

***New Horizons* SOC to Instrument Pipeline ICD**

March 2023

SwRI® Project 05310

Document No. 05310-SOCINST-01

Contract NASW-02008



SOUTHWEST RESEARCH INSTITUTE®
Space Science and Engineering Division
6220 Culebra Road, San Antonio, Texas 78228-0510
(210) 522-2600 • FAX (210) 647-4325

New Horizons SOC to Instrument Pipeline ICD

SwRI Project 05310

Document No. 05310-SOCINST-01

Contract NASW-02008

Prepared by: Joe Peterson	08 November 2013
Revised by: Brian Carcich	August, 2014
Revised by: Tiffany Finley	March, 2016
Revised by: Tiffany Finley	October, 2016
Revised by: Brian Carcich	December, 2016
Revised by: Tiffany Finley, Brian Carcich, PEPSSI team	April, 2017
Revised by: Tiffany Finley, PEPSSI team	September, 2017
Revised by: Jillian Redfern, SDC Team	June 2019
Revised by: Jillian Redfern, LEISA Team	March 2020
Revised by: Joel Parker, ALICE Team, SWAP Team, REX Team	March 2023

Contributors:

ALICE specifics prepared by: Maarten Versteeg, Joel Parker, Andrew Steffl
LEISA specifics prepared by: George McCabe, Allen Lunsford
LORRI specifics prepared by: Hal Weaver, Howard Taylor
MVIC specifics prepared by: Cathy Olkin
PEPSSI specifics prepared by: Stefano Livi, Matthew Hill, Lawrence Brown,
Peter Kollman
REX specifics prepared by: Ivan Linscott Brian Carcich
SDC specifics prepared by: David James
SWAP specifics prepared by: Heather Elliott



SOUTHWEST RESEARCH INSTITUTE®

Space Science and Engineering Division
6220 Culebra Road, San Antonio, Texas 78228-0510
(210) 522-2600 • FAX (210) 647-4325

General Approval Signatures:

Approved by: _____ Date: _____
Hal Weaver, JHU/APL, Project Scientist

Approved by: Leslie Young via email Date: 1 December 2005
SwRI, Deputy Project Scientist

Approved by: John Andrews via email Date: 4 December 2005
SwRI, SOC Project Manager

**Instrument-specific Signatures:
(Science Theme Teams members)**

Approved by: _____ Date: _____
Jeff Moore, NASA, GGI Science Theme Team Lead
(applies to MVIC, LORRI)

Approved by: /s/ Dale Cruikshank via email Date: 6 March 2006
NASA, COMP Science Theme Team Lead
(applies to LEISA)

Approved by: Randy Gladstone via email Date: 6 December 2005
SwRI, ATM Science Theme Team Lead
(applies to ALICE, REX)

Approved by: /s/ Fran Bagenal via email Date: 23 January 2006
CU, P&P Science Theme Team Lead
(applies to SWAP, PEPSSI, SDC)



New Horizons SOC to Instrument Pipeline ICD

(Instrument PIs & Scientists)

Approved by: Alan Stern via email
SwRI, ALICE Instrument PI 5 December 2005

Approved by: Dave Slater via email
SwRI, ALICE Instrument Scientist 2 December 2005

Approved by: Alan Stern via email
SwRI, ALICE Instrument PI 5 December 2005

Approved by: Dennis Reuter via email
NASA, LEISA Instrument Scientist 5 December 2005

Approved by: _____ Date: _____
Andy Cheng, JHU/APL, LORRI Instrument PI

Approved by: _____ Date: _____
Hal Weaver, JHU/APL, LORRI Instrument Scientist

Approved by: Alan Stern via email
SwRI, ALICE Instrument PI 5 December 2005

Approved by: Dennis Reuter via email
NASA, LEISA Instrument Scientist 5 December 2005



Approved by: _____ Date: _____
Ralph McNutt, JHU/APL, PEPSSI Instrument PI

Approved by: _____ Date: _____
Stefano Livi, JHU/APL, PEPSSI Instrument Scientist

Approved by: /s/ Len Tyler via email Date: 8 March 2006
Stanford, REX Instrument PI

Approved by: /s/ Ivan Linscott via email Date: 8 March 2006
Stanford, REX Instrument Scientist

Approved by: _____ Date: _____
Mihaly Horanyi, CU, SDC Instrument PI

Approved by: _____ Date: _____
Mihaly Horanyi, CU, SDC Instrument Scientist

Approved by: Dave McComas via email Date: 12 December 2005
SwRI, SWAP Instrument PI

Approved by: Heather Elliot via email Date: 8 March 2006
SwRI, SWAP Instrument Scientist

Release to Document Control:

Approved by: _____ Date: _____
TBD, SwRI Project CM



SOUTHWEST RESEARCH INSTITUTE®
Space Science and Engineering Division
6220 Culebra Road, San Antonio, Texas 78228-0510
(210) 522-2600 • FAX (210) 647-4325

TABLE OF CONTENTS

	Page
1. Scope	1
2. Signatures, Authorship, and Responsibility.....	1
3. Applicable Documents	1
3.1 New Horizons Documents.....	1
4. Introduction	2
4.1 The New Horizons SOC.....	2
4.2 The SOC Pipeline and Definition of Data Levels	2
5. Interface Description	4
6. Requirements.....	5
6.1 Level 2 (output) Files	5
6.2 Pointing Information	6
6.3 The Code	6
6.4 Calibration Data.....	7
6.5 Documentation	7
6.6 Error Conditions and Integration Troubleshooting.....	7
6.7 PDS Archiving.....	7
6.8 Configuration Management.....	8
6.9 Pipeline Updates	8
6.10 Acceptance Review	8
6.11 Longevity.....	8
7. ALICE Instrument description	9
7.1 Overview	9
7.2 Level 1 Data Specifics.....	10
7.3 Level 2 Data Specifics.....	18
8. LEISA Instrument description.....	23
8.1 Overview	23
8.2 Level 1 Data Specifics.....	24
8.3 Level 2 Data Specifics.....	27



SOUTHWEST RESEARCH INSTITUTE®

Space Science and Engineering Division
6220 Culebra Road, San Antonio, Texas 78228-0510
(210) 522-2600 • FAX (210) 647-4325

New Horizons SOC to Instrument Pipeline ICD

9.	LORRI Instrument description	34
9.1	Overview	34
9.2	Level 1 Data Specifics	35
9.3	Level 2 Data Specifics	37
9.4	References	56
10.	MVIC Instrument description	57
10.1	Overview	57
10.2	Level 1 Data Specifics	61
10.3	Level 2 Data Specifics	65
11.	PEPSSI Instrument description	74
11.1	Overview	74
11.2	Introduction to PEPSSI Data	79
11.3	Level 1 Data Specifics	83
11.4	Level 2 Data Specifics	84
12.	REX Instrument description	105
12.1	Overview	105
12.2	Level 1 Data Specifics	107
12.3	Level 2 Data Specifics	115
13.	SDC Instrument description	121
13.1	Overview	121
13.2	Level 1 Data Specifics	121
13.3	Calibration	125
13.4	Level 2 Data Specifics	129
14.	SWAP Instrument description	133
14.1	Overview	133
14.2	Electronics and Flight Software	135
14.3	SWAP Data Types	135
14.4	Level 1 Data Specifics	138
14.5	Level 2 Data Specifics	143
14.6	Operations	150



New Horizons SOC to Instrument Pipeline ICD

14.7	Observation Examples	151
14.8	Updates to the Level 2 Calibrated SWAP Data Pipeline	154
14.9	Gain and Deflector Angle Sweep Tests	156
14.10	SWAP Science Goals	156
14.11	SWAP Publications Examples	158
14.12	Packet Description	159
15.	Acronyms and abbreviations	166



SOUTHWEST RESEARCH INSTITUTE®

Space Science and Engineering Division
6220 Culebra Road, San Antonio, Texas 78228-0510
(210) 522-2600 • FAX (210) 647-4325

REVISION HISTORY

- Draft 1: October 2005
- Draft 2: February/March 2006
- Initial Issue: February 2007. Updated Sections 10 (MVIC), 11 (PEPSSI), 13 (SDC), and 14 (SWAP).
- Draft 3: August, 2014. Updated formatting of sections for PDS deliveries.
- Draft 4: October 2016. Updated formatting of document for PDF/A-1a.
- Draft 5: April 2017. Updated instrument sections for PDS deliveries. Cleaned up graphics. Redid SWAP equations. Updated PEPSSI, REX, MVIC, other instrument calibrations.
- Draft 6: September 2017. Updated PEPSSI and REX calibration specifics. Added information on MVIC Housekeeping parameters.
- Draft 7: June 2019: Updated SDC information per reviewers comments.
- Draft 8: March 2020: Updated description of LEISA calibrated extensions 5 and 6 to match reality. Updated formatting issues.
- Draft 9: March 2023: Section 4 - added more details about data processing levels, and made terminology (raw/uncalibrated/calibrated, Level 1 and 2) more consistent among all sections. Updated Alice description for HCCR data, REX text and table for side-dependent constants, additional descriptions of PEPSSI calibration and PHA data, and SWAP evolution of the HISTOGRAM Level 2 Data Product. A few typo and format fixes.

**SOUTHWEST RESEARCH INSTITUTE®**

Space Science and Engineering Division
6220 Culebra Road, San Antonio, Texas 78228-0510
(210) 522-2600 • FAX (210) 647-4325

1. SCOPE

This document specifies the interfaces between the *New Horizons* Science Operations Center (SOC) and the instrument pipeline, which process data from *New Horizons* SOC Data Level 1 to Level 2. The purpose is to define the various aspects of the interfaces in sufficient detail to establish a clear understanding between the SOC and the instrument team to allow for a parallel pipeline development.

2. SIGNATURES, AUTHORSHIP, AND RESPONSIBILITY

Representatives from each instrument team and the SOC contributed to this document. Specifically, sections 1 through 6 were written by the document author, whereas the bulk of the instrument-specific sections (7 through 14) were written by representatives of the instrument described. Each instrument team has a person or persons responsible for their section. If changes are made to that section, only the person(s) responsible need to sign the new revision. If, however, changes are made to sections 1 through 6, all parties need to sign. The title(s) of the person(s) responsible for each instrument section are given in the signature section above.

The *New Horizons* project maintains a copy of the signed version of this document under Configuration Management (CM); from 2008 through 2017, a fork of this document has been edited many times as a support document for PDS data set deliveries, and these edits have now been incorporated into the release as of September 2017.

3. APPLICABLE DOCUMENTS

3.1 New Horizons Documents

- *New Horizons* SOC Data Pipeline Guide (SwRI Doc. No. 05310-SOCPL-G-01)
- *New Horizons* SOC Level 1 Data Formats (SwRI Doc. No. 05310-SOCL1DATA-01)
- *New Horizons* SOC Pipeline User Manual (SwRI Doc. No. 05310-SOCPLUM-01)
- *New Horizons* SOC Technical Reference Document (SwRI Doc. No. 05310-SOCTRD-01)
- *New Horizons* Data Management and Archiving Plan (SwRI Doc. No. 05310-DMAP-01)
- *New Horizons* Longevity Plan (APL Doc. No. 7399-9009)
- Definition of the Flexible Image Transport System (FITS)
 - The FITS SUPPORT OFFICE is based out of GSFC.

4. INTRODUCTION

4.1 The New Horizons SOC

The *New Horizons* SOC is part of the ground system that processes data returned from the *New Horizons* planetary spacecraft. Data downlinked from the spacecraft in raw packetized form is retrieved from the DSN by the Mission Operations Center (MOC), located at the Johns Hopkins Applied Physics Laboratory (APL), and is processed through *New Horizons* Level 0 (CODMAC Level 1), which then is retrieved by the SOC along with navigation and related ancillary data. Those Level 0 data - low-speed housekeeping and high-speed science data – are in the form of Supplemented Telemetry Packets (STPs), which consist of standard Consultative Committee for Space Data Systems (CCSDS) telemetry packets with prepended blocks of supplementary data. Navigation data, including spacecraft and planetary ephemerides, and spacecraft pointing data, are relayed to the SOC in SPICE format and are used in the SOC data processing pipeline. The *New Horizons* SOC primary location is the Tombaugh Science Operations Center (TSOC) located at SwRI-Boulder; the Contingency Science Operations Center (CSOC) is at APL.

The SOC generates the higher level (more refined) data products used by the instrument teams and science teams. In addition, the SOC performs archiving of data to the Planetary Data System (PDS).

4.2 The SOC Pipeline and Definition of Data Levels

The science data processing component of the SOC is called the SOC pipeline. The SOC pipeline is divided into two main parts: the Level 1 pipeline segment and the Level 2 (instrument) pipeline segment. Pipeline processing is carried out sequentially, such that results of the Level 1 pipeline are provided as inputs to the instrument Level 2 pipeline segment. The SOC pipeline creates PDS standard data products in Flexible Image Transport System (FITS) format, which is described in the document referenced herein (Definition of the Flexible Image Transport System).

The Level 1 pipeline converts the Level 0 data to Level 1 data (CODMAC 2, referred to in this ICD as “engineering”, “eng”, or “uncalibrated” data). The Level 1 pipeline code is written by the SOC. More information about the formats of Level 1 data can be found in the “*New Horizons* SOC Level 1 Data Formats” document (SwRI Doc. No. 05310-SOCL1DATA-01).

The Level 2 pipeline (also called the “instrument pipeline” or the “calibration pipeline”) transforms *New Horizons* SOC Level 1 data into Level 2 science data (CODMAC 3 and 4, also referred to in this ICD as “science”, “sci”, or “calibrated” data). The Level 2 pipeline code for each instrument is written by the appropriate instrument team.

Higher-level “derived” data products (CODMAC Level 5) are produced by the Science Team and delivered to the SOC.

The SOC forwards the Level 1, Level 2, and derived data to the PDS archiving process. *Throughout this ICD “Level 1” and “Level 2” will always refer to the New Horizons SOC processing level unless explicitly referred to as “CODMAC Level”.* “Raw” usually means either data as they are stored on or telemetered from an instrument, or the data between downlink from the spacecraft to input to the SOC Level 1 pipeline.

Upon getting packets (housekeeping and science data) from the MOC, the data are decommutated in the SOC and written to an SQL database. Housekeeping from the database and science data are associated by MET time and other methods, such as by using meta data inserted in the high-speed telemetry. Data for an observation are combined to create the Level 1 uncalibrated data file. A PDS detached header file (a separate file containing the PDS header with a pointer to the data file itself) is also created for each file. The header of the Level 1 file contains most of the necessary information about the particular observation needed by the instrument pipeline (an exception is detailed pointing, which will be calculated during calibration). The instrument pipeline segment creates the Level 2 calibrated data file from the contents of the Level 1 file and calibration data provided by the instrument team. Level 1 and Level 2 science data files are stored in FITS format. Lastly, the SOC archives pipeline data products to the PDS.

SOC pipeline processing is automated under the control of the Master Data Manager (MDM), which is a collection of scripts that control the flow of the pipeline. While manual execution of the program is permitted, normal operation of the SOC pipeline is not directed by manual requests or user inputs. Pipeline segments are called by the MDM when data from the MOC or from a previous pipeline step is available.

The hardware platform used for the SOC as implemented for launch and early mission began with an Intel Pentium 4 processor running at 3.2GHz with 4GB RAM and a 146GB SCSI hard disk. In the case of the primary SOC (TSOC), located in Boulder, Colorado, two of these machines were used (one for pipeline processing and the other for data storage). The secondary (backup) SOC – the CSOC – used similar hardware, but only one server handling both of these roles.

This hardware was refreshed twice during the mission, so far, and one more refresh is planned on the way to the KBO. The current hardware (used at both the SOC and the CSOC) are rack-mounted servers with Intel(R) Xeon(R) CPU E5-1620, 3.60GHz processors, integrated 19TB RAID arrays, and 32GB RAM. The operating system used in all cases is Linux.

5. INTERFACE DESCRIPTION

The SOC pipeline code will call the Level 2 pipeline code by executing a separate process.

SOCINST-1 The name of the executable or script shall be:

```
[instrument]_level2_pipeline
```

where “[instrument]” is the full instrument name (i.e. alice, leisa, lorri, mvic, pepssi, rex, sdc, or swap) in lower case.

The parameters (all are character strings) passed to the Level 2 code will follow the executable name and will be in the following order (note that “full path,” when stated below, means a file specification containing the filename and all directories in the file’s path):

- **in_file - Location of input (Level 1) file (in_file)**
The SOC will provide the full path of the Level 1 input file to the Level 2 pipeline code.
- **in_pds_header - Location of input (Level 1) detached PDS header**
The SOC will provide the full path of the Level 1 PDS header to the Level 2 pipeline code.
- **calibration_dir - Location of calibration data and temporary file storage**
Data provided by the instrument team that is needed for calibration shall be supplied by the instrument team. The SOC will provide the root directory containing these files (and potentially, subdirectories) to the Level 2 pipeline code so it references the correct location. The structure of the directories under this directory is up to the instrument team.
- **temp_dir - Location for temporary storage used by the code**
This is a place where the instrument pipeline code may write files for temporary use. The contents of this directory will be erased upon completion of the instrument pipeline.
- **out_status - Location of status file**
The Level 2 pipeline, upon completion, may write a short machine readable status file used to communicate the results of the processing to the main SOC pipeline. The location (full path) of this file will be supplied by the SOC.
- **out_file - Location of output (Level 2) file**
This is the file (full path) to which the output will be written. The SOC will provide this to the Level 2 pipeline code.
- **out_pds_header - Location of output (Level 2) detached PDS header**
This is the file (full path) to which the Level 2 PDS header will be written. The SOC will provide this to the Level 2 pipeline code.

6. REQUIREMENTS

This section describes the various requirements that the instrument team shall follow with regard to the Level 2 pipelines.

6.1 Level 2 (output) Files

SOCINST-2 The Level 2 data files shall be FITS files.

SOCINST-3 There shall be exactly one Level 2 file produced for each Level 1 file (in any given run of the instrument pipeline). The headers will be mostly the same, except for optional additional keywords inserted in the Level 2 files (this could include, for example, refined pointing). In other words, typically, the Level 2 header keywords will be a superset of the Level 1 header keywords.

The filename of the Level 2 file will be supplied by the SOC, and it will be similar to the Level 1 filename.

SOCINST-4 Level 2 filenames shall be a maximum length of 27.3 ASCII characters, as PDS requires a “27.3” ASCII character limit on the filenames.

SOCINST-5 The format of the Level 2 filenames shall be as follows:

- Level 2: `[instrument]_[MET]_[ApID]_sci.fit`

where “[instrument]” is the first three letters of the instrument name (i.e. ali, lei, lor, mvi, pep, rex, sdc, or swa) in lower case, “[MET]” is the 10-digit MET time, either from the instrument itself (low-speed data) or from the instrument interface card (high-speed data), and “[ApID]” is the hexadecimal value of the ApID (Application ID of the main packet used in the data product) for this observation.

- (No longer used)

Level 2: `[instrument]_[MET]_[ApID]_sci_[version].fit`

Note: This naming convention with “version” in the filename is no longer used as of March 2016, and appears in the archive only in the first two datasets (Post-Launch Checkout, Jupiter flyby).

The instrument/mode strings above will be derived from the ApID of the data, and these filenames will be supplied to the instrument pipeline (see the interface description above).

Whereas the Level 1 files will be in the same “units” as the data coming from the spacecraft/instrument (i.e. same binary representation - this is partly to avoid any round-off or conversion loss), the Level 2 files shall express values in physical units useful for scientific interpretation.

- Level 1: `[instrument]_[MET]_[ApID]_eng.fit`

As a note, prior versions again had a format that included a version number, as shown below, however the version was removed as it caused some confusion.

- (No longer used)

Level 1: `[instrument]_[MET]_[ApID]_eng_[version].fit`

Note: This naming convention with “version” in the filename is no longer used as of March 2016, and appears in the archive only in the first two datasets (Post-Launch Checkout, Jupiter flyby).

Whereas the Level 1 files only contain the header and data itself, the Level 2 files will contain (when appropriate) two additional "panes" (FITS extensions):

- Error (specifies error bars on the numbers – defined by the instrument team)
- Quality (Indicates the quality of the data – defined by the instrument team)

6.2 Pointing Information

The pointing information included in the Level 1 files will be mostly non-instrument specific (except for bore-sight vector where applicable). It also may not cover the time granularity needed for calibration in the Level 2 pipelines (see the “*New Horizons* SOC Level 1 Data Formats” document 05310-SOCL1DATA-01 for specifics). Therefore it is expected that the Level 2 pipelines may have to make use of SPICE. It is therefore the responsibility of the Level 2 pipelines to provide this functionality. SPICE kernels will be available from the SOC.

6.3 The Code

The SOC defines the interface the code uses to access the required data. This interface includes the directory structure on the disk where the Level 1 data file can be found as well as the path (specific to each instrument) where instrument-team-supplied calibration files and other data will be stored and accessed. Also, the filename of the output file is supplied.

SOCINST-6 The pipeline code shall be written in a language that meets the SOC's longevity requirements (see Section 6.11). More information on this can be found in the “*New Horizons* SOC Pipeline Guide” (SwRI Doc. No. 05310-SOCPL-G-01). The languages allowed are as follows:

- C
- C++
- Fortran
- IDL
- Python
- Java
- Perl

SOCINST-7 Each instrument team's Level 2 code shall complete processing within 1 hour, to ensure the pipeline can complete processing within 1 day.

In other words, if the time to compute Level 2 data is so extreme that it jeopardizes the completion of each daily run of the pipeline (so ample idle time is not available between runs), the situation will need to be re-evaluated. It is expected that a daily run of the entire pipeline be complete within a few hours of its start. This schedule gives most of the day for users to access new data before the next run is initiated.

The maximum time allowed for execution of an instrument's Level 2 pipeline is recommended to be less than **5 minutes** (for each input file processed). The predicted actual maximum time is negotiated and specified in the instrument-specific sections.

SOCINST-8 Any needed 3rd party libraries also shall meet the longevity requirements. Specifically, source code should be available and must be provided with the code unless already available to the SOC.

6.4 Calibration Data

The code will most likely need calibration data in addition to the Level 1 data files themselves. These data can be anything required. The SOC will provide a directory where these files will be placed on the SOC pipeline system, and the instrument pipeline code will be able to access them there.

The combination of the Level 1 file (and detached PDS label) and the data provided must be sufficient to produce each Level 2 file. If housekeeping information (instrument or spacecraft) is needed, these must be already in the header (or tables) of the Level 1 file. If continuously varying values (e.g. temperature over many seconds, etc.) are needed, a FITS table will be written into the Level 1 file with this information.

6.5 Documentation

All code and data files will be accompanied by thorough documentation. The code itself should have clear and appropriate comments throughout. Error conditions will be documented in the code as well (see Section 6.6 for more on this topic).

Documentation and code will be written assuming that it will be read by someone years from now who is unfamiliar with the system. Understanding of the documentation should not rely on special scientific knowledge or specific domain knowledge.

6.6 Error Conditions and Integration Troubleshooting

SOCINST-9 If there are any reasons the code might abort processing, these shall be defined, and the resulting action should be specified. Also, if such an abort happens, the reason should be noted in the status file written ("out_status" file described above).

SOCINST-10 A contact person shall be specified who will be responsible for helping the SOC operators when unexpected errors occur. This person should be able to help with debugging and should also be available to respond and help in two days or less for consultation during the pipeline integration process.

6.7 PDS Archiving

SOCINST-11 The format of the Level 1 and Level 2 (output) files shall meet PDS v3.8 requirements, as they will be archived to the PDS. This format includes size requirements set forth in the *New Horizons* Data Management and Archiving Plan (05310-DMAP-01).

SOCINST-12 PDS detached labels for the Level 2 files shall be generated by either the instrument pipeline code or by the SOC using a translation table (from FITS to PDS keywords). Which method is appropriate will be determined on an instrument by instrument basis. If

generated by the SOC, “in_pds_header” and “out_pds_header” can be ignored in the instrument pipeline code.

6.8 Configuration Management

All code, documentation, and calibration files will be put under configuration management at the SOC. SOCINST-13 The necessary keywords shall be inserted into the Level 2 headers by the Level 2 pipeline code to specify the version of the code and data used to produce the Level 2 files. This method ensures that data are traceable back to the correct code version.

SOCINST-14 In addition, the Level 2 pipeline code shall insert, using header keywords, the calibration files used and the versions thereof (if applicable).

6.9 Pipeline Updates

Updates to the instrument pipeline (including code, documentation, and calibration data) are to be delivered to SOC personnel for integration; all such updates will require appropriate documentation and will fall under SOC CM. The code will be checked in to the SOC configuration management after regression tests are run. Any special instructions or changes should be communicated to the SOC personnel, and a file containing release notes (called “RELEASE_NOTES”) should accompany the update. The SOC personnel will notify the instrument team when the new update is in place and active.

6.10 Acceptance Review

SOCINST-15 The instrument pipeline (including code, documentation, and calibration data) shall be subject to an acceptance review.

6.11 Longevity

SOCINST-16 The code and all third party libraries and data files used shall meet the longevity requirements as specified in the *New Horizons* Longevity Plan (#7399-9009). Also, development, documentation, and testing of the pipeline shall adhere to these requirements.

7. ALICE INSTRUMENT DESCRIPTION

7.1 Overview

ALICE (also PERSI-ALICE or P-ALICE) is a UV spectrograph that is sensitive to Ultraviolet (UV) light in the range of 520-1870 Å. The instrument consists of a telescope section with an off-axis parabolic mirror, and a spectrograph section that includes a diffraction grating and a microchannel plate (MCP) detector. The MCP detector is an electro-optical device sensitive to extreme and far ultraviolet light and energetic particles. It consists of a photo cathode-coated MCP surface that converts UV photons to electrons, a Z-stack configuration of three MCPs to amplify the signal, and a two-dimensional double delay-line (DDL) readout anode. The first (x) dimension provides the spectral location of the detected photon and the second (y) dimension provides one-dimensional spatial information. The DDL detector system outputs to the command-and-data-handling (C&DH) electronics the pixel location for each detected photon event. The pixel location (X, Y) is given as a pair of coordinates, spectral (X axis) and spatial (Y axis). The events are processed by the C&DH electronics. The C&DH is also the controller of ALICE; it receives commands from the spacecraft, acquires data from the MCP detector system, and returns telemetry to the spacecraft. Science data telemetry generation is performed by the detector hardware but the C&DH also controls this function. ALICE has three acquisition modes (see Section 7.2.1 below) in which the spectral/spatial data from the detector are processed by the C&DH subsystem.

All following descriptions assume a nominal operating instrument. The instrument hardware also provides a basic, hardware-controlled, default pixel list acquisition mode, which is activated when the instrument control hardware detects multiple successive software failures. This mode is called the 'State Machine Mode' (SMM); for a description of this mode, the reader is referred to the instrument C&DH hardware description. Once this mode is activated any software control of the acquisition operation is disabled, until the instrument C&DH is reset either by a power cycle reset or by a hardware S/C reset.

Data recorded on *New Horizons* are sent to the ground via the Deep Space Network. From there the data are sent to the Mission Operations Center (MOC) at the Applied Physics Laboratory (APL). The Science Operations Center (SOC) at the Southwest Research Institute (SwRI) in Boulder retrieves new data from the MOC daily. The data from the MOC are in a packetized form. Software pipelines at the SOC convert the data from these raw packets into FITS (Flexible Image Transport System) format files with scientifically useful and calibrated data. The initial processing sorts the packets into FITS files (as images and data tables) with useful header keywords. These keywords include the mode of the observation, exposure and timing information, and basic pointing information of the instrument boresight. The initial processing also adds relevant housekeeping telemetry (temperatures, voltages, etc.) in a binary table as an extension to the FITS file. The calibration pipeline then performs basic scientific calibration on these data.

7.2 Level 1 Data Specifics

The term “RAW” (in uppercase) refers to ALICE packetized raw data, to distinguish from ALICE compressed raw data.

7.2.1 Level 1 Data Format

ALICE High Speed data frame formats:

Science data frames consist of “raw” science data of 32768 16-bit words in size, consisting of 1 word (16 bits) of frame ID header information and a 32767-word data block. Science data frames are generated by the acquisition hardware and sent to the spacecraft via the dedicated LVDS interface. Data are transferred at a rate of 1 Mbit/sec, resulting in a frame transfer time of 557 milliseconds. The spacecraft tags the received data with a receive time (referred to as ‘collect time’) and stores the data in the Solid State Recorder (SSR). Note that these science frames are not identical to the CCSDS (Consultative Committee for Space Data Systems) telemetry packets that are used to transfer the science data to the ground. The generation of telemetry (TM) packets from the frames stored in the SSR is handled by the spacecraft and multiple CCSDS TM packets are used to transfer a single science frame.

To identify the Science frames, a single 16-bit word header is inserted into each frame. This header is generated by the acquisition hardware and includes the information listed in Table 7-1. The complete header word of the most recently generated science frame is included in the housekeeping TM packet to allow for correlation between the two data streams (low and high speed) after the data are received by the ground system.

Table 7-1: ALICE Science Packet Telemetry Header

Field	Size in bits	Description
Packet contents (msb)	1	0 - pixel list or count rate 1 - histogram
Memory	1	0 - ping (side A), 1 - pong (side B)
Last block	1	0 - intermediate block 1 - last block (acquisition cycle terminated)
HW acquisition	1	HW controlled acquisition (hardware “limp along”, a.k.a. “State Machine Mode”), overrides software control after 8 consecutive watchdog timeouts (see C&DH hardware documentation).
Block number	12	Least significant 12 bits of the block number

The contents of the remainder of the data frames (the 32767-word data block) depend on the acquisition mode:

- **Pixel List** Each word in the data block from a pixel list exposure describes a detector event, which can be either a photon event or a time hack. Photon events are indicated by

bit 15 having a value of zero and time hacks are indicated by bit 15 having a value of one. For a photon event, the remaining 15 data bits encode the location of the detected event consisting of a 10-bit encoded spectral location (X) and a 5-bit encoded spatial location (Y). The time hack is used to provide temporal information about the photon events. The acquisition hardware will generate and insert time hacks in the frame on a periodic basis; the frequency of the time hacks is configurable (by command) for each acquisition in a range of 4 – 512 ms. For a time hack, the remaining 15 bits contain an incrementing counter that counts the number of 4 ms periods. This value allows for verification and correction of timing in case of lost frames or packets.

- **Histogram** Each word in the data block from a histogram exposure is a 16-bit “counter” giving the number of photon events detected at each specific X, Y location on the detector. The format of the detector is 1024x32 pixels, which are the spectral and spatial dimensions respectively, i.e., there are 1024 spectral elements along the X -axis and 32 spatial elements along the Y -axis, giving a total of 32768 values (however, the first word always contains the header word). The counters are stored row-wise, corresponding to the spectral dimension indexing most quickly. These counters saturate at a maximum value of 65535 to indicate a completely filled counting bin, meaning that the counters do not wrap around. In addition some special data words (header cross-identification and pulse height distribution) are overlying the lower left-hand corner of the actual array in a block of 4 (spatial) x 32 (spectral) words. This usage doesn’t affect the science data contents because detector events do not occur in this region. In this over-written block, the first row contains the header cross-identification word, the second and third rows contain the 64 words of pulse height information, and the fourth row is filled with zeros. A “pulse height” is the amplitude of a photon event, and this pulse height distribution (PHD) shows the number of events in a 64-bin distribution with 6-bit resolution; the value in each PHD bin gives the number of events that occurred with the particular amplitude associated with that bin. These PHD counters also saturate at a value of 65535. So a single photon event is counted both in the spectral/spatial array and in the pulse height list. The PHD is used as a diagnostic for the health and behavior of the detector.
- **High Cadence Count Rate (HCCR)** Each byte in the data block from a HCCR exposure describes the total number of photon events detected by ALICE in a given (commandable) time sample period. HCCR acts as a UV photometer, and does not preserve wavelength or spatial information, so it provides a low downlink data volume collection mode. This total detector analog countrate can be captured at a very fast cadence, with time samples as short as 20 ms. The FSW update (loaded on the spacecraft in August 2021) to provide the HCCR allows it to be executed via the nominal ALICE pixellist science acquisition commanding, by extending the range of the existing mode parameter to select the multiplier of the base sample period, providing sample times from 0.02 s to 4.96 s. Since the number of counts in a sample period are given in single 8-bit bytes, the maximum number of counts in a sample is 255 and any higher incident count rate is reported as 255; so the range of measureable incident count rates ranges from 51 Hz (at 4.96 s sample period) to 12,750 Hz (at 0.02 s sample period). A standard P-Alice pixellist science frame, on which the HCCR data are based, is 65536 bytes. The HCCR data has an

extended frame header of 12 bytes, which provides metadata that allows the HCCR data to be identified uniquely from a regular pixel list acquisition (two hack entry values of 0x789A and 0x7ABC that can't occur in a regular pixel list frame), the sample period multiplier (in units of 20 ms), the hack counter (4 ms resolution, 18.6 hr rollover), and the number of samples in the HCCR packet. The rest of the HCCR science frame holds the HCCR data, up to 65524 values, padded with zeros. A HCCR uses the ALICE_START_COUNTRATE command that is based on (and has the same command OpCode as) the ALICE_START_PIXELLIST command, and the command's existing MODE_SELECT parameter is extended to allow values of 8-255, which identify the HCCR acquisition and indicates the sample period multiplier (1-248). Because the HCCR mode is enabled via a pixel list acquisition, enabled only with a special parameter value, HCCR data appear to the spacecraft as a regular pixel list frame, and doesn't require any special handling by the spacecraft systems that is any different of that for other ALICE data.

For test purposes the instrument can fill the memory with known deterministic patterns so the interfaces to the spacecraft and ground can be verified. The instrument software allows for the generation of 5 different test patterns.

7.2.1.1 Histogram FITS file

The Primary Data Unit (PDU) in the FITS file is a 2-D raw histogram frame (also referred to as an "image") consisting of 1024x32 16-bit integer numbers. Note that the ALICE instrument data are unsigned 16-bit integers (giving values of 0 to 65535). Following the PDU is the first Extension Data Unit (EDU), which contains the 64-element pulse height distribution (PHD) that was acquired together with the histogram. When the Level 1 pipeline saves the PHD to this extension, it then zeros out that part of the histogram array (the 4 (spatial) x 32 (spectral) pixels in the lower left corner). The second data extension contains a 117 column by t row binary table, where t is the time of the exposure in seconds, of housekeeping values recorded during the observation (the housekeeping report rate is programmable and typically varies from 1-1/30 Hz).

Table 7-2: ALICE Level 1 FITS File Data Units: Histogram

FITS File Data Units	Description
Primary Data Unit (PDU)	Raw Histogram image (uncorrected counts)
Extension #1	Pulse Height Distribution (PHD)
Extension #2	Binary Housekeeping Table

7.2.1.2 Pixel list FITS file

Upon receiving a pixel list frame, the Level 1 processing creates a ground-calculated "reconstructed histogram" from the received pixel list data and places it in the primary data unit of the FITS file; this enables an easy quick-look inspection of the pixel list data (e.g., using most FITS viewers that by default typically display the data in the primary data unit). The pixel list data itself can be hard to interpret, so this reconstructed histogram image is desirable to enable the scientist to determine, e.g., data quality, whether the target was in the field of view, etc... The first data extension contains the raw pixel list data set, which includes the full stream of photon events and time hacks. The second extension contains the derived count rate derived from the pixel list data. Each count rate bin shows the number of events that occurred between successive time hacks. The resolution of this count rate data set is determined by the hack rate used for the pixel list acquisition. The length of this vector is variable depending on the source flux and the hack rate. The third data extension contains a 117 column by t row binary table, where t is the number of housekeeping values recorded during the observation.

Table 7-3: Alice Level 1 FITS File Data Units: Pixel List

FITS File Data Units	Description
Primary Data Unit (PDU)	Reconstructed Histogram image (uncalibrated counts)
Extension #1	Raw pixel list
Extension #2	Count rate vector from pixel list data (sampled at time hack rate)
Extension #3	Binary Housekeeping Table

7.2.1.3 High Cadence Count Rate FITS file

The HCCR Level 1 data FITS files have the 8-bit values of the number of detected events within each time sample period. The HCCR section of the FITS header (delimited by "HCCRSEP") contains information about the sample period, saturation value, and some statistics of the values. The first data extension contains a 117 column by t row binary table, where t is the number of housekeeping values recorded during the observation.

Table 7-4: Alice Level 1 FITS File Data Units: High Cadence Count Rate

FITS File Data Units	Description
Primary Data Unit (PDU)	Number of counts in each time sample period
Extension #1	Binary Housekeeping Table

7.2.2 Data Sources (High/Low Speed, CCSDS, ITF)

ALICE data are transferred via CCSDS packets that are packetized by the spacecraft from the Solid State Recorder.

The spacecraft will packetize the ALICE High Speed Telemetry data into CCSDS packets before sending the data to the ground. Different ApIDs (packet Application Process Identifiers; also Application IDs) are used to distinguish the histogram and pixel list data packets. For the ALICE science frames the spacecraft will either use the packetized (“RAW”) format or the lossless compressed format to transfer the data. In either case the result will be data encoded in CCSDS telemetry packets. The packet ApIDs listed in Table 7-5 are used to distinguish between the different ALICE CCSDS packets.

The packet time as listed in the CCSDS packet represents the time at which the packetization operation was performed. The second time field contains the frame collection time as sent from a spacecraft perspective, meaning this represents the time at which the spacecraft received the frame from the instrument. As the instrument immediately sends the frame at the completion of the acquisition this corresponds to the time at which the acquisition was completed. For each data type there are two ApIDs, one for each C&DH (Command and Data Handling) system on the spacecraft.

Note: In the Table 7-5, the ApIDs for HCCR data are effectively “pseudo” ApIDs. As discussed in Section 7.2.1, the spacecraft doesn't identify HCCR data as different from Pixel List data, and it is processed through the instrument -> spacecraft -> downlink with the Pixel List ApIDs. The data only acquire the special HCCR ApIDs in the ground pipeline processing where, once the data are identified as HCCR and the ApID is put in the FITS header and filename

Table 7-5: ALICE Science ApIDs

ApID CDH-1	ApID CDH-2	Data type
0x4B0	0x4B4	Compressed (Lossless) ALICE Pixel list Data (DT-1)
0x4B1	0x4B5	Packetized (RAW) ALICE Pixel list Data (DT-1)
0x4B2	0x4B6	Compressed (Lossless) ALICE Histogram Data (DT-2)
0x4B3	0x4B7	Packetized (RAW) ALICE Histogram Data (DT-2)
0x4C0	0x4C1	Compressed (Lossless) ALICE HCCR Data (DT-3)
0x4C4	0x4C5	Packetized (RAW) ALICE HCCR Data (DT-3)

Packetized “RAW” telemetry:

The nominal data transfer method for the ALICE pixel list science data is packetized (RAW) data. Each CCSDS science packet can transfer a segment of up to 480 data bytes. In order to transfer a full ALICE frame of 32768 words (16-bits), 137 science packets are needed; the first 136 packets will all be full size segments of 480 bytes, the last packet will transfer the remaining 256 bytes. The grouping flags of the packets indicate the start and end segment within a complete frame transfer.

