ACCUM, one array per distribution (NUM_DISTSx48x12x10 elements). For convenience, for the case with GROUPING = 1, a full 48x12x10 array is still provided, but with only the central values containing non-zero counts. This allows Fine data products with different formats to be mixed in a transparent way. In order to determine which elements of the energy and theta tables in the supporting data to use when working with these data, it is necessary to use the starting energy step and deflection step. In other words, the correct energies for each distribution will consist of the 48 elements of the 96 energies ranging from the starting energy step (0 to 48) to the starting energy step + 47. Similarly, the correct theta angles for each distribution will consist of the 12 elements of the 24 theta angles ranging from the starting deflection step (0 to 12) to the starting deflection step + 11.
GEOM_FACTOR	FLOAT (NOVARY)	Full sensor geometric factor [cm^2 s sr eV/eV]
DE_OVER_E_FINE	FLOAT (NOVARY)	Energy resolution of fine distributions
ACCUM_TIME_FINE	FLOAT (NOVARY)	Accumulation time for each sample
ENERGY_FINE	FLOAT (NOVARY)	96-element array of energies (eV) covered by the distribution
THETA_FINE	FLOAT (NOVARY)	96x24-element array of theta angles (instrument coordinates) covered by the distribution with the attenuator open
THETA_ATTEN_FINE	FLOAT (NOVARY)	96x24-element array of theta angles (instrument coordinates) covered by the distribution with the attenuator closed
G_THETA_FINE	FLOAT (NOVARY)	96x24-element array of relative sensitivity as a function of theta with the attenuator open

G_THETA_ATTEN_FINE	FLOAT (NOVARY)	96x24-element array of relative sensitivity as a function of theta with the attenuator closed
PHI_FINE	FLOAT (NOVARY)	10-element array of phi angles (instrument coordinates)
G_PHI_FINE	FLOAT (NOVARY)	10-element array of relative sensitivity as a function of phi with the attenuator open
G_PHI_ATTEN_FINE	FLOAT (NOVARY)	10-element array of relative sensitivity as a function of phi with the attenuator closed
NUM_DISTS	INTEGER (NOVARY)	Number of fine 3d distributions in the file

Table 15: Contents for swia.calibrated.onboard_svy_mom calibrated data files

Field Name	Data Type	Description
EPOCH	EPOCH	Spacecraft event time for this data record (UTC Epoch time from 01-Jan-0000 00:00:00.000 without leap seconds), one element per set of moments (NUM_MOM elements)
TIME_TT2000	TT2000	UTC time from 01-Jan-2000 12:00:00.000 including leap seconds), one element per ion distribution (NUM_MOM elements)
TIME_MET	DOUBLE	Mission elapsed time for this data record, one element per set of moments (NUM_MOM elements)
TIME_UNIX	DOUBLE	Unix time (elapsed seconds since 1970-01-01/00:00 without leap seconds) for this data record, one element per set of moments (NUM_MOM elements)
ATTEN_STATE	INTEGER	Attenuator state (1 = open, 2 = closed, 3 = cover shut), one element per set of moments (NUM_MOM elements)
TELEM_MODE	INTEGER	The telemetry mode (1 = ‘Sheath’, 0 = “Solar Wind”) that defines which distribution the moments are calculated from (Coarse for Sheath, Fine for Solar Wind), one element per set of moments (NUM_MOM elements)
QUALITY_FLAG	FLOAT	Quality flag based on whether the bulk of the distribution is covered by the angular field of view and the energy sweep (exact definition TBD, 0 = bad, 1 = good), one element per set of moments (NUM_MOM elements)
DECOM_FLAG	FLOAT	Quality flag based on whether the attenuator state or telemetry mode is ambiguous (0 = bad, 1 = good), one element per set of moments (NUM_MOM elements)

DENSITY	FLOAT	Onboard density moment (particles per cc), calculated assuming all ions are protons, one element per set of moments (NUM_MOM elements)
PRESSURE	FLOAT	Onboard pressure tensor (eV per cc) calculated assuming all ions are protons, in instrument coordinates, one array per set of moments (NUM_MOM*6 elements)
VELOCITY	FLOAT	Onboard velocity vector (km/s) calculated assuming all ions are protons, in instrument coordinates, one array per set of moments (NUM_MOMx3 elements)
VELOCITY_MSO	FLOAT	Onboard velocity vector (km/s) calculated assuming all ions are protons, in MSO coordinates, one array per set of moments (NUM_MOMx3 elements)
TEMPERATURE	FLOAT	Onboard temperature vector (km/s) calculated assuming all ions are protons, in instrument coordinates, one array per set of moments (NUM_MOMx3 elements)
TEMPERATURE_MSO	FLOAT	Onboard temperature vector (km/s) calculated assuming all ions are protons, in MSO coordinates, one array per set of moments (NUM_MOMx3 elements)
NUM_MOM	INTEGER (NOVARY)	Number of moment samples in the file