Note the ‘#’ marks in the following tables refer to the third hexadecimal digit of the ApID; valid values for that digit are as indicated in the table captions and in Table 7-5.

Table 7-6: ALICE CCSDS Packetized (RAW) Science Packet; # = 1, 3, 5, 7

Parameter	Bit Count	Size in bytes	Description
Primary Header			
PH_VER_NUM_4B#	3	2	Version Number, fixed value = 0; designates a source packet
PH_PKT_TYP_4B#	1		Type Indicator, fixed value = 0; designates a telemetry packet
PH_SH_FLG_4B#	1		Secondary Header Flag, fixed value = 1; designates presence of secondary header
PH_APP_ID_4B#	11		Application Process identifier, see Table 7-5
PH_SEQ_FLG_4B#	2	2	Grouping Flags: 1 – first segment 0 – intermediate segment 2 – last segment
PH_SEQ_CNT_4B#	14		Source Count, continuous sequence count of all generated packets (per ApID) (modulo 16384)
PH_PKT_LEN_4B#	16	2	Number of bytes (secondary header + data bytes – 1): 511 or 287
Secondary Header			
SH_PACKET_TIME_4B#	32	4	Spacecraft MET at time the Telemetry packet is constructed
SH_COLLECT_TIME_4B#	32	4	Spacecraft MET at time the high-speed science data was collected
Data			
ERROR_STATUS	64	8	Information from the SSR forward error correcting code, not important for simple decoding
SSR_HEADER	128	16	Information from the SSR storage administration, not important for simple decoding
DATA_BLOCK	3840	480	Data bytes, all packets except the last packet 480 bytes, the last one is 256 bytes

Lossless Compressed telemetry:

The nominal data transfer method for the ALICE histogram science data is lossless compressed data. When applied to pixel list data, the ‘FAST’ algorithm results in negligible compression rates, and occasionally in a 1% expansion, therefore lossless compression will generally not be used with pixel list data. The spacecraft uses the so called ‘FAST’ algorithm to compress the image data. The ‘FAST’ algorithm uses one-dimensional correlation between successive data elements to remove redundancy. Data are encoded in blocks of 16 successive science values, the first value of such a block is sent in full 16 bits, the remainder of the block is encoded using successive differences, using an adaptive coding mechanism.

Table 7-7: ALICE CCSDS Compressed (Lossless) science packet; # = 0, 2, 4, 6

Parameter	Bit Count	Size in bytes	Description	
Primary Header				
PH_VER_NUM_4B#	3	2	Version Number, fixed value = 0; designates a source packet	
PH_PKT_TYP_4B#	1		Type Indicator, fixed value = 0; designates a telemetry packet	
PH_SH_FLG_4B#	1		Secondary Header Flag, fixed value = 1; designates presence of secondary header	
PH_APP_ID_4B#	11		Application Process identifier, see Table 7-5	
PH_SEQ_FLG_4B#	2	2	Grouping Flags: 1 – first segment 0 – intermediate segment 2 – last segment	
PH_SEQ_CNT_4B#	14		Source Count, continuous sequence count of all generated packets (per AplD) (modulo 16384)	
PH_PKT_LEN_4B#	16	2	Number of bytes (secondary header + data bytes – 1): maximum 265	
Secondary Header				
SH_PACKET_TIME_4B#	32	4	Spacecraft MET at time the Telemetry packet is constructed	
SH_COLLECT_TIME_4B#	32	4	Spacecraft MET at time the high-speed science data was collected	
Data				
COMP_UNIT_INDICATOR	16	2	Compression Unit (CU) Indicator The compression unit indicator (CUI) is sequential with the MSB set if it is the first packet within a CUI.	
FIRST CU only	IMAGE_LINE_SIZE	16	2	Width of an image line
	IMAGE_LENGTH	16	2	Total image length
	X_POSITION	16	2	Bottom left start position (windowed data)
	Y_POSITION	16	2	Bottom left start position (windowed data)
	COLOR_INDICATOR	16	2	Color channel of data, 0 = not applicable for ALICE

Parameter	Bit Count	Size in bytes	Description
FAST_DATA	8 ... 2048	1..256	Lossless (FAST) compressed ALICE science data; total packet data size is limited to 256+2 bytes, so for the first CU only 246 data bytes are used.
<p>Description of the FAST encoded data:</p> <ul style="list-style-type: none"> - Each compressed observation consists of an integer number of Compression Units (CUs). - Each CU represents a fixed number of samples (2048). - Each CU starts with a 16 bit initial value, this value is byte aligned. - In the first CU of an observation the third byte contains the block length (16). - Each Compression Unit contains an integer number of Blocks (= 2048/block length = 128). - Each Block starts with a 5 bit number that indicates the number of bits used to encode the successive differences. - Each Block continues with the differences, the least significant bit of these small numbers is used as a sign indicator, so the differences can be positive and negative. <p>Note that even after an initial value (start of a CU) a difference is encoded to calculate the first data sample, so the initial delta after the initial value will be zero.</p>			

7.2.3 Definition of an "Observation"

An observation will be a single histogram image or one frame of a pixel list series. Each observation will be written to a separate FITS file. A pixel list resulting from a single exposure command may therefore produce many such frames, each of which will be saved as a separate FITS file.

7.2.4 S/C Housekeeping Needed in Level 1 Files (for Calibration)

Spacecraft housekeeping that may be needed in the ALICE pipeline include any temperature sensors on the spacecraft around the ALICE instrument and the spacecraft-measured instrument bus voltage and power consumption on the different busses.

Spacecraft measured temperatures related to ALICE (ApId 0x00D and 0x08D):

T_A.CDH_TEMP_ALICE_BRACK_BASE_00D
T_A.CDH_TEMP_ALICE_1_00D
T_A.CDH_TEMP_ALICE_2_00D

PDU parameters related to ALICE (ApId 0x009, 0x00a, 0x089 and 0x08a):

ALICE_LVPS_A_VOLT_009
ALICE_LVPS_A_CURR_009
ALICE_LVPS_B_VOLT_009
ALICE_LVPS_B_CURR_009
ALICE_ACT_A_VOLT_00A
ALICE_ACT_A_CURR_00A
ALICE_ACT_B_VOLT_009
ALICE_ACT_B_CURR_009

Note that these temperatures currently are not used in the pipeline processing, but may be used in the future as the code and calibrations are revised.

7.3 Level 2 Data Specifics

“Calibrated” data as used here refers to CODMAC level 3 data.

7.3.1 Algorithm for Pipeline

Overview: The ALICE calibration pipeline that is run at the SOC applies various calibrations to Level 1 ALICE data to convert the data from units of counts to flux units (photons/s/cm²). Three types of operations can be performed. They are, in order of application to the data: deadline correction, dark correction, effective area correction. These are described in more detail below.

7.3.1.1 Deadline Correction

The ALICE detector electronics require a finite time to process an event pulse. As a result, if event pulses arrive too close together in time, the latter pulse(s) will not be recorded, resulting in an effective decrease in the sensitivity of the instrument that is a function of the count rate. The deadline correction time constant for the high-speed electronics, i.e. the science data, is 18 microseconds. The deadline correction time constant for the low-speed electronics, i.e. the housekeeping data, is 4 microseconds. At input rates below 50 kHz, the detector electronics is non-paralyzable (i.e. a fixed deadline per event that is not re-triggerable). To calculate the detector output count rate, the following formula is used:

$$C_{out} = C_{in} / (1 + C_{in} \tau)$$

where C_{out} is the output (i.e., detected) count rate and C_{in} is the input count rate. At a count rate of 1 kHz, the deadline correction factor (τ) is approximately 1.02, while at 20 kHz, the deadline correction factor is approximately 1.56.

7.3.1.2 Dark Correction

The ALICE detector electronics register events even when the aperture door is closed and the detector is not illuminated by FUV photons. The spatial distribution of these dark counts is approximately uniform across the detector. However, there is some low-level 2-D structure to the dark counts. ALICE observations made with the aperture door closed are summed together to create a "superdark". This superdark image is then scaled to the exposure time of an ALICE science observation and subtracted from the data.

During in-flight commissioning, these dark counts were measured at a rate of approximately 94 Hz across the entire detector. The primary source of dark counts is the spacecraft RTG. Dark exposures are made throughout the mission to monitor the background event rate and detector performance.

7.3.1.3 Effective Area

The sensitivity of ALICE to UV photons varies as a function of wavelength. It is convenient to think of the ALICE sensitivity in terms of the effective area of the instrument. For a point source located at infinity, effective area is defined as the area of the surface that intercepts incident photons at the same rate as is detected by the ALICE instrument. Dividing the observed count rate by the effective area yields the incident flux of photons. In general, effective area depends on the geometric size of the instrument aperture, reflectivities of the optical surfaces, sensitivity and quantum efficiency of the detector, etc.

The ALICE effective area curve is based on observations of stars that have been previously observed by IUE and HST. In wavelength regions outside those covered by the IUE/HST observations, the pre-flight effective area estimate has been linearly scaled so that it matches the effective area derived at longer wavelengths.

7.3.1.4 Calibration File set

As of October, 2016, the current set of calibrations files to use with all ALICE flight data are

```
AEFFFILE= 'pa_aeff_007.tab'      / Detector effective area file
DARKFILE= 'pa_dark_002.fit'     / Detector dark count rate file
```

Future calibration files follow this same format but with updates to the version. All other calibration files provided with PDS data sets are there for historical purposes, to support tracking of instrument performance.

7.3.2 Level 2 (Calibrated) Data Format

7.3.2.1 Histogram FITS File

The primary data unit in the FITS file is a 2-D calibrated histogram frame consisting of 1024x32 array of 32-bit floating-point numbers. The units of the histogram image are photons/s/cm². The first data extension in the FITS file is a 1024x32 array of 32-bit floating numbers containing the uncertainty in the histogram image. The second data extension is a 1024x32 element array containing the wavelength for each pixel in the histogram image. The third data extension is the 64-element PHD, identical to that in the Level 1 data. The fourth data extension is an array containing the number of photon events per housekeeping sampling interval. The fifth data extension is the 141 column by t row housekeeping row as in the Level 1 data.

Table 7-8: ALICE Level 2 Calibrated FITS File Data Units: Histogram

FITS File Data Units	Description
Primary Data Unit (PDU)	Calibrated Histogram image (photons/sec/cm ²)
Extension #1	Uncertainties in histogram data values
Extension #2	Wavelength Image (Angstroms)
Extension #3	Pulse Height Distribution (PHD)
Extension #4	Count rate vector from HK
Extension #5	Binary Housekeeping Table

7.3.2.2 Pixel List FITS File

The primary data unit in the FITS file is a 2-D calibrated reconstructed histogram image consisting of a 1024x32 array of 32-bit floating-point numbers. The units of the histogram image are photons/s/cm². The first data extension in the FITS file is a 1024x32 array of 32-bit floating numbers containing the uncertainty in the reconstructed histogram image. The second data extension is a 1024x32 element array containing the wavelength for each pixel in the reconstructed histogram image. The third data extension contains a binary table of 5 columns and rows for each photon event. The five columns are: the X (spectral) position of each photon event, the Y (spatial) position of the photon event, the wavelength of the photon event, the cumulative number of elapsed time hack intervals (starting from 0 at the beginning of the file), and the deadtime correction factor events occurring during the time hack in question. The temporal resolution of this pixel list data set is determined by the hack rate used for the pixel list acquisition. The length of this vector is variable depending on the source flux and the hack rate. The fifth data extension contains the binary housekeeping table.

Table 7-9: ALICE Level 2 Calibrated FITS File Data Units: Pixel List

FITS File Data Units	Description
Primary Data Unit (PDU)	Reconstructed Calibrated Histogram image (photons/sec/cm ²)
Extension #1	Uncertainties in reconstructed histogram data values
Extension #2	Wavelength Image (Angstroms)
Extension #3	Binary Pixel List Table (X , Y , wavelength, time hack #, MET)
Extension #4	Count rate vector from pixel list data (sampled at time hack rate)
Extension #5	Binary Housekeeping Table

7.3.2.3 High Cadence Count Rate Fits File

The HCCR data FITS files have a blank Primary Data Unit (PDU), since the data are in a binary table which cannot be in the PDU. The first data extension contains a binary table of 4 columns and rows for each time sample period. The four columns are: the mission elapsed time (MET), the count rate in Hz, the deadtime-corrected countrate in Hz, the Gaussian (\sqrt{N}) uncertainty in that countrate, and the raw counts (i.e., the same values as in the Level 1 data).

Table 7-10: ALICE Level 2 Calibrated FITS File Data Units: High Cadence Count Rate

FITS File Data Units	Description
Primary Data Unit (PDU)	empty
Extension #1	Count Rate Table (MET, count rate, count rate error, raw counts)
Extension #2	Binary Housekeeping Table

7.3.3 Scientific Units

For Histogram, units are photons/s/cm² (histogram), angstroms (wavelength array), and counts (PHD array).

For Pixel List, units are photons/s/cm² (generated histogram), angstroms (wavelength array), counts, pixel location, and angstroms, and seconds (pixel list array), and counts per second (count rate array).

7.3.4 Additional FITS and PDS Keywords Added

This section is an example of the Mike pipeline keyword block added to the FITS header:

```
COMMENT =====
COMMENT
MIKE_BEG= 'Feb 15 16:12:57 2005'      / START MIKE KEYWORD BLOCK
MIKE_VER= '2.0 [2005 Feb 15]'        /Version of Mike pipeline code
K_MODE = 'ACQMODE'                   / Keyword containing the mode name
K_ETIME = 'EXPTIME'                  / Keyword for the effective exposure time
FILE_IN = 'test/ali_0000006498_0x4b3_eng_1.fit' / Input file for processing
FILE_OUT= 'test/test_his.fit'        / Output file after processing
DIR_CAL = 'cal/'                     / Directory of calibration data
DIR_DONE= ''                         / Directory to put raw data after processing
BADFILE = ''                         / FITS file of bad pixel mask array
BADFLAG = -1                         / Bad pixel mask flag
BADVALUE= -666                       / Bad pixel value
DEADFILE= 'deadtime/ra_dead_002.txt'  / Deadtime correction file
DEADFLAG= 1                          /
DEADTYPE= 'FUN'                      / Correct using FUNction or lookup table (LUT)?
DEADCORR= 'TOTAL'                   / Correct by TOTAL or each PIXEL count rate?
BIASFILE= ''                         / Bias image filename
BIASFLAG= -1                         / Bias correction flag
DARKFILE= 'dark/ra_dark_001.fit'     / Dark image filename
DARKFLAG= -1                         / Dark correction flag
FLATFILE= 'flat/ra_flat_001.fit'     / Flat field image filename
FLATFLAG= -1                         / Flat field correction flag
FLATNORM= 'AVERAGE '                / How to normalize flat field
WCALFLAG= 0                          / Wavelength calibration flag
WCALPRO = 'alice_wavecal'           / IDL program to perform wavelength calibration
WCALPARS= 'T_DELECC'               / keywords for parameters to use for wave cal
AEFFFLAG= 1                          /
AEFFPRO = 'alice_aeff'              / IDL program to get effective area
AEFFPARS= 'T_DELECC'               / keywords for parameters to get effective area
LOG_FILE= 'test/log.out'            / Filename to save log file (default = append to
LOG_MAIL= ''                        / address (if any) to e-mail log file
MIKE_ERR= 1                          /
MIKE_END= 'Tue Feb 15 16:12:57 2005' / END MIKE KEYWORD BLOCK
COMMENT
COMMENT =====
```

7.3.5 *Hardware/OS Development Platform*

Dell Linux, Redhat 7.2; Apple G5 Power PC and PowerBook G4, OS X v10.4

7.3.6 *Language(s) Used*

IDL

7.3.7 *Third Party Libraries Required*

IDL Astro

7.3.8 *Predicted Execution time*

A few seconds per file.

7.3.9 *Contact/Support Person(s)*

Joel Parker and Maarten Versteeg

8. LEISA INSTRUMENT DESCRIPTION

8.1 Overview

The Linear Etalon Imaging Spectral Array (LEISA) is an infrared imaging spectrometer. The detector is a 256x256 pixel array (so $256 \times 256 = 65,536$ pixels in the array). Spectral separation is done with a wedged optical etalon filter placed in close proximity to the detector array. The filter is made of two pieces, a high spectral resolution ($\lambda/\Delta\lambda=560$) segment and a low spectral resolution ($\lambda/\Delta\lambda=240$) segment, bonded together. The detector-filter assembly is located at the plane of focus of the Ralph telescope where a 2-D image is recorded that contains the infrared spectrum of the scene. The layout for the filter assembly is shown in Figure 8-1. The wavelength range of the sensor is 1.25-2.5 μm for the low resolution segment and 2.1-2.25 μm for the high resolution segment. The wavelength of transmission of the filter varies along one axis and is relatively constant in the other. Lines of nearly constant wavelength are aligned with the row direction of the detector array. The number of pixel-limited spectral channels is the number of rows of the detector, excluding a number of rows (4) obscured by opaque adhesive at the bond joint between the two filter segments.




Figure 8-1: Layout for LEISA Linear Variable Filter Assembly; displayed samples (columns) increase right to left, and lines (rows) increase bottom to top.

The LEISA detector array is a Rockwell PICNIC HgCdTe device. It is read out by the Ralph electronics. The signal is converted to a 12-bit value using the middle 12 bits of a 16 bit analog to digital (A/D) converter. There are two data transfer modes, one in which both signal and reset level data are returned (raw mode), and the other in which the reset level is subtracted from the signal level and only the difference is returned (Correlated Double Sample (CDS) mode). Note: analog saturation occurs when the digitally converted signal contains 13 bits so that, in principle, the dynamic range is 0 to 8192. This is accounted for in the pipeline.

One read out of all 65,536 detector pixels is called a frame. The “frame” is converted to an “image” by sorting the read-out pixel order so that it is consistent with the optical definition (see Section 8.2.2). An image cube is obtained from a series of N images, with N determined by the duration of the scan multiplied by the frame rate. The frame rate is adjustable between 0.25 and 8 Hz in 1 ms steps. LEISA is normally operated in a scanning mode, with the target moving

through the image plane. In the normal scan mode, the integration time is set from the measured spacecraft scan rate so that the image moves 1 row per integration time. Slicing the image cube along one row gives a scanned image of the target over a few spectral resolution elements. Co-registering each image (removing motion and optical distortion and correcting for the spectral “smile”) yields an IR spectrum of the target.

Data recorded on *New Horizons* is sent to the ground via the Deep Space Network. From there the data are sent to the Mission Operations Center (MOC) at the Applied Physics Laboratory (APL). The Science Operations Center (SOC) retrieves new data from the MOC daily. The SOC software pipelines convert the data from the MOC archives into FITS (Flexible Image Transport) files with scientifically useful and calibrated data. The SOC first sorts the packets into image cubes of raw (12-bit) sensor counts with useful header keywords. These keywords include the recording mode of the observation, timing information and basic pointing information of the instrument boresight. The raw processing also gathers housekeeping (H/K) telemetry from the Ralph instrument into a table. Once the raw processing is complete, the SOC produces a calibrated data set for each observation.

8.2 Level 1 Data Specifics

8.2.1 Data Format

Dataset

The SOC stores the LEISA data cubes in Band Interleaved by Line (BIL) order, i.e. images are stored sequentially. To re-order LEISA frames as received from the spacecraft to images, the SOC does the following to each frame of data:

1. de-interlace by quadrant
2. reverse the Y direction
3. rotate 180 degrees

In the resulting sequence of images, the (0,0,0) element of the data cube corresponds to the location of wavelength 2.5 μm on the LEISA filter at the minimum X axis location in the first image.

The SOC (Level 1) data product is a FITS format data file and a PDS detached label file. Ancillary data for an observation is placed in the primary header of the FITS data file. The 256x256xN data cube is stored in the primary data unit as an array of integers. The first FITS extension is a binary table of Ralph housekeeping data.

Outline of the (Level 1) FITS file:

- Primary Header Data Unit (HDU) - 12 bit image counts
 - Primary Header (FITS + pointing + observation keywords)
 - 256 x 256 x N integer point array
- Extension 1 - Binary table of Ralph housekeeping
 - Ext. Header (keywords + binary table definition)
 - Ext. Binary Table (115 x S binary table of Ralph housekeeping data)

* [N is the number of data frames in the observation, S is the number of seconds in the observation]

** [In the case of raw mode, images alternate between read and reset signal levels]

Image Data

The primary data unit contains the digital spectral image data. Values recorded by the instrument with 12 bit precision are stored as 16 bit integers.

Housekeeping Data

Housekeeping data generated by the Ralph instrument are stored in Extension 1 as a binary table. The first field in each row of the table is mission elapsed time (MET). Table entries are sorted by increasing MET. The time interval between each table entry is fixed, one second per entry, unless there is missing data.

Pipeline Processing

The limits of an observation are established by the SOC using information in each telemetry packet of an observation sequence.

8.2.2 Data Sources (High/Low Speed, CCSDS, ITF)

Ralph housekeeping data are transmitted in the form of CCSDS (Consultative Committee for Space Data Systems) packets. One housekeeping packet is produced by the Ralph instrument each time a spacecraft PPS (pulse per second) signal is received. State information is gathered, time tagged, and written to the low speed bus in CCSDS packet form. The CCSDS packets are transmitted during the next DSN pass.

LEISA frame data are transmitted in the form of CCSDS packets produced by the spacecraft compression/packetization routines from the data written to the high-speed bus. The LEISA detector has 4 output channels, one for each array quadrant. The first 4 elements of the data stream are the first pixels from each quadrant. The second 4 elements are the second pixels from each quadrant, and so on. One 'line' of data in this order is 4 x 128 pixels long, and is not the same as a line in the final image. During the observation, frame data are written to the spacecraft high speed bus. These data are not automatically transmitted. A compression/packetization routine is scheduled some time later that converts the Ralph sensor counts to a packetized form.

The data can be packetized without compression, with lossless compression, or with lossy JPEG compression. The packetization routine can also process a sub-frame area of interest, or a more complicated sliding subframe that tracks the image target as the scanning observation proceeds. In raw mode, spacecraft operations team only downlinks uncompressed data. Regardless of windowing or compression, the SOC Level 1 data processing reassembles the data into a full 256x256xN data cube.

8.2.3 Definition of an "Observation"

Science Operation

An observation is a sequence of frames. The number of frames per observation is variable. Pixel values are recorded with 12 bit precision. One image contains 65536 pixels x 1.5 bytes/pixel = 98304 bytes of data, however, data are stored as 16 bit values in the SOC data files.

For normal science observations the Ralph electronics use the measured spacecraft scan speed to set the frame rate to minimize smearing by compensating for the spacecraft motion relative to the target. Reset levels are stored temporarily by the electronics and subtracted from read levels (CDS mode). The frame rate is set at the beginning of the scan. The difference in read and reset levels is transferred to the spacecraft. Alternately, the instrument can be forced to use a forced frame rate value.

Un-subtracted Read-out

There is a voltage offset for each LEISA quadrant to assure the sensor signal will be in the correct range of the A/D. The offset values are set from a table when Ralph is powered on. Raw mode data are used to evaluate the offset values. In raw mode, reset levels are not subtracted from read levels. Both read and reset are transferred. The number of values in one data frame of raw mode (131,072) is twice that of CDS mode. Read and reset values are interleaved by data line and the number and order of the pixel elements in a line are the same as for CDS mode readout. The reset of a pixel occurs after the integrated signal is read, so read levels correlate with reset levels recorded in the preceding frame.

8.2.4 Housekeeping Needed in Level 1 Files (for Calibration)

Most of the H/K values are used for engineering troubleshooting and not needed for data processing. Housekeeping data that are important to further processing (see Table 8-1) are stored in header keywords.

Table 8-1: LEISA Housekeeping Level 1 Keywords

Keyword	Description
SIDE	Instrument hardware side
DETECTOR	Always LEISA for LEISA data
FILTER	Always WEDGE for LEISA data
LEI_OFFx	Value used to set voltage offsets for the four LEISA quadrants. x=1-4
LEI_RATE	Time between LEISA readouts (ms)

8.2.5 Science Data and/or Housekeeping Requirements

Other important information is determined by the SOC while processing the raw observation data. These values are also stored as keywords in the FITS header.

Table 8-2: LEISA SOC-Processed FITS Keywords

Keyword	Description
MET510	The MET of the Ralph housekeeping packet that marks the start of an observation, used to determine the observation start time and frame rate
TRUE510	Whether the 0x510 packet is real (YES) or assumed from a gap (NO)
SCANTYPE	Always LEISA for LEISA data
LEI_MODE	RAW for un-subtracted readout mode SUBTRACTED for CDS mode
STARTMET	Actual start time of first integration, in MET (s)
EXPTIME	LEISA exposure time (s). Same as RALPHEXP and as LEI_RATE / 1000. There is zero dead time between frames so the frame rate is exactly 1/EXPTIME

SPICE and SPICE Kernels

The SOC maintains an archive of SPICE kernels that describe the position and attitude of the spacecraft as any time in the mission. The kernels are used to calculate many values that describe the instrument pointing during each observation. The names of the SPICE kernels used to process the observation are stored as header keywords with the prefix SPCK. All SPICE-based keywords start with SPC; the kernel-related SPICE keywords are the only ones that start with SPCK.

8.3 Level 2 Data Specifics

8.3.1 Algorithm for Pipeline

There are six processing steps applied to the digital LEISA data to produce the calibrated output:

1. Validate digital image file
2. Preprocess raw mode data
3. Process A/D rollover pixels
4. Convert digital counts to radiometrically calibrated values
5. Compute pointing data
6. Construct FITS file

Validate digital image file

The input file is validated to assure the data are ready for further processing. Checks are for valid mission, instrument, mode, and image array size. The values of important keywords are validated and collected in this step.

Preprocess raw mode data

If the data readout was in raw mode, the reset values are subtracted from the read values in this step and the rest of the processing is the same, regardless of readout mode.

Process A/D rollover pixels

There are two instances where the reset subtracted data value will be off by exactly 4096 counts. The first is in cases of low signal when the subtraction of the reset count results in a small negative number because of noise in the data. Small negative numbers are returned as large positive numbers because 13 bits are used in the subtraction, but only 12 bits are stored in memory. The second instance is when the subtraction of the reset count results in a number greater than 4095 (12 bits). Again, because only 12 bits are stored, the stored number will be low by 4096 counts. These cases are corrected on a case by case basis.

If a file identifying rollover pixels for the observation exists (see next paragraph), the identified pixels are corrected for rollover. If no file exists, any subtracted count greater than 3850 is considered to have rolled over, and 4096 is subtracted from the raw count value. This is done in initial processing because the observations, particularly at Pluto, will not be capable of producing signal counts this high.

During initial image analysis by the Ralph team, each observation is analyzed in detail. A file identifying rollover pixels is generated which identifies the pixels that are deemed to need rollover correction. The case where the read count is higher than 4095 can be detected by analyzing surrounding pixels and by watching the target scan through the array. These are also included in the rollover file. Once this file is installed on the SOC, processing of the calibrated data for the observation will automatically use the rollover file instead of the default processing.

Convert digital counts to radiometrically calibrated values

Radiometrically calibrated values are given as $\text{ergs/sec/cm}^2/\text{\AA}/\text{sr}$. The conversion involves multiplying the measured digital counts/sec $((\text{measured digital counts})/(\text{integration time}))$ in each pixel by calibration gains for that pixel. The measured digital counts have been corrected for additive calibration offsets and, as mentioned in the step above, have had the reset value subtracted. The calibration gains have been determined from a combination of pre-launch measurements, in-flight stellar and Jupiter system flat field observations and, for observations within the Pluto system and beyond, a flat-field measurement made using Pluto itself.

The flat field changes slightly as the mission progresses so a different flat field can be defined for an individual observation or a range of observations. The actual flat field used in the processing is included in the output FITS file. Also included in the output FITS file are the calibration coefficients and the calibration offsets.

Compute pointing data

The pointing for each pixel of each image is computed using the timing information from the observation, reconstructed ephemeris and attitude files, and knowledge of the optical distortion of the instrument. One array is generated giving the Cartesian pointing vector of each pixel in the LEISA array. This is a function only of the optical distortion of the system. A second array is generated giving the rotation quaternion of the instrument boresight into the J2000 reference frame for the middle of each exposure. By rotating the pointing vector of a pixel by the quaternion for the image frame, the J2000 pointing vector of each pixel can be derived.

Construct FITS files

A FITS file is constructed to store all the calibrated image data and related processing data.

8.3.2 Dataflow Block Diagram

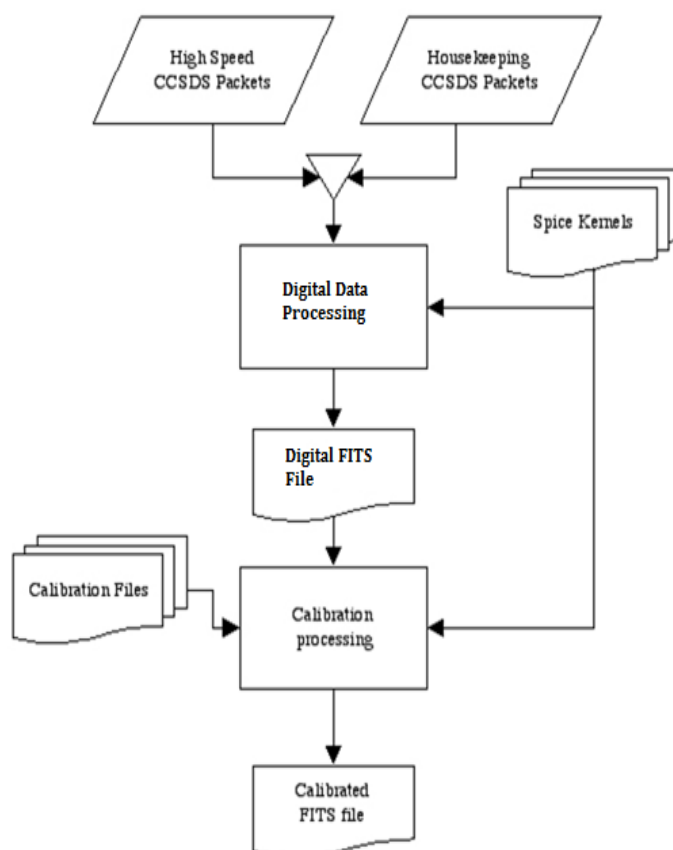


Figure 8-2: Digital and Calibrated data pipeline processing flow

8.3.3 Data Format

Calibrated Dataset

The calibrated LEISA data are stored in Band Interleaved by Line (BIL) order, exactly as the digital count data are stored. The resulting images from LEISA have the (0,0,0) element of the 3-D array corresponding to the location of wavelength 2.5 μm on the LEISA filter at the minimum X axis location in the first image.

The calibrated data product is a FITS format data file and PDS detached label file. Ancillary data for an observation are placed in the primary header of the FITS data file. The 256x256xN data cube is stored in the primary data unit as an array of floating point numbers. The FITS extensions are outlined below.

Outline of the calibrated FITS file:

- Primary HDU - Calibrated image data
 - Primary Header (FITS + pointing + observation keywords)
 - 256 x 256 x N floating point array
- Extension 1 - Center wavelength and filter width for each pixel
 - 256 x 256 x 2 floating point array
- Extension 2 - Cartesian pointing vector for each pixel
 - 256 x 256 x 3 floating point array
- Extension 3 - Flat field correction for each pixel
 - 256 x 256 x 1 floating point array
- Extension 4 - Radiometric gain and offset for each pixel
 - 256 x 256 x 2 floating point array
- Extension 5 - Error estimates for each pixel
 - 256 x 256 x N floating point array
 - N.B. all error estimates are zero; see *Calibration and data characteristics* below
- Extension 6 - Data quality flags for each pixel
 - 256 x 256 x N Integer array
- Extension 7 - Ephemeris time and quaternion for each frame
 - 5 x N floating point array
- Extension 8 - Binary table of Ralph housekeeping
 - Ext. Header (keywords + binary table definition)
 - Ext. Binary Table (115 x S binary table of Ralph housekeeping data)

*[N is the number of data image frames in the observation, S is the number of seconds in the observation]

For a description of the contents of the FITS extension, see the above section describing the SOC calibration processing.

Calibrated Image Data

The Image Data Unit of the Level 2 file contains data expressed in physical units useful for scientific interpretation. The instrument pipeline converts the data values of digital counts to radiance units, ergs/sec/cm²/Å/sr.

Calibration and data characteristics

- i) *Error estimates.* The error estimates for each pixel are calculated as $(RN^2 + \text{digital count} * 11)^{1/2} / 11$, where RN is the system read noise in electrons and 11 is the number of electrons per digital count.
- ii) *Fixed-pattern noise.* There is often fixed-pattern noise (FPN) in the data. This may be removed by generating an average of several run-up images containing background sky, and subtracting that average from every image in the scan. The number of frames to use will vary by observation, so this operation is not part of the automated pipeline.
- iii) *Scattered light.* Although it is known that scattered light is present at some level in the images, no photometric modeling of the effect has presently been applied to the data. Stray light modeling is being pursued, and it is expected that later versions of the pipeline will account for this effect.

Data quality flags

The data quality flag bits are set for known problem(s) with the corresponding pixel in the image; multiple bits may be set for any one pixel. A quality flag value of zero indicates a good pixel.

Table 8-3: LEISA Data Quality Flag Value Descriptions

Quality Flag Value (bit)	Description
1	Defect in one of the calibration files
2	Flat field out of bounds
4	Known CCD defect
32	Bad pixel not in any of above categories

8.3.4 Extra FITS Extensions (planes) and Their Definitions

See above

8.3.5 Scientific Units

Radiometric units: ergs/sec/cm²/Å/sr for extended sources, ergs/sec/cm²/Å for point sources.

8.3.6 Additional FITS and PDS Keywords Added

See above

8.3.7 Hardware/OS Development Platform

The software for processing the digital and calibrated data files has been developed on the SOC computers, running GNU/Linux/i686 Version 2.6.17-1.2142_FC4.

8.3.8 Language(s) Used

The software for processing the digital data files is written in Python.

The software for processing the calibrated data files is written in C, with a Perl script wrapper.

8.3.9 Third Party Libraries Required

Python SPYCE interface library

Python MySQLdb interface library

Independent JPEG Group's JPEG software

CSPICE processing library

CFITSIO processing library

8.3.10 Calibration Files Needed (with Quantities)

The instrument software package will include additional datasets needed for calibrating the data, as describe in the above sections. The instrument pipeline maintains version control on calibration datasets, and calibration procedures. Some calibration files are associated with specific observations, and some are associated with a range of observations. The calibration files are partitioned by time (MET), and the METs are the subdirectory names under the calibration top-level directory (e.g. on the SOC in the MET subdirectory will be pipeline/level2/lei/cal/METxxxxxxx/; in the PDS data sets it will be CALIB/METxxxxxxx/). The algorithm, for choosing which partition of calibration files to use, is as follows:

- Find the MET subdirectory name with the highest MET value less than or equal to the MET of the data to be calibrated
 - E.g. if the data MET is 250, and the MET subdirectory paths are cal/100/, cal/200/, cal/300/, then directory path chosen is cal/200/.
 - If such a directory is found, then use the calibration files in it
- If no MET subdirectory is found, look for the cal/default/ path
 - If such a directory is found, then use the calibration files in it
- If no cal/default/ path is found, use the calibration files in cal/initial/

The estimated number of calibration files needed for the mission is 300, totaling 3 GBs.

8.3.11 Memory Required

500MB

8.3.12 Temporary File System Space Needed

500MB

8.3.13 Predicted Size of Output File(s)

Up to 500MB

8.3.14 Predicted Execution time

Processing time for raw data files is approximately 3 seconds per image frame.

Processing time for calibrated data files is approximately 0.1 seconds per frame

8.3.15 Contact/Support Person(s)

Allen Lunsford 301-246-0286

Dennis Reuter 301-286-2042

Donald Jennings 301-286-7701

8.3.16 Maintenance Schedule (Code/Data Updates, Documentation)

The LEISA Processing System (LPS) is installed by extracting files from an archive. The sub-directory structure for the software package is created during extraction. Symbolic links to external directories may be substituted for default directory references. The shell for execution of the LPS is tcsh/csh. Changes will be made to shell initialization files; new elements are appended. Instructions for configuring the shell environment are given. A guide to installation and setup is included with the LPS package.

An initial period of testing and refinements is expected. Pieces of the software are tested separately during development. LPS modules are re-tested upon installation to the SOC. Sample datasets are provided to verify the function of software. Integrated testing of the instrument pipeline under the control of the SOC MDM is performed in accordance with the SOC. The instrument software engineer is available exclusively to the SOC to support the integration of pipeline software.

Changes to calibration datasets are made as needed. A facility is provided by the SOC so that software changes are reversible. A LEISA team member will be available to assist SOC operators in responding to unexpected errors in the instrument pipeline. Persons supporting the LEISA Instrument Pipeline software are listed above.

The LEISA instrument pipeline is developed at GSFC. The first fully functional version (v0) of the software was tested on the GSFC computer system. Documentation including instructions for installation and setup were delivered to the SOC with Version 1 of the LPS. Updates of software after Version 1 are performed on an as needed basis.

9. LORRI INSTRUMENT DESCRIPTION

9.1 Overview

The LOnG Range Reconnaissance Imager (LORRI) is a narrow angle (FOV=0.29°), high resolution (IFOV=5 μ rad), Ritchey-Chrétien telescope with a 20.8 cm diameter primary mirror, a focal length of 263 cm, and a three lens field-flattening assembly. A 1024 x 1024 pixel (optically active region), back-thinned, backside-illuminated CCD detector (model CCD 47-20 from E2V) is located at the telescope focal plane and is operated in standard frame-transfer mode. LORRI does not have any color filters; it provides panchromatic imaging over a wide bandpass extending approximately from 350 nm to 850 nm. The LORRI telescope has a monolithic silicon carbide structure, built by SSG Precision Optonics, Inc., is designed to maintain focus over the entire operating temperature range (-125 C to +40 C) without a focus adjustment mechanism. The instrument is shown in Figure 9-1, and a detailed description of the design and fabrication of LORRI can be found in the paper by Conard, et al., "Design and fabrication of the *New Horizons* Long-Range Reconnaissance Imager" in SPIE proceedings 5906-49, 2005. A detailed discussion of the performance of LORRI, as measured during calibration testing before launch, can be found in the paper by Morgan et al., "Calibration of the *New Horizons* Long-Range Reconnaissance Image" in SPIE proceedings 5606-49, 2005.




Figure 9-1: Cutaway Views of LORRI

LORRI is a supplemental instrument on *New Horizons* and is not needed to meet the baseline scientific objectives of the mission. Nevertheless, LORRI adds significant capabilities to *New Horizons*, including the highest available spatial resolution (50 m/pixel at the Pluto closest approach distance of 10,000 km) and redundancy for the primary optical imager, MVIC on Ralph.

The exposure time for LORRI is adjustable in 1 msec increments from 0 ms to 29,967 msec. However, exposure times will normally be limited to ≤ 150 msec to prevent image smear associated with spacecraft motion during observations. Initially, the shortest useful exposure time was expected to be ~ 40 msec owing to frame transfer smear associated with the transfer of charge from the active CCD region to the storage region, during which time the active region remains exposed to the image scene because LORRI has no shutter, but an improved frame transfer smear removal algorithm was developed that now permits exposure times as little as 1 msec. The LORRI exposure time can be commanded to a specific value, or LORRI can be operated in “auto-exposure” mode, in which the LORRI flight software sets the exposure time automatically based on the signal level in a previous image. In auto-exposure mode, the algorithm used to set the exposure time depends on several adjustable parameters that are stored in an onboard table. The optimal values for these table parameters vary with the type of scene being observed, which means that new table loads may be required prior to some observations. Although the LORRI auto-exposure mode worked well during ground testing, no decision has yet been made on whether it will be used in-flight during encounter observations.

LORRI can also be operated in “rebin” mode, in which case the signal in a 4 x 4 pixel region is summed on-chip to produce an active region that is effectively 256 x 256 pixels covering the entire 0.29° FOV. The main purpose of this mode is to provide high sensitivity acquisition of a Kuiper Belt object (KBO), which requires an exposure time of ~ 10 sec. Although LORRI rebin mode may never be used for science observations, the LORRI pipeline is still required to calibrate rebinned images.

9.2 Level 1 Data Specifics

9.2.1 Data Format

The Level 1 image data are organized in a FITS file. The primary header and data unit (HDU) is used to store the reconstructed image from telemetry. Additional data are stored in the extensions of the file. The two tables below contain a description of the layout for the extensions for Level 1 data.

As described previously, LORRI operates in two binning modes: 1x1 and 4x4. For the 1x1 binning mode, the Level 1 image dimensions are 1028x1024 where columns 0 through 1023 are the optically active region of the CCD and the remaining columns (1024-1027) are from optically inactive region (dark columns) of the CCD and represent a temperature-specific measurement of the bias value. For the 4x4 binning mode, the Level 1 image dimensions are 257 x 256 where columns 0 through 255 are optically active and column 256 for the dark column.

Table 9-1: Level 1 FITS file extension layout (ApID => Packet Application ID)

FITS File Storage Location	Description
Primary HDU	Reconstructed image from telemetry
First Extension	histogram from image descriptor packet (ApID 0x611)
Second Extension	Instrument housekeeping from first 34 pixels
Third Extension	Matching image descriptor

9.2.2 Data Sources (High/Low Speed, CCSDS, ITF)

The LORRI high-rate data are delivered to the Instrument Interface card over a low-voltage differential signal (LVDS) interface and is then transferred to the SSR through the spacecraft high-speed PCI bus by the C&DH software. The image data are stored directly on the SSR and CCSDS (Consultative Committee for Space Data Systems) packets are generated by command to the C&DH software as is described in the table below. The ApID (packet Application Process Identifier; also Application ID) from which the image originated is part of the filename, so this mapping may provide some assistance in decoding the filenames retrieved from the SOC.

Table 9-2: Low Rate Instrument Telemetry Description

ApID	MNEMONIC	Description
0x0601	LORRI_MEM_DMP	Memory Dump
0x0602	LORRI_MEM_CKSM	Memory Checksum
0x0603	LORRI_CMD_ECHO	Command Echo
0x0604	LORRI_ALARM	Alarm
0x0605	LORRI_STAT	Status
0x0606	LORRI_MON	Monitor Limits
0x0607	LORRI_BOOT	Boot Status
0x0608	LORRI_MAC_DMP	Macro Dump
0x0609	LORRI_MAC_CKSM	Macro Checksum
0x0610	LORRI_PARM	Parameters
0x0611	LORRI_IMG_DES	Image Descriptor

Table 9-3: LORRI high-speed telemetry description

ApID	C&DH side	Binning mode	Compression type
0x630	1	1x1	lossless
0x631	1	1x1	packetized
0x632	1	1x1	lossy
0x633	1	4x4	lossless
0x634	1	4x4	packetized
0x635	1	4x4	lossy
0x636	2	1x1	lossless
0x637	2	1x1	packetized
0x638	2	1x1	lossy
0x639	2	4x4	lossless
0x63A	2	4x4	packetized
0x63B	2	4x4	lossy

9.2.3 Definition of an “Observation”

Each LORRI image is an “observation.”

9.2.4 Housekeeping Needed in Level 1 Image Files (for Calibration)

No special requirements other than pointing

9.2.5 Level 1 Science Data and/or Housekeeping Requirements

No special requirements

9.3 Level 2 Data Specifics

9.3.1 Algorithms for Pipeline Calibration Process

The calibration of LORRI images potentially involves all of the following steps:

- 1) Bias subtraction
- 2) Signal linearization
- 3) Charge transfer inefficiency (CTI) correction
- 4) Dark subtraction
- 5) Smear removal
- 6) Flat-fielding
- 7) Absolute calibration (*N.B. this step does not change pixel values; see Section 9.3.1.4 below*)

Ground testing has demonstrated that the linearization, CTI, and dark subtraction steps will not be needed, so they are not described below. Nevertheless, the LORRI pipeline architecture will be maintained to allow these additional steps to be incorporated quickly, if in-flight data suggest they are needed.

The LORRI pipeline software consists of a series of IDL routines that implement the above processing steps. In general, the IDL routines have the following naming convention: *lorri_function.pro*, where “function” refers to the specific task performed by that routine. (The “pro” extension will be omitted below when discussing specific routines.) Each routine typically has several command line arguments and keywords that specify the input and output files and, possibly, parameters for tailoring the routine for particular circumstances. The routines that perform the bias subtraction, the smear removal, and the flat-fielding are described below. No special routines are provided to perform the absolute calibration. Instead, the absolute calibration is performed using keywords provided in the FITS header, as described further in Section 9.3.1.4.

9.3.1.1 Bias Subtraction

If an image has an associated “dark” image (i.e., an image taken with the same exposure time but without any illumination) then the debiased image is simply the difference of those two images. This was usually the case during on-ground testing when images taken of a scene were immediately followed by images taken with the scene blocked (i.e., an obstruction was placed in the optical path to block the illumination). However, in-flight images may often be taken without accompanying darks either because of limitations on downlink bandwidth, or because a decision is made to take more target images at the expense of concurrent darks. In either case, the same pipeline routine will be used to debias the image (*lorri_debias*), but the algorithm employed is different in each case and different reference files are required.

If in-flight data indicate that bias images are stable over time, many bias images will be combined (after filtering out clearly discrepant pixels) to produce a “super-bias” image. Then the median value of the inactive region of the image (i.e., the median of a 1024 row by 4 column region) is subtracted from the super-bias image to produce a “delta-bias” image. The IDL procedure that produces the delta-bias image is called *lorri_delta_bias*, but this routine is *not* part of the standard LORRI calibration pipeline; rather, it is an ancillary routine used to produce a calibration reference file.

The delta-bias image will exhibit the pixel-to-pixel variation in the bias and will oscillate about zero. The bias subtraction for any new image is then a two-step process:

- 1) The median signal level in the inactive region of the image is subtracted from each pixel’s value to remove the overall bias level, and
- 2) The delta-bias image is subtracted from the image created in the previous step to remove the pixel-to-pixel variation and produce the final, debiased image.

Ground calibration testing showed that the overall bias level in step (1) above depends on the signal level in the last few columns of the active region of the CCD. The effect is produced by amplifier undershoot, which means that the bias level recorded by the pixels in the inactive region is smaller than the actual bias level. The magnitude of the effect depends on the signal

level in the active region and on the column number in the inactive region and can be as large as ~12 DN. Thus, prior to computing the median signal in the inactive region (step 1 above), the intensities of all the pixels must be corrected for amplifier undershoot. This correction step is incorporated into the *lorri_debias* procedure.

If the in-flight bias images vary significantly in time, separate bias images (i.e., 0 ms exposures) must be taken for each science image obtained. In this case, the bias subtraction proceeds exactly as performed during ground calibration testing, with the bias removal achieved by simple subtraction of the bias image from the science image. There are several drawbacks to this approach: (1) more images must be taken, which affects the data volume that must be stored on the on-board solid-state recorder, (2) more data must be downlinked, which may not be possible because of limited downlink bandwidth and/or the cost associated with the extra Deep Space Network (DSN) support required, (3) the signal-to-noise ratio (SNR) may be degraded because the bias subtraction no longer involves a high SNR reference file, and (4) fewer science images can be obtained because they have been displaced in the observing timeline by extra bias images.

Some improvements were put into place just after the Pluto encounter that adjusted how the bias subtraction was performed in the presence of missing data. The general steps of the algorithm were not changed but were altered to consider only pixels from the image that were populated with valid data. The pixels corresponding to missing data are excluded from subsequent processing steps and their values are set to 0 DN.

9.3.1.2 Smear Removal

LORRI does not have a shutter, so the target being observed illuminates the active region of the CCD whenever LORRI is pointed at the scene. In particular, the CCD continues to record the scene as the charge is transferred from the active portion to the storage area, and this results in a smearing of the observed scene. Fortunately, this smear can be removed to high accuracy using the correction algorithm described below.

When bright objects are observed, the readout smear makes the Level 1 image difficult to use for analysis purposes. In the image of Jupiter in Figure 9-2, the Level 1 image is on the left and the Level 2 calibrated image with readout smear (aka frame transfer smear) removed is on the right.

The need for the readout smear removal arises from the operation of the frame transfer CCD used in LORRI, where first the image zone is flushed, then an exposure is taken, and finally the image is transferred into the storage zone. Hence a pixel of the Level 1 image is exposed to the scene radiance from the corresponding geometrical element of the scene, but it is also exposed to the radiances of all the scene elements in the same image column during the image transfers. Thus the Level 1 image is the superposition of the scene radiance and the signal acquired during frame transfers, which is called readout smear

$T_{exp} = 2 \text{ ms}$, Jupiter diameter = 517 pixels, 2007 January 24




Figure 9-2: Demonstration of Smear Removal

The readout smear is removed as follows. Let $P_{i,j}^{meas}$ = measured image array in DN where i,j are the column and row indices, respectively. Let the exposure time be written T_{exp} , with the transfer times for the frame scrub T_{f1} and the frame storage T_{f2} and with N the number of rows (which is 1024 for 1×1 images and 256 for 4×4). Let $T_{fav\bar{g}}$ be the average of T_{f1} and T_{f2} to define the constant

$$A = \frac{T_{exp}}{T_{exp} - \frac{T_{fav\bar{g}}}{N}}$$

Finally we define the $N \times N$ constant matrix

$$\varepsilon_{k,j} = \begin{cases} T_{f1}/T_{fav\bar{g}} & \text{for } k < j \\ 1 & \text{for } k = j \\ T_{f2}/T_{fav\bar{g}} & \text{for } k > j \end{cases}$$

with $k, j = 1, \dots, N$, and we calculate the $N \times N$ matrix:

$$\lambda_{i,j}^{(1)} * T_{exp} = A \left[P_{i,j}^{meas} - \frac{AT_{fav\bar{g}} \sum_k P_{i,k}^{meas} \varepsilon_{k,j}}{N(T_{exp} + AT_{fav\bar{g}})} \right]$$

The desmeared image is then

$$P_{i,j}^{\text{desmear}} = A \left[P_{i,j}^{\text{meas}} - \frac{AT_{\text{avg}} \left[\sum_k P_{i,k}^{\text{meas}} \varepsilon_{k,j} + \frac{E_{i,j}}{A} \right]}{N(T_{\text{exp}} + AT_{\text{avg}})} \right]$$

with

$$\frac{E_{i,j}}{A} = T_{\text{avg}} \left[\sum_k \lambda_{i,k}^{(1)} \varepsilon_{k,j} - \frac{1}{N} \sum_l \sum_k \lambda_{i,k}^{(1)} \varepsilon_{k,l} \varepsilon_{l,j} \right]$$

In-flight tests have verified desmear by this technique using observations of Jupiter obtained at exposure times as short as 1 ms.

The value for T_{avg} is dependent on the desired exposure time and has been determined empirically using in-flight data. The following Table 9-4 provides the appropriate values at different exposure times.

Table 9-4: Value for T_{avg} from T_{exp}

Desired Exposure Time (msec)	Value for T_{avg} (msec)
1	7.1
2	8.75
3	9.65
6	10.5
Nominal	10.7

It should be noted that when the Level 1 data are saturated, the resulting readout smear correction will be inaccurate. The algorithm relies on an accurate accumulation of charge in all rows of each column and if the Level 1 data are clipped for lack of dynamic range to capture that integrated signal, the effect of readout smear cannot be completely and properly removed.

In the presence of missing data, calculation of the readout smear is not possible since all rows are not available. A best effort can be made by estimating the missing pixel values using simple interpolation on a per-column basis. This is only done in an attempt to provide an estimate of the readout smear. The estimated pixel values are not used in any other way. If the missing data are bound by valid pixel data, linear interpolation is used. If missing data are at the edge of the detector, the nearest pixel value in the row is used. These values are used in place of the measured values indicated in the formulas above. After the readout smear is removed, the pixel locations of missing data are set to 0 DN and the appropriate value in the quality flag is set to represent missing data.

This correction algorithm has been implemented in the IDL routine *lorri_desmear*.

9.3.1.3 Flat-Fielding

Flat-fielding refers to the process of removing the pixel-to-pixel sensitivity variations in the image. An exposure obtained by illuminating the LORRI aperture uniformly with light is called a “flat-field” image. During ground calibration testing, flat-fields were obtained by using an “integrating sphere” to provide uniform illumination. The light source was a xenon arc lamp with a spectrum similar to that of the sun. The absolute intensity of the input illumination was measured using a calibrated photodiode. For the panchromatic case, which is the one most relevant for flat-fielding LORRI images, the light from the xenon lamp was unfiltered. Flat-field images were also obtained by passing the light through bandpass filters centered at five different wavelengths spanning the range over which LORRI is sensitive, prior to injection into the reference sphere, in order to estimate the sensitivity of the flat-fields to the spectral distribution of the source. The spatial patterns in the flat-field images change fairly dramatically with wavelength. However, the variation in panchromatic flat-fields caused by differences in the spectral distribution of the illumination source should be much less significant. Indeed, panchromatic flat-field images produced using a tungsten lamp were virtually indistinguishable from those produced by the xenon lamp. Flat-fields were obtained at four different telescope temperatures (at standard laboratory room temperature, and at the lowest, nominal, and highest temperatures predicted for in-flight conditions), but no significant temperature variations in the flat-field images were detected.

The flat-field reference file used in the LORRI pipeline was produced by averaging 100 flat-field images taken at room temperature using the xenon arc lamp as the light source, debiasing and desmearing the average image as described earlier, and normalizing the intensities in the active region to a median value of 1. If “S” (units are DN) is an image of a target that has already been desmeared and debiased, and if “FF” is the reference flat-field image, then the flat-fielded (i.e., photometrically-corrected) target image (“C”; units are DN) is given by:

$$C = S/FF$$

The flat-fielding correction is implemented in the LORRI pipeline by the routine *lorri_flatten*.

If in-flight measurements indicate that the LORRI flat-field characteristics are different than those measured during ground calibration tests, new reference flat-field images must be obtained. Although LORRI has two internal reference lamps (sometimes referred to as “cal lamps”), the illumination pattern is highly non-uniform and, thus, not very suitable as a secondary flat-field standard. Various test measurements will be performed during the early portion of the mission to determine if scattered sunlight can serve as a suitable secondary flat-field standard. After a Jupiter encounter, smeared images of Jupiter might also prove to be useful as a secondary flat-field standard. In any case, there will be an attempt to monitor the flat-field characteristics of LORRI over time, and the reference flat-field image used by the LORRI pipeline will be updated as necessary to maintain an accuracy better than 1% in the correction of the pixel-to-pixel sensitivity variation, except possibly near the center of the field where image ghosts may compromise the quality of the reference flat-field (see further discussion below).

During ground calibration tests, intensity artifacts caused by optical ghosts were observed near the center (roughly covering a 200 x 200 pixel region) of the flat-field images. Ray tracing of the

optical system indicates that the intensity of the ghost image should be less than ~1% of the intensity produced by the direct illumination, but measurements indicated that ghost intensities have an amplitude of ~5-7% of the direct intensity for panchromatic illumination. The ghost intensity is scene-dependent with most (~80%) of the ghost signal arising from regions outside the nominal field-of-view of LORRI. There is a suspicion that at least some of the ghost signal is an artifact of the test conditions, and the reference flat-fields currently used by the pipeline do *not* include the ghost signal produced by the out-of-field light. Any flat-field data taken in-flight will be carefully scrutinized to search for any effects attributable to optical ghosts. Depending on those results, further modifications to the reference flat-fields may be required. There is also the possibility that different flat-field reference images may be required depending on the scene being imaged (i.e., a ghost subtraction step may be required prior to application of the flat-field correction under some circumstances).

9.3.1.4 Absolute Calibration (Conversion from corrected DN to physical units)

The calibration software pipeline does not perform per-pixel conversion from DN to physical units, because that conversion requires knowledge of the spectral distribution (i.e. color) of the target. Instead, the pixel values are left in their **corrected (calibrated) DN units** from the flat-fielding, and the absolute calibration step adds “photometry” keyword divisors to the FITS header (see Section 9.3.9). Each divisor scales per-pixel count rate values from DN/s/pixel to physical units according to its assumed spectral type and spatial distribution (diffuse vs. point source) for the target.

Photometry keywords are provided for targets having spectral distributions similar to Pluto, Charon, Pholus, Jupiter, and the Sun. The units adopted for the **radiance of diffuse targets** (a.k.a. "intensity") are ergs/cm²/s/sr/Å. The units adopted for the **irradiance of point (i.e., unresolved) targets** (a.k.a. "flux") are ergs/cm²/s/Å. Tables providing the values for the photometry keywords at the time of launch are given below. The latest (i.e., current) values of the photometry keywords are provided in the header of the calibrated image FITS file for the image being analyzed.

The absolute calibration is achieved by specifying a keyword (e.g. RPLUTO) in the header of the calibrated image file that allows the user to convert a count rate (“C/TEXP” in DN/s/pixel, where “C” is the flat-fielded signal in a pixel, **in DN units**, and “TEXP” is the exposure time) for a resolved source into a radiance value (“I” in ergs/cm²/s/sr/Å) at LORRI’s pivot wavelength (specified by the FITS keyword PIVOT; see below), assuming that the spectrum of the target is identical to the globally-averaged spectrum of Pluto. The relevant formula is:

$$I = C/TEXP/RPLUTO$$

Similarly, the keyword RSOLAR allows the conversion of the count rate for a resolved source into a radiance value at the pivot wavelength assuming that the target has a solar-like spectral distribution:

$$I = C/TEXP/RSOLAR$$

Finally, the keyword RPHOLUS allows the conversion of the count rate for a resolved source into a radiance value at the pivot wavelength assuming that the target has a spectral distribution

identical to that of the centaur object 5145 Pholus, which may be a good analog for the reddest regions on Pluto:

$$I = C/\text{TEXP}/\text{RPHOLUS}$$

Table 9-5 provides estimates for these radiance sensitivity keywords for 1x1 mode images. In-flight calibration observations of photometric standard stars will be used to verify these values and to monitor them over time, and the SOC pipeline will have latest values. For 4x4 mode images, use the 1x1 value times 19.2 (1.2 times a factor of 16 to account for the 4x4 pixel grouping).

Table 9-5: Estimates for Radiance Sensitivity Keywords in 1x1 Mode

Keyword	Value [(DN/s/pixel)/(ergs/cm ² /s/sr/Å)]
RSOLAR	2.664 x 10 ⁵
RPLUTO	2.575 x 10 ⁵
RCHARON	2.630 x 10 ⁵
RJUPITER	2.347 x 10 ⁵
RPHOLUS	3.243 x 10 ³

Conversions for other source spectral distributions must be derived by integrating the product of the LORRI monochromatic response function and the other source spectrum, as described in the LORRI Space Science Reviews and SPIE papers (for references see Section 9.4).

The pivot wavelength (PIVOT) is given by:

$$\text{PIVOT} = \sqrt{\frac{\int P\lambda d\lambda}{\int P d\lambda/\lambda}}$$

where “P” is the LORRI system quantum efficiency (i.e., fraction of photons detected) at wavelength “λ”. The current best estimate for the LORRI pivot wavelength is 6076 Å.

For unresolved sources (e.g., stars), the absolute calibrated flux (also called “irradiance”) at the pivot wavelength can be determined using keywords that are defined analogously to the photometry keywords discussed above for resolved sources. In the case of an unresolved source having a spectral distribution identical to that of a globally-averaged Pluto spectrum, the observed count rate integrated over the LORRI PSF (“CINT/TEXP” in DN/s, where CINT is the total number of flat-field corrected counts integrated over the image and “TEXP” is the exposure time) can be related to the flux (“F” in ergs/cm²/s/Å) by:

$$F = \text{CINT}/\text{TEXP}/\text{PPLUTO}$$

Similarly, the flux at the pivot wavelength for an unresolved target having the same spectral distribution as the sun is given by:

$$F = \text{CINT}/\text{TEXP}/\text{PSOLAR}$$

And the flux at the pivot wavelength for an unresolved target having the same spectral distribution as 5145 Pholus is given by:

$$F = \text{CINT}/\text{TEXP}/\text{PPHOLUS}$$

Table 9-6 provides estimates for these irradiance sensitivity keywords for 1x1 mode images. In-flight calibration observations of photometric standard stars will be used to verify these values and to monitor them over time, and the SOC pipeline will have latest values. For 4x4 mode images, use the 1x1 value times 16 to account for the 4x4 pixel grouping.

Table 9-6: Estimates for Irradiance Sensitivity Keywords in 1x1 Mode

Keyword	Value [(DN/s)/(ergs/cm ² /s/Å)]
PSOLAR	1.066 x 10 ¹⁶
PPLUTO	1.030 x 10 ¹⁶
PCHARON	1.052 x 10 ¹⁶
PJUPITER	9.386 x 10 ¹⁶
PPHOLUS	1.297 x 10 ¹⁶

Synthetic photometry techniques can be used to convert the fluxes derived in the manner described above to fluxes at other wavelengths, and then into standard *UBVRI* magnitudes in the Landolt (1992) photometric system, which is essentially identical to the Johnson *UBV* system combined with the Kron-Cousins *RI* system. The results described in the LORRI calibration paper can be used to derive fluxes for targets whose spectral distributions do not match the three cases discussed above.

We provide below some examples showing how to convert from engineering units to physical units, for both diffuse and point targets.

Consider a diffuse target whose spectrum is similar to that of Pluto. You should then use the RPLUTO photometry keyword from the header of the calibrated image file to convert a count rate (“C/TEXP” in DN/s/pixel, where “C” is the flat-fielded signal in a pixel and “TEXP” is the exposure time) into a radiance value (“I” in ergs/cm²/s/sr/Å) at LORRI’s “pivot” wavelength (specified by the FITS keyword PIVOT for the formal definition of the pivot wavelength):

$$I = \text{C}/\text{TEXP}/\text{RPLUTO}$$

Similarly, the photometry keywords RSOLAR, RCHARON, RJUPITER, and RPHOLUS should be used to convert count rates into radiance values at the pivot wavelength assuming that the target has, respectively, solar-like, Charon-like, Jupiter-like, or Pholus-like spectral distributions.

For LORRI, the pivot wavelength is 6076.2 Å, and we don't expect this to change, at least not significantly. Since the solar flux (F_{solar}) at a heliocentric distance of 1 AU at the pivot wavelength is 176 erg/cm²/s/Å, the value for the radiance can be converted to I/F (where $\pi * F = F_{\text{solar}}$) using:

$$I/F = \pi * I * r^2 / F_{\text{solar}}$$

where "r" is the target's heliocentric distance in AU.

In the case of an unresolved target having a spectral distribution identical to that of a globally-averaged Pluto spectrum, the observed count rate integrated over the LORRI PSF ("CINT/TEXP" in DN/s, where CINT is the total number of flat-field corrected counts integrated over the image and "TEXP" is the exposure time) can be related to the flux ("F" in ergs/cm²/s/Å; not to be confused with "F" in I/F) by:

$$F = \text{CINT/TEXP/PPLUTO}$$

When observing point targets, it is more common to convert the absolute flux to a magnitude in a standard photometric system. The following equation can be used to transform a measured value of the irradiance (aka "flux") of an unresolved target to a magnitude in the standard V band:

$$V = -2.5 \log S + \text{PHOTZPT} + \text{CC} + \text{BC}$$

where "V" is the visual magnitude in the Johnson photometric system, PHOTZPT is the "stellar photometry keyword", which is the "zero point" of the LORRI instrumental magnitude system, "S" is the integrated net signal rate from the target in DN/s, "CC" is the color correction (i.e., correction for the spectral distribution of the target), and "BC" is the aperture correction (in case the flux is not integrated over the entire stellar image; a careful analysis of the flux versus aperture size for a bright star in the field can then be used to determine the value of BC for the aperture selected for the photometry).

In-flight photometry of stars in the open galactic cluster M7 yield the following:

$$\text{PHOTZPT} = 18.94$$

Table 9-7: Color correction coefficient for various targets

Spectral Type	CC
O, B, A stars	-0.06
F, G stars	0
K stars	+0.4
M stars	+0.6
Pluto	-0.037
Charon	-0.014
Jupiter	-0.138
Pholus	+0.213

The following reference flux information is provided for convenience and was gathered from several sources. The UBV are in the Johnson system, RI are in the Landolt-Kron-Cousins system, and JHK_sK are in the UKIRT system.

The fluxes for Vega are from the model STScI absolute calibrated spectrum. At near-IR wavelengths, the model underestimates the actual Vega flux by about 5-6% owing to the excess flux from the Vega dust disk. Note also that Vega has $U=B=V=0.03$ (i.e., not 0).

Table 9-8: Fluxes for Vega

Band	Center (Å)	Vega Flux (ergs/cm ² /s/Å)
U	3600	3.05×10^{-9}
B	4400	6.74×10^{-9}
V	5500	3.54×10^{-9}
R	6500	2.11×10^{-9}
I	8000	1.12×10^{-9}
J	12200	3.18×10^{-10}
H	16540	1.11×10^{-10}
K _s	21570	4.10×10^{-11}
K	21790	3.97×10^{-11}

9.3.1.5 Pointing Information

Pointing information for the LORRI boresight (center of the LORRI field-of-view, which is pixel [511,511]) is included in the FITS header in both the Level 1 and the Level 2 calibrated image files. An example of this information follows:

```
SPCBLRA = 233.4199004768138 / [degrees] Boresight RA, EME J2000
SPCBLDEC= -17.96897170490819 / [degrees] Boresight DEC, EME J2000
SPCEMEN = 283.935414259362 / [degrees] EME J2k North Clk Angle, CW from UP
```

There are also WCS (World Coordinate System) and SIP (Simple Imaging Polynomial for distortion) keywords in the FITS header that define the pointing of each pixel; these are typically interpreted by existing application software.

9.3.1.6 Conversion of instrument housekeeping items to engineering units

The LORRI-specific housekeeping items reported in the Level 1 FITS file are in units of counts or DN. To make these values more useful for data analysis, they have been converted to engineering units (volts, amps, degrees Celsius) and reported at the tail end of the header of the primary HDU of the calibrated FITS file. Because the contents of the Level 1 header are duplicated in the calibrated file, a different set of tag names are used for the values that have been converted to engineering units. The new tags are reported after the comment that reads "LORRI Level 2 Calibrated telemetry items".

9.3.2 Instrument Characterization

There are several characteristics of the instrument that are related to the radiometric calibration of LORRI that will be useful when analyzing the calibrated image data. They are the quantum efficiency and spectral responsivity, each as a function of wavelength. There is a table for each of these in the CALIB/ or DOCUMENT/ directory of the PDS archive data sets, but a graph for each is reproduced in the figures below. **N.B. there has been some confusion about the ordinate units in some published versions of this plot; the units of $(DN\ s^{-1}) / (W\ cm^{-2}\ sr^{-1})$, both here and in the pre-prints of LORRI-related papers in PDS data sets produced ca. December, 2015, are correct.**




Figure 9-3: LORRI Spectral Response vs. Wavelength, adapted from Cheng et al. (2008)




Figure 9-4: LORRI Quantum Efficiency vs. Wavelength

9.3.3 Geometric Distortion

The geometric distortion on the LORRI camera is quite small at less than 2 pixels in the corners of the detector. There are two models used by the LORRI team. The first model was developed by Bill Owen and has been used on many previous NASA missions. The second model is called the Simple Imaging Polynomial (SIP) distortion model and is commonly used within the astronomical community. Although these are considered separate models, the forward transformation for Bill Owen's model is mathematically equivalent to the reverse transformation for the SIP model. Both models have been thoroughly documented in the LORRI Instrument Kernel (IK). The coefficients for these models are updated periodically using in-flight observations. Because of their transient nature, the coefficients for these models are best captured in the LORRI IK, which is included as a regular part of the Planetary Data Systems (PDS) delivery schedule for the mission.

9.3.4 Special Processing

After the data have been calibrated, additional processing steps are likely to be required. Obvious examples of this are ghost removal and stray light processing. At present, there have been no algorithms developed for public release because they are highly scene dependent. Individual images must be analyzed to understand the structure of the effects to determine an appropriate method for its removal. In the example below, a cutout from a calibrated image is presented to illustrate the effect of stray light from Jupiter's disk, which is just out of the field of view. The circular structure is an example of the ghost pattern. The image on the right demonstrates the processed version of that image. The gradient from the stray light has been removed, as well as the majority of the effects of the ghost.




Figure 9-5: Example of Special Processing of Calibrated Data

9.3.5 Dataflow Block Diagram




Figure 9-6: LORRI Pipeline Processing Dataflow

9.3.6 Data Format

The calibrated image data are organized in a FITS file. The primary header and data unit (HDU) is used to store the calibrated image that results from the calibration pipeline. The first extension is the error estimate image, followed by the second extension containing the data quality image. Table 9-9 below contains a description of the layout for the extensions for calibrated data.

For 1x1 binning mode, the calibrated image dimensions are 1024x1024, and for 4x4 mode, the dimensions are 256x256 pixels. In both situations, these pixels correspond to the optically active pixels from the Level 1 image mentioned previously.

Table 9-9: Calibrated FITS file extension layout

FITS File Storage Location	Description
Primary HDU	Calibrated image
First Extension	Error image
Second Extension	Data Quality Image

9.3.7 Extra FITS Extensions (planes) and Their Definitions

LORRI calibrated FITS files have 3 extensions. The debiased, desmeared LORRI image is written into the primary HDU as a 2-dimensional, 32-bit real image. The unit for each data value

is photometrically-corrected DN. The estimated errors in these corrected DN values are stored as a 2-dimensional, 16-bit real image in the first extension. A data quality image is stored in the second extension as a 2-dimensional, 16-bit integer image.

The error in the photometrically-corrected signal is estimated from:

$$\sigma = \frac{\sqrt{(P_{meas}/g) + (RN)^2 + (fP_{meas})^2}}{FF}$$

where “σ” is the 1-sigma error in the corrected signal for a particular pixel (DN), “P_{meas}” is the observed signal in that pixel (DN, after bias subtraction but before smear removal), “g” is the electronics gain (22 e/DN), “RN” is the electronics noise (1.3 DN), “f” is the estimated error in the reference flat-field image (0.005), and “FF” is the value of the reference flat-field image at the relevant pixel. The above formula neglects any noise contributed by the bias and smear removal steps, but those errors are generally expected to be small compared to the other sources of error.

The data quality image is used to flag pixels that have known artifacts and may need special consideration when performing scientific analysis. The pixel value in the quality flag image represents the sum of all quality flags present for that pixel. This pixel value can also be described as the result of the bitwise ‘OR’ of each quality flag value. The list of data quality values and their descriptions are listed in Table 9-10.

Table 9-10: Quality flag value descriptions

Quality Flag Value	Bit position in 2-byte word	Description
0	n/a	Good pixel
1	0	Defect in reference deltabias image (set if 0 or NaN)
2	1	Defect in reference flatfield image (set if 0 or NaN)
4	2	Permanent CCD defect; pixel>0 in dead_ground_NxN.fit map
8	3	Hot Pixel identified; pixel>0 in hot_ground_NxN.fit map
16	4	Saturated pixel in Level 1 data (A/D value of 4095)
32	5	Missing Level 1 data (assume fill value of 0)
64 and higher	6-15	unused at present

9.3.8 Scientific Units

Following the convention adopted by the *New Horizons* Principal Investigator, the unit used for a calibrated data product is “photometrically-corrected DN”. The procedure given above must be

completed to obtain absolute calibrated data products. The units adopted for the radiance (aka “intensity”) of diffuse targets are $\text{ergs/cm}^2/\text{s}/\text{sr}/\text{\AA}$. The units adopted for irradiance (aka “flux”) of point (i.e. unresolved) targets are $\text{ergs/cm}^2/\text{s}/\text{\AA}$. Wavelengths are quoted in angstrom units.

9.3.9 Additional FITS and PDS Keywords Added

Listed below are the keywords and sample values for those keywords that have been added to the FITS header and are stored with the primary HDU of the output calibrated image FITS file.

```

COMMENT *****
COMMENT *** LORRI Level 2 software name and version info ***
COMMENT *****
L2_SWNAM= 'lorri_level2_pipeline' /Level 2 calibration software
L2_SWVER= 'untagged' /software version tag
COMMENT *****
COMMENT *** LORRI Level 2 software logic flow control flags ***
COMMENT *****
IMGSUBTR= 'OMIT ' / image subtraction step
BIASCORR= 'PERFORM ' / bias subtraction step
SLINCORR= 'OMIT ' / signal linearization step
CTICORR = 'OMIT ' / charge transfer inefficiency step
DARKCORR= 'OMIT ' / dark subtraction step
SMEARCOR= 'PERFORM ' / smear removal step
FLATCORR= 'PERFORM ' / flat-fielding step
GEOMCORR= 'OMIT ' / geometric correction step
ABSCCORR= 'PERFORM ' / absolute calibration step
COMPERR = 'PERFORM ' / compute error estimate
COMPQUAL= 'PERFORM ' / compute quality flags
COMMENT *****
COMMENT *** LORRI Level 2 Reference Filename ***
COMMENT *****
REFDEBIA= 'sap_006_combined_100img_1x1.fit' / debias image filename
REFFLAT = 'cflat_grnd_SFA_20050309_v2.fit' / flat field image filename
REFDEAD = 'dead_ground_1x1_synthetic.fit' / dead pixel image filename
REFHOT = 'hot_ground_1x1_synthetic.fit' / hot pixel image filename
REFSUBIM= ' ' / subtraction image filename
COMMENT *****
COMMENT *** LORRI Level 2 Absolute Calibration Parameters ***
COMMENT *****
PIVOT = 6076.20019531 / LORRI pivot wavelength. units=angstroms
RSOLAR = 266400.000000 / Conv to radiance for solar source
RPLUTO = 257500.000000 / Conv to radiance for pluto source
RPHOLUS = 324300.000000 / Conv to radiance for 5145 pholus source
RCHARON = 263000.000000 / Conv to radiance for charon source
RJUPITER= 234700.000000 / Conv to radiance for jupiter source
PPLUTO = 1.03000005170E+16 / Conv to irradiance for pluto source
PSOLAR = 1.06600003807E+16 / Conv to irradiance for solar source
PPHOLUS = 1.29700002225E+16 / Conv to irradiance for 5145 pholus source
PCHARON = 1.05199994793E+16 / Conv to irradiance for charon source
PJUPITER= 9.38600033786E+15 / Conv to irradiance for jupiter source
PHOTZPT = 18.9400000000 / Zero point for visual magnitude, V
COMMENT *****
COMMENT *** LORRI Level 2 Calibrated telemetry items ***
COMMENT *****
EPU_P5VO= 5.04305504857 / EPU +5 voltage. units=Volts
EPU_P5CU= 0.143143000000 / EPU +5 current. units=Amps

```

```

FPU_P15V=      15.0005851594 / FPU +15 voltage. units=Volts
FPU_P15C=      0.049382700000 / FPU +15 current. units=Amps
FPU_P6_V=      6.05666080780 / FPU +6 voltage. units=Volts
FPU_P6_C=      0.152152000000 / FPU +6 current. units=Amps
FPU_HTRC=      0.000000000000 / FPU heater current. units=Amps
EPU_25PV=      2.50943456804 / EPU +2.5 voltage. units=Volts
RINGTEMP=     -66.8836898878 / Intermediate ring temp. units=celsius
MFOOTTMP=     -61.8964797242 / Mounting foot-top temp. units=celsius
M2MNTTMP=     -66.8836898878 / M2 mirror mount temp. units=celsius
RADTEMP =     -88.9564863168 / Radiator temp. units=celsius
BAFATEMP=     -62.9653774259 / Baffle-aft temp. units=celsius
BAFFTEMP=     -70.8007466057 / Baffle-forward temp. units=celsius
M1SUPTMP=     -67.2398321861 / M1 mirror support temp. units=celsius
M1MIRTMP=     -66.5275372052 / M1 mirror temp. units=celsius
CCDTEMP =     -79.5485000000 / CCD temperature. units=celsius
M1VFTMP=     -66.0128183000 / M1 V/F temperature. units=celsius
M2VFTMP=     -66.3287025000 / M2 V/F temperature. units=celsius
FPUBTEMP=     29.5499120000 / FPU board V/F temp. units=celsius
STEMPCVR= 'ENABLE ' / Temperature conversion enable
SCLMP2PE= 'OFF ' / Cal lamp 2 power enable
SCLMP1PE= 'OFF ' / Cal lamp 2 power enable
SSOURCE = 'CCD ' / Image source
SFORMAT = '1X1 ' / Image format
SEXPMODE= 'MANUAL ' / Exposure mode
PDUNAME = 'Level 2 LORRI image' /

```

9.3.9.1 Reading FITS file contents using IDL

The main method for accessing the various extensions and headers from the FITS file within IDL rely on a third-party library known as the Goddard Astron library. From within IDL, one can load the primary HDU from a fits file using the following command:

```

IDL> calimg=readfits('lor_0035015237_0x630_sci_1.fit', hdr )
IDL> help, calimg
CALIMG      FLOAT   = Array[1024, 1024]

```

The return value of this function (“calimg”) is a two dimensional array containing the image data from the primary HDU and its type depends on the data that is read from the file. In the case of Level 1 data, it will be a 16-bit integer array and for calibrated data, it will be a 32-bit floating-point array. The first argument in the call to readfits() is the name of the FITS file to be read. The second argument is an ASCII string variable that will contain the FITS header for the primary HDU upon completion of the function.