Table 16: Contents for swia.calibrated.onboard_svy_spec calibrated data files

Field Name	Data Type	Description
EPOCH	EPOCH	Spacecraft event time for this data record (UTC Epoch time from 01-Jan-0000 00:00:00.000 without leap seconds), one element per energy spectrum (NUM_SPEC elements)
TIME_TT2000	TT2000	UTC time from 01-Jan-2000 12:00:00.000 including leap seconds), one element per ion distribution (NUM_SPEC elements)
TIME_MET	DOUBLE	Mission elapsed time for this data record, one element per energy spectrum (NUM_SPEC elements)
TIME_UNIX	DOUBLE	Unix time (elapsed seconds since 1970-01-01/00:00 without leap seconds) for this data record, one element per energy spectrum (NUM_SPEC elements)

NUM_ACCUM	INTEGER	Number of four second accumulations per energy spectra, one element per spectrum (NUM_SPEC elements)
ATTEN_STATE	INTEGER	Attenuator state (1 = open, 2 = closed, 3 = cover shut), one element per energy spectrum (NUM_SPEC elements)
DECOM_FLAG	FLOAT	Quality flag based on whether the attenuator state or telemetry mode is ambiguous (0 = bad, 1 = good), one element per energy spectrum (NUM_SPEC elements)
SPECTRA_COUNTS	FLOAT	48-element array of counts calculated by summing over all angles in the coarse 3d ion distributions, one array per energy spectrum (NUM_SPECx48 elements)
SPECTRA_DIFF_EN_FLUX	FLOAT	48-element array of differential energy fluxes [eV/(cm^2 s sr eV)] computed from the SPECTRA_COUNTS array, the full sensor geometric factor GEOM_FACTOR, the accumulation time ACCUM_TIME, and the number of accumulations per distribution NUM_ACCUM, one array per energy spectrum, one array per energy spectrum (NUM_SPECx48 elements)
GEOM_FACTOR	FLOAT (NOVARY)	Full sensor geometric factor [cm^2 s sr eV/eV]
DE_OVER_E_SPECTRA	FLOAT (NOVARY)	Energy resolution of energy spectra
ACCUM_TIME_SPECTRA	FLOAT (NOVARY)	Accumulation time for each sample, nominally 12*64*1.7 ms to account for the summation over 2 energy and 6 deflection steps for each element of the P1 product that is used to make the coarse distributions, then summation over the 16x4 angle elements of the coarse distributions to make the spectra
ENERGY_SPECTRA	FLOAT (NOVARY)	48-element array of energies (eV) covered by the distribution
NUM_SPEC	INTEGER (NOVARY)	Number of energy spectra in the file

6.2 Document Product File Formats

Documents are provided in either Adobe Acrobat PDF/A or plain ASCII text format. Other versions of the document (including HTML, Microsoft Word, etc.) may be included as well.

6.3 **PDS Labels**

PDS labels are ASCII text files written, in the eXtensible Markup Language (XML). All product labels are detached from the digital files (if any) containing the data objects they describe (except Product_Bundle). There is one label for every product. Each product, however, may contain one or more data objects. The data objects of a given product may all reside in a single file, or they may be stored in multiple separate files. PDS4 label files must end with the file extension ".xml".

The structure of PDS label files is governed by the XML documents described in Section 6.3.1.

6.3.1 XML Documents

For the MAVEN mission PDS labels will conform to the PDS master schema based upon the 1.1.0.0 version of the PDS Information Model for structure, and the 1.1.0.0 version of the PDS Schematron for content. By use of an XML editor these documents may be used to validate the structure and content of the product labels.

Examples of PDS labels required for the SWIA archive are shown in Appendix C (bundle products), Appendix D (collection products), and Appendix E (basic products).

6.4 **Delivery Package**

Data transfers, whether from data providers to PDS or from PDS to data users or to the deep archive, are accomplished using delivery packages. Delivery packages include the following required elements:

1. The package which consists of a compressed bundle of the products being transferred.
2. A transfer manifest which maps each product's LIDVID to the physical location of the product label in the package after uncompression.
3. A checksum manifest which lists the MD5 checksum of each file included in the package after uncompression.

SWIA archive delivery packages (including the transfer and checksum manifests) for delivery to PDS are produced at the MAVEN SDC.

6.4.1 The Package

The directory structure used in for the delivery package is described in the Appendix in Section F.1. Delivery packages are compressed using either [zip, or tar/gzip] and are transferred electronically using the ssh protocol.

6.4.2 Transfer Manifest

The "transfer manifest" is a file provided with each transfer to, from, or within PDS. The transfer manifest is external to the delivery package. It contains an entry for each label file in the package, and maps the product LIDVID to the file specification name for the associated product's label file. Details of the structure of the transfer manifest are provided in Section F.2.

The transfer manifest is external to the delivery package, and is not an archive product. As a result, it does not require a PDS label.

6.4.3 Checksum Manifest

The checksum manifest contains an MD5 checksum for every file included as part of the delivery package. This includes both the PDS product labels and the files containing the digital objects which they describe. The format used for a checksum manifest is the standard output generated by the md5deep utility. Details of the structure of the checksum manifest are provided in section F.3.

The checksum manifest is external to the delivery package, and is not an archive product. As a result, it does not require a PDS label.

Appendix A Support staff and cognizant persons

Table 17: Archive support staff

SWIA team			
Name	**Address**	**Phone**	**Email**
Jasper S Halekas	Space Sciences Laboratory, 7 Gauss Way, University of California, Berkeley, CA 94720	510-643-4310	jazzman@ssl.berkeley.edu

UCLA			
Name	**Address**	**Phone**	**Email**
Dr. Steven Joy PPI Operations Manager	IGPP, University of California 405 Hilgard Avenue Los Angeles, CA 90095-1567 USA	+001 310 825 3506	sjoy@igpp.ucla.edu
Mr. Joseph Mafi PPI Data Engineer	IGPP, University of California 405 Hilgard Avenue Los Angeles, CA 90095-1567 USA	+001 310 206 6073	jmafi@igpp.ucla.edu

Appendix B Naming conventions for MAVEN science data files

This section describes the naming convention used for science data files for the MAVEN mission.

Raw (MAVEN Level 0):

mvn_<inst>_<grouping>_l0_< yyyy><mm><dd>_v<xx>.dat

Level 1, 2, 3+:

mvn_<inst>_<level>_<descriptor>_<yyyy><mm><dd>T<hh><mm><ss>_v<xx>_r<yy>.<ext>

Code	Description
<inst>	3-letter instrument ID
<grouping>	Three-letter code: options are all, svy, arc for all data, survey data, archive data. Primarily for P&F to divide their survey & archive data at Level 0.
<yyyy>	4-digit year
<mm>	2-digit month, e.g. 01, 12
<dd>	2-digit day of month, e.g. 02, 31
<hh>	2-digit hour, separated from the date by T. OPTIONAL.
<mm>	2-digit minute. OPTIONAL.
<ss>	2-digit second. OPTIONAL.
v<xx>	2-digit software version: which version of the software was used to create this data product?
r<yy>	2-digit data version: is this a new version of a previous file, though the same software version was used for both? (Likely to be used in the case of retransmits to fill in data gaps)
<descriptor>	A description of the data. Defined by the creator of the dataset. There are no underscores in the value.
.<ext>	File type extension: .fits, .txt, .cdf, .png
<level>	A code indicating the MAVEN processing level of the data (valid values: l1, l2, l3)

Instrument name	<instrument>
IUVS	iuv
NGIMS	ngi
LPW	lpw
MAG	mag
SEP	sep
SWIA	swi
SWEA	swe
STATIC	sta
P&F package	pfp

Appendix C Sample Bundle Product Label

This section provides a sample bundle product label.

Appendix D Sample Collection Product Label

This section provides a sample collection product label.

Appendix E Sample Data Product Labels

This section provides sample product labels for the various data types described in this document.

Appendix F PDS Delivery Package Manifest File Record Structures

The delivery package includes two manifest files: a transfer manifest, and MD5 checksum manifest. When delivered as part of a data delivery, these two files are not PDS archive products, and do not require PDS labels files. The format of each of these files is described below.

F.1 Transfer Package Directory Structure

[Insert a description of the directory structure contained in the delivery package.]

F.2 Transfer Manifest Record Structure

The transfer manifest is defined as a two field fixed-width table where each row of the table describes one of the products in the package. The first field defines the LIDVID of each product in the package. The second field defines the file specification name of the corresponding product label in the package. The file specification name defines the name and location of the product relative to the location of the bundle product.

F.3 Checksum Manifest Record Structure

The checksum manifest consists of two fields: a 32 character hexadecimal (using lowercase letters) MD5, and a file specification from the root directory of the unzipped delivery package to every file included in the package. The file specification uses forward slashes (“/”) as path delimiters. The two fields are separated by two spaces. Manifest records may be of variable length. This is the standard output format for a variety of MD5 checksum tools (*e.g.* md5deep, etc.).