The same function may be used in order to read any of the extensions listed in the files. For example, to read the data quality image from the calibrated FITS file, one would use a statement such as:

```

IDL> quality=readfits('lor_0035015237_0x630_sci_1.fit', hdr2, exten_no=2)
IDL> help, quality
QUALITY     UINT    = Array[1024, 1024]

```

In this example, the ASCII string variable “hdr2” contains the FITS header associated only with the second extension and has no portion of the header from the primary HDU.

9.3.10 Hardware/OS Development Platform

The pipeline software was developed in a variety of environments with the commonality of UNIX-style operating systems. There are no dependencies on the endian properties of the environment.

9.3.11 Language(s) Used

IDL

9.3.12 Third Party Libraries Required

There are two third party IDL libraries that are needed by the calibration pipeline software:

- 1) Goddard Astron library, which contains routines needed to read and write FITS files, the format used by the Level 1 data files. Because this library is provided by the SOC for use by many instruments, we will not be delivering this library, but will rely on the version provided to us.
- 2) IDLUSR, a collection of useful IDL routines made available for public release at APL. Information about this library can be found at fermi.jhuapl.edu.

9.3.13 Calibration Files Needed (with Quantities)

There are currently five categories of reference files needed to perform the calibration process. The reference image categories are the delta-bias, flat-field, dead pixel, hot pixel and desmear e-matrix. Because the LORRI instrument can produce images in either 1024 x 1024 mode or 256 x 256 mode, there are two varieties of each of these images. The filenames associated with these images will be obvious by inspection, although no formal file naming convention has been adopted.

There are two ASCII description files in the calibration directory that don't qualify as calibration files but are related to the operation of the pipeline. The first is a configuration file that details all of the configuration parameters for the pipeline (“default_config.txt”). The other file is a description of the housekeeping items that are stored in the first 34 pixels (51 bytes) of the Level 1 image data (“binary_lorri_image_hdr.txt”). These values can be used to validate the FITS header tags that were produced by associating the high-speed image data with the low-speed telemetry values. The values in the first 34 pixels are guaranteed to be correctly associated with a particular image (provided they were not compressed in a lossy fashion) because the LORRI ASE put them in place prior to the transfer of the image data to the SSR. As such, they represent a valuable check of the telemetry processing performed on the ground after receipt.

The following is a table of the types of files in the calibration directory:

Table 9-11: File Types in Calibration Directory

Description	Quantity	1x1 filesize	4x4 filesize
delta bias	2	~ 8 MiB	~0.5MiB
flat field	1	~ 8 MiB	~0.5MiB
hot pixel map	1	~ 8 MiB	~0.5MiB
dead pixel map	1	~ 8 MiB	~0.5MiB
desmear e-matrix	1	~ 8 MiB	~0.5MiB
pipeline configuration file	1	~5KiB	
LORRI header description	1	~4KiB	

9.3.14 Memory Required

~ 100 MiB

9.3.15 Temporary File System Space Needed

None

9.3.16 Predicted Size of Output File(s)

Table 9-12: Size of Output Files

Image dimensions	Binning	binmode	Expected File Size
1024 x 1024	1x1	0	~ 10.5 MiB
256 x 256	4x4	1	~ 700 KiB

9.3.17 Predicted Execution time

Less than 5 seconds per image.

9.3.18 Contact/Support Person(s)

Level 1 data support: Howard Taylor, John Hayes, and Hal Weaver

Level 2 calibrated data support: Howard Taylor and Hal Weaver

9.3.19 Maintenance Schedule (Code/Data Updates, Documentation)

As in-flight calibration data are collected and analyzed, certain aspects of the calibration pipeline will require updates, either in the form of updated reference files, or updated code for bug fixes or future improvements.

9.4 References

A. F. Cheng, H. A. Weaver, S. J. Conard, M. F. Morgan, O. Barnouin-Jha, J. D. Boldt, K.

A. Cooper, E. H. Darlington, M. P. Grey, J. R. Hayes, K. E. Kosakowski, T. Magee1 E.

Rossano, D. Sampath, C. Schlemm, H. W. Taylor, "Long-Range Reconnaissance Imager on *New Horizons*", *Space Sci. Rev.*, Volume 140, Numbers 1-4, pp. 189-215, 2008. **N.B. The published version of this paper has an error in the LORRI responsivity plot; refer to the PDS labels of pre-prints of this paper provided with LORRI data sets after 2015.**

S. Conard, F. Azad, J. Boldt, A. Cheng, K. Cooper, E. Darlington, M. Grey, J. Hayes, P. Hogue, K. Kosakowski, T. Magee, M. Morgan, E. Rossano, D. Sampath, C. Schlemm, and H. Weaver, "Design and fabrication of the *New Horizons* Long-Range Reconnaissance Imager," in *Astrobiology and Planetary Missions*, G. R. Gladstone, ed., Proc. SPIE 5906, 2005.

F. Morgan, S.J. Conard, H.A. Weaver, O. Barnouin-Jha, A.F. Cheng, H.W. Taylor, K.A. Cooper, R.H. Barkhouser, R. Boucarut, E.H. Darlington, M.P. Grey, I. Kuznetsov, T.J. Madison, M.A. Quijada, D.J. Sahnou, and J.M. Stock, "Calibration of the *New Horizons* Long-Range Reconnaissance Imager," in *Astrobiology and Planetary Missions*, G. R. Gladstone, ed., Proc SPIE 5906, 2005.

10. MVIC INSTRUMENT DESCRIPTION

10.1 Overview

The Ralph instrument consists of two sets of focal planes: MVIC (Multispectral Visible Imaging Camera) a visible, near-IR imager and LEISA (Linear Etalon Imaging Spectral Array), a short-wave IR spectral imager. This section only relates to the MVIC part of the Ralph instrument. LEISA is described in a different section (Section 8). There are 7 separate CCD arrays in the MVIC focal plane. The MVIC telemetry is communicated via a low-speed interface and the imaging data uses a high-speed interface.

Figure 10-1 shows a model of Ralph in the spacecraft coordinate system. The MVIC detector package is the light blue box on the +Y face of the instrument.




Figure 10-1: A model of Ralph in the spacecraft coordinate system. The Ralph aperture points in the $-X$ direction, the normal to the radiator is in the $+Z$ direction and the SIA points in the $+Y$ direction.

Data recorded on *New Horizons* is sent to the ground via the Deep Space Network. From there the data are sent to the Mission Operations Center (MOC) at the Applied Physics Laboratory (APL). The Science Operations Center (SOC) at the Southwest Research Institute (SwRI) in Boulder, Colorado, retrieves new data from the MOC daily. The data are in a raw form (packetized). The Level 1 and 2 SOC software pipelines convert the data from these raw packets into FITS (Flexible Image Transport) files with scientifically useful and calibrated data. The Level 1 processing sorts the packets into images (in the case of MVIC) with useful header keywords. These keywords include the mode or filter of the observation, timing information and

basic pointing information of the instrument boresight. The Level 1 processing also adds relevant housekeeping telemetry (temperatures, voltages, etc.) in a binary table as an extension to the FITS file. The Level 2 processing performs the basic scientific calibration, and is equivalent to CODMAC Levels 3/4.

Before we get into a description of the MVIC calibration, we will describe the image formats for each of the CCD arrays that comprise MVIC. There are 8 detectors in the Ralph instrument, and seven of those detectors are part of MVIC (the yellow and blue ones shown in Figure 10-2. The boresight information is summarized in Figure 10-2.

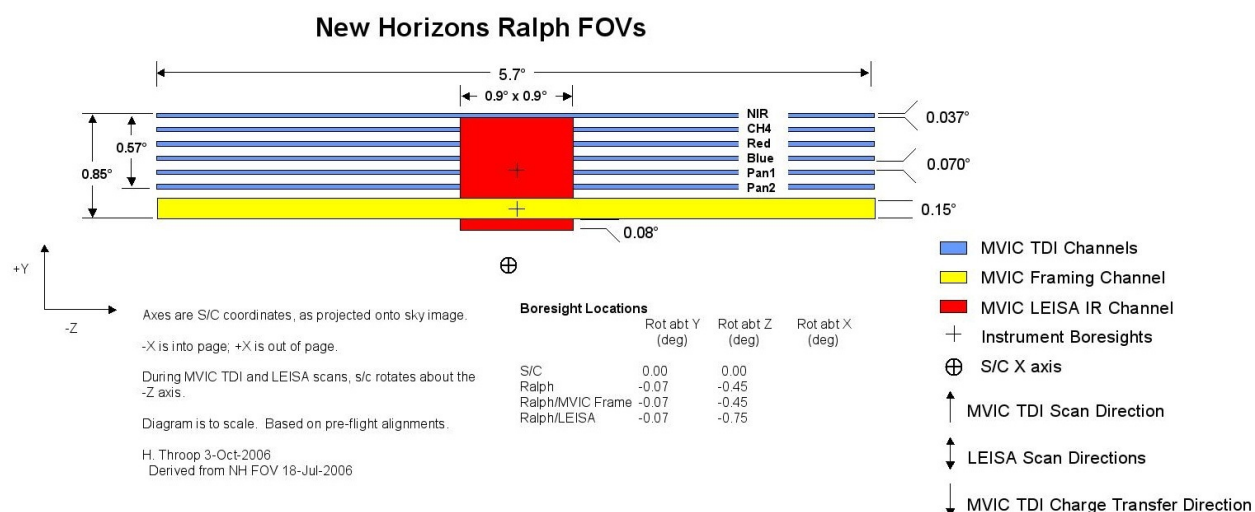


Figure 10-2: Ralph Fields-Of-View. MVIC focal planes are shown as blue (TDI) and yellow (Framing Channel) boxes.

The "pan frame" array (yellow box in Figure 10-2) has 5024x128 pixels. The first and last 12 pixels in each row are not optically active. When observing with the pan frame array, multiple images are recorded in each observation. Correspondingly, we store these images together in one FITS file. The first three letters of these files are "mpf" standing for MVIC pan frame. The images are stored as an image cube (3-dimensions). The first dimension is the column number, the second dimension is the row number and the third dimension is the image number within that observation.

The Pan Frame camera was originally the prime instrument for optical navigation. For optical navigation it is easiest to analyze data from a frame transfer array rather than a TDI array. During the proposal and development phase, the LORRI instrument was descopable which means that it wasn't guaranteed to be used for optical navigation, however, it was the preferred instrument because of the smaller iFOV and deeper imaging. So the purpose for the Ralph Pan Frame array was optical navigation.

The remaining 6 arrays (blue boxes in Figure 10-2) are operated in a time-delay integration mode (TDI). Each of these arrays have 5024x32 pixels. The first and last 12 pixels in each row are not optically active. To take an observation with the TDI arrays, the spacecraft is scanned and (typically) the charge is moved through each of the 32 rows at a rate that matches the spacecraft's scan rate. Using this method an arbitrarily long image in the row direction can be built up. For the "Pan 1" and "Pan 2" (panchromatic -- unfiltered) detectors, the resulting FITS files are standard 2-dimensional images. The first three letters of these files are either "mp1" or "mp2" corresponding to MVIC Pan 1 or MVIC Pan 2. Only one of these is operated at any time. The color arrays (NIR, CH4, Red and Blue) are operated together. The data for each filter is stored in separate FITS files. The first three letters for MVIC Color TDI are mc# where mc0, mc1, mc2, and mc3 correspond to Red, Blue, NIR and CH4, respectively.

In Table 10-1 the variable "Ni" stands for the number of images in a pan frame observation. When we command a pan frame observation, we always take multiple images. The "Nr" in the table is the number of rows in an observation. This number is determined by the length of time that we are recording data and the rate that we clock the rows in TDI mode.

Table 10-1: Observation Modes and their filename prefixes and data dimensions

Detector	Prefix for FITS file	Dimensions of data in FITS file
Pan Frame	mpf	3 (5024 x 128 x Ni)
Pan 1	mp1	2 (5024 x Nr)
Pan 2	mp2	2 (5024 x Nr)
Red	mc0	2 (5024 x Nr)
Blue	mc1	2 (5024 x Nr)
NIR	mc2	2 (5024 x Nr)
CH4	mc3	2 (5024 x Nr)

The Level 1 FITS file has a Primary Data Unit (PDU) that contains the raw data number (DN) pixel values formatted for the image (or image cube), plus 2 Extension Data Units (EDUs). EDUs 1 and 2 are binary tables. The housekeeping (HK) binary table in EDU 1 is image dependent as its content is a time series of housekeeping information (i.e. long scans will have more entries). "Nt" is the number of time entries for how many times it is sampled. The window mismatches binary table in EDU 2 is dependent on how many downlinked windows were merged to make the image, and how many inter-window pixel mismatches were found when merging those downlinked windows; "Nw" is the number of rows in the table. See PDS label description for WINDOW_MISMATCHES extension details; see also Table 10-2 and Table 10-3.

Table 10-2: Level 1 Pan Frame Image Data Format

FITS Data Unit	Dimension	Description
Primary	5024 x 128 x Ni (pixels x pixels x Ni)	Raw (DN) image cube
Extension 1	115 x Nt (bytes x Nt)	Binary table extension
Extension 2	61 x Nw (bytes x Nw)	Binary table extension

Table 10-3: Level 1 TDI Image Data Format

FITS Data Unit	Dimension	Description
Primary	5024 x Nr (pixels x pixels)	Raw (DN) image cube
Extension 1	115 x Nt (bytes x Nt)	Binary table extension
Extension 2	57 x Nw (bytes x Nw)	Binary table extension

The Level 2 FITS file has a Primary Data Unit (PDU) that contains the bias-subtracted, flattened image (or image cube in the case of pan frame) plus 2 extensions. Extension 1 is an array with the estimated per pixel error of the bias-subtracted, flattened, image (or image cube) (see Section 10.3.1.4). Extension 2 is an array with a data quality flag for each pixel of the bias-subtracted, flattened, image (or image cube) (see Section 10.3.1.5). See Table 10-4 and Table 10-5.

Table 10-4: Level 2 Pan Frame Image Data Format (pixels x pixels x Ni)

FITS Data Unit	Dimension	Description
Primary (PDU)	5024 x 128 x Ni	Bias-subtracted and flattened image cube
Extension 1	5024 x 128 x Ni	1-sigma error per pixel in Primary image
Extension 2	5024 x 128 x Ni	Data quality flag of Primary image

Table 10-5: Level 2 TDI Image Data Format (pixels x Nr)

FITS Data Unit	Dimension	Description
Primary (PDU)	5024 x Nr	Bias-subtracted and flattened image
Extension 1	5024 x Nr	1-sigma error per pixel in Primary image
Extension 2	5024 x Nr	Data quality flag of Primary image

10.2 Level 1 Data Specifics

10.2.1 Data Format

The Level 1 MVIC files are FITS files. Details of the dimensions of the FITS images and header keywords are given in the following sections.

10.2.2 Data Sources (*High/Low Speed, CCSDS, ITF*)

MVIC telemetry data are recorded using the low-speed interface and images are recorded with the high-speed interface. There are three Ralph-MVIC data formats used in flight:

- A. MVIC Panchromatic TDI (Time-Delay Integration) Format
- B. MVIC Color TDI Format
- C. MVIC Panchromatic Frame Format

The MVIC Panchromatic TDI format has two modes (either “Pan 1” or “Pan 2” indicating which detector was used).

Table 10-6 lists the MVIC ApIDs and their corresponding data types. For each data type there are two ApIDs, one for each C&DH (Command and Data Handling) system on the spacecraft. Compression is handled by the spacecraft C&DH system (compressed Lossless, compressed Lossy, or uncompressed (packetized)).

There is not a different ApID for “Pan1” and “Pan2” observations; however the SOC-generated filenames have different prefixes mp1 and mp2. This information is also stored in the low-speed housekeeping data (FITS keyword MODE). This information is inserted into the FITS header of the Level 1 FITS file. A value of 3 means that the observation was taken with the “Pan 1” array and a value of 4 indicates the “Pan 2” array was used. For more information, see the Ralph User’s Manual (SwRI Doc 05310.05-UM-01).

Table 10-6: ApID and Data Type for MVIC Data

ApID	Data Type
0x530	MVIC Panchromatic TDI Lossless (CDH 1)
0x53f	MVIC Panchromatic TDI Lossless (CDH 2)
0x531	MVIC Panchromatic TDI Packetized (CDH 1)
0x540	MVIC Panchromatic TDI Packetized (CDH 2)
0x532	MVIC Panchromatic TDI Lossy (CDH 1)
0x541	MVIC Panchromatic TDI Lossy (CDH 2)
0x533	MVIC Panchromatic TDI 3x3 Binned Lossless (CDH 1)*
0x542	MVIC Panchromatic TDI 3x3 Binned Lossless (CDH 2)*
0x534	MVIC Panchromatic TDI 3x3 Binned Packetized (CDH 1)*
0x543	MVIC Panchromatic TDI 3x3 Binned Packetized (CDH 2)*
0x535	MVIC Panchromatic TDI 3x3 Binned Lossy (CDH 1)*
0x544	MVIC Panchromatic TDI 3x3 Binned Lossy (CDH 2)*
0x536	MVIC Color TDI Lossless (CDH 1)
0x545	MVIC Color TDI Lossless (CDH 2)
0x537	MVIC Color TDI Packetized (CDH 1)
0x546	MVIC Color TDI Packetized (CDH 2)
0x538	MVIC Color TDI Lossy (CDH 1)
0x547	MVIC Color TDI Lossy (CDH 2)
0x539	MVIC Panchromatic Frame Transfer Lossless (CDH 1)
0x548	MVIC Panchromatic Frame Transfer Lossless (CDH 2)
0x53a	MVIC Panchromatic Frame Transfer Packetized (CDH 1)
0x549	MVIC Panchromatic Frame Transfer Packetized (CDH 2)
0x53b	MVIC Panchromatic Frame Transfer Lossy (CDH 1)
0x54a	MVIC Panchromatic Frame Transfer Lossy (CDH 2)

*The Panchromatic TDI Binned data are binned 3x1 in the cross-track direction with 2 out of 3 along-track lines discarded. This format is obsolete and was never used in flight.

10.2.3 Definition of an “Observation”

For this ICD and as is consistent with the “*New Horizons* Spacecraft to RALPH Interface Control Document (APL Doc 7399-9201)”, an “observation” is a coherent sequence of data-taking operations, with data reported over the high-speed telemetry interface conducted automatically after being initiated by telecommands from the C&DH system. The observation may end automatically, or may run until a telecommand ends it.

All observations consist of one or more “frames.” A “frame” is the amount of high-speed image data between returns of the +frame signal to the high state. All frames consist of a sequence of “words”. A word is a pixel and is 12 bits in length.

10.2.4 Housekeeping Needed in Level 1 Files (for Calibration)

There are some housekeeping values that are kept track of to monitor the health of MVIC and to assess the data quality. These keywords are monitored during Ralph observations. The Ralph telemetry will be sent each second during an MVIC observation and these keywords will be stored as a binary table in EXTENSION 1 of the Level 1 FITS file. They will be used to set the data quality flag if the value exceeds its yellow limit. As is usual, a yellow limit indicates the monitored parameter is outside its expected operational range while a red limit indicates the parameter is far enough outside its expected operational range that the system is not performing correctly. The red and yellow limits are summarized in Table 10-7. Conversion factors for the data are provided in the instrument catalog file as part of the PDS delivery.

Table 10-7: RALPH/MVIC Housekeeping Limits

Mnemonic	Yellow Low Limit	Yellow High Limit	Red Low Limit	Red High Limit	Meaning
RALPH_HK2.POS_12V	11.75	12.25	11.5	12.5	Positive 12 Volt monitor
RALPH_HK2.NEG_12V	-12.25	11.75	-12.5	11.5	Negative 12 Volt monitor
RALPH_HK2.POS_5V	4.75	5.25	4.5	5.5	Positive 5 Volt monitor
RALPH_HK2.NEG_5V	-5.25	-4.75	-5.5	-4.5	Negative 5 Volt monitor
RALPH_HK2.POS_30V	29.75	30.25	29.5	30.5	Positive 30 Volt monitor
RALPH_HK2.MVIC_TEMP	-144	26	-150	30	MVIC temperature
RALPH_HK2.MVIC_VRD	1.9	2.1	1.8	2.2	Reset Drain Bias Reading
RALPH_HK2.MVIC_VOD	1.9	2.1	1.8	2.2	Output Drain Bias
RALPH_HK2.MVIC_VOG	1.9	2.1	1.8	2.2	Output Gate Bias

10.2.5 Raw Science Data and/or Housekeeping Requirements

There are no additional raw science data requirements beyond those already specified. Ralph housekeeping data are stored in a binary table extension. Ralph specific keywords in the Level 1 file are given in Table 10-8.

Table 10-8: RALPH/MVIC Level 1 Keywords

Keyword	Description
MET510	The MET (Mission Elapsed Time) of the Ralph housekeeping packet (ApID 0x510) that marks the start of an observation, used to determine the observation start time (STARTMET) and frame rate
TRUE510	Whether the 0x510 packet is real (YES) or assumed from a gap (NO)
RALPHEXP	Ralph per-row (SCANTYPE=TDI) or per-frame (SCANTYPE=FRAMING) exposure time in seconds
MODE	Instrument mode (0=Invalid; 1=Pan Frame; 2=Color TDI; 3=Pan1 TDI; 4=Pan 2 TDI; 5=Pan Bin (mode not used) 6=Invalid; 7=LEISA Sub; 8=LEISA Raw; 9-17=Invalid; 18=Idle; 19=Abort; 20-31=Invalid)
SIDE	Instrument hardware side (0 or 1) Note: Side 0=B, Side 1=A
DETECTOR	CCD detector use (RED; BLUE; NIR; CH4; PAN1; PAN2; FRAME)
FILTER	Detector-dependent filter (RED; BLUE; NIR; CH4; CLEAR)
SCANTYPE	Scan type (FRAMING; TDI)

For the MVIC pan frame data, the standard keywords describing the pointing (example: RA/Dec) and mid-exposure time are repeated for each image in the image cube.

10.3 Level 2 Data Specifics

10.3.1 Algorithm for Pipeline

There are four processing steps applied to the Level 1 MVIC data to produce the Level 2 output:

1. Remove bias and flat-field pattern
 - a. This converts pixel values from raw DN to calibrated DN
2. Add calibration divisors to the FITS header to be used to convert pixel values to physical units
 - a. **Pixel values remain in calibrated DN units**
 - b. Multiple divisors are provided; which to use depends on the image source (target) spectrum and geometry (point source or resolved object)
3. Calculate error for each pixel and construct the error array in a new extension
4. Construct data quality extension.

There will not be any correction for scattered light in the Level 2 products. A complete assessment of the scattered light field will be made in flight, and corrections will be implemented if necessary. There is no correction for cosmic rays in the Level 2 product.

We do not apply the Level 2 calibrations to the non-optically active pixels of the detector to maintain our high-speed header data that is encoded in some of these pixels.

Geometric and motion distortion (for the TDI arrays) corrections are higher level products and not part of the Level 2 pipeline, nor part of any PDS data sets NH-*-MVIC-2-* or NH-*-MVIC-3-* as resampling is required.

10.3.1.1 Remove bias and flat-field pattern

First, the removal of the bias level and flat-field pattern for the framing array will be discussed, and then we will address modifications for the TDI modes.

For the “Pan Frame” data, we use the shielded pixels on both sides of the array to compute the row-by-row bias. For column numbers 12 to 2511 (the active area on the left half of the chip), the bias is the median of pixels in columns 2 to 11 (inclusive and zero-based). Similarly, the bias for pixels from column numbers 2512 to 5011 (the active area on the right half of the chip) comes from the shielded pixels in columns 5012 to 5021 (inclusive and zero-based). We do not include the pixels closest to the edge of the array as they contain the high-speed header information. In the FITS headers, the bias levels for left and right are indicated with the keywords BIASL F_{xx} and BIASR T_{xx} , respectively, where ‘xx’ maps to the slice number starting with 00, 01, etc.

For the TDI mode, only pixels on one side of the image would be usable as bias level. The others are used for charge injection and do not represent the bias level. However, if we only use pixels from one side of the array we do not get a good model of the bias level. Instead, optically-active pixels from 2015 in-flight images pointed at dark sky with EXPTIME=0.6 seconds are used to compute a bias level for each TDI array. The two electronic sides have different BIAS levels.

Table 10-9 summarizes the integer bias value for each TDI array per side. The value is captured in the Level 2 FITS keyword BIASLEVL. The bias value may be different for TDI images taken with different scan rates; however that dependence is not included here.

Table 10-9: TDI Bias Levels Determined In-Flight

Detector	SIDE 0	SIDE 1
Pan1	25	25
Pan2	25	25
RED	25	23
BLUE	24	23
NIR	25	24
CH4	24	24

The correction for the bias level and flat-field pattern are described here. Let B be the bias level, R is the raw image, and F is the normalized flat-field image:

We calculate the corrected image, R_f :

$$R_f = \frac{R - B}{F} \quad (10-1)$$

The name of the flat field file used will be stored in the header keyword FLATNAME.

The naming convention for each of the flats is unique. Each file name will contain the name of the array as well as the date:

(i.e. “flat_mvic_pan1_20050612.fits” or “dark_mvic_colorBlue_20050612.fits” or “mc0_20061106.fits” for the red channel flat made on Nov 11, 2006).

For our purposes, the dark current is expected to be negligible.

For the TDI exposures, the flat field image is one-dimensional since each pixel in a given column is clocked through all 32 rows before being read out. This one-dimensional flat field has the dimensions of a single row in the TDI image. Instead of dividing the TDI image by a two-dimensional normalized flat, it is produced by dividing each row by the normalized (one-dimensional) flat. There will be 6 different flat-field one-dimensional arrays: 4 color and 2 panchromatic. There is only one 2-D flat field image (for the framing array).

The function that constructs the flattened, bias-subtracted image (Eq. 10-1) is called **mvic_flatten**. The function will determine what imaging mode was used (pan1, pan frame, etc.), use the appropriate flat field image, and produce an image with the bias subtracted, and flat field pattern removed. This image (or image cube, in the case of the pan frame detector) is stored in the Primary Data Unit (PDU) of the Level 2 FITS file.

10.3.1.2 Geometric Correction

Geometric correction is not included in the Level 2 pipeline. Instead, this correction can be implemented by third party software (e.g. USGS ISIS). There are two sources of geometric distortion: the first is instrumental, due to optical effects in the instrument. The second source of distortion is motion distortion, which only affects TDI observations and is caused by the cross-track drift of the boresight within the deadband during the scan. The first has been modeled with fifth order Legendre polynomials fitting stellar calibration images. The standard deviation of the residuals is less than 0.1 pixel. The maximum distortion at the edge of the field (near column numbers 12 and 5000) is about 2.5 pixels.

TDI images also need correcting for motion distortions. The c-kernel provides the spacecraft pointing information. The spacecraft scan can be oriented in any direction in space. From the c-kernel, we get the pointing at the beginning and end of the scan and construct a line between those two points. This is our nominal scan path. Deviations from the linear path are the motion distortion. At the mid-exposure time for each pixel, we query the c-kernel for the spacecraft pointing and compare that to our nominal scan path. We determine the cross-track error from the nominal scan path and linearly interpolate the pixels in that row to correct for the motion distortion at the mid-exposure time of the row.

10.3.1.3 Conversion to Physical Units: Calibrated DN; radiance; flux

The calibration software pipeline does not perform per-pixel conversion from DN to physical units, because that conversion requires knowledge of the spectral distribution (i.e. color) of the target. Instead, **the pixel values are left in their corrected (calibrated) DN units** from the flat-fielding, and the absolute calibration step adds “photometry” keyword divisors to the Level 2 FITS header. Each divisor scales per-pixel count rate values from DN/s/pixel to physical units (radiance) according to its assumed spectral type and spatial distribution (diffuse vs. point source) for the target.

Conversions for other source spectral distributions must be derived by integrating the product of the MVIC detectors' spectral response functions and the other source spectrum; the details are beyond the scope of this document, but the process is described in the MVIC Icarus paper listed in the References in Section 10.3.17. The MVIC detectors' spectral response functions are provided with MVIC PDS data sets delivered after December, 2015, along with the default source spectra (Pluto, Solar, etc.) as described in the Icarus paper.

Radiometric calibration keywords are provided for targets having spectral distributions similar to Pluto, Charon, Pholus, Jupiter, and the Sun. Tables providing the values for the photometry keywords at the time of launch are given below. The latest (i.e., current) values of the photometry keywords are provided in the header of the calibrated image FITS file for the image being analyzed.

The absolute calibration for each target is given by a keyword (i.e., RPLUTO) in the header of the Level 2 file that allows the user to convert a count rate (“C/TEXP” in DN/s/pixel, where “C” is the flat-fielded signal in a pixel and “TEXP” is the exposure time in seconds) for a resolved source into a radiance value (“I” in $\text{ergs/cm}^2/\text{s}/\text{sr}/\text{\AA}$) at MVIC’s pivot wavelength (specified by

the FITS keyword PIVOT; see below), assuming that the spectrum of the target is identical to the globally-averaged spectrum of Pluto. The relevant formula is:

$$I = C/TEXP/RPLUTO \quad (10-2)$$

The keyword RPHOLUS allows the conversion of the count rate for a resolved source into a radiance value at the pivot wavelength assuming that the target has a spectral distribution identical to that of the centaur object 5145 Pholus, which may be a good analog for the reddest regions on Pluto:

$$I = C/TEXP/RPHOLUS \quad (10-3)$$

Similarly, the keywords RSOLAR, RJUPITER and RCHARON provide the divisors for resolved targets with Solar-, Jupiter- or Charon-like spectra. Table 10-10 gives the values of the radiometric keywords (ie, RSOLAR, RJUPITER...) for resolved objects.

*Table 10-10: The values of the conversion factor [DN/s/pixel]/[ergs/cm²/s/Å/sr] for resolved objects. The first column is the FITS keyword in the MVIC Level 2 header. The values are derived using stellar calibrations taken on Side 0 electronics except for Pan 1. *Pan1 radiometric keywords are derived from data taken on Side 1, due to excessive noise on Side 0 discovered during instrument commissioning. Only Side 1 Pan 1 observations were taken during the Pluto encounter.*

	Red	Blue	NIR	CH4	Pan 1*	Pan 2	Pan Frame
RSOLAR	31710.05	8114.32	42993.80	10475.01	88449.55	96276.94	100190.64
RJUPITER	33642.48	8033.69	69827.44	24969.52	75954.84	82676.51	86037.34
RPHOLUS	32633.10	8404.07	41713.33	10426.00	88748.05	96601.86	100528.77
RPLUTO	31675.77	8227.81	43312.17	10541.14	85082.49	92611.91	96376.62
RCHARON	31619.96	8092.69	42989.39	10474.49	87928.24	95709.50	99600.13

The pivot wavelength (PIVOT) is given by:

$$PIVOT = \sqrt{\frac{\int P \lambda d\lambda}{\int P d\lambda/\lambda}} \quad (10-3)$$

where “P” is the MVIC system responsivity (quantum efficiency (i.e., fraction of photons detected) multiplied by the filter throughput) at wavelength “λ”. The pivot wavelength is passband specific and is stored in the FITS file header with keyword PIVOT.

Table 10-11: MVIC Pivot Wavelengths in μm

Red	0.624
Blue	0.492
NIR	0.861
CH4	0.883
Pan 1	0.692
Pan 2	0.692
Pan Frame	0.692

For unresolved sources (e.g., stars), the calibrated flux (also called “irradiance”) at the pivot wavelength can be determined using keywords that are defined analogously to the photometry keywords discussed above for resolved sources. In the case of a source having a spectral distribution identical to that of a globally-averaged Pluto spectrum, the observed count rate integrated over the MVIC PSF (“CINT/TEXP” in DN/s, where CINT is the total number of flat-field corrected counts integrated over the image and “TEXP” is the exposure time) can be related to the flux (“F” in $\text{ergs/cm}^2/\text{s}/\text{\AA}$) by:

$$F = \text{CINT}/\text{TEXP}/\text{PPLUTO} \quad (10-4)$$

Similarly, the flux at the pivot wavelength for a target having the same spectral distribution as Charon (PCHARON), the sun (PSOLAR), and Pholus (PPHOLUS) are given in the FITS header. Note: PPLUTO = $\text{RPLUTO}/((19.806 \times 10^{-6})^2)$, because the IFOV of an MVIC pixel = 19.806 μradian . The same conversion applies to the other conversion factors.

*Table 10-12: The values of the conversion factor $[\text{DN/s}]/[\text{erg/cm}^2/\text{s}/\text{\AA}]$ for unresolved objects. The first column is the FITS keyword in the MVIC Level 2 header. The values are derived using stellar calibrations taken on Side 0 electronics except for Pan 1. *Pan1 radiometric keywords are derived from data taken on Side 1, due to excessive noise on Side 0 discovered during instrument commissioning. Only Side 1 Pan 1 observations were taken during the Pluto encounter.*

	Red	Blue	NIR	CH4	Pan 1*	Pan 2	Pan Frame
PSOLAR	8.0836E+13	2.0685E+13	1.0960E+14	2.6703E+13	2.2548E+14	2.4543E+14	2.5541E+14
PJUPITER	8.5762E+13	2.0480E+13	1.7801E+14	6.3653E+13	1.9363E+14	2.1076E+14	2.1933E+14
PPHOLUS	8.3189E+13	2.1424E+13	1.0634E+14	2.6578E+13	2.2624E+14	2.4626E+14	2.5627E+14
PPLUTO	8.0748E+13	2.0974E+13	1.1041E+14	2.6872E+13	2.1689E+14	2.3609E+14	2.4568E+14
PCHARON	8.0606E+13	2.0630E+13	1.0959E+14	2.6702E+13	2.2415E+14	2.4398E+14	2.539E+14

10.3.1.4 Error Propagation

The standard deviation of each pixel is estimated and the resulting 2-D array of errors is put into extension 1 of the Level 2 FITS file. The gain, g , and the DN value of each pixel of the flattened image (Primary in the FITS file) are used to determine the photon noise. The photon noise and the read noise, RN , will be used to calculate the standard deviation per pixel in DN. The error equation is:

$$\sigma = \frac{\sqrt{(P_{meas} * g) + (RN)^2 + (f * g * P_{meas})^2} / g}{FF} \quad (10-5)$$

Where $P_{meas} = (Raw - Bias)/Flat$, and f =error in the Flat.

The gain (58.6 electrons per DN) and read noise (30 electrons) values used to calculate the standard deviation are entered in the file header with the keywords GAIN and READNOI. We also include the error propagation due to uncertainty in the flat field pattern. Other sources of error such as the uncertainty in the bias level are not included.

10.3.1.5 Data Quality Flags

The data quality flags are set to a non-zero value if there is a problem with the data, and zero if the data does not display any problems. Table 10-13 lists the factors that will cause the data quality flag to be set to its appropriate value. The data quality information is written into extension 2 of the Level 2 FITS file.

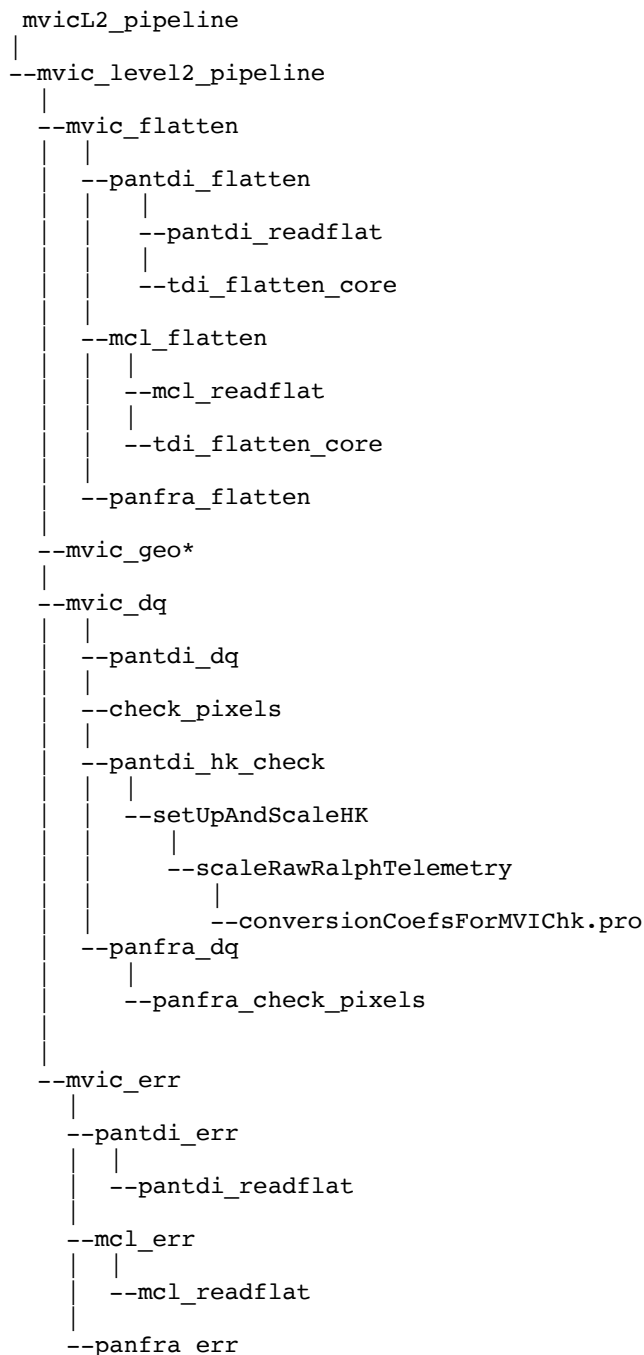
Table 10-13: Data Quality Flag Descriptions

Quality Flag Value	Description
0	Good pixel
1	Housekeeping keyword out of yellow limits, see Table 10-7.
2	Defect in one of the reference calibration files
4	Permanent CCD defect (e.g., dead pixel)
8	DN level in non-linear regime of detector
16	Zero-value pixel
32	Bad pixel not in any of above categories
-1	Missing data

If a housekeeping keyword exceeds its yellow limit at any time during the exposure the data quality flag for all the pixels in the image are set to 1.

10.3.2 Dataflow Block Diagram

Here is a diagram that shows how the data flows through the different IDL procedures and functions that comprise the Level 2 pipeline.



* mVIC_geo is where the geometric and motion distortion corrections would have been placed in the pipeline. They are not implemented in the current pipeline and this function acts as a pass-through.

10.3.3 Data Format

All of the MVIC Level 2 FITS files have a Primary data unit and two extensions. The bias-subtracted, flattened data are stored in the Primary data unit. The first extension has the error array for the Primary data. The second extension has the data quality array for the Primary data.

The image data for the TDI observations is stored in a two-dimensional array. The number of columns in the array is always 5024. The number of rows in the array depends on the duration of the observation and the scan rate.

The image data for the pan frame observations is stored in a three-dimensional array (an image cube). The number of columns is always 5024, the number of rows is always 128.

Geometric and motion distortion corrections are not part of the Level 2 pipeline, nor part of any PDS data sets NH-*-MVIC-2-* or NH-*-MVIC-3-* as resampling is required.

10.3.4 Scientific Units

We will be using cgs units for the Level 2 output. The flux per pixel will have the units of ergs/s/cm²/Angstrom or ergs/s/cm²/sr/Angstrom, for unresolved and resolved sources, respectively, after applying the divisors in the FITS keyword described in Section 10.3.1.3.

10.3.5 Additional FITS and PDS Keywords Added

Additional FITS keywords, beyond those keywords already described in Section 10.3.1.3, that have been included in the Level 2 product include:

```
SOCL2VER= '2.3      '           /Version number of the Level 2 software
PIXSIZE  =          13.0000 /Pixel size in microns
READNOI  =          30.0000 /Readnoise in electrons
GAIN     =          58.6000 /Gain in electrons/DN
PIXFOV   =          19.8065 / Plate scale in microradians/pixel
CALDIR   = 'cal/    '           /Directory for calibration files
PIVOT    = '0.692   '           /Pivot wavelength (microns)
FLATNAME= 'mxx_flat_20yymmdd.fits' /flat file used
BIASLEVL=          /median bias level, units=DN
```

The radiometric keywords summarized in Table 10-7 and Table 10-12 are filter and source dependent.

The values of the radiometric keywords are source-dependent as discussed above.

10.3.6 Extra FITS Extensions() and Their Definitions

The FITS extensions for the calibrated data are the error and data quality maps, which are described above.

10.3.7 Hardware/OS Development Platform

The Level 2 software was developed on a PowerBook G4 running Mac OS X version 10.3.4. Future migrations will be backwards compatible.

10.3.8 Language(s) Used

The MVIC Level 2 software was written using IDL.

10.3.9 Third Party Libraries Required

The code requires the “IDL Astronomy User’s Library” which can be found on the internet.

10.3.10 Calibration Files Needed (with Quantities)

The following calibration files are used.

- Flat field images for each array
- Bad pixel files for each array
- File with the acceptable levels of the housekeeping keywords that we are monitoring

[There will be multiple generations of each over time during the mission. These files will be archived with the observations.]

10.3.11 Memory Required

Not specified.

10.3.12 Temporary File System Space Needed

None.

10.3.13 Predicted Size of Output File(s)

The size of the files are dependent on the number of rows (TDI) or frames (PAN). A Level 2 file will be approximately four to five times as large as the corresponding Level 1 file.

10.3.14 Predicted Execution time

Not specified.

10.3.15 Contact/Support Person(s)

Cathy Olkin
Dennis Reuter

10.3.16 Maintenance Schedule (Code/Data Updates, Documentation)

The code will be updated on the schedule mandated by the mission. Calibration files will be updated throughout the mission as necessary. Necessary criteria include new calibration information (such as new flat-field observations, new geometric distortion correction or inflight stellar radiometric calibration) such as might be available after an annual checkout.

10.3.17 References

Howett, C.J.A., A.H. Parker, C.B. Olkin, D.C. Reuter, K. Ennico, W.M Grundy, A.L. Graps, K.P. Harrison, H.B. Throop, M.W. Buie, J.R. Lovering, S.B. Porter, H.A. Weaver, L.A. Young, S.A. Stern, R.A. Beyer, R.P. Binzel, B.J. Buratti, A.F. Cheng, J.C. Cook, D.P. Cruikshank, C.M. Dalle Ore, A.M. Earle, D.E. Jennings, I.R. Linscott, A.W. Lunsford, J.Wm. Parker, S. Phillippe, S. Protopapa, E. Quirico, P.M. Schenk, B. Schmitt, K.N. Singer, J.R. Spencer, J.A. Stansberry, C.C.C. Tsang, G.E. Weigle II, A.J. Verbiscer. Inflight Radiometric Calibration of *New Horizons*’ Multi-spectral Visible Imaging Camera (MVIC), Icarus 287, 140-151. DOI:10.1016/j.icarus.2016.12.007.

11. PEPSSI INSTRUMENT DESCRIPTION

11.1 Overview

11.1.1 PEPSSI Investigation

PEPSSI (Pluto Energetic Particles Spectrometer Science Investigation) is a particle telescope and a time-of-flight (TOF) spectrometer that measures ions and electrons over a broad range of energies and pitch angles. Particle composition and energy spectra are measured for H to Fe from ~ 30 keV to ~ 1 MeV (but not all species are uniquely separated) and for electrons from ~ 30 keV to 700 keV.

The PEPSSI instrument traces its heritage back to the MESSENGER Energetic Particle Sensor (EPS) instrument. EPS/PEPSSI was developed with the support of a NASA Planetary Instrument Definition and Development (PIDDP) grant aimed at designing a low-mass, low-power sensor that can measure energetic pickup ions produced near planets and comets (Andrews *et al.*, 1998; McNutt *et al.*, 1996). The overall PEPSSI instrument weighs 1.5 kg and uses maximum power of 1.4 W.

Figure 11-1 shows the placement of PEPSSI on the spacecraft and the PEPSSI fields-of-view (FOV).




Figure 11-1: Location of PEPSSI on the New Horizons spacecraft. The lightly shaded area denotes PEPSSI Field-of-View (FOV)

The science goals of the PEPSSI instrument are:

1. Determine the escape rate of Pluto's atmosphere
2. Measure the interaction of the solar wind with Pluto's ionosphere
3. Determine the source and nature of energetic particles found near Pluto

11.1.2 PEPSSI Sensor Description

PEPSSI is a compact particle telescope with a time-of-flight (TOF) section and a solid-state detector (SSD) array (see Figure 11-2). A mechanical collimator defines the acceptance angles for the incoming ions and electrons. A cutaway view of the assembly is shown in Figure 11-3. The TOF section is axially symmetric; entrance and exit apertures are 6 mm wide with an azimuthal opening angle of 160° . The entry apertures are covered by a thin polyimide/aluminum/ polyimide foil. The stop foil is a polyimide/ palladium/ polyimide foil. The foils are mounted on high-transmittance stainless-steel grids. The foil thickness and composition is a compromise to minimize the energy threshold, secondary electron production, and scattering of particles in the foil while blocking UV from the direct Sun and Lyman- α background. PEPSSI measures the ion TOF using secondary electrons generated as the ion passes through the entrance and exit foils in the spectrometer. Total energy is measured by the SSD array. Each of the six SSDs has two pixels, three of the SSDs are dedicated for ion measurement. The other three have one pixel covered with $\sim 1 \mu\text{m}$ Al absorber, to block low energy ions and permit measurements of electrons. The fan-like collimator together with the internal geometry defines the acceptance angles. The FOV is 160° by 12° with six angular sectors of 25° each; the total geometric factor is $\sim 0.15 \text{ cm}^2 \text{ sr}$. As an ion passes through the sensor, it is first accelerated by the potential of $\sim 3 \text{ kV}$ on the front foil prior to contact with that foil. The ion generates secondary electrons at the foils, which are then electrostatically steered to well-defined separate regions on a single micro channel plate (MCP), providing “start” and “stop” signals for the TOF measurements (from 1 ns to 320 ns). The segmented MCP anode, with one start segment for each of the six angular entrance segments, allows determination of the direction of travel even for lower-energy ions that do not give an SSD signal above threshold.

The combination of measured energy and TOF provides unique particle identification by mass and particle energy depending on the range: for protons from $\sim 30 \text{ keV}$ to $\sim 1 \text{ MeV}$; for heavy (CNO) ions from $\sim 80 \text{ keV}$ to $\sim 1 \text{ MeV}$. Lower-energy ($>3 \text{ keV}$) ion fluxes are measured by TOF only, but without the SSD signal, providing velocity spectra at these energies as well. Molecular ions, expected from Pluto's atmosphere, will break up in the foil prior to their full detection, but will be detected as high-mass events. Internal event classification electronics determine the mass and produce an eight-point energy spectrum for each of four species for six arrival directions. Energetic electrons are measured simultaneously in the dedicated electron pixels in the range from $\sim 30 \text{ keV}$ to 700 keV . Only protons with energies $> 300 \text{ keV}$ (expected to be very rare at Pluto) can penetrate the absorbers on these pixels, and even those would be eliminated by on-board MCP coincidence requirements and ground comparisons with the simultaneously measured ion flux.




Figure 11-2: Schematic drawings of the PEPSSI sensor




Figure 11-3: A Cut-Away View of the PEPSSI FOV

The instrument is equipped with an internal degraded Am-241 alpha source that provides a falling spectrum towards the lowest energies. The alpha source irradiates sectors 4 and 5 and can be used to track the time dependent instrument efficiency.

11.1.3 PEPSSI Electronics Description

Extensive uses of miniaturization and custom electronics in the design allow PEPSSI to weigh less than 1.5 kg and consume less than 1.4 W. PEPSSI is made up of six modular 4"x4" slices. They consist of:

- 1) Energy board;
- 2) High Voltage Power Supply (HVPS);
- 3) TOF board;
- 4) Digital processing board;
- 5) Common event processor board; and
- 6) Low Voltage Power Supply (LVPS) board.

Figure 11-4 shows the exploded view of PEPSSI with each board labeled. A brief description of the functionality of each board is highlighted below.

Energy board: The energy board is the interface between the SSDs and the signal conditioning electronics. It houses the sensor, MCP anodes, charge amplifiers, pulse shapers, etc. In addition, it also outputs the pulse signal from the 6 start anodes and 1 stop anode.

HVPS board: The HVPS board contains the high voltage (HV) drive circuitry, HV transformer, and its control circuitry. It provides HV up to -2900 V for the sensor electrostatic lens and MCP bias. In addition, the HVPS also needs to provide bias voltage over the range of 0 to -200 V with <10 mV ripple.

Digital processing board: The digital processing board provides valid event logic functions, which include channel enables, programmable coincidence window, event packet generation and rate counters for event statistics. It provides the logic to distinguish between electrons, ions and directionality.

Common event processor board: This board contains PEPSSI's main processor (RTX2010RH), the Filed Programmable Gate Array (RT54SX72S), and various memory modules (SRAM, EEPROM, PROM).

LVPS board: This board converts primary spacecraft power into multiple low voltage outputs used by PEPSSI. It provides highly efficient power conversion into two digitals (+5, +2.5V) and four analogs (+5, -5, +15, and -15) outputs.




Figure 11-4: Expanded view of the PEPSSI sensor.

11.2 Introduction to PEPSSI Data

The PEPSSI instrument can operate in two modes: Normal and Diagnostic. On the spacecraft, each event generates a PHA record. This record is classified by event type: Electron, High-Energy Ion (or “Hi-Ion” or “Triple”), or Low-Energy ion (or “Low-Ion,” “Double,” or “TOF-only”). In diagnostic mode, events are not classified; alternatively, all events are “diagnostic events”. Events of a given type are further classified into “Rate Boxes” by their energy and/or time of flight (TOF). Thus each event has a type, a rate box, and a detector in which it occurred. Instead of detector number, we will often use arrival direction (or sector) since there is a one to one relation between them (see Figure 11-5). A six character string is used to specify each possible classification (or Rate) of an event. The construction of this string is (type)(rate box)S(arrival sector). The arrival sector numbering is shown in Figure 11-5. The “type” string is:

- B – Hi-Ions (possesses Energy and TOF)
- R – Electrons (energy only, no TOF); note: cosmic ray ions dominate the R rate away from Jupiter.
- L – Low-Ions (TOF-only, no-energy).





Figure 11-5: PEPSSI Layout Labeling

For High Energy Ions, the “Rate Boxes” are determined by areas in the TOF vs Energy plane (see Figure 11-6). These correspond to different particle species and different energies.

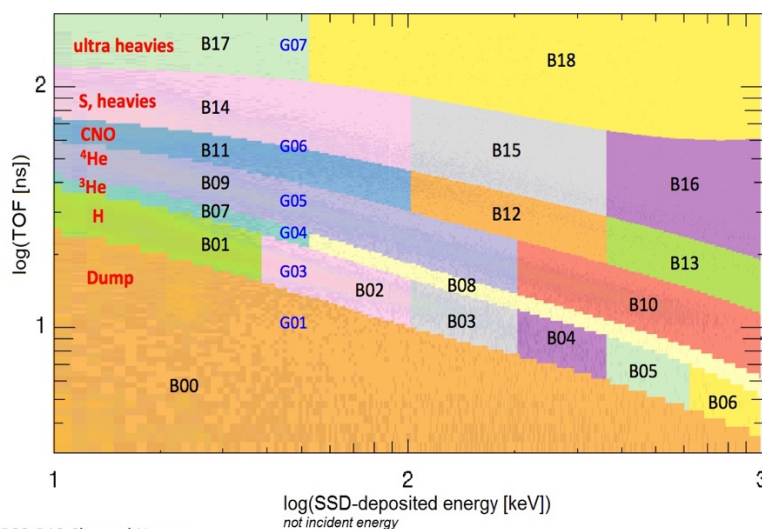
For Electrons, the “Rate Boxes” are determined by energy ranges.

For Low Energy Ions, the “Rate Boxes” are determined by TOF ranges *and* by which heavy ion discriminators fired. Note that because of the way the PEPSSI electronics work, frequently the arrival sector is unknown or uncertain for the Low-Ion measurements.



a)

PEPSSI TOF (ns) vs. E (keV) Post-2007/144 16:30 (as of 2013-08-14)



B00-B18 Channel Names
G01, G03-G07 Priority Groups (by species)

b)

Figure 11-6: PEPSSI Rate Boxes on the TOF vs Energy Plane. Normal Mode. (a) Boxes until 2007 DOY 144; (b) boxes after this time.

At a few points in the mission the “Rate Box” definitions for the Hi-Ions and/or the Low-Ions have been changed by uploading new “Look Up Tables” (LUTs) to the instrument. The LUTs are documented in the PDS data:

- For Hi-Ions, FITS images which specify the Rate Box for each “pixel” in the TOF-Energy plane of possible measured data numbers are presented in FITS image files in the calibration directory.
- For Low-Ions and Electrons, the Rate Boxes are actually just ranges in a one-dimensional space and are given in the headers of the Level 2 files.

Some examples:

B02S04 – High-Energy Rate Box #2, arriving from the sector 4 direction. Protons with deposited energy in channels 95-169, in analog-to-digital units (ADUs).

L06S03 – Low-Ions Rate # 6 arriving from the sector 3 direction. Ions with TOF indices from 45-79 ADUs and for which the heavy ion discriminator H0 fired but H1 didn’t.

R00S05 – The 0th electron rate, arriving from the sector 5 direction. Electrons with Energy channels in the range 720-1023 ADUs. (Note that cosmic ray ions dominate the “electron” channels away from Jupiter.)

There are two counters for each Rate that are incremented whenever a corresponding event occurs. The N2 counter is accumulated for a certain time interval (programmable down to 1 second during the Pluto encounter but typically 15-60 seconds during cruise), then recorded and zeroed. The N1 counter is accumulated for some multiple of the N2 interval (usually 10 minutes during cruise), then recorded and zeroed. A certain number of PHA events are kept according to a complex priority scheme and telemetered along with the Rate data. (Note: if the multi-hit, i.e. anti-coincidence, flag is set for an event, the event is *not* counted. This is programmable, but the “don’t count multi-hit events” rule was true outside of Bad Time Intervals for the whole Jupiter phase).

Various housekeeping and status data and certain global hardware and software counters are also present in the data at Levels 1 and 2.

PEPSSI Level 1 is an internal format that is not used outside the SOC. It is nearly identical to the Level 2 format apart from extra meta data added to the FITS headers of Level 2 data.

PEPSSI Level 2 data represents, with 3 exceptions, the raw data taken from the spacecraft telemetry. It has merely been reformatted for ease of use. No data has been added, removed or altered with the following 3 exceptions:

- a. Instrument Status information has been calibrated to physical units where applicable (see discussion in Section 11.4 below).
- b. For clarification, a “DT” column has been added to the Rate tables to indicate the integration time over which the count data was accumulated. This information is not available in the spacecraft telemetry until after a flight software change on November 29 (DOY 334), 2016. Before that, it is calculated from a collection of spacecraft command files. This “DT” value may be inaccurate during Bad Time Intervals (BTIs), see below.

- c. For ease of use, we have added a column giving the deduced “Rate Box” of High-Ion PHA and Electron PHA events to the Level 2 PHA data. While this can, in principle, be calculated from the Level 2 quantities and the RATEBOXDEFINITIONPLANES.FIT file available in the CALIB/ directory of the PDS archive, the procedure is complex enough that we have found it convenient to perform this calculation in advance and include the information in the Level 2 files.

PEPSSI Level 3 data presents the data in a format that should be convenient for scientific analysis. All of the calibration parameters needed to convert Level 2 data to Level 3 data are present in the headers of the Level 3 data files. The formulas used to calculate the calibrated quantities are also present in the Level 3 headers.

Rate data are presented in physical flux units with uncertainties as well as counts per second.

For the Doubles (“L”) rates, only time of flight information is available, so we do not know the actual mass/particle species of the measured events. Therefore, we present multiple calibrations for various “assumed” particle composition. That is, we calibrate the rate as if the flux were exclusively composed of a particular species. If more than one species calibration is present for a given energy channel, the column name will have a character appended to it (see Table 11-1). There will also be a column with no species letter which uses the “best guess” species for that column. The “no suffix letter” column will duplicate the appropriate lettered column. Some of these species (such as CNO or Ions) are only present as labels of the Triples rates but are retained here for completeness).

Table 11-1: PEPSSI Species Letter Codes

Species	Code Letter
Protons	P
Helium (He ⁺)	H
Alphas (He ⁺⁺)	A
CNO	B
Ions	I
Oxygen	O
Nitrogen	N
Singly charged molecular Nitrogen	L
Doubly charged molecular Nitrogen	M
Oxygen +6	Q
Sulfur	S
Electrons	E
Rate (i.e. not a flux)	R

PHA data are presented with calibrated TOF and deposited energy. Further calibrated incident energies are given for assumed ion species. There are also some “quick look” PHA images and Rate spectrograms in the Level 3 data to allow for a simple overview of each “day's” (see below) observations.

IMPORTANT NOTE: the quick-look PHA images are only presented to give a single, uniform overview of data throughout the entire mission and to allow users to “sanity check” their understanding of the actual data in the FLUX and PHA extensions. They do not represent a best presentation of the data and should not be used for scientific analysis.

No diagnostic mode data are present in the Level 3 data. BTI data *are* present in the Level 3 data but should be ignored. Level 3 data for the “Launch” phase is present for quick-look purposes, but, apart from the deposited energy calibration (which is well known), the calibrations are performed with dummy values, as will be evident from examination of the header information.

11.2.1 Effect of Spacecraft Spin

Because *New Horizons* is not a pure “spinning” spacecraft, PEPSSI data does not have specific adaptations available for data taken while the spacecraft is spinning. For event data, the timestamp and sampling period are too coarse to accurately correlate any given event with a specific spacecraft attitude. For rate data, since the data accumulation periods rarely “line up” precisely with the start and end of a spin, there will typically be some spin-modulation of the measured rates due to sample aliasing.

For example, N2 counter accumulation data timestamped at 15s sampling intervals while the spacecraft is spinning at ~12s per revolution (5RPM) would sample the same 90 degrees of sky (3s of a revolution) twice per sampling interval, but would sample the other 270 degrees of sky only once per sampling interval.

11.3 Level 1 Data Specifics

The SOC Level 1 data product is nearly identical to the Level 2 product apart from additional metadata added to the Level 2 product which is described below. The Level 1 product is not delivered to the PDS.

11.3.1 Data Sources (High/Low Speed, CCSDS, ITF)

PEPSSI is low-speed only. Data will be from CCSDS (Consultative Committee for Space Data Systems) packets.

11.3.2 Definition of an “Observation”

PEPSSI doesn't make specific “Observations.”

11.3.3 Header with Keywords

This section has intentionally left blank. The keywords are described in the Level 2 documentation.

11.3.4 Spacecraft Housekeeping Needed in Level 1 Files (for Calibration)

Spacecraft attitude is necessary to calculate PEPSSI pointing in the ecliptic coordinate system, and with respect to Jupiter, Pluto, and Charon. An accuracy of 1 degree is adequate.

11.3.5 Raw Science Data and/or Housekeeping Requirements

No special requirements

11.4 Level 2 Data Specifics

11.4.1 Algorithm for Pipeline

11.4.1.1 Level 1 Step

The PEPSSI Level 1 pipeline converts a stream of CCSDS packets into FITS tables and writes them to a FITS file with each HDU corresponding to a PEPSSI instrument packet type and the columns of the table corresponding to the fields in the packet.

The extensions/packet types are:

- N1
- N2
- (STATUS) N2_STATUS
- D_N1
- D_N2
- PHA_ELECTRON
- PHA_LOW_ION
- PHA_HIGH_ION
- PHA_DIAG

For historical reasons, the STATUS packet type is placed in an extension labelled N2_STATUS.

11.4.1.2 L2 Data Format

The contents of the L2 files are detailed in the following sections. The ordering of the extensions is *not* guaranteed, programs accessing the data should search for the desired extension by name (in the EXTNAME keyword). There will always be only one extension of each type (EXTNAME) and not all types will necessarily be present. An extension will only be present if there is data of that type taken during the time period covered by that file. The coverage of a single L2 “daily” file begins at an integer multiple of 86,400 ephemeris (TDB, Barycentric Dynamical Time) seconds from a time near a nominal midnight UTC time at launch (January, 2006). All events and rate measurements with timestamps in that file occur before the 86,400th TDB second following. All event records in the file occur *between* those two TDB times, but

there are rarely if ever event records with timestamps *at* those times. The values and keyword values in the FITS header dealing with time are as follows:

ET – The ephemeris time of the beginning of coverage of this file. Note: Ephemeris time is not in the FITS header or PDS label. This corresponds to a time near 00:00:00 (midnight) UTC at the start of the mission (January, 2006), but slides back with respect to 00:00:00 UTC as the mission progresses through UTC leap second events.

MET, METEND – Integer MET values approximating the beginning and ending of the coverage period of this file, expressed in terms of the spacecraft clock. Since these values are truncated to the previous integer value and do not include spacecraft clock partition information, they should be used with caution, if at all. The filename contains a truncated version of the beginning MET keyword value. Note that, because the integer MET values are rounded *down*, the beginning MET usually represents a time *not* actually contained in this file. Usually, two consecutive-in-time files have beginning MET values that differ by a multiple of 86,400; however, because the spacecraft clock runs slower (MET seconds are longer) than TDB seconds, sometimes that [difference modulo 86,400] will be 86,399s. The PDS keyword values START_TIME and STOP_TIME are millisecond-rounded UTC times based on these rounded MET values, and the PDS keyword values SPACECRAFT_CLOCK_START_COUNT and SPACECRAFT_CLOCK_STOP_COUNT are in turn based on START_TIME and STOP_TIME, and include sub-millisecond artifacts of the sequence of roundoffs.

STARTMET, STOPMET – Floating-point MET values of the first and last record (in any EXTENSION) in this file. These values may be useful for quickly answering data availability questions. $STARTMET \geq MET$ and $STOPMET < METEND$.

Other time related keywords in the header are artifacts from the processing pipeline and should be considered as “for instrument team use in validation only”.

The Primary HDU contains no data, only informational header keywords identifying mission info, observational start time, and information about the file creation (date, software version, etc.).

The available extensions are:

D_N1, D_N1_STATUS, D_N2, D_N2_STATUS, N1, N1_STATUS, N2, N2_STATUS,
PHA_DIAG, PHA_ELECTRON, PHA_LOW_ION, PHA_HIGH_ION

The different extensions are described below. For detailed information such as the data-type of different columns, see the FITS header in the data file.

Extension names beginning with “D_” (and the PHA_DIAG extension) represent data taken in diagnostic mode (ApIDs 0x695 through 0x698). Diagnostic data are taken for purposes of calibration and understanding the instrument at a very low level. Diagnostic data are complex and often taken under unusual conditions. It is unlikely that the general user will be able to use diagnostic data in any meaningful way. In diagnostic mode, a PHA event is generated whenever any of the instrument detectors sees an event. All events are counted except for those with the

multi-hit flag set (by default, the multi-hit setting can be changed). This means that there are no “Low-Ion” events because, by definition, these events are required to have no SSD fire, which is mutually exclusive with the diagnostic mode requirement that every diagnostic event is initiated with an SSD fire. Many of the columns will contain “fill values” for certain types of events. The “Fill” values for invalid Energy and TOF PHA data are: Energy – 1023 and TOF – 2047. In the PHA_DIAG extension, what would normally be an electron event, for instance, will have a Time of Flight value of 2047.

11.4.1.3 PHA_HIGH_ION

When an ion enters the PEPSSI detector, if it has enough energy, we measure an Energy and a Time of Flight (TOF). Since we don’t have enough bandwidth to telemeter all of our events, we use a round robin priority scheme to decide which PHA events to discard and which to telemeter. All of the events are counted in the N1 and N2 Rate data. The Rate data can be used to remove the priority group effects (to a large extent) from the PHA data by weighting events by their respective rates. See the L3 documentation below for more on rate-weighting.

The PHA_HIGH_ION extension contains one row for each high energy ion event that was not discarded by the priority scheme. The columns are:

ET	Ephemeris Time (s past J2000 epoch) of start of accumulation interval
MET	Mission Elapsed Time
ApID	Which ApID was this event telemetered in: N1 – 0x691, N2 – 0x692, N3 – 0x693
Cross_Talk_Indicator	Did more than one detector fire?
H0	Did Heavy Ion Discriminator 0 fire?
H1	Did Heavy Ion Discriminator 1 fire?
Ion_Channel	Detector Channel (0-8)
Raw_Energy	Energy deposited (ADU)
Raw_TOF	Time of Flight (ADU)
Start_Anode	Bits 0-5 are set if that start anode fired.

Notes for PHA High Ion data:

- **IMPORTANT NOTE:** The timing of a PHA event is, in general, *not* known to 1 second precision. The “time tag” of a PHA event only represents the start time of the accumulation interval of the Rate packet with which it was telemetered. So, in normal cruise operation, for example, “N1 PHA” event arrival times are known only to the nearest 10 minutes, N2 events to the nearest minute, and N3 events to within 2 hours. See discussion of Level 3 PHA data for more details.
- See Figure 11-5 for a diagram of the detector numbering. 9 of the 12 detectors are dedicated to Ions (the other 3 are dedicated to electrons) and are configured accordingly.

- **IMPORTANT NOTE:** Only High Ion event data following the TOF(E) tracks (going from long TOF and small energy to short TOF and high energy, roughly aligned with the rate boxes shown in Fig 11-6) are valid. Events with TOF<5ns and deposited energy >1MeV or <40keV are artifacts.
- The event data are not weighted by the instrumental efficiency (different to the rate-weighting that is included). Similar to raw count rates, event data therefore cannot be used as-is to determine energy spectra in physical units.
- Mission Elapsed Time starts about 19-JAN-2006-18:09:05.184.
- The H0 and H1 ion discriminators were found to be of limited usefulness.
- Events with the Cross_Talk_Indicator value set are discarded from the rate counters and are not usually used in analysis.
- The Start Anode column consists of a single byte. The individual bits 0-5 indicate whether the corresponding Start Anode (0-5) registered an event. See Figure 11-5 for start anode layout. Note that, unfortunately, the numbering of the anodes is reversed from the numbering of the incoming angle sectors.
- Energy and TOF are given in raw “Analog to Digital Units” (ADU).
- The electronics of the Start Anodes are such that, while a given event may have a known TOF, the exact Start Anode information may be uncertain or completely unknown. Thus, a valid event may show more than one Start Anode, or none.

11.4.1.4 PHA_ELECTRON

The PHA_ELECTRON data are very similar to the PHA_HIGH_ION data except that the TOF-related values aren't present since electrons aren't detected by that part of the instrument (i.e. they only have solid state detector (SSD)-related values:

ET	Ephemeris Time (s past J2000 epoch) of start of accumulation interval
MET	Mission Elapsed Time
ApID	Which ApID was this event telemetered in
Cross_Talk_Indicator	Did more than one detector fire?
Electron_Channel	Detector Channel (0-2)
Raw_Energy	Energy deposited solid state detector (ADU)

Notes for PHA_ELECTRON data:

1. Only Sectors 0, 2, and 5 have electron detectors (0, 1, and 2, respectively) associated with them.
2. See PHA_HIGH_ION notes for other relevant info.

11.4.1.5 PHA_LOW_ION

The PHA_LOW_ION events are from low energy ions that register in the TOF part of the detector but do not trigger the SSDs. Hence “Low Ions” have TOF data (and associated quantities like Start Anode) but no Energy data:

ET	Ephemeris Time (s past J2000 epoch) of start of accumulation interval
MET	Mission Elapsed Time
ApID	Which ApID was this event telemetered in
H0	Did Heavy Ion Discriminator 0 fire?
H1	Did Heavy Ion Discriminator 1 fire?
Raw_TOF	Time of Flight (ADU)
Start_Anode	Bits 0-5 are set if that start anode fired.

- **IMPORTANT NOTE:** Low Ion events with TOF<20ns are instrumental artifacts.
- The event data are not weighted by the instrumental efficiency (different to the rate-weighting that is included). Similar to raw count rates, event data therefore cannot be used as-is to determine energy spectra in physical units. Most notably, the event-counts peak around 100ns in the low ion event data is not a peak in intensity.

See PHA_HIGH_ION notes for other relevant info.

11.4.1.6 PHA_DIAG

In diagnostic mode, the following columns are present:

ET	Ephemeris Time (s past J2000 epoch) of start of accumulation interval
MET	Mission Elapsed Time
ApID	Which ApID was this event telemetered in?
Cross_Talk_Indicator	Did more than one detector fire?
Fired	0 - electron event 1 - ion event
H0	Did Heavy Ion Discriminator 0 fire?
H1	Did Heavy Ion Discriminator 1 fire?
Ion_Channel	Detector Channel (0-8) or (0-2 if electron)
Raw_Energy	Energy deposited solid state detector (ADU)
Raw_TOF	Time of Flight (ADU)
Start	Did a start anode fire?
Start_Anode	Bits 0-5 are set if that start anode fired.
Stop	Did the stop anode fire?

Notes for Diagnostic PHA Data:

1. The “Fired” flag indicates whether the event is an Ion or an Electron event. This determines which Rate counter gets incremented as well.
2. Raw Energy frequently has the 1023 fill value in diagnostic mode.
3. Raw TOF frequently has the 2047 fill value in diagnostic mode.
4. Diagnostic mode event data before the November 29, 2016 (DOY 334) software change did not correctly perform the round robin priority scheme so that some rate boxes were not properly populated with event data.

11.4.1.7 N1 and D_N1

The N1 and N2 (and D_N1 and D_N2) extensions contain several types of “Rate” data. The Rate data are accumulated in histograms which are then dumped at set intervals. During cruise: for N1 data, usually the histograms are accumulated for 600 seconds. For N2 data, the accumulation time is usually 60 seconds except for the first hour of the day when it is 15 seconds. These values will be different during the Pluto encounter or certain testing events. The DT column will indicate the accumulation time for a given row or Rate data. The DT column may not be accurate during BTIs. Before November 29, 2016 (DOY 334), the DT value was not present in the instrument telemetry and must, therefore, be deduced from archived instrument command files. Sometimes, it is difficult to accurately deduce the DT value; in that case, the DT value will be given as -1.

B Rates: The number of high energy ion events in the various Hi-Ion “Rate Boxes”.

C Rates: The contents of various hardware counters

HK Rates: Various housekeeping quantities such as power levels and discriminator thresholds

J Rates: Software counters that represent overall quantities like total number of Electron Events.

L Rates: The number of low energy (TOF-only) ion events in the various Lo-Ion Boxes.

R Rates: The number of electron events in the various electron “Rate Boxes”.

The N1 and D_N1 data are identical in format; the D_N1 data are taken when the instrument is in diagnostic mode. The definitions of some of the Rate Boxes are different in diagnostic mode and normal mode (i.e. the Rate Box number is the same, but its definition is different depending on the mode). The events being counted are triggered with different rules, as well, see the discussion of PHA_DIAG data for more detail.

We describe some of the rates in more detail below. A detailed description of each rate is also available in the FITS header as a comment to the keyword defining a Rate column. Example:

```
TTYPE12 = 'B01S03 ' / B01S03: Protons (60-94) Energy ADUs Sector: 3
```

means that column 12 contains rate B01S03 which is nominally Protons with energy between 60 and 94 ADUs incident from Sector 3.

B Rates: The TOF vs Energy plane is divided into 19 “Rate Boxes” as shown in Figure 11-6. Each high energy ion is classified into a Rate Box and further its incident SSD is used to classify its sector (and therefore direction), resulting in a Rate, or histogram cell designation of the form BnnSnn. The B boxes at any given point in the mission can be found in the Rate Box Definition Planes file in the CALIB/ directory.

Important Notice: B rates in sectors 4 and 5 are contaminated or dominated by counts from the internal alpha source of the instrument and are therefore only useful for diagnostic purposes of the instrument. Note that the alpha source also affects channels that nominally do not measure alpha particles. This is because the edges of the alpha “track” in TOF x Energy space (same space as used for the channel definitions) reach to the tracks of other species. Even though only a small fraction of alphas reaches the other channels, their contamination becomes significant if the foreground is as low as in the interplanetary medium. B rates in the affected sectors (like B01S04, B01S05, etc) should therefore not be used for science. Make sure to not average over all sectors. Before 2007 DOY 144 the alpha source was measured with B06-08 S04-05, after that day the rates were B07-10 S04-05.

Background or electronic noise dominates the following B channels after 2007 DOY 144: B00, B07, B08, B17, and B18. The B00 channel is a dump channel. B07 and B08 were designed to be He-3 channels, but there is no sign of this isotope so they contain only the wings of the He-4 distribution. The B17 and B18 channels are 'ultra heavies' but have shown no evidence of responding to such particles. Also B06, nominally measuring high-energy protons, is unreliable. Beyond these channels, the channels B14, B15, B16 are nominally measuring sulfur but are dominated by noise in the interplanetary medium far from Jupiter. Before 2007 DOY 144, the similar issues with background and noise as described above exist but the channel definitions are different. Background or electronic noise dominates the dump channels B00 and B18. The B05 channel measuring high energy protons is at background, except in near-Jupiter environment. The channels measuring sulfur respond well during this period.

C Rates:

These are a collection of total counting rates:

C00D00 Detector 00 Singles
C01D01 Detector 01 Singles
C02D02 Detector 02 Singles
C03D03 Detector 03 Singles
C04D04 Detector 04 Singles
C05D05 Detector 05 Singles
C06D06 Detector 06 Singles
C07D07 Detector 07 Singles
C08D08 Detector 08 Singles
C09D09 Detector 09 Singles
C10D10 Detector 10 Singles

C11D11	Detector 11 Singles
C12S05	Anode 00 Singles Sector: 5
C13S04	Anode 01 Singles Sector: 4
C14S03	Anode 02 Singles Sector: 3
C15S02	Anode 03 Singles Sector: 2
C16S01	Anode 04 Singles Sector: 1
C17S00	Anode 05 Singles Sector: 0
C18	Electron Events
C19	Ion Events
C20	H0 Singles
C21	H1 Singles
C22	Starts
C23	Stops
C24	Valid TOFs

Notes:

-“Singles” means a single “start” event, without a coincident “stop” event, on the named detector, or a “stop” event without a coincident “start” event.

HK Rates:

These are various housekeeping values:

HK00	Peak Discriminator Level 0
HK01	Peak Discriminator Level 1
HK02	Peak Discriminator Level 2
HK03	Peak Discriminator Level 3
HK04	Peak Discriminator Level 4
HK05	Peak Discriminator Level 5
HK06	Peak Discriminator Level 6
HK07	Peak Discriminator Level 7
HK08	Peak Discriminator Level 8
HK09	Peak Discriminator Level 9
HK10	Peak Discriminator Level 10
HK11	Peak Discriminator Level 11
HK12	HVPS Level
HK13	BIAS Level

HK14	Start Constant Fraction Discriminator Power
HK15	Stop Constant Fraction Discriminator Power
HK16	Stop Constant Fraction Discriminator Threshold
HK17	Start Constant Fraction Discriminator Threshold
HK18	Start Anode 0
HK19	Start Anode 1
HK20	Start Anode 2
HK21	Start Anode 3
HK22	Start Anode 4
HK23	Start Anode 5
HK24	Heavy Discriminator 0
HK25	Heavy Discriminator 1
HK26	Command word A mirror
HK27	Command word B mirror
HK28	Event parameter A mirror
HK29	Event parameter B mirror
HK30	Event parameter C mirror
HK31	Bus read address (from most recent bus read)
HK32	Bus read value (resulting data value)
HK33	Unknown0
HK34	Unknown1

Notes:

- HK33 and HK34 are only retained because they're present in the telemetry. They're just spare values at this time.

J Rates:

These are software counter totals:

J00	Electron Events
J01	Hi-E Ion Events
J02	Low-E Ion Events
J03	Electron Discards
J04	Hi-E Ion Discards
J05	Diagnostic Events
J06	Diagnostic Discards

L Rates:

The ADU ranges defining the L Rates are given in the Level 2 data file headers.

Important Notice: The L rates in sectors 4 and 5 have a strong contribution from the internal alpha source of the instrument and should therefore ignored for most purposes.

Each L rate is associated with sectors 0-5, as it is the case for the B rates. Counts, where the instrument was not able to track the sector in which they occurred, are accumulated in sector “unknown”. On 2014 DOY 178, electronic thresholds in the instrument were optimized to ensure that most counts occur in sectors 0-5.

R Rates:

R rates nominally represent electron events. Far from Jupiter they measure cosmic rays. The ADU ranges defining the R Rates are given in the header of the Level 2 files.

11.4.1.7.1 Rate Box Definitions

For Electrons and Low-Ions, the rate box definitions are simple ranges in Energy and TOF in ADUs which can be found in the Level 2 headers. For Hi-Ions, the Rate Boxes are regions in the TOF-Energy plane (see Figure 11-6). The precise specification of the rate boxes is complex and this is why we include rate box classifications in the Level 2 PHA data. However, we also provide the file RATEBOXDEFINITIONPLANES.FIT in the CALIB/ directory of PDS data sets.

Through the Jupiter mission phase and into the Pluto Cruise mission phase, PEPSSI has had a few sets of Rate Box definitions. There are separate rate box definitions for Normal and Diagnostic mode. The RateBoxDefinitionPlanes.fit file contains an extension for each version of the “B” Look Up Table for each mode (i.e. normal or diagnostic). The Rate Box classification data are presented as an image containing the Rate Box id number (0-18) in the Time of Flight (TOF) vs Energy plane. TOF and Energy are presented in raw analog to digital units (ADUs). Thus, TOF ranges from 0-2047 and Energy from 0-1023. The Rate Box classification of any event can be determined by looking up the corresponding pixel in the classification image. In diagnostic mode, Energy-only detections have a fill value of 2047 for TOF, and the Rate Box classification image assumes these fill values. The time range for each set is given in the header of the respective HDU.

11.4.1.8 N2 and D_N2

N2 (and D_N2) are identical to their N1 counterparts except that they are typically sampled much more frequently (every 15 or 60 seconds) and only some of the L and C rates are present. Which specific L and C rates are present at N2 cadences has changed over the course of the mission.

11.4.1.9 (D)_N(1/2)_STATUS

All the STATUS extensions contain the same quantities for their respective coverage periods.

STAT00	STATINT	Status interval (seconds)
STAT01	MACBLCKS	Number of macro blocks fre
STAT02	TLMVOL	Telemetry volume produced
STAT03	WTCHADDR	Memory watch address
STAT04	WTCHMEM	Watched memory (pg. no)
STAT05	WTCHDATA	Watched memory
STAT06	PEPSWVER	Software version number
STAT07	ALARMID	Latest Alarm Id
STAT08	ALARMTYP	Latest alarm type
STAT09	ALARMCNT	Count of alarms
STAT10	CMDEXEC	Commands executed
STAT11	CMDREJCT	Commands rejected
STAT12	MACEXEC	Macro commands executed
STAT13	MACREJCT	Macro commands rejected
STAT14	MACROID	Id of most recent macro executed
STAT15	MACROLRN	Macro learn mode
STAT16	MONRESP	Monitor response
STAT17	WRITEENB	Memory write enable
STAT18	HVPSCURR	HVPS current
STAT19	HVPSVOLT	HVPS voltage
STAT20	BIASCURR	Bias current
STAT21	BIASVOLT	Bias voltage
STAT22	PEPSTAT	PEPSSI status word
STAT23	DVOLTP5	+5V digital voltage
STAT24	AVOLTN5	+5V analog voltage
STAT25	VOLTP2	+2.5V voltage
STAT26	VOLTN5	-5V voltage
STAT27	VOLTP15	+15V voltage
STAT28	VOLTN15	-15V voltage
STAT29	DCURRP5	+5V digital current
STAT30	ACURRP5	+5V analog current
STAT31	CURRP2	+2.5 volt current
STAT32	CURRN5	-5V current

STAT33	CURRP15	+15V current
STAT34	CURRN15	-15V current
STAT35	PRIMCURR	Primary current
STAT36	LVPSTEMP	LVPS temperature
STAT37	ENGYTEMP	Energy temperature
STAT38	HVPSTEMP	HVPS temperature
STAT39	ADDR12C	I2C read command address
STAT40	RSLT12C	I2C read command result

As of KEM-1 release 6.0, new STATUS quantities have been added to support a flight software update. At this time, they are not meaningful for the currently released data.

11.4.2 Bad Time Intervals (BTIs)

Various instrument conditions can make the PEPSSI data difficult or impossible to use for scientific purposes. Powering down, ramping the high voltage power up or down, running in diagnostic mode, etc. will all make the PEPSSI data unusable for standard analysis. The PEPSSI team delivers to the SOC a “PEPSSI_BTI_(*).txt” file¹ containing a comma-separated table of “Bad Time Intervals” (BTIs), which should be used to exclude data that are unsuitable for science analysis. It should be noted that the entire “Launch” phase of PEPSSI data are classified as a BTI. The columns in the PEPSSI_BTI_(*).txt file are:

- INTERVAL_IDENTIFIER – always “BTI” in this file, to indicate a Bad Time Interval. In other uses, the instrument team may include Good Time Intervals (GTI) in this file format.
- START_TIME – beginning of the interval in Year-DayOfYearTHour:Minute:Second format.
- END_TIME – end of the interval in Year-DayOfYearTHour:Minute:Second format.
- COMMENT – an optional comment.

These data are put in the file document/pep_bti.tab that is delivered to the PDS. Only START_TIME and END_TIME values are included in the pep_bti.tab file.

11.4.3 L3 Data Format

The L3 Files contain calibrated scientific data in an easily accessible form. There are three basic types of data in the L3 files: Quick-Look, flux-calibrated Rate Data, and calibrated PHA data. As with the L2 files, each file contains one UTC days’ worth of data. No Diagnostic mode data are present in the L3 files.

The Level 3 files are meant to be, as much as possible, self-documenting, and correspond to CODMAC Level 3 products. All calibration constants, calibration formulas, and physical units should be present in the FITS header in an easily readable format. It should be possible, albeit with a lot of work, to reproduce the Level 3 files independently from the Level 2 files using the information in the Level 3 headers.

¹ As of this writing, the most recent PEPSSI_BTI_(*).txt file is named: PEPSSI_BTI_GoodTimeIntervalsBadTimeIntervals_2022_149.txt

11.4.3.1 Primary HDU: Rate Weighted 2-D Histogram

The image in the primary array of the L3 file is a rate-weighted 2-D histogram of the PHA data for that day binned in calibrated deposited energy. It represents a “best available” overview of the day’s most detailed high energy ion data.

The priority scheme distorts ion abundances, so we correct for that by using a “rate-weight” rather than a single count. For each period of 600 seconds, we divide the counts reported in the N2 rate box by the number of PHA events observed in that rate box. This is the weight those events are then assigned in constructing the histogram (see Figure 11-8 and Figure 11-9 for comparison of weighted and unweighted histograms). The two axes: energy deposited in the SSD and time of flight, are simple linear calibrations of the measured values. The calibration parameters are reported in the primary FITS header.

11.4.3.2 Quick Look Spectrograms

IMPORTANT NOTE: These quick-look spectrograms do **not** represent best available views of the data. They are **only** presented to allow a single, uniform overview of general detector activity over the course of the whole mission. They rarely represent the best binning of the data. Users might also attempt to recreate these spectrograms themselves, as they are produced with a minimum of calculation, as a check on their understanding of the main data extensions.

The extensions: SPEC_Protons, SPEC_Helium, SPEC_Heavies, SPEC_Electrons, and SPEC_LowIon, contain quick-look spectrograms of their respective species. These spectrograms present counts/second N2 data, averaged over 60 second intervals and summed over all incidence directions (i.e. “Sectors”). During cruise, 60 seconds is, except for the first hour of the day, the default accumulation interval of N2 data (Before 2007 day 42, the default N2 accumulation interval was 30 seconds). The N2 accumulation interval has been changed from time to time. The FLUX extension contain the detailed record of the N2 setting in the DT column. The x-axis of the spectrograms is hour of day. On the y-axis each pixel represents a different rate (e.g. B00, L01, R02, etc.). Nominal deposited energies of the rate boxes (or, in the case of Low-Ions, nominal time of flight bins) are given in the FITS header.

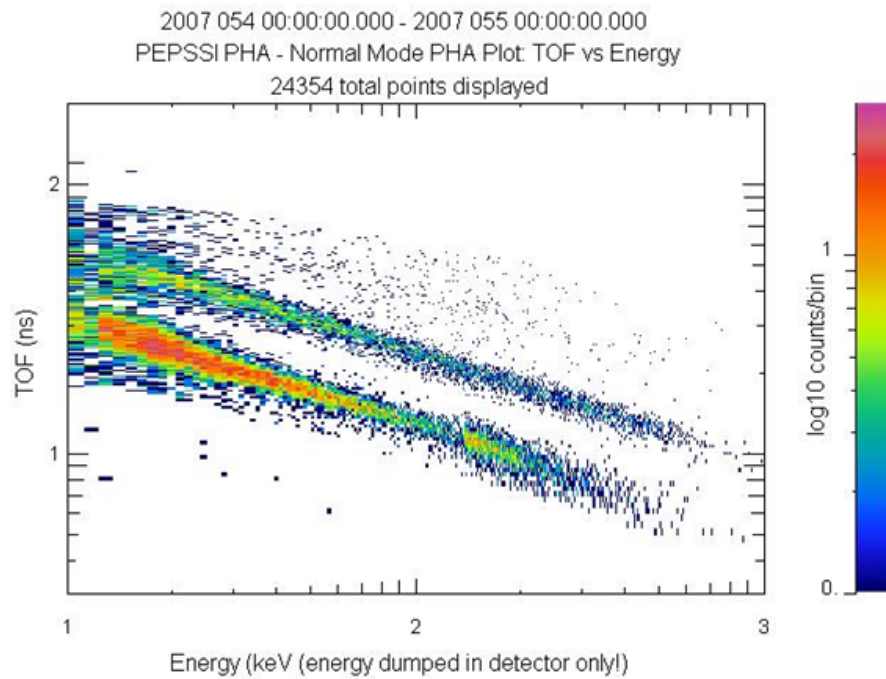


Figure 11-8: 2D PHA Histogram: No weighting. Note artifact in high energy protons (Box B03).




Figure 11-9: 2D PHA Histogram with Rate-Weighting applied.

11.4.3.3 FLUX, FLUX_N1A, FLUX_N1B

These HDUs contain calibrated fluxes, uncertainties, and raw counts/sec rates for all of the High and Low Energy Ion and Electron N1 and N2 Rate data. It also contains counts/sec rates for the various hardware and software rate counters from the N1 and N2 data. There is also an accumulation time column (**DT**) and three timing columns.

We do not provide calibrated fluxes for channels that are dominated by electronic noise or by background from other energies or species than what the channels nominally measure. These channels and their issues were described in Sec 11.4.1.7.

IMPORTANT NOTE: The UTC **YEAR** and **DAY_OF_YEAR** columns are only included for convenience in plotting. **DO NOT USE THEM** for precise timing as there could be leap-second ambiguities in them. Use the ephemeris time (**ET**) column if precision is important.

Some High Energy (aka “Triples” or “B”) Rate Boxes mix oxygen and sulfur. Separate calibrations are supplied assuming that the ions are purely oxygen or sulfur. The Rate Box name is modified by appending an “O” or an “S”, respectively, in the FITS table column name.

Even though the Rate Boxes associated with sulfur show counts in the interplanetary medium, these are mostly instrumental background counts that have no physical meaning, even after converting them to fluxes.

IMPORTANT NOTE: For Low Energy (aka “Doubles” or “L”) Rate Boxes, the particle species is not known without ambiguity since the data are based on TOF-only measurements. Therefore there are no fully calibrated fluxes provided. Instead, we provide corrected count rates (aka “efficiency*flux” values) that only account for the geometry factors of the instrument sectors and for the time-dependent portion of the efficiency (see 11.4.3.3.2). The corrected rates are also normalized to the energy range that the boxes are sensitive to. These energy ranges, as well as the ambient energies of each box, are under the assumption of various species, with suffix letters as given in Table 11-1.

For each High and Low Energy Rate Box, we also present a single “best guess” rate using whichever species the team judges to have been the “most likely” species for that rate in a column with no suffix letter in its name. This rate duplicates the “suffixed” rate.

All of the quantities used in the calibration of the flux measurement, including their uncertainties are included in keywords in the FITS header. A description of the calibration procedure follows.

As of KEM-1 release 6.0, The energy calibrations are much more accurate and include efficiency calculations to adjust for the time dependent changes to instrument sensitivity after power cyclings.

The FLUX extension contains N2 rate data. The FLUX_N1A extension contains N1 rate data for the “Triples”/High Energy/”B” rates. The FLUX_N1B extension contains N1 rate data for all of the other quantities collected in the N1 data accumulation interval. The separation of the N1 data into two extensions was required due to FITS format limitations on the allowed number of columns.

As of KEM-1 release 6.0, uncalibrated channels (e.g. B00,B17,B18,L00,L02) which were not present in the L3 data are now “carried over” from the L2 data and converted to counts per second.

11.4.3.3.1 FLUX Calibration Procedure

We calculate the differential intensity j ($1/\text{cm}^2 \text{ sr s keV}$) in terms of the counts C , time coverage T (s), geometric factor G ($\text{cm}^2 \text{ sr}$), upper and lower energy bounds E_{hi} and E_{lo} (keV), and detection efficiency η :

$$j = \frac{C/T}{G \Delta E \eta},$$

where $\Delta E = E_{\text{hi}} - E_{\text{lo}}$. We assume Poisson statistics for C , no error in T , absolute errors in G , E_{hi} , E_{lo} and relative error in η . So, formally we quote the counts as $C = C \pm \sqrt{C}$, the energies as $E = E \pm \delta E$, and the geometry factor is $G = G \pm \delta G$. We could also write $\eta' = \eta \pm \delta \eta$ but choose instead to define the relative error $\varepsilon \equiv \frac{\delta \eta}{\eta}$. Starting with the minimum and maximum efficiencies that describe a two sigma confidence band, η_{lo} and η_{hi} , we determine our best efficiency using the geometric mean $\eta = \sqrt{\eta_{\text{hi}} \eta_{\text{lo}}}$ and subsequently determine the relative error $\varepsilon = \sqrt{\frac{\eta_{\text{hi}}}{\eta_{\text{lo}}}}$. So, we can quote the efficiency and relative error as $\eta = \eta_{\pm}^{\times} \frac{\delta \eta}{\eta}$ or $\eta = \eta_{\pm}^{\times} \varepsilon$, meaning that the actual efficiency is between η/ε and $\eta\varepsilon$, to one sigma confidence.

With these given errors the formal error in j is given by:

$$\delta j = j \sqrt{\frac{1}{C} + \left(\frac{\delta G}{G}\right)^2 + \frac{(\delta E_{\text{hi}})^2 + (\delta E_{\text{lo}})^2}{\Delta E^2} + \left(\frac{\delta \eta}{\eta}\right)^2} = j \sqrt{\frac{1}{C} + \left(\frac{\delta G}{G}\right)^2 + \frac{(\delta E_{\text{hi}})^2 + (\delta E_{\text{lo}})^2}{\Delta E^2} + \varepsilon^2}$$

Here we have assumed that all errors $\delta x_1, \dots, \delta x_N$, are uncorrelated and have used the general expression for the error in a function $f=f(x_1, \dots, x_N)$:

$$(\delta f)^2 = \sum_{i=1}^N \left(\frac{\partial f}{\partial x_i} \delta x_i \right)^2$$

A “pseudo-code” version of the actual calculation code used is given in COMMENT keywords in the FITS header of the Level 3 files. The actual calibration parameters used in the production of a file are also given in the keywords for each flux value. NOTE: At the time of this writing, the constant uncertainties used for calibration are set to zero, so that only the Poisson error is calculated for the FLUX Uncertainty values in the Level 3 files.

Note on the correlation of the errors in E_{lo} and E_{hi} :

To zeroth order, we can treat the errors on E_{lo} and E_{hi} as uncorrelated. One pragmatic reason is that it is a conservative assumption; if we are wrong then we are *overstating* the errors (at most by a factor of $\sqrt{2}$). Also, there are times when the uncertainty in E_{lo} or E_{hi} will be quite uncorrelated. For example the E_{lo} could depend on our understanding of the SSD threshold, while the E_{hi} could depend on our estimate of how fast the spectrum will be falling, both very different things. The most problematic case (from an error propagation point of view) would be where we believe we know the passband width ΔE very well but we are not sure of the absolute energies. We think most of these potential cases are taken care of in the channel-to- $E_{deposit}$ calibration (which establishes the scale) whereas most of the potential cases of uncorrelated errors in E_{lo} and E_{hi} occur in the $E_{deposit}$ -to- $E_{incident}$ calibration.

11.4.3.3.2 Derivation and Explanation of Calibration Table Values

In the headers of the L3 files, we have supplied values to convert the instrument specific data (e.g., count rates) into physical instrument-independent units (e.g., differential intensity), as well as computing the physical quantities themselves. It must be stressed that these are not final values and calibration work by the instrument team is ongoing.

The calibration quantities are energy pass-band ($\Delta E \equiv E_{hi} - E_{lo}$, lower and upper limit of the energies of the particles measured), measurement efficiency (η , the fraction of valid incident particles that are actually measured), the geometry factor (G , the measurement of the physical detector size and solid angle subtended by the field of view). These values are all given and applied with uncertainties in the Level 3 files.

The calibration quantities for the B channels are now physically meaningful and their derivation is explained below. The fluxes might still change in the course of calibration efforts of the team.

The R channels nominally measure electrons and the calibration quantities assume this. This is applicable in the Jupiter environment. At all other times, the R channel count rates were found to be dominated by penetrating particles consistent with cosmic rays. The current calibration does not account for this, so that the provided R channel intensities near Pluto and in the interplanetary medium can only be used to track relative changes of cosmic rays, not to measure electrons.

The L channels do not distinguish ion species since they solely rely on TOF measurements without SSD information. Attempts to deconvolve their response to different species to back out their respective intensities are ongoing. We currently do not provide detection efficiencies for the L channels. Formally, the detection efficiencies of each sector are set to 1. Even though this formally yields units of flux, this value is not a final calibrated value as the energy and sector dependent efficiency has not been applied. The different L channel sectors (like L01S00, L01S01, etc.) show electronic cross talk that makes it difficult to determine the true intensity distribution as a function of direction. Comparisons of different sectors are therefore not meaningful, independent if count rates or the provided formal fluxes are used. Best use of the L channels is to

focus on a single sector with high count rates (SUnknown before 2014 DOY 178, S00 afterwards) and use it throughout.

Generally we distinguish deposited, incident, and ambient energies. Deposited energies are measured by the instrument's SSDs. The particles have incident energy before entering the instrument aperture that is kept at $\approx 3\text{kV}$ to repel electrons. The ambient energy is the energy that the particles have far away from the spacecraft. The L3 files provide ambient energies. We determine this ambient energy based on the measured TOF or deposited energy by calculating the energy loss in the start foil (for the L rates), as well as in the stop foil and SSD dead layer (for the B and R rates) and the $\approx 3\text{kV}$ accelerating potential. This calculation is based on tabulated values of differential energy loss. The deposited energy is calculated similarly but based on tabulated values of "electronic"/ionization energy loss since this is what is measured by the SSD. The electronic thresholds of the SSDs determine the lowest deposited energy that each of PEPSSI's sectors can measure. These thresholds differ between the sectors but the calibration accounts for this. The lowest energy of each sector and channel was determined from PHA_HIGH_ION event data accumulated through the mission.

The geometry factors are derived based on the formalism described in Sullivan, 1971, Nucl. Instr. and Methods, using a simplified instrument geometry and Monte-Carlo technique.

The PEPSSI flight unit was calibrated pre-launch to convert the measured deposited energy- and TOF-values to physical units. For a long time there were no reliable efficiencies available because there was only time to measure two efficiency points that were of limited use. Significant effort was put into deriving the missing efficiencies.

The current helium efficiencies are based on a combination of in-flight measurements and calibrations with the engineering model (EM). Non helium-efficiencies are based on a combination of in-flight measurements and theory. For the helium calibration, we used a degraded alpha source that was measured both with the PEPSSI EM and an independently calibrated SSD.

The theoretical approach is similar to what is used for other PEPSSI-like "Puck" instruments as RBSPICE or JEDI. It assumes that the efficiency is proportional to the secondary electron yield of ions passing the start and stop foils. This yield is usually proportional to the differential energy loss of the ions in the foil material that faces the MCP. This was confirmed for protons and helium passing the start foil (Smidths et al., 1992, Nucl. Inst. and Methods; Baragiola et al., 1978, Phys. Rev. B).

The second efficiency factor is from particles that scatter after the start foil in a way that they miss the stop foil and/or the SSD. However, intercalibration of RBSPICE with HOPE showed that theory overestimates the scattering factor. Additionally, we found that neglecting the scattering factor makes the theoretical helium efficiency similar to the measured efficiency. The used theoretical efficiency as we use it is therefore only a product of the electron yields in the start and stop foil. This is not critical since we only provide efficiencies for the B channels that have high enough energies that scattering is not important in any case.

Both the theory as the measurements with the EM only provide the energy dependence of the efficiency, not absolute values. This scaling is done based on in-flight measurements in

diagnostic mode. In this mode, all particles hitting the SSDs are counted (instead of requiring a coincidence with both start and stop counts). SSDs are 100% efficient at energies well above their threshold. In an environment where the counts are dominated by a single ion species it is therefore possible to use the SSDs as an independent, calibrated measurement that can be compared with the SSD-TOF-coincidence measurement. Such an “environment” is usually given in sectors 4 and 5 that are irradiated by an internal alpha source that is usually much stronger than the radiation from the environment. Before Jupiter we found that the counts of other sectors were dominated by protons (around 2007 DOY 39). Around Jupiter they were dominated by sulfur (around 2007 DOY 90). To reduce the contribution from electrons, photons, and electronic noise, we only consider SSD counts that are in coincidence with at least a start or a stop count. The absolute efficiencies of other species (like oxygen) are interpolated (along a relative theory efficiency axis) based on the simultaneously measured absolute helium and sulfur efficiencies.

PEPSSI's efficiency is changing on a timescale of about 3 months after each time it was switched off. This is tracked by the channels and sectors that measure the internal alpha source. All intensities were corrected for this, even though there may be residual trends or artifacts from this process.

On top of this, the efficiencies change because electronic thresholds or rate box definition tables were changed. We therefore distinguish several calibration epochs that are accounted for in the calibration.

PEPSSI's sectors have different efficiencies. This was tested in-flight by pointing each sector to the same direction relative to the solar wind flow and comparing the response. For the B channels we use the test that was performed around 2007 DOY 287, where the high-energy count rates were high enough to accumulate enough counts for a significant result. For the L channels we use the test that was performed around 2015 DOY 8, after the electronic thresholds were optimized to distinguish directions well (see Section 11.4.1.7). The result is applied to the entire mission but might not be reliable for the L channels before 2014 DOY 178.

Note that even though significant effort was put into providing absolute intensities that there are still uncertainties. Comparisons between different energies (energy spectra) or species (abundance fractions) need to consider what effect artifacts in the calibration can have.

Note that the PEPSSI instrument was specifically engineered to make low rate measurements, which is a given in the interplanetary and Pluto environment, but not near Jupiter. For example, the CPU was selected for its low power consumption, which means that there is an upper limit to the total number of events that can be processed. Therefore the user has to be aware that saturation of the rates can take place, especially within about $100R_J$ of Jupiter. This saturation does not have to be uniform across different rates and over time. It is possible during high rates of triple coincidence ions to impede the processing of electrons. None of this is not accounted for by the currently provided efficiencies.

11.4.3.4 PHA Data

The three PHA extensions: PHA_ELECTRON, PHA_LOW_ION, and PHA_HIGH_ION contain the PHA event data. As in the L2 data, each row represents a single PHA event. Events with the multi-hit (cross talk) flag set have been excluded. Quantities of limited usefulness (such as Heavy Ion Discriminator triggers) have been excluded. Because of the difficulty of removing priority scheme biases from non-N2 PHA data, only N2 (ApID = 0x692) PHA data are present in the L3 files.

Calibrated Deposited Energy and/or Time of Flight values are given. The linear calibration constants and formulas are in the FITS headers. A **Speed** column is calculated from the Time of Flight assuming a 6.0cm flight path.

The Rate Box classification for each event is given in the **Rate_Box** column. Special cases follow:

- **R00A and R00B:** R00 was originally the "dump bin" channel for electrons. Early in the mission it had a low-energy and a high energy band ("A" and "B" respectively). Since we raised the threshold, there is now only the high energy band in practice, that is, we will never see electrons in the R00A band. "R00A" and "R00B" are used in the L3 PHA data column ("Rate_Box") that indicates how the flight software categorized the event. The actual channel in N1 or N2 (L2) or FLUX (L3) extensions only exists as "R00".
- **SAll:** SAll means all Sectors (0-5) averaged together.

The PHA_HIGH_ION extension contains additional columns:

- The **H_Incident_Energy**, **He_Incident_Energy**, **O_Incident_Energy**, and **S_Incident_Energy** columns contain the calculated Incident energy assuming that the event is of that (H, He, O, or S) species.
- The **Rate_Normalized_Weight** column has removed Priority Group artifacts from the PHA data by the procedure described in the Primary HDU section above. This column is usually used in making histograms of the High Energy Ion PHA data.

As of KEM-1 release 6.0, more information has been included in the L3 PHA data. Multi-hit data is now included as are the non-N2 PHA data. New columns are:

- Low Energy Ion PHA Events now contain a column (calculated on the ground) that reports what number energy bin they would have been classified into ('Lo_Bin').
- Rate Normalized Weight and Rate Normalized Weight with 1-day integration are now included for High Energy Ion PHA Events. The instrument team now feels these quantities are understood well enough to include.
- 'Raw' energy and time of flight values are included (in instrument units). These values are just copied from the raw data files, for ease of use.

11.4.4 Memory Required

1 GB memory and a 3 GHz Pentium is sufficient for processing.

11.4.5 Temporary File System Space Needed

The Level 2 files require up to 50MB per days' worth of data.

11.4.6 Predicted Size of Output File(s)

Level 2 files are in the range 1MB – 30MB. Level 3 files are typically 5-6 times larger than the corresponding Level 2 file, and can range in size up to 600MB.

11.4.7 Predicted Execution time

Less than a minute per file, typically. The L1 to Pre-L2 conversion takes a few seconds per file. The entire Jupiter phase takes 40 minutes to convert from Pre-L2 to L2 and L3 files on a Red Hat Linux machine with 4 4GHz Xeon processors. It's not known how much parallelization was actually responsible for the speed.

11.4.8 Contact/Support Person(s)

Matthew Hill, Larry Brown

11.4.9 Maintenance Schedule (Code/Data Updates, Documentation)

As calibration data are collected during flight, the Level 2 pipeline code requires updates either to calibration files or to code for bug fixes or enhancements.

12. REX INSTRUMENT DESCRIPTION

12.1 Overview

The primary purpose of the Radio Experiment (REX) system is to investigate open questions regarding the atmospheric and ionospheric structure, surface conditions, and planetary radii of both Pluto and Charon.

The REX instrument is unique among the suite of instruments comprising the *New Horizons* payload in that it is physically and functionally incorporated within the spacecraft telecommunications subsystem. REX consists of both a ‘Flight Element’ carried on-board the *New Horizons* spacecraft, and a ‘Ground Element’ comprising the existing Earth-based NASA Deep Space Network transmitting and operations facilities, which radiate an uplink signal that is sampled and measured by REX on-board the spacecraft.

REX is designed to fulfill the mission objectives by performing the following distinct experiments:

1. A radio occultation experiment, designed to detect and measure the atmosphere, ionosphere and size of Pluto, and of Charon where detectable. The spacecraft was maneuvered so that Pluto, and then Charon, passed through the line of sight from Earth to the spacecraft. That line of sight is coincident with the uplink signal path from the Ground Element, and so Pluto and Charon effect changes in the signal received and measured by REX. The experiment detects phase differences between the uplink signal as measured by REX, and a modeled signal which would be expected if there were no occultations. Subsequent analysis of the behavior of those phase differences over time, combined with knowledge of the ephemerides of Pluto, Earth and the spacecraft, is used to deduce atmospheric temperature and pressure profiles down to the surface of Pluto (and of Charon, should it be found to have a sensible atmosphere), as well as electron density profiles in the ionosphere of Pluto (and possibly Charon). Analysis of the on-board-measured uplink signal strength as the spacecraft goes in and out of occultation is used to deduce the chord lengths of the line of sight path across the disks of Pluto and Charon.
2. A gravity experiment, designed to measure the independent gravitational fields of Pluto and possibly of Charon. Around the time of the occultation experiment, but with the line of sight to Earth well away from the limbs of Pluto and Charon and above any atmosphere or ionosphere, accurate tracking of Doppler shifts in the received frequency of the uplink signal, again compared to a model of what would be expected if there were no body-induced gravity fields nearby, is used to deduce gravitationally-induced changes in velocity along the spacecraft’s flight path.
3. A radiometry experiment, designed to measure the spatially averaged surface emission brightness at a wavelength of 4.2 cm (7.182 GHz, the nominal operating frequency of the *New Horizons* radio) of Pluto and of Charon. The dark-side emissions are measured during the occultation interlude. The day-side emissions are measured as is operationally feasible.

In those first two experiments the on-board REX system hardware downconverts and samples an uplink signal from the Ground Element (Earth-based Deep Space Network; DSN) for the purpose of analyzing the frequency and phase behavior of that signal compared to an on-board oscillator. In the third experiment the on-board hardware measures (integrates) the radiometric emission of a source.

Note that this section's focus, and the PDS archive data sets produced from REX data, comprise samples taken and measurements made *on-board* the spacecraft *by REX* either of a one-way uplink signal from the Ground Element (DSN), or of a 4 cm-wavelength (7.2GHz) radiation from a planetary surface, from an astronomy source, or from cold sky.

As of about a year before the Pluto Encounter, almost all REX observations using one-way uplink data to-date were instrument functional tests, instrument characterizations, and instrument calibrations, so tuning profiles of the one-way uplink data to analyze those observations are not provided. The only exceptions are Lunar Occultations in May of 2011 and January of 2012; Tracking and Navigation Files (TNFs; TRK-2-34 files) with uplink tuning profiles will not be covered here, but have been delivered with PDS data sets containing the lunar occultation data; these data are described by PDS labels and by DOCUMENT/TNFSIS.LBL in those data sets. The TNFs for the Pluto encounter including the Pluto and Charon occultations are provided as part of the Pluto phase PDS delivery.

Finally note that Earth-spacecraft-Earth non-coherent (a.k.a. one-way mode) tracking data, with measurements made *at the Ground Element (DSN)* of signals sent from *New Horizons*, are *not processed* by the SOC pipeline, are *outside the scope* of this document, and are *not included* in any REX PDS data sets.

The heart of the REX instrument is an Actel Field Programmable Gate Array (FPGA) that takes samples of the downconverted & digitized intermediate frequency (IF) receiver output and generates wideband radiometer and narrowband sampled signal data products. The REX hardware also includes an analog-to-digital converter (ADC) and other direct interface components, and by extension all of the RF telecommunications system hardware along the uplink (receive) path from the High Gain Antenna (HGA) to the input to the ADC.

Stanford is responsible for the FPGA design and system analysis. APL is responsible for the design of the telecommunications system and incorporating the REX FPGA system therein.

The interfaces to the REX FPGA (see Figure 12-1) include a 30 MHz clock signal from the Ultra-Stable Oscillator (USO), the secondary power connections, the command and telemetry data interfaces to the Uplink Card, the high-speed data interface to the Instrument Interface Card, a 1 PPS signal for data framing, and the interface to the ADC, where the wideband IF signal from the Uplink Card is sampled.

The input to the REX FPGA is normally the uplink signal from the DSN after being filtered by a 4.5 MHz bandpass filter (not shown) and digitized by the ADC at a sample rate of 10 Msamples/s. The Input Select function, commandable via uplink, allows the FPGA to process any of seven predetermined digital signals for testing the FPGA functionality (see the ROF Status Byte section below).




Figure 12-1: Electrical Interfaces to the REX FPGA

REX is part of the redundant spacecraft telecommunication subsystems and signal paths that use the single HGA in common. The two REX cards are designated as Sides A and B (also Channels A and B). Sides A and B can be operated simultaneously, to increase SNR, using uplink signals with RCP and LCP (Right and Left Circular Polarization), respectively. Normally Side A communicates with spacecraft CDH1 (Command and Data Handling), and B with CDH2, but that can be switched if required by spacecraft events. There are two USOs, and each REX side is referenced to a separate USO, and that must be considered when using the data. The USOs are also cross-strapped so either can provide timing to both sides in the event of a single USO failure. Execution of the command to do so would be a one-time, irreversible event, and as of the end of 2015 that has not occurred.

12.2 Level 1 Data Specifics

After REX receives a power-on command, on the next 1PPS strobe from the spacecraft clock the REX FPGA starts generating a continuous stream of data containing In-Phase & Quadrature-Phase value pairs as well as integrated radiometer values. This stream of data is divided into fixed-length units called REX Output Frames (ROFs) at a rate of one ROF per 1.024s (s implies seconds on the spacecraft clock in this section). The ROFs are stored on the spacecraft solid state recorder (SSR), and eventually played back via the High-Speed Telemetry interface to the DSN and arrive at the SOC as raw telemetry packets.

REX continues to generate ROFs until turned off. Each ROF also contains Time Tags that may be used to verify that a sequence of ROFs is a contiguous set. If REX is turned off while writing data to the SSR, the data at the end of the last ROF in a contiguous sequence of ROFs *may* comprise random bits. It is up to the user to assess where this has occurred; in general the last ROF of a sequence can be ignored. Depending on the mode (e.g. compression) of communication with C&DH, similar issues have been seen in the first ROF in a sequence, so those should be examined as well before use.

The In-Phase, Quadrature-phase, Radiometer, and Time Tags are all multi-byte values, and the individual bytes for all quantities are interleaved throughout the ROF.

12.2.1 Level 1 Data Format

The SOC Level 1 pipeline decommutates each ROF from telemetry and places it into the Primary Data Unit (PDU) of an individual FITS file. Each PDU is stored as a one-dimensional image of 5088 bytes: the first 5082 bytes are the ROF; the last 6 bytes in the PDU are spare.

The SOC Level 1 pipeline also looks in the telemetry for packets corresponding to the time of the ROF, and places them in Extension Data Units (EDUs). Specifically, data from spacecraft housekeeping ApIDs (packet Application Process Identifiers; also Application IDs) 0x004, 0x016, 0x084 and 0x096 as well as from Thruster packets are placed in EDUs 1 through 5, respectively, of the Level 1 FITS files.

12.2.1.1 PDU Content

Each ROF contains the items listed in Table 12-1, with the individual bytes in an interleaved format:

Table 12-1: REX Output Frame Contents

Item	Item Description	Count/ROF	Bytes/item	Total bytes
ID byte	ROF Identifier = 0xB7	1	1	1
Status byte	Input select	1	1	1
Radiometry	40-bit integrated power; reset each ROF	10	5	50
Time tag	24-bit accumulator	10	3	30
I&Q value pair	In-phase & Quadrature @ 16 bits	1250 (pairs)	(2 + 2 =) 4	5000

12.2.1.1.1 ROF ID byte

The ID byte is the first byte in the ROF and should always have the same value; see Table 12-2:

Table 12-2: REX Output Frame ID byte value

Unsigned Decimal	Hexadecimal	Binary
183	B7	1011 0111

12.2.1.1.2 ROF Status byte

The ROF Status byte is the fourth byte in the ROF. In bit positions 6, 5 & 4 (bit position 7 is the most significant bit = $128 = 2^7$; bit position 0 is the least significant bit = $1 = 2^0$) it contains the three bits that make up the Input Select setting for the ROF; all other bits are normally zero, otherwise corruption should be suspected. When Input Select is set to any of its non-zero values, the ADC output is replaced as the FPGA input with a predetermined 10Msample/s Test Pattern signal as described in Table 12-3. In that case, the output of the REX FPGA should be deterministic and known, and may be compared bit-for-bit against the expected output as a limited check on the health of the FPGA as well as that of the Input Select system. In general during normal operations, at power-on the first ROF status is default, the next few ROF's are [All zeros] or another Test Pattern, and subsequent ROFs have a default status.

Table 12-3: REX Input Select (ADC vs. Test Patterns) & Status Byte values

Input Select (binary)	Status Byte (binary)	Input Select description
000	0000 0000	ADC output from NH receiver system (default)
001	0001 0000	Impulse: 2 samples (200ns) of value 128 at the start of each ROF, followed by zeros
010	0010 0000	Low-Frequency Square Wave: +/-256 @ 610.3515625 Hz
011	0011 0000	Mid Frequency Square Wave: +/-256 @ 19.53125 KHz
100	0100 0000	Pseudo-Random Number (PRN) of value +/-1 @ 10 MHz
101	0101 0000	Pseudo-Random Number full scale @ 10 MHz
110	0110 0000	Hi Frequency Square Wave: +/-256 @ 78.125 KHz)
111	0111 0000	All zeros

12.2.1.1.3 Integrated Radiometry values

The details of how incoming power is used as radiometry are given in Tyler et al., 2008.

The FPGA integrates the incoming power from its input signal by squaring and summing the ~10Msamples/s that compose its input. Ten accumulating radiometry values are stored in each ROF, and the FPGA resets the value to zero at the start of each ROF. Each radiometry value comprises 40 bits, or 5 bytes, as an unsigned integer, and the bytes are in MSByte-first order, interleaved with I&Q values.

The time interval between radiometry values is one-tenth of a ROF or 102.4ms. In each ROF, REX stores an integrated radiometry value at the start time of that ROF (and not 102.4ms after its start), so the tenth, or last, radiometry value associated with a ROF (i.e. the one that represents a full 1.024s of integration time) is actually stored as the first radiometry value in the following ROF.

12.2.1.1.4 Time Tag values

REX places ten incrementing time tags in each ROF. The first time tag of the first ROF after a start command is zero, and following time tags increment by one. The time tag is not reset at the start of each ROF. Each time tag value comprises 24 bits, or 3 bytes, as an unsigned integer, and the bytes are in MSByte-first order, interleaved with I&Q values. Each increment of the time tag represents 102.4ms. The rollover time is about a fortnight and a half (16.7Msamples) and will never be reached in practice.

The time tags can be used both to identify any breaks in a sequence of ROFs, and to determine the time between any two ROFs within a sequence.

12.2.1.1.4.1 Time tag anomalies in ROF sequences

The normal sequence for time tags is to start at zero in the first ROF and increment ten times per ROF, so the first time tag of the second ROF is 10, that of the third ROF is 20, etc. The time tags continue to increment until either REX is turned off or the high-speed allocation for the observation in the Solid-State Recorder (SSR) is full. In practice, the first few and last few

ROFs in a sequence do not always show simple zero starts and clean finishes, respectively; the first time tag may be 10, or it may continue from the previous ROF sequence before resetting to zero; the time tags in the last ROF often start correctly for the first few values but then drop to zero for the final values. This can be due to the compression issue mentioned elsewhere, but has also been observed with uncompressed data. These time tag anomalies indicate data corruption in just those starting and ending ROFs, but there is no indication of corruption elsewhere within ROF streams. REX commanding ensures there are always adequate ROFs before and after any observation, so discarding starting and ending ROFs in a sequence based on simple inspection of time tags is the way to handle this issue.

12.2.1.1.5 I & Q value pairs

Each ROF contains 1250 pairs of In-Phase (I) & Quadrature-Phase (Q) values. Each I value and each Q value comprises 16 bits or two bytes as a two's-complement signed value.

The process of down conversion from 10 Msamples/s is accomplished by heterodyning to zero frequency the uplink carrier signal centered initially at the 2.5MHz Intermediate Frequency (IF) center frequency, followed by use of time-invariant baseband filters to reduce the bandwidth. The details are too extensive to include here, but are explained in detail in Tyler et al. (2008).

12.2.1.2 Method 1 - values' bytes' layout within ROF: Interleaving

In each ROF, the bytes of the ID, Status, Radiometry and Time Tag values are interleaved with the I&Q value pairs, but none of the values start or end on other than a byte boundary. There are two methods to describe such an arrangement: describe the layout of the bytes of the data values as a sequence of bytes from the ROF; describe the layout of the ROF as built out of a sequence of bytes extracted from the de-interleaved data values (ID, Status, Radiometry, Time Tag, I&Q). Both methods will be described here, the former first as it lends itself more easily to writing computer code to build the data values by extracting bytes from the ROF. Indeed, IDL(tm) and Python routines to de-interleave ROF data exist each with less than two dozen statements.

Table 12-4: Data values' bytes' positions and offsets within an ROF. For example, there are 10 Radiometry values, each comprises 5 bytes (1st column, Radiometry row). The first Radiometry value comprises the 7th (2nd column), then 10th, 13th, 16th, and 19th bytes (3rd column) of the ROF. The second Radiometry value comprises the 515th (7+508; 4th column), then 518th, 521st, 524th, 527th bytes of the ROF.

Item, Item Size (bytes), Item Count	First (MSB) byte of first value (1-based)	Offset to succeeding byte(s) within an Item	Byte Offset(s) to successive Items	First (MSB) byte of first value (0-based)
ID byte, 1, 1	1	N/A	N/A	0
Status byte, 1, 1	4	N/A	N/A	3
Radiometry, 5, 10	7	3	508	6
Time Tag, 3, 10	22	3	508	21
I, 2, 1250	2	1	6 & 4	1
Q, 2, 1250	5	1	6 & 4	4

The following sections use the notation in Table 12-5 to specify and/or locate the various quantities' bytes within the ROF (Method 1), or vice versa (Method 2):

Table 12-5: Notation used in this section

ROF[N]	Nth byte of ROF (N=1 to 5082), interpreted as an unsigned integer (Range is 0 to 255)
R[I]	Ith Radiometry value; for I=1 to 10
T[J]	Jth Time Tag value; for J=1 to 10
I[K] or Q[K]	Kth In-phase (I) or Quadrature-phase (Q) value, for K=1 to 1250
X[M:N]	Bits M through N of multi-byte quantity X. E.g. R[10][31:24] second byte of R[10]
Y[MSB]	The MSByte, as an unsigned integer, of a 16-bit signed IQ value Y; Y[MSB] = Y[15:8]
Y[LSB]	The LSByte, as an unsigned integer, of the a 16-bit signed IQ value Y; Y[LSB] = Y[7:0]
O _i	A useful offset quantity equal to (508 * (i-1)) ; for i = 1 to 10

12.2.1.2.1 Method 1 - Layout of ID & Status bytes

The ID and Status bytes are the first and fourth bytes in the ROF, respectively, and can simply be obtained from the ROF as such as there are no following bytes or values:

$$\text{ID byte} = \text{ROF}[1]$$

$$\text{Status byte} = \text{ROF}[4]$$

12.2.1.2.2 Method 1 - Layout of Radiometry bytes

The first byte of the first radiometry value is the seventh byte of the ROF, and the following four bytes of that first radiometry value are each offset by three bytes from the previous byte. The order is MSByte-first. So, the first radiometry value of a ROF may be calculated from the following formula:

$$R[1] = (((\text{ROF}[7] * 256 + \text{ROF}[10]) * 256 + \text{ROF}[13]) * 256 + \text{ROF}[16]) * 256 + \text{ROF}[19])$$

The following radiometry values' bytes in the same ROF are each offset 508 bytes from the previous radiometry value's bytes. So, more generally:

$$R[i] = (((\text{ROF}[7+O_i] * 256 + \text{ROF}[10+O_i]) * 256 + \text{ROF}[13+O_i]) * 256 + \text{ROF}[16+O_i]) * 256 + \text{ROF}[19+O_i])$$

12.2.1.2.3 Method 1 - Layout of Time Tags' bytes

The first byte of the first time tag value is the 22nd byte of the ROF, and the following two bytes of that first time tag value are each offset by three bytes from the previous byte. The order is MSByte-first. So, the first time tag value of a ROF may be calculated from the following formula:

$$T[1] = (\text{ROF}[22] * 256 + \text{ROF}[25]) * 256 + \text{ROF}[28]$$

The following time tag values' bytes in the same ROF are each offset 508 bytes from the previous radiometry value's bytes. So, more generally:

$$T[i] = (\text{ROF}[22+O_i] * 256 + \text{ROF}[25+O_i]) * 256 + \text{ROF}[28+O_i]$$

12.2.1.2.4 Method 1 - Layout of I & Q values' bytes

The bytes that are used to store the I & Q values alternate in sequence (I[1], Q[1], I[2], Q[2], I[3], Q[3], ..., I[1250], Q[1250]) as 16-bit MSByte-first two's complement signed integers, that cover all of the bytes not used by the other values (ID, Status, Radiometry, Time Tags) above.

Specifically, Table 12-6 describes how ROF bytes are used to calculate I[K] and Q[K] for K=1 to 1250:

Table 12-6: I & Q values from ROF

KMOD	$KMOD = (K+123) \text{ MOD } 125$	
j	$j = (K+123-KMOD) / 125$	
Oj	$Oj = 508 * (j-1)$	
K=1	$I [K] = IQ(\text{ROF}[2],$	$\text{ROF}[3])$
	$Q [K] = IQ(\text{ROF}[5],$	$\text{ROF}[6])$
K>1 & KMOD<4	$I [K] = IQ(\text{ROF}[8+KMOD*6+Oj],$	$\text{ROF}[9+KMOD*6+Oj])$
	$Q [K] = IQ(\text{ROF}[11+KMOD*6+Oj],$	$\text{ROF}[12+KMOD*6+Oj])$
K>1 & KMOD>3	$I [K] = IQ(\text{ROF}[15+KMOD*4+Oj],$	$\text{ROF}[16+KMOD*4+Oj])$
	$Q [K] = IQ(\text{ROF}[17+KMOD*4+Oj],$	$\text{ROF}[18+KMOD*4+Oj])$

Where Function IQ(MSByte,LSByte) is defined as (with MSByte & LSByte interpreted as unsigned 8-bit - i.e. 1-byte - integers)

$$IQ(\text{MSByte}, \text{LSByte}) = 256 * \text{MSByte} + \text{LSByte}$$

if MSByte is between 0 and 127 inclusive. Otherwise, it is defined as

$$IQ(\text{MSByte}, \text{LSByte}) = 256 * \text{MSByte} + \text{LSByte} - 65536$$

12.2.1.3 Method 2 - Layout of ROF

The tables below have one cell per ROF byte, and indicate which data values (ID, Status, Radiometry, &c) each byte contributes to. The order of bytes in these table is left-to-right and down i.e.

ROF[1]	ROF[2]	ROF[3]
ROF[4]	ROF[5]	
ROF[6]	ROF[7]	
...	...	

The first six ROF bytes contain the ID & Status bytes and the first I&Q pair:

ID	I [1] [MSB]	I [1] [LSB]
Status	Q [1] [MSB]	Q [1] [LSB]

The next 508-byte "chunk" (again, ROF byte order is left-to-right then down) contains one Radiometry value, one Time Tag, and 125 I&Q pairs:

R [1] [39:32]	I [2] [MSB]	I [2] [LSB]
R [1] [31:24]	Q [2] [MSB]	Q [2] [LSB]
R [1] [23:16]	I [3] [MSB]	I [3] [LSB]
R [1] [15:8]	Q [3] [MSB]	Q [3] [LSB]
R [1] [7:0]	I [4] [MSB]	I [4] [LSB]
T [1] [23:16]	Q [4] [MSB]	Q [4] [LSB]
T [1] [15:8]	I [5] [MSB]	I [5] [LSB]
T [1] [7:0]	Q [5] [MSB]	Q [5] [LSB]
I [6] [MSB]	I [6] [LSB]	
Q [6] [MSB]	Q [6] [LSB]	
I [7] [MSB]	I [7] [LSB]	
Q [7] [MSB]	Q [7] [LSB]	
I [8] [MSB]	I [8] [LSB]	
Q [8] [MSB]	Q [8] [LSB]	
...	...	
I [126] [MSB]	I [126] [LSB]	
Q [126] [MSB]	Q [126] [LSB]	

The rest of the ROF comprises 9 more chunks of 508 bytes per chunk, essentially identical to the one above, incrementing the R[] & T[] indices by one per chunk, and incrementing the I[] & Q[] indices by 125 per chunk. Each chunk except the last contains 125 I&Q pairs; N.B. the tenth chunk ends after its 504th byte and after its 124 I&Q pair which is the 1250th, and last, I&Q pair of the ROF.

12.2.2 Data Sources (High/Low Speed, CCSDS, ITF)

REX data are in the high-speed stream and come to the SOC in CCSDS (Consultative Committee for Space Data Systems) packets.

12.2.3 Definition of an “Observation”

One REX Output Frame (ROF), as defined above, is an observation.

12.2.4 Data and housekeeping in Level 1 FITS files' PDU and EDUs

The Level 1 pipeline puts the raw ROF into the Primary Data Unit (PDU; a.k.a. EDU “0”), and puts ROF-derived information (I/Q pairs, Radiometry, Time Tags), plus several types of housekeeping (HK) packets, as FITS binary tables (BINTABLEs) into Extension Data Units (EDUs) 1 through 8 (see Table 12-7). The pipeline attempts to find the closest HK packet to the observation time. If no packet is available, the EDU data are not present and the corresponding Extension Header Unit (EHU) will indicate a zero-sized EDU.

Table 12-7: REX Level 1 (and Level 2 calibrated; see Section 12.3.2 below) FITS PDU and EDU descriptions: Numbers; Names; Calibration relevance; Descriptions. Level 1 extensions to be calibrated have “Yes” in “Cal” column. The PDU is a.k.a. EDU 0. The PDU and Extensions 3-8 may typically be ignored by PDS users.

#	<u>EXTNAME</u>	<u>Cal</u>	<u>Description</u>
0	N/A	No	Raw ROF (REX Output Frame)
1	I AND Q VALUES	Yes	1250 (In-, Quadrature-phase) value pairs
2	RADIOM. AND TIME	Yes	Radiometer and Time Tag values (10 each)
3	HOUSEKEEPING_0X004	No	Phase Locked Loop HK from ApID 0x004
4	HOUSEKEEPING_0X016	No	REX HK from ApID 0x016
5	HOUSEKEEPING_0X084	No	Phase Locked Loop HK from ApID 0X084
6	HOUSEKEEPING_0X096	No	REX HK from ApID 0x096
7	THRUSTERS	No	Thruster HK from ApID 0x124
8	SSR_SECTOR_HEADERS	No	SSR sector header information

12.2.4.1 FITS keywords added to PDU (a.k.a. EDU 0 or DU 0)

The Status byte value is added to the PHDU (FITS header of the PDU).

12.2.4.2 FITS BINTABLE specifics

12.2.4.2.1 Extension Data Unit (EDU) 1 - I & Q values

The In-phase and Quadrature-phase values are each combined from two ROF bytes as described above and stored as signed 16-bit values in their respective columns (TFORM_n = I; TFORM_n is a keyword in EDU BINTABLE header and declares the format of table column n; a value of I indicates a 16-bit integer; refer to FITS references given in Section 3 for more detail).

The BINTABLE comprises 1250 rows, and each row comprises two columns laid out over four bytes.

12.2.4.2.2 EDU 2 – Radiometry and Time Tags

The Radiometry values are each combined from five ROF bytes as described above, and stored within signed 64-bit integers in the first column of the BINTABLE (TFORM1 = K).

The Time Tag values are each combined from three ROF bytes as described above, and stored within signed 32-bit integers in the second column of the BINTABLE (TFORM2 = J).

The BINTABLE comprises 10 rows, and each row comprises two columns laid out over twelve bytes.

12.2.4.2.3 EDUs 3 through 8 – Housekeeping

The remaining BINTABLES contain housekeeping information and are there for historical reasons and for diagnostics performed by the science team, and will not be described here. PDS users may safely ignore them.

12.2.5 Note for PDS users reading FITS files

There are many FITS libraries that render trivial the reading of data from any PDU or EDU in a FITS file.

If however the PDS user, for whatever reason, still desires to write their own access routines, then we recommend they should instead use the detached PDS labels as a data map. PDS labels provide, in both machine- and human-readable form, a simple arithmetic path to interpret the bytes in these FITS files.

12.2.6 Raw Science Data and/or Housekeeping Requirements

Radio receiver housekeeping (ApIDs 0x004 and/or 0x084 noted above).

12.3 Level 2 Data Specifics

12.3.1 Calibration Algorithms

The conversion of Level 1 REX data to Level 2 calibrated data is concerned with several data streams from REX:

- (1) the REX filter output, comprising 16-bit samples at 1250 samples (complex) per ROF, i.e., 1250 In-phase samples per ROF and 1250 Quadrature-phase samples per ROF, and
- (2) the Radiometer output, comprising 40-bit samples at a rate of 10 samples per ROF, and
- (3) the Time Tags, 10 per ROF, and
- (4) the Quality flags, 10 per ROF. While not a data stream *per se* the Radiometer calibration pipeline applies heuristics to per-ROF Radiometer and Time Tag data to assess the reliability of the calibrated Radiometer values. See Section 12.3.1.5.

12.3.1.1 Calibrating the REX filter output: In-phase & Quadrature-phase values

The conversion of the I/Q samples from the raw DN to calibrated physical units (milliVolts) involves applying a gain-independent scaling, since the FIR process producing the samples has no adjustable parameters.

The algorithm takes each I or Q raw `FilterValue` that represent a single filter output value and does the following:

- Scales the filter value by the ratio of the ADC input range ($\pm 1V = \pm 1000mV$) to its output range ($\pm 2^{13}$; the FIR output values are signed twos-complement integers):

$$mV_{IorQ} = (1000 / (2^{13})) * FilterValue$$

12.3.1.2 Calibrating the REX Radiometry

The formula for converting the Raw REX radiometer data to power, in units of dBm, is dependent on RCP vs. LCP for `Rbase` and `Ro`, and is as follows:

$$dBm = Rbase + 10 * \log_{10}(4.5e6 * RAW) + dBstep(AGC - AGCoffset) + Ro$$

where

`Rbase` = -176.852 (RCP) and -177.177 (LCP), base dBm values

4.5e6 = Bandwidth of the REX instrument, Hz

`RAW` = Increase in raw (40-bit) 64-bit integer value `ROF` (raw FITS BINTABLE), scaled as needed to represent the equivalent increase for one `ROF` (1.024s),

- There are ten raw (accumulating) samples stored in column 1 of EDU 2 in the raw file. The variable `RAW` in the equation above applies to the increase of each raw sample, except the first, from the previous sample, with that increase scaled by a factor of 10.
- The first sample in an `ROF` represents the total tenth accumulating sample for the entire *previous* `ROF`. Calculating the increase from the previous sample would involve extracting the last sample from the previous `ROF` in a separate FITS file; the design of the SOC pipeline calibration process treats FITS files as independent entities, and so cannot guarantee the existence of that previous `ROF` FITS file. Instead, the SOC calibration uses the first sample as is, without scaling because it already represents the increase over one `ROF`, as an expedient to avoid a more complicated process.
- If any `RAW` sample is zero, then the power value will be set to -999.0.

`dBstep` = -0.475dB, Radiometer response per -1 step in the gain word (AGC above)

`AGC` = Gain word setting from which REX Side, A or B, generated the data to be calibrated.

- A state table of commanded, per-Side AGC settings as a function of time, derived from SeqGen SSF output files, is stored in the Uplink DataBase (UplinkDB) on the SOC, and in TABLEs `AGCGAINA` and `AGCGAINB` for PDS data sets, and
- When creating the Level 1 FITS file, the pipeline retrieves the appropriate value based on the observation time of the FITS file and the side, and puts it into the FITS header as

keyword AGCGAIN (Side A has ApIDs 0x7b0, 0x7b1, 0x7b6, 0x7b8; Side B has ApIDs 0x7b2, 0x7b3, 0x7b7, 0x7b9). For observations before the state table, use 167 or 163 for Side A or B, respectively. ApID 0x7b4 is general housekeeping data that applies to both sides, and 0x7b5 is for incomplete playbacks.

- The SOC calibration retrieves the value from the FITS header when calibrating the data.

Table 12-8: REX Side-Dependent Constants

Name	Description	Side A Value (0x7b0; 0x7b1, 0x7b6, 0x7b8)	Side B Value (0x7b2, 0x7b3, 0x7b7, 0x7b9)
AGCOffset	REX Side-dependent constant; gain word offset	167	163
Ro	REX Side-dependent response, mW/RAW, dB	-101.030	-104.547
Rbase	REX Base values, dBm	-176.852	-177.177

12.3.1.2.1 Radiometry calibration

The Rbase and Ro dBm constants were set based on Cold Sky and radio source radiometer observations taken throughout the mission. The gain-related coefficients are based on gain linearity tests where the gain word is stepped through several values while receiving a fixed-amplitude signal from the DSN. See the draft of Linscott, et al, 2017 provided in the PDS dataset for full details of the calibration.

12.3.1.3 Calibrating the REX Time Tags

The time tags are 24-bit integers, stored as 32-bit integers in the FITS BINTABLE, that increment ten times per ROF frame of 1.024s and represent the nominal (see Note 1) time since the first contiguous ROF frame in a sequence, so the formula to convert from the 24-bit Time Tag value TTraw to seconds is

$$T_s = T_{Traw} * 0.1024 \quad \text{*** N.B. relative spacecraft seconds}$$

Note 1: In the unlikely case of REX taking data for more than about three weeks, this will represent the time from the last Time Tag rollover. **N.B. the initial and final ROFs between power-on and power-off may have corrupt time tags and other data due to spacecraft data handling issues, and should generally be ignored; such products are identified in the ERRATA.TXT of PDS data sets.**

12.3.1.4 Converting REX Time Tags to absolute time

The absolute time of any single calibrated REX Time Tag in a FITS file is the difference, between it and the first time tag in the file, added to the START_TIME UTC value from the PDS label.

12.3.1.5 Setting Quality flag bits

The raw REX data stored, then compressed or packetized, and then transmitted by the spacecraft can be corrupted by several processes along the way. However, since all samples are represented by integers made up of bits, the SOC pipeline cannot tell from any single sample whether that value is valid or not. As noted above in Sections 12.2.1.1.4.1 and 12.3.1.3, the first and last few ROFs should generally be discarded. The input source may come from a Test Pattern (see Section 12.2.1.1.2); also, the gain word could change, and because there is no record of the gain word in effect within the raw ROF data themselves, the SOC pipeline uses modeled commands to set the gain word. Also, because the first raw Radiometer sample in a ROF is the accumulation of radiometry measurements from the previous ROF, the in-ROF status byte does not give a completely reliable indication of the input source for that first sample. To deal with these uncertainties, a third column has been added to the calibrated [Radiometer and Time Tag] BINTABLE in EDU 2; that column is called [Quality flag] and comprises 32-bit integers. The low six bits of each integer are independent flags indicating the result of several heuristic tests applied to the data in the ROF; Table 12-9 below gives the meaning of each bit. A bit is set (1) if its heuristic indicates there may be an issue with the calibration of the corresponding Radiometer value on the same row. All higher bits are set to zero, so if there are no detected issues with the calibration then the Quality flag integer value will be zero. The bits are independent and the Quality flag integer value will be the sum of the values of the bits; e.g. if the [All Zeros] Test Pattern is active, then Quality flag bit 0 (raw Radiometer sample is zero, bit value=1), bit 1 (possible corruption detected because all Radiometer samples are zero, bit value=2), and bit 4 (Input Source is a Test Pattern, bit value=16) will all be set, and the quality flag will be (1+2+16)=19. The heuristics for bit 3 (near-in-time gain word change, bit value=8) and for bit 2 (near-in-time Input Source change, bit value=4) are not yet implemented.

12.3.2 Level 2 Calibrated FITS file data format

The calibrated data from each ROF are stored in a single FITS file. The data layout in the PDU and EDUs is similar to that of a Level 1 data FITS file (see Table 12-7); the PDU containing the raw ROF data and BINTABLES in EDU 3 through 8 containing HK data are identical; the only differences are described as follows.

1. In the EDU 1 BINTABLE, the I&Q pairs have been calibrated to units of mV.
2. In the EDU 2 BINTABLE, the radiometry power values are derived from rates – differences between successive accumulating values - instead of the accumulating values themselves. The radiometry values have been calibrated to dBm over a 4.5MHz bandwidths. The Time Tags have been calibrated to s. A Quality flag column has been added to express the result of heuristics assessing the validity.
3. All EDU 1 and EDU 2 values stored in the FITS file are 32-bit floating-point or integer values.

12.3.2.1 Extension Data Unit (EDU) 1: I & Q values

The first extension of the Calibrated FITS file is a FITS BINTABLE, containing the calibrated I&Q value pairs in units of milliVolts.

The BINTABLE comprises 1250 rows, and each row comprises two columns laid out over eight bytes.

12.3.2.2 EDU 2: Radiometry & Time Tags; also Quality flags

The second extension of the Calibrated FITS file is also a FITS BINTABLE, containing three columns: Radiometry values in dBm; Time Tags in seconds; Quality flags as integers, no units.

The EDU 2 BINTABLE comprises 10 rows, and each row comprises three columns laid out over twelve bytes.

A Quality flags is a four-byte integer, with each of the low five bits indicating the result of an assessment whether the quality of the corresponding radiometry value is suspect.

If a bit is set (1), then the radiometry value corresponding to that Quality flag integer (i.e. on the same table row) should be ignored, or used with care; if a bit is clear (0), then that assessment does not indicate any issue with the radiometry value. See Table 12-9.

Table 12-9: Quality flag bits; red italicized items are not yet implemented

Bit	Value	Description
0	1	The scaled raw radiometry sample is zero; the calibrated value will be -999.0
1	2	A check of this ROF's data indicates there may be corruption in this ROF
<i>2</i>	<i>4</i>	<i>The ROF is within 5s of a change in the input source to or from a Test Pattern</i>
<i>3</i>	<i>8</i>	<i>The ROF is within 5s of a change in the AGC gain word</i>
4	16	The input source to this ROF is a Test Pattern and the calibrated value is not real

12.3.2.3 Scientific Units

The units of the calibrated values, after applying the scaling factors if present, are as follows:

```
Filter outputs (IQ pairs):      Voltage (mV)
Radiometry:                    Power (dBm over the 4.5MHz REX bandwidth)
Time Tag:                       Relative spacecraft time (s)
Quality flag:                   N/A
```

12.3.2.4 Additional FITS and PDS Keywords

12.3.2.4.1 Keywords added to Level 1 and Level 2 Calibrated PHDUs

```
FSTATUS = '0x00      ' / Raw status byte from TLM
AGCGAIN =              167 / AGC Gain
AGCSRC = 'ULCMD     ' / Source of AGC Gain info
AGCPROV = '15188.ssf:PERX_P_REX_THERMSCAN$8REX_8CMD,167' / Provenance
```

12.3.2.4.2 Provenance added to Level 2 Calibrated PHDU

```
RADRAD__ = '*** Radiometry formula and coefficients:'
RADRADIO= 'Power, dBm=Rbase+10*log10(Bandwidth*RAW)+dBstep*(AGC-AGCOF)+Ro'
RADRBASE=          -176.8520 / Base dBm value
RADBNWDW=          4.5000 / Bandwidth, MHz
RADDBSTP=          -0.4750 / Radiometry response, -raw REX units/AGC, dB
```

```

RADAGC = 167.0000 / Gain word setting for observation
RADAGCOF= 167.0000 / Base gain word setting
RADRO = -101.0300 / Rad. resp, raw REX units/mW @ AGC=AGCOF, dB
RADIQ__= '*** I&Q formula:'
RADIANDQ= 'mV = (1000 / (2^13)) * RAW = KIQ * RAW'
RADKIQ = 0.1221 / REX IQ response, mV/count
RADTT__= '*** Time Tag formula:'
RADTIMTG= 's = 0.1024 * RAW = DT * RAW'
RADDT = 0.1024 / Time tag scaling, s/countRADRAD__= '***

```

12.3.3 Hardware/OS Development Platform

PC/Linux

12.3.4 Language(s) Used

Python (Telemetry to Level 1); C (Level 1 to Level 2 Calibrated)

12.3.5 Third Party Libraries Required

cfitsio (where C used) or Pyfits (where Python used)

12.3.6 Calibration Files Needed (with Quantities)

AGC gain table, in the Uplink Database on the SOC, and in DOCUMENT/AGCGAIN{A,B}.* in PDS data sets. All other calibration factors are in the source code and listed above.

12.3.7 Memory Required

< 128MB

12.3.8 Temporary File System Space Needed

None.

12.3.9 Predicted Size of Output File(s)

< 70 Kbyte

12.3.10 Predicted Execution time

Less than a second per ROF

12.3.11 Contact/Support Person(s)

Ivan Linscott

12.3.12 Maintenance Schedule (Code/Data Updates, Documentation)

None planned.

12.3.13 References

[LINSOTTETAL2017] Linscott, I.R., D.P. Hinson, G.L. Tyler, S.A. Stern, L.A. Young, R.A. Beyer, M.K. Bird, K. Ennico, G.R. Gladstone, C.B. Olkin, M. Pätzold, P.M. Schenk, D.F. Strobel, M.E. Summers, H.A. Weaver, W.W. Woods, "REX Radiometer Calibration at 4.2 cm on *New Horizons*," in draft, 2017.

13. SDC INSTRUMENT DESCRIPTION

13.1 Overview

The mission of the Venetia Burney Student Dust Counter (SDC) is to analyze the size and distribution of dust particles along the New Horizon's trajectory to the Kuiper Belt. The SDC instrument consists of the front-end analog electronics, the digital interface electronics, the detector panel, and the intraharness.

Each particle impact on 1 of the 12 active SDC detectors (see Section 13.1.1 below) will be a candidate for a science event. The impact crater causes a depolarization signal in the Polyvinylidene Fluoride (PVDF) detector film dependent on the size and speed of the particle. This signal gets converted to a digital number via the electronics. If the amplitude is above the value at which the threshold is currently set, then the signal is stored in memory as a science event along with other relevant housekeeping data.

These depolarization signals are measured in charge (Q) produced (Note that SDC reports charge in number of electrons. Even though this is not strictly charge, the number of electrons will from here on be referred to as the charge.) The charge from an impacting particle depends on the particles mass and velocity. Because the unit of the raw data is data number (DN), a calibration curve from data number to charge (DN=>Q) is needed. This curve is a function of box temperature and detector channel. For SDC, this curve was produced pre-flight and is checked during the mission with internal calibration procedures. The DN=>Q calibration curves are shown in Figure 13-2. The Level 2 calibrated files are derived from the Level 1 files through these curves.

13.1.1 Notes

Each detector has an independent set of electronics called a channel.

There are a total of 14 channels; there are also two inactive detectors mounted on the back of the panel, out of the path of any dust, for estimating the rate of detectable non-dust events such as spacecraft thruster firings and vibrations.

Channel 11 failed before launch; it can still generate spurious data, which are processed by the pipeline, but they should be ignored.

13.1.2 Box Temperature

The box temperature is the average of the two temperatures recorded on the analog side of the electronics box of SDC; these are stored in the housekeeping packet with ApID 0x004; see the next section.

13.2 Level 1 Data Specifics

The Level 1 data are unprocessed telemetry. At the SOC and PDS, all levels of data are recorded in FITS format. The SDC team uses IDL for our data processing and hence would like to be able to load these FITS files into IDL as structures/arrays, etc. To do this we typically use an IDL fits reader which can be found in the Goddard IDL library. Specifically we use `mrdfits.pro`. If this is used, please note that a `"/unsigned"` flag must be given as the data are all unsigned integers.

The Level 1 data FITS file consists of housekeeping and science data. Some of these data are not used in the calibration process to produce the calibrated data. It's stated in the PDS label files which telemetry points are and are not used by the calibration process.

In addition to the IDL functions for FITS files, generic programs such as fv can also be found. If opened in this program, the Level 1 data tables are displayed in Figure 13-1 below.

Index	Extension	Type	Dimension	View								
<input type="checkbox"/> 0	Primary	Image	0	Header	Image	Table						
<input checked="" type="checkbox"/> 1	DATA	Binary	6 cols X 3 rows	Header	Hist	Plot	All	Select				
<input type="checkbox"/> 2	HOUSEKEEPING_SDC	Binary	9 cols X 1 rows	Header	Hist	Plot	All	Select				
<input type="checkbox"/> 3	HOUSEKEEPING_0X004	Binary	37 cols X 1 rows	Header	Hist	Plot	All	Select				
<input type="checkbox"/> 4	HOUSEKEEPING_0X00D	Binary	8 cols X 1 rows	Header	Hist	Plot	All	Select				
<input type="checkbox"/> 5	HOUSEKEEPING_0X00A	Binary	5 cols X 1 rows	Header	Hist	Plot	All	Select				
<input type="checkbox"/> 6	THRUSTERS	Binary	28 cols X 268 rows	Header	Hist	Plot	All	Select				

Figure 13-1: Primary Data Unit (PDU) and Extension Data Unit layout in FITS file

13.2.1 Data Format

The data in the FITS file are stored as a binary table extension. There are five tables in the Level 1 file. These tables and their columns are :

DATA –

- 1) Copy Number – Not used in calibration
- 2) Channel ID – Detector number (0-13) [Channel 10 has an electrical issue and is not used for science. Channels 6 and 13 are reference detectors. These detectors cannot detect real dust as they are covered. For all higher level data products the channel IDs are incremented by one and become 1-14.]
- 3) Zero Fill – Not used in calibration
- 4) Threshold – First note that this DN scale is reversed. This means 65535 is a small event while 0 is a very large event. This reverse scale is also true for the Magnitude described below. The threshold value is the maximum (highest DN but smallest signal) magnitude (see next item) for accepted hits. Hits above (smaller amplitude) the threshold are rejected at the instrument level. These thresholds are adjustable and vary from channel to channel.
 - i. Note that it is SOMETIMES possible for a slightly smaller amplitude hit to come in just above this value; this is a timing effect due to the way the instrument software work. An event that triggers a measurement occurs before that measurement. The peak of the pulse from a dust event that is

below (larger amplitude) threshold triggers a measurement, but that measurement occurs with a slight delay after its trigger event. For some events, between the time of the trigger and the time of the measurement the pulse may have decayed to very slightly *above* (smaller amplitude) the threshold, with the result that the measured and recorded Magnitude (next item) is *above* (smaller amplitude) the threshold.

- 5) Magnitude – The size of the hit in DN [Note that this scale is also reversed. This means that 65535 is a small event while 0 is a very large event.]
- 6) Time Stamp – The time the hit was recorded in Mission Elapsed Time (MET)

HOUSEKEEPING_SDC –

- 1) MET – Mission Elapsed Time
- 2) PanTemp A-D – Temperatures recorded on the panel of SDC
- 3) BoxTemp 1-4 – Temperatures recorded on the electronics box of SDC

HOUSEKEEPING_0X004 – Values used in Calibration from this table:

- 1) MET – Mission Elapsed Time
 - a. For a given Event MET (e.g. from an event row in the DATA table), to select data that applies to that event, use the row in this table with the MET that is closest to the Event MET.
- 2) CDH_PNL_A-D_TEMP – Temperatures recorded on the panel of SDC (Note these are the same as above in HOUSEKEEPING_SDC)
- 3) CDH_ANA_A-B_TEMP – Temperatures recorded on analog side of SDC electronics box
 - a. These are the sources for the two temperatures that are averaged (after conversion to Celsius) to calculate the box temperature used in the calibration.
- 4) CDH_ANA_DCDC_TEMP – Temperature recorded on DCDC
- 5) CDH_ANA_DCDC_TEMP – Temperature recorded on the FPGA

HOUSEKEEPING_0X00D – Values used in Calibration from this table:

- 1) MET – Mission Elapsed Time for the columns in this table
- 2) CDH_TEMP_SDC_ELEC – Electronics box temperature as recorded by the spacecraft
- 3) CDH_TEMP_SDC_DET – Detector temperature as recorded by the spacecraft
- 4) RALPH temperatures in this table are expected, but can be ignored for SDC calibration

HOUSEKEEPING_0X00A

- 1) MET - Mission Elapsed Time for the columns in this table
- 2) SDC_LVPS_VOLT – Voltage of SDC recorded by the spacecraft
- 3) SDC_LVPS_CURR – Current of SDC recorded by the spacecraft

4) RALPH values in this table are expected, but can be ignored for SDC calibration

THRUSTERS – Values used in this Table

- GC1_DATA_VALID_MET – MET of Thruster Fire

13.2.2 Data Sources (High/Low Speed, CCSDS, ITF)

- GC1_RCS_FIRE_MINOR_1-24 – Tells whether one of the thrusters fired

SDC data are low-speed CCSDS (Consultative Committee for Space Data Systems) packets only.

13.2.3 Definition of an “Observation”

One observation is one collection of events in one CCSDS packet.

13.2.4 Housekeeping Needed in Level 1 Files (for Calibration)

See Section 13.2.1.

13.2.5 Raw Science Data and/or Housekeeping Requirements

From launch to the end of the first month, HK packet METs should be within 1 minute of a dust observation

From the end of the first month until Jupiter, HK packet METs should be within 1 hour

From Jupiter until the end of the mission HK packet METs should be within 1 day

For the redundant points, such as temperatures, only one of these needs to satisfy this requirement.

13.2.6 Notes about Level 1 Data

- 1) The scale in DN is “backwards” on a 0-65535 scale. In other words, a very large hit represent a number near 0. A small hit registers as a number close to 65535.
- 2) The threshold can be tuned and represents the maximum DN of a detectable hit. HOWEVER, it is possible due to the way the electronics work, that you might get a hit with a slightly higher DN (smaller hit) than the threshold. Usually this is no more than a few tens of DN higher than the threshold.
- 3) SDC has on-board flight rules for autonomously turning a channel off, and there is also commanding throughout the mission that turns the instrument on and off. The user needs to know when the channel was on/off. This information (commanded and autonomous on/off events) is in a separate file named `sdc_on_off_times.dat`.

- a. In PDS data sets, these data are available as files `sdc_on_off_times_vNNNN.lbl` and `sdc_on_off_times_vNNNN.tab`, where NNNN is the version number.
 - b. There is also a PDS product, `sdc_on_off_cmd_vNNNN.lbl`, which contains only those on/off events that were commanded. This commanded-only product was created as a result of a data set lien generated at the December, 2016 peer review.
- 4) The maximum number of recorded hits in one second on a given channel for SDC is in general 3. The way the timing works it is possible to get up to 5 hits/second. However, if more than one hit is recorded in one second (instrument wide) this is considered a coincident event and will be flagged. The science processing interprets such an event as s/c noise and flags it.
 - 5) Hits within 1 second of thruster firings are also flagged as s/c noise.

13.3 Calibration

The data calibration is a three-step process:

1. the telemetry is stored as raw DN (Section 13.2, above);
2. each DN value representing the size of a hit is converted into charge (Section 13.3.1, below);
3. each charge is converted into mass via the ground calibration results and an assumed particle velocity (Section 13.3.2, below).

Note that each event (hit) is converted to mass regardless of whether or not it is believed to be noise.

13.3.1 Pre-Flight Calibration Procedure - Charge

In a temperature controlled environment, the electronics from the end of the PVDF to the DN in the raw data were calibrated, at each of 4 calibration box temperatures and for each of the 14 channels. This was done by injecting 19 (actually 21; see below) fixed-amplitude charge pulses 100 times into a channel and recording the DN value each time. From those recorded values, the average DN (DN_{avg}) and its standard deviation (SIG) at each charge pulse amplitude, box temperature and channel were calculated. Then, for each box temperature and channel, a 9th order polynomial fit of Q (DN_{avg}) was derived. Finally, these 3 sets of values (the polynomial coefficients, DN_{avg} , and SIG) were stored in a matrix. This matrix contains all information required to calculate the charge equivalent to a DN as a function of box temperature and channel (detector), as well as the uncertainty in that calculated charge value.

13.3.1.1 Charge Calibration File

The calibration file contains the calibration values described above as a matrix of floating point values with dimension (4 × 14 × 3 × 19) representing values for the 4 calibration box temperatures, the 14 channels, and the 3 types of calibration values (coefficients, DNavg & SIG). The zero-based indices have the following meanings:

First Index – 4 Calibration Box Temperatures:

- 0) 49.9deg
- 1) 40deg
- 2) 34.25deg
- 3) -7.1deg

Second Index – 14 Detector Channels

- 0) First channel
- 1) Second channel
- ...
- 13) Last (fourteenth) channel

Third Index – 3 types of data; selects which array of values are retrieved via the Fourth Index:

- 0) Coeffs – Charge as a function of DN, polynomial coefficients (only the 1st 10 are used)
- 1) DNavg - Average DN recorded during Calibration at this Tbox & Channel
- 2) SIG - Standard deviation of the corresponding DNavg value (Third Index = 1)

Fourth Index – Meaning is dependent on the Third Index; see also Note 1 below

- For the coefficients of the polynomial (C_n ; third index = 0), for an in-flight hit on the selected channel (Second Index), and at the calibration box temperature (First Index), the charge (Q) as a function of in-flight hit DN is calculated as follows:

$$\log_{10}(Q(DN)) = C_0 + C_1*DN + C_2*DN^2 + C_3*DN^3 + \dots + C_9*DN^9$$

- 0) Zeroth order coefficient, C_0
- 1) First-order coefficient, C_1
- N) Nth-order coefficient, C_n
- For DNavg & SIG data types (third index = 1 & 2)
 - 0) the index of each charge pulse test ordered by increasing charge (decreasing DN)

Note 1: We injected charge pulses at 21 different values, but some of these were too small to record, and no channel had more than 19 recordable values at any box temperature. Also, there are only 10 coefficients in the 9th-order polynomial. So, although the matrix can hold up to 19 coefficients, average DN's or standard deviations per box temperature and channel, only the derived/recorded values are stored in the matrix, and any unused matrix values are set to zero.

This does not affect the polynomial evaluation, but when using the DNavg and SIG values one should ignore zero values.

Thus from this matrix you can get 3 things: Fit coefficients, Average DNs, and standard deviations. So, for example, to get the fit coefficients for a box temperature of -7.1 degrees on the first channel you want (-7.1, first channel, Coeff, *) => CALARRAY[3, 0, 0, *] (IDL notation). See Figure 13.2 for a plot of Charge vs DN represented by the Fit Coefficients.

For more details about this calibration procedure see *Horanyi, et al., "The Student Dust Counter on the New Horizons Mission", Space Sci. Rev., Volume 140, Numbers 1-4, pp. 387-402, 2008.*




Figure 13-2: Calibration curves for SDC. All 14 channels are shown for reference.

13.3.2 Calibration – Mass

13.3.2.1 Pre-Flight and other ground-based calibrations

The mass can be derived from the charge. It was discovered by J.A. Simpson and A.J. Tuzzolino (1985) [S-T] that a particle impacting a 28 μm PVDF film (such as those on SDC) will produce a charge given by the equation:

$$N[e^-] = c \times (m[\text{g}])^a \times (v[\text{km/s}])^b$$

In this equation N is the charge in equivalent number of electrons, m is the mass in grams, v is the detector-relative particle speed in km/s; see Table 13-1 for constants a, b, and c.

Additional data were considered from tests on flight spare detectors built at the Laboratory for Atmospheric and Space Physics (LASP). Initial ground calibration results updated the constant **c** for the three-parameter model above; James et al., 2010 [JAMESETAL2010] updated all of the constants for the three-parameter model, and also proposed a four-parameter model that includes temperature dependency (T in Celsius; d is another constant in Table 13-1):

$$N[e^-] = (c + d \times T[\text{Celsius}]) \times (m[\text{g}])^a \times (v[\text{km/s}])^b$$

The *New Horizons* pipeline code switched to using this four-parameter model in 2013, and redelivered past PDS data (Launch and Jupiter mission phases) based on the four-parameter model in 2014.

13.3.2.2 Mass Production

On the NH SDC instrument we have two measurements per event: the charge, N[e⁻]; the temperature of the detector, T[Celsius]. Thus to find either mass or velocity we must assume the other. In the pipeline we assume a Keplerian velocity for the hypothetical dust particle that produced each event, and use that to determine a mass for the event.

Note that each event is converted to mass regardless of whether or not it is believed to be noise.

Thus one can simply use the number of electrons produced and the assumed spacecraft-relative speed of the dust particle calculated through SPICE to determine the mass of the impacting particle.

Rearranging the equation for the four-parameter model above yields

$$m[\text{g}] = \{N[e^-] / ((c + d \times T[\text{Celsius}]) \times (v[\text{km/s}])^b)\}^{1/a}$$

The *New Horizons* pipeline code started using this equation in 2013. In 2014 the project redelivered past Launch and Jupiter mission phase SDC data to PDS based on this four-parameter model; see Table 13-1 for constants a, b, c, and d.

13.3.2.3 Constants

Table 13-1: Constants for the mass-charge-velocity relationships; Simpson and Tuzzolino, 1985, and James et al., 2010.

	Simpson and Tuzzolino three-parameter	Initial ground calibrations three-parameter	James et al., 2010 three-parameter	James et al., 2010 four-parameter
a	1.3 ± 0.1	1.3 ± 0.1	0.9343 ± 0.0009	1.052 ± 0.004
b	3.0 ± 0.1	3.0 ± 0.1	2.4 ± 0.1	2.88 ± 0.06
c	3.8x10 ¹⁷	5.63x10 ¹⁷	(6.7 ± 0.3.2)x10 ¹²	(1.2 ± 0.1)x10 ¹⁵
d	0	0	0	(6.7 ± 0.3.2)x10 ¹²

13.4 Level 2 Data Specifics

13.4.1 Algorithm for Pipeline

Pre-flight calibration of the electronics box was performed to find the relationship between charge in and DN out. This was done at 4 electronics box temperatures for all 14 channels. Fits were established from these data and the coefficients were stored in a matrix (see Sections 13.3.1 and 13.3.2 above).

The box temperature is derived from two raw sensor values in the HOUSEKEEPING_0X004 table: fields (columns) CDH_ANA_A_TEMP and CDH_ANA_B_TEMP. The conversion from sensor DN to temperature is via this polynomial:

$$T_{\text{Box, Celsius}} = 63.06613 - 0.36194 \text{ DN} + 0.001984 \text{ DN}^2 - 7.45\text{E-}6 \text{ DN}^3$$

For example, if the CDH_ANA_A_TEMP and CDH_ANA_B_TEMP DNs are 100 and 150, respectively, then the calculated temperatures are as follows:

Symbol	Formula	Result
T _{ANA_A}	$63.06613 - 0.36194 \times 100 + 0.001984 \times 100^2 - 7.45\text{E-}6 \times 100^3$	39.2621 Celsius
T _{ANA_B}	$63.06613 - 0.36194 \times 150 + 0.001984 \times 150^2 - 7.45\text{E-}6 \times 150^3$	28.2714 Celsius
T _{Box}	$(T_{\text{ANA_A}} + T_{\text{ANA_B}}) / 2$	33.7668 Celsius

The code for Level 2 data uses the channel number and the in-flight electronics box temperature (T_{Box}) to find the correct coefficients in the matrix. Those coefficients are then used in a polynomial function, with the raw DN as the independent value, to calculate the corresponding charge, and then converted to mass using the equation the Mass Production equation above. For in-flight box temperatures other than the four calibration temperatures represented by the first index of the calibration matrix, the in-flight charge is interpolated (or extrapolated) in box temperature from two calculated charges, which in turn use the two calibration box temperatures nearest to T_{Box}. Finally, in like manner using the standard deviations calibration matrix (SIG), for the nearest DN calibration measurements (DN_{avg}) and nearest two calibration temperatures, as an analog for the 1-sigma combined uncertainty of the calibration charge pulse measurement and of the calibration and in-flight DN measurement, the +/- 1-sigma masses (M_{sigplus} & M_{sigminus}) are calculated.

13.4.2 Dataflow Block Diagram

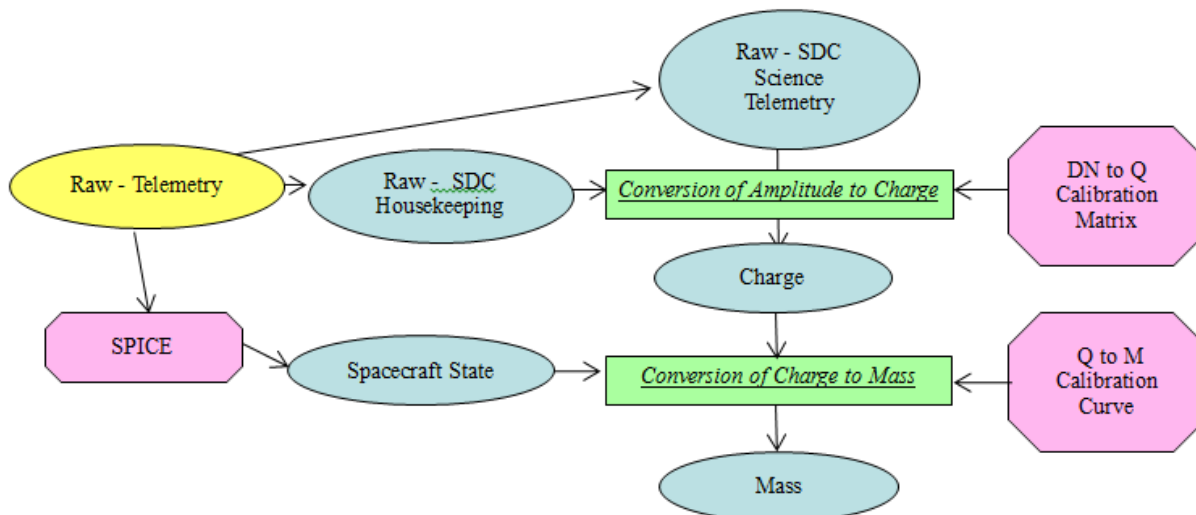


Figure 13-3: SDC Pipeline Processing Dataflow

13.4.3 Data Format

The calibrated FITS file has an empty Primary Data Unit (PDU; FITS keyword NAXIS=0), and a on Extension Data Unit (EDU) with a binary table with each row representing a detected event. Each row comprises science data with the event (hit) magnitude expressed in units of charge (equivalent number of electrons) and in units of mass (grams), the channel number (PVDF detector) and time of the event, and other quantities (see Section 0). Quality flags signal whether or not any of the housekeeping values were out of the standard operating range when the hit occurred. The quality flags also tell whether or not the data was extrapolated or interpolated from our pre-flight calibration curve, and whether there were any spacecraft thruster firings near in time to the event that may have induced the event.

Note that for scientific convenience in calibrated data, the channels are labeled 1-14 instead of 0-13.

13.4.4 Extra FITS Extensions (planes) and Their Definitions

The table in the calibrated FITS file is

1) CALIBRATED_DATA –

- a. UTC Time
- b. MET - Time in Mission Elapsed Time (MET)
- c. Channel – [1-14]
- d. Charge [Number of Electrons]
- e. Mass [grams]
- f. Mass_Thrsh - The threshold in mass [grams]
- g. M_sigplus - Mass Plus sigma [grams]
- h. M_sigminus - Mass Minus sigma [grams]
- i. Coincident_quality_flag - This flag indicates whether this event shares a time tag with another event.
 - "OK" - No other event shares the same time tag
 - "COIN" - At least one other event has the same time tag and thus we believe the event to be a result of noise.
- j. Thruster_quality_flag - Because we are susceptible to thruster firings (i.e. a thruster fire can cause false hits) a flag has been created to identify events we believe were caused by a thruster.
 - "OK" - No thruster firing occurred near this event
 - "TF" - A thruster firing occurred within 1 second of this event and thus we believe the event was possibly caused by a thruster firing
- k. IMP_VEL - Theoretical impact speed used to convert Charge to Mass [km/s]

13.4.5 Scientific Units

Charge - Number of electrons produced from impact.

Mass – Grams of impacting particle.

13.4.6 Additional FITS and PDS Keywords Added

13.4.7 Hardware/OS Development Platform

Intel, Linux or Windows

13.4.8 Language(s) Used

IDL

13.4.9 Third Party Libraries Required

JPL Astro Library downloaded from NASA at Goddard.

13.4.10 Calibration Files Needed (with Quantities)

IDL .sav file consisting of a table for fit coefficients. (<1MB)

13.4.11 Memory Required

N/A

13.4.12 Temporary File System Space Needed

N/A

13.4.13 Predicted Size of Output File(s)

Less than 1kB

13.4.14 Predicted Execution time

A few seconds

13.4.15 Contact/Support Person(s)

Level 1: Marcus Piquette, David James

Level 2: Mihaly Horanyi, Marcus Piquette

13.4.16 Maintenance Schedule (Code/Data Updates, Documentation)

We do have on-board calibration capabilities for the instrument and a place to insert these changes built into the code. Currently this simply multiplies by 1, but it the capability to adjust the values by some specified function remains.


14. SWAP INSTRUMENT DESCRIPTION

14.1 Overview

Solar Wind Around Pluto (SWAP) instrument is designed to measure the properties of solar wind ions for the *New Horizons* mission to Pluto. The bulk (thermal) solar wind ion distribution is typically Maxwellian. For most of the long journey to Pluto we expect to encounter bulk solar wind cold ion distributions that are nearly Maxwellian since the density and temperature of the solar wind decrease with increasing distance from the Sun. One notable exception is when the *New Horizons* spacecraft encountered Jupiter’s magnetosheath. Ion distributions are known to be hot in sheath regions. Since there have been no prior in situ measurements near Pluto, we do not know if it has a well-developed sheath region.

The SWAP instrument is an electrostatic instrument. The SWAP electro-optics control the energy passband of ions entering the instrument. The electro-optics has three parts: the Retarding Potential Analyzer (RPA), the Electrostatic Analyzer (ESA), and the deflector (DFL). Figure 14-1 shows a cross section of the instrument. The RPA consists of four grids with the inner two having a positive voltage, which repels ions with energies less than the corresponding potential energy (qV) (top right and left of Figure 14-1). The Electrostatic Analyzer has two parts, which are concentrically spaced, an inner dome and an outer spherical shell at ground (Figure 14-1). Only ions with a limited range of energies pass through the ESA to reach the detector.

The SWAP instrument is mounted on the $-Z_{sc}$ side of the spacecraft and the normal to the center of the aperture is aligned with $+Y_{sc}$ (Figure 14-2). Figure 14-3 shows the instrument being mounted to the spacecraft.



TA003358_B

Figure 14-1: Diagram of SWAP electro-optics. Two ion trajectories are drawn: one having energy greater than the RPA voltage and one less. [Figure 8 from McComas et al., 2008]




Figure 14-2: Diagram of New Horizons spacecraft highlighting the location of the SWAP instrument on the left. On the right is a diagram of the SWAP instrument with the spacecraft axes labeled.

The deflector is used to adjust the field-of-view (FOV). That is if the solar wind, which is highly collimated (spanning only a few degrees), enters at the bottom of the RPA, the voltage on the deflector could be set so that only ions that are not part of the solar wind beam enter the instrument. This would allow pickup ions, which occur over a wide range of angles, to be studied. In the inner heliosphere the pickup ions have substantially lower fluxes than the solar wind. The SWAP deflector can be used to bring the solar wind into the field of view if the solar wind beam is slightly above the top of the nominal field of view. Operating the deflector affects the energy of the ions that can enter the ESA. The RPA voltage is adjusted to compensate such that the same energy ions enter the ESA as did prior to the deflector voltage change. The deflector voltage can be automatically varied based on the commanded angle. The voltage settings for the ESA, RPA, deflector, and the amount the RPA should be adjusted to compensate for the deflector setting are all specified using lookup tables, which allow many instrument operation changes to be made by uploading new tables without having to make any software changes. Additional information on the electro-optical design is given in the introduction of Section 3 and in Section 3.1 in the *McComas et al.* [2008] instrument paper. The Channel Electron Multiplier (CEM) detector design is also described in Section 3.1.

The SWAP instrument has two kinds of voltage scans (also called sweeps): coarse and fine. The sweeps are performed in pairs. The two possible pair options are either two back-to-back coarse scans or a coarse scan followed by a fine scan. The voltage settings are predefined with onboard voltage tables. In coarse scans large voltage steps are taken with the ESA and RPA holding the ratio of the two voltages fixed. In the fine scans we also hold the RPA and ESA at constant ratio, but take smaller steps. Our voltage tables allow us to vary the ratio between the RPA and ESA voltages, but typically this ratio is held constant as much as possible. For high ESA voltages we cannot set the RPA to a high enough voltage to keep the RPA and ESA voltage ratio fixed because the highest RPA voltage is 2000 V and the highest ESA voltage is 4000 V. A given step number refers to a pair of RPA and ESA voltages. In the inner heliosphere we set the RPA to ESA voltage ratio high in order to narrow the passband slightly by removing ions at the low energies. All fine scans are approximately centered about the step number where the peak counts are observed in the coarse scan. To determine plasma properties from the detected count rates as a function of step number, the following calibration information is necessary: the ESA and RPA response functions, angular response function, the instrument solid angle, the detector gain, and the effective area.

Onboard there is one ESA table with 1024 steps and 4 RPA tables with 1024 steps each. For a given sweep we use the ESA table and one of the RPA tables. The different RPA tables can be used for coarse and fine scans, but for Jupiter operations we used the same RPA table for both the coarse and fine scans. In a scan/sweep the same step number is used in the software to reference rows in the ESA table and the chosen RPA table. The coarse scans use every 16th step in the 1024 voltage step lookup table where a step refers to a RPA and ESA voltage pairing. A fine scan consists of 64 steps with the coarse step at which the peak counts were detected in the middle of the fine scan.

14.2 Electronics and Flight Software

The instrument electronics are described in section 3.4 of the *McComas et al.* [2008] instrument paper. In subsequent sections of this document some information related to the flight software is provided, but further details are provided in sections 3.6 of the instrument paper.

14.3 SWAP Data Types

There are six types of SWAP science and engineering data: real-time science (0x584), summary (0x585), histogram (0x586), housekeeping, messages, and memory dump. Housekeeping, messages, summary and memory dump provide engineering data and the other two modes contain science data. Real-time data provide the most detailed science measurements since they contain the full count rate distribution as a function of energy (speed). For summary and science histogram modes, the full distribution is not recorded. Instead, parameters are derived from the count rate distribution stored by SWAP. These derived parameters require less memory than storing the whole distribution.



Figure 14-3: Picture of SWAP being mounted to the spacecraft.

The summary and science histogram modes are primarily used during the cruise phase of the mission.

The real-time science data contain the full count rate energy distribution for the primary, secondary and coincidence rates. The full distribution is desired because in bow shock and sheath regions plasma distributions may not be Maxwellian. The shape of the distribution will provide valuable information about what physical processes are occurring. In real-time mode the instrument can take measurements at a rate of 1 Hz, which is crucial for studying plasma boundaries and shocks.

Summary data consist of parameters related to the average speed, temperature, and density. The summary data are designed to study the bulk solar wind. The peak of the count distribution is related to the density, the bin location of the peak is related to the speed, and the distribution width is related to the temperature and speed combined. Along with the average values, the variance, maximum and minimum values of the peak counts, width of the peak, and energy of the peak are also recorded. The summary data 0x585 packet is included in the Level 1 PDS data sets (CODMAC Level 2) only; they are not included in the Level 2 calibrated data sets (CODMAC Level 3) as it was determined that the solar wind speed, density and temperature data provided in the summary packets were not suitable for science studies of the solar wind. Since early 2017, the *New Horizons* Project delivers CODMAC Level 4 PDS data sets based on the summary data.

The histogram data are designed to study pickup ions. The pickup ion distribution has a characteristic shape once it is normalized by the average solar wind energy (or speed). The histogram data conserve storage space by adding up all the counts detected in given bins. The accumulation time for the histogram is variable. The bins for the histograms are not energy bins, but are bins relative the average solar wind energy (E_{sw}). The steps for the fine scan are roughly centered on the coarse scan step where the peak counts were observed allowing the energy of the solar wind to be more precisely determined in the fine scan. The energy found in the fine scan is then used to place the counts determined in the coarse ESA scan into a new large 1-D histogram array. The coarse scan count rate data array is placed into the larger histogram array such that the bin with the maximum counts in the fine scan is placed into histogram bin 1024. There are 1024 possible voltages in a given onboard voltage table. To understand the histogram data further information about the scans is necessary. Onboard there is one ESA table with 1024 steps and 4 RPA tables with 1024 steps each. For a given sweep we use the ESA table and one of the RPA tables. The different RPA tables can be used for coarse and fine scans, but for Jupiter operations we used the same RPA table for both the coarse and fine scans. In a scan/sweep the same step number is used in the software to reference rows in the ESA table and the chosen RPA table. The coarse scans use every 16th step in the 1024 step table where a step refers to a RPA and ESA voltage pairing. A fine scan consists of 64 steps with the coarse step at which the peak counts were detected in the middle of the fine scan. There are 2048 bins because the peak count rate in a fine scan could occur at any step in the fine scan, and the step in the coarse scan containing the peak counts is always placed in bin 1024.

14.4 Level 1 Data Specifics

14.4.1 Data Format

There are separate files for summary, histogram, and real-time data, and corresponding housekeeping data are placed in each file. Note that not all kinds of packets will be generated every day. For example, during commissioning there may only be housekeeping and memory dump packets, and during cruise there will be housekeeping, summary, and histogram packets. All the packet types have a CHKSUM parameter. This parameter is calculated onboard and is also calculated on the ground to check the data. For the real-time science data (0x584) there are black and white images of the coincidence spectrogram array where the y-axis is energy bin number and the x-axis is time bin number.

Housekeeping packets are included as extensions in the Level 1 files. Housekeeping time tags are compared to the time tags in the three types of science packets to insure that each of the three Level 1 files includes all the housekeeping data required for processing. Since the summary and histogram packets are assembled over an extended interval, the beginning and ending times of the measurements (included in the summary and histogram packets) are used to determine the correct housekeeping measurements.

Real-time science packets can occur at a rate as high as 1 Hz where each packet contains a set of counts, voltages, etc. Observations for 2 energy steps are stored in one packet. One observation occurs in the 1st half second and a 2nd observation occurs in the 2nd half second. The 1st and 2nd half second measurements correspond to two different steps in a given sweep. Each step consists of an RPA and ESA voltage pairing, and 64 such pairs complete either a coarse-coarse scan or a coarse-fine scan. In a coarse-coarse scan two 64 step (32 packets) coarse scans are done back to back. In a coarse-fine scan a 64 step (32 packets) coarse scan and then a 64 step (32 packets) fine is performed. Both a complete coarse-coarse, and a coarse-fine scan have 64 packets. There is a parameter called SWAP_RT.sec64_ST, which is included in every real-time science packet and has the value 1 at the start of a pair of scans (a set of 64 packets) and is zero otherwise. We use this parameter to insure that a 64 second cycle (pair of scans) is not split across a day.




Figure 14-5: Real-time Level 2 (raw) file displayed using FV.

The SWAP Level 1 data are arranged in a binary table such that the columns are instrument parameters and measurements, and rows are measurement times. The FITS format has a binary table that allows for columns and rows. Figure 14-5 is a picture of what the Level 1 real-time files look like using FITS Viewer (FV). As mentioned earlier the housekeeping data are also in a table extension. The histogram counts and the number of samples in each bin are stored as image extensions. The zeroth extension contains only the primary header, the first extension holds the real-time data, the second extension holds the housekeeping data, and the third extension holds the thruster data.

SWAP data are in CCSDS (Consultative Committee for Space Data Systems) packets packetized by the spacecraft from the low-speed bus. Note that on the *New Horizons* mission, every instrument also outputs a non-packetized portion of telemetry to the S/C. This portion is also called the “instrument state” and these data are incorporated into the general spacecraft housekeeping data and not into the SWAP packets. All the bits in the SWAP packets were defined for the MOC at APL in EXCEL spreadsheets in the form required by the mission. The *New Horizons* SOC used the same bit level format description (APL EXCEL spreadsheets) for all the parameters in our packets to decode our Level 1 data.

14.4.2 Definition of an “Observation”

A complete histogram observation consists of one histogram type 1 packet and 63 histogram type 2 packets. A complete set of real-time science measurements consists of a full 64-second cycle. This is described in detail in section 14.4.1. One summary packet constitutes a complete measurement. Summary packets are included only in Level 1 (CODMAC Level 2) PDS data sets; Level 2 (CODMAC Level 3) PDS data sets do not contain data from summary packets. Housekeeping data are required for all science measurements since the housekeeping data are key to interpreting the data and determining error flags.

14.4.3 Housekeeping Needed in Level 1 Files (for Calibration)

Unlike some of the other instruments all housekeeping data for SWAP are included into the Level 1 files as an extension.

14.4.4 Level 1 Science Data and/or Housekeeping Requirements

In addition to the complete housekeeping packets, summary, real-time, histogram, and thruster fire packets are included into our Level 1 files. The thruster data format for the Level 1 files was reformatted to reduce space. In the calibrated (Level 2) data the thruster data has been arranged by thruster name and time. The numbers in the table indicate the duration of the thruster firings.

14.4.5 Evolution of the HISTOGRAM Level 1 Data Product (0x586/7)

The flight software for the histograms has been updated a couple of times to make these data product more useful. Originally, there were only normalized coincidence count rate histograms (2048 array) typically spanning about 1 day. The first flight software update to improve the coincidence count rate histogram data product to store total counts for each of the 64 bins in the coarse scans without normalizing by the location of the peak count rate in the last 64 bins of the 2048 element array. We refer to this histogram product as non-normalized histograms, and these also spanned about 1 day. The second flight software update eliminated the normalized

histograms and reallocated the space to store many non-normalized histograms within the same space to increase the time cadence of the non-normalized histograms to ~30 minutes. The number of changes made to the flight software to implement this was minimized, which has led to a more convoluted format for these new high time resolution non-normalized histograms. Table 14-1 summarizes the evolution of the histogram data products and file formats. Each of these histograms data products are described more fully below. Note all histogram count rates are derived from coincidence measurements. The histogram data products only store data rates summed up when the deflector voltage is zero because the deflector voltage reduces the sensitivity, the majority of the time the deflector voltage is zero, and we do not want to sum data with different levels of sensitivity. For all of the histogram products, the energies for the energy bins for the coarse scans can be found in the 0x584 data taken during the same timeframe as the histograms.

Original Histograms

In the original normalized histogram data product, the 64 step coarse scan count rate arrays were placed into a very large array with 2048 bins (extension 1/histogram) based on the location of the energy bin that had the highest count rate during the fine scan. Note that the fine scans are taken right after the coarse scans and consist of another set of 64 measurements closely spaced and centered on the energy step with the highest count rate in the coarse scan. An additional array was used to keep track of the number of times samples (extension 0/primary) that were put into the count rate array. The average count rate could then be calculated by dividing the count rate array (extension 1) by the number of samples array (extension 0) element by element. The normalized histograms turned out not to be useful because while spinning the solar wind beam could go in and out of the FOV. When the solar wind peak was outside the FOV, the coarse scan count rates would be misplaced because the energy bin with the peak count rate did not reflect the count rate at the peak of the solar wind. The time information for this histogram product was stored as MET values in the header for extension 0 in the BEGTIME and ENDTIME parameters.

Addition of Non-normalized Histograms

A prior flight software load addressed the lack of utility of the normalized histograms by storing in the last 64 bins of the 2048 count rate array (extension 1) the total counts for the 64 step coarse scans without making any adjustments based on which bin had the highest count rates. For these non-normalized histograms, the average count rates can be determined by dividing all the count rates in last 64 bins (extension 1) by the value in the SMPLSCNT parameter in the header of extension 0.

Addition of High Time Resolution Non-normalized Histograms

The latest flight software update increases the time resolution of the non-normalized histograms by reallocating all of the original space occupied by the normalized count rate array and the number of samples array. For these new high time resolution non-normalized histograms, the file

format is quite different with several 64 element histogram count rate arrays stored in a large array in time order with each one spanning a time range of ~30 minutes. Up to 47 individual total histogram count rate arrays each with 64 elements can be stored in this data product as a large array with dimensions 64 by 47 in extension 0 (primary).

The start times for each of these histograms is stored in the array **TIMESTAMP** in extension 1 and has 47 elements. If the **TIMESTAMP** value is zero then the corresponding count rate array is not valid and should contain zeros. The end time for each count rate histogram is given by the start time of the subsequent histogram except for the last valid count rate histogram. The end time for this last valid histogram is given by the **ENDTIME** parameter value in the header of extension 0.

The average count rate for all, but the last valid set of 64 in the count rate distributions is found by dividing all elements of those rate arrays by 29. To calculate the average count rates for the last valid 64 element histogram count rate array, all elements of that array are divided by the value in the **SMPLSCNT** quantity in the header of extension 0 after subtracting 29 times the value of **nmax-1** where **nmax** is the number of valid sets of 64 count rate arrays (divide by $SMPLCNT - ((nmax - 1) * 29)$).

Table 14-1: Summary of the types of histogram products and the data format.

Types of Histograms	Description	First File	Method to Find Average Rates	Time Range
Normalized Histograms	Total of coarse scan count rates after being placed into bins organized by which energy bin in the fine energy scan the peak count rate was detected.	Launch configuration	Count rate array (extension 1) divided by the number of samples array (extension 0) element by element.	BEGTIME and ENDTIME parameters in the header of extension 0.
Normalized Histograms and Non-normalized Histograms	The normalized histograms in the first 1984 bins with the last 64 bins containing total count rates in coarse scan bins without any shifting based on which energy bin the peak count rate was detected in.	swa_0087555680_0x586_sci	For the normalized histograms, the average is the first 1984 bins of the count rate array (extension 1) divided by the first 1984 bins of the number of samples array (extension 0). For the non-normalized histograms, the last 64 bins of the count rate array	The BEGTIME and ENDTIME parameters in the header of extension 0.

			(extension 1) are divided the value for SMPLSCNT parameter in the header for extension 0).	
High Time Resolution Non-normalized Histograms.	<p>Several 64 element histogram count rate arrays stored in a large array in time order with each one spanning a time range of ~30 minutes.</p> <p>Up to 47 individual total histogram count rate arrays each with 64 elements can be stored in this data product as a large array with dimensions 64 by 47 in extension 0.</p> <p>Note if the corresponding TIMESTAMP element value is 0 then that 64 element count rate histogram array is not valid and should contain zeros.</p>	swa_0476064032_0x586_sci	The average count rate for all, but the last valid set of 64 in the count rate distributions are found by dividing those all elements of those rate arrays by 29. To calculate the average count rates for the last valid 64 element histogram count rate array, all elements of that array are divided by the value in the SMPLSCNT quantity in the header of extension 0 after subtracting 29 times the value of nmax-1 where nmax is the number of valid sets of 64 count rate arrays (divide by SMPLCNT-((nmax-1)*29)).	The start times for each of these histograms is stored in the array TIMESTAMP in extension 1 and has 47 elements. If the TIMESTAMP value is zero then the corresponding count rate array is not valid and should contain zeros. The end time for each count rate histogram is given by the start time of the subsequent histogram except for the last valid count rate histogram. The end time for this last valid histogram is given by the ENDTIME parameter value in the header of extension 0.

14.4.6 Change to Energy Sweep Voltage Table

Over time we noticed that the power that the SWAP instrument drew began to increase and get close to the yellow power limits while stepping the electrostatic analyzer (ESA) voltage to the highest values in our ESA voltage sweep tables. To make sure the power draw from the SWAP instrument remained within specification at all times, we changed the voltage stepping values such that we would not sweep the ESA to such high voltages. Therefore, we performed a commanded table load to change the values in the SWAP ESA voltage table such that we no longer collect measurements at such a high voltages. This means we no longer collect measurements up to 7.8 keV. This is an acceptable solution scientifically because we were typically observing background noise at those highest energy steps. The solar wind and interstellar pickup ions were not being observed at such high energies. The last SWAP observations that extend up to 7.8keV occur on early July 10, 2021. The first observations after restricting the energy range occur on early August 9, 2021. We kept the energy range on our plotting the same for continuity. Such that now you will see a grey band at the highest energies on our plots starting on August 9,2021 and beyond that date because we longer take measurements at those energies.

14.5 Level 2 Data Specifics

14.5.1 Data Format

The SWAP calibrated (Level 2) pipeline requires the following input information, SWAP Level 1 files which include all the housekeeping data, SWAP calibration information and engineering factors, orbit and attitude information, and spacecraft information such as thruster firings. In Figure 14-6 we show a general schematic for our Level 2 (calibrated) real-time data files. The

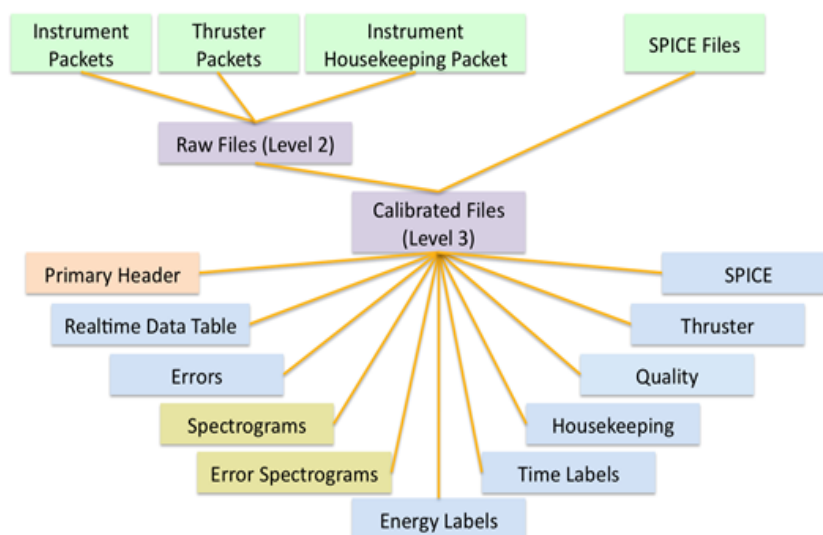


Figure 14-6: General Schematic for SWAP Pipeline. In green are inputs to the pipeline. Data files are in purple. The primary extension holds the main header (orange). In blue are the data tables, and in yellow are the image extensions. There are 3 spectrogram extensions and 3 error spectrogram extensions for our three count rate signals: PCEM, SCEM, and coincidence.

main input to the calibration pipeline are the SPICE kernels, and the Level 1 files, which include the real-time data, thruster firings, a few minor calculations performed using SPICE in the header and the housekeeping data.

The SWAP Level 2 (calibrated) analysis software has four parts corresponding to the different data (packet) types. There are simple algorithms for converting the raw data to engineering units. For example in the raw data the RPA voltage is stored as a step number for the Digital to Analog Converter (DAC), and not the actual voltage. We first make all such engineering conversions. Most information in our calibrated (Level 2) files is stored in table format (blue blocks in diagram). The tables contain all the data (science and housekeeping) converted to engineering units, and counts are converted to count rate (Hz). There are extensions for housekeeping data, thruster firing data, and quality flags. In the real-time data there are additional image extensions for spectrograms derived from high voltage science real-time measurements. There are three count rate 2-D arrays for the Secondary Channel Electron Multiplier (SCEM), Primary Channel Electron Multiplier (PCEM), and coincidence (COIN) signals stored as images, and three corresponding count rate error image extensions. These errors will be based on counting statistics. In addition to the 2-D arrays, axis information is also necessary for the spectrograms. The axis information for the spectrograms is contained in two tables one with the energy per charge (E/q) and one with the time tags for each sweep. In Figure 14-7 we show a picture of what the real-time files look like when opened using the FITS viewer FV.

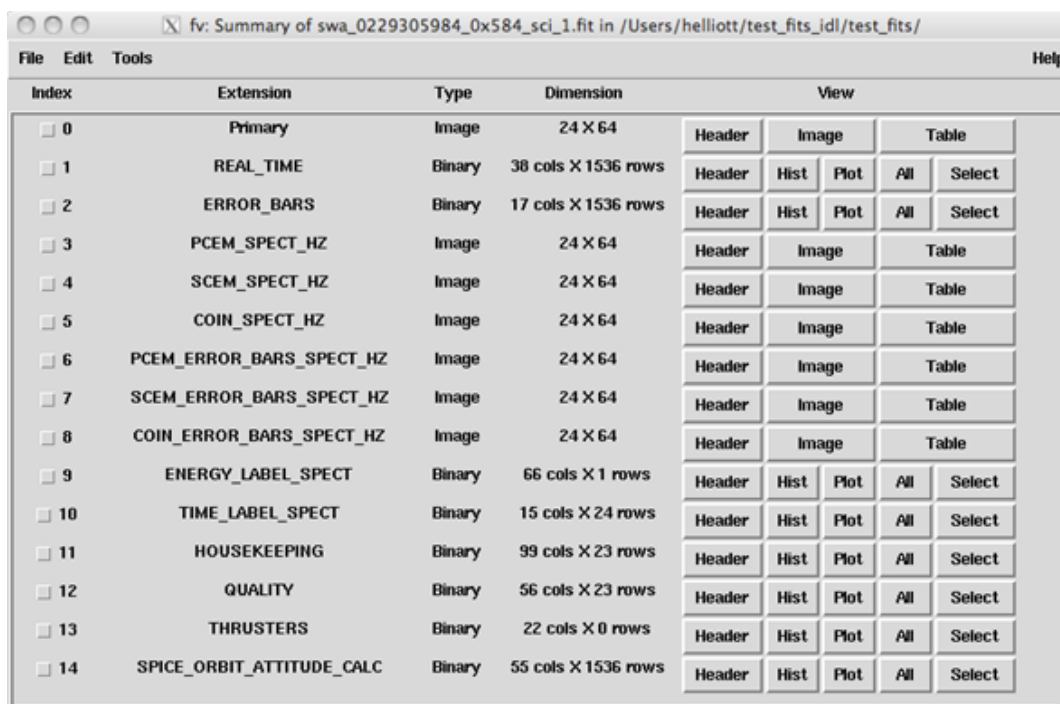


Figure 14-7: Picture of what the real-time calibrated (level 3) files look like using FV. The numbers listed under index are the extension numbers. All tables are Binary type and all images are image type.

14.5.2 Further Algorithm details for Pipeline

In this section we describe the algorithm for SWAP Level 2 (calibrated) processing. As mentioned above, the first step in our processing is to convert all raw values over to engineering units. These conversion factors are also stored in the command and telemetry spreadsheets used in the APL GSEOS system. The details of the housekeeping processing are not discussed further since processing of the housekeeping data consists of simple conversion factors. Analysis of ground calibration data provides critical information used to process the SWAP data, and is consequently a crucial input to our software. A description of the type of calibration information used in the pipeline is given in the calibration document.

14.5.3 Real-time science data processing




Figure 14-8: Examples of a coincidence (right) and error coincidence (left) spectrogram arrays.

The identifier in the SOC filenames for real-time data is the packet ApID (Application Process Identifier; also Application ID) 0x584. Our real-time high voltage science (HVSCI) analysis begins by determining the count rates in Hz as a function of energy for each measurement. A spectrogram is created by sorting the data into sweeps to build a 2-D array where the y-axis is energy per charge and the x-axis is time. A spectrogram spans the time range in HVSCI mode in a given daily Level 1 file. Spectrograms are created for each of the 3 count signals (PCEM, SCEM, and coincidence). Corresponding count rate error spectrograms are created based on

counting statistics for each of the three signals as described in the next paragraph. The x-axis time information is provided in the TIME_LABEL_SPECT extension along with a column indicating if a background has been removed. The background is mentioned in Section 14.5.10 and described in detail in the calibration document. Also in the TIME_LABEL_SPECT extension is a column indicating the plan and sweep used since the energy bins are different for different plans and sweeps. The y-axis energy labels for a given sweep and plan number are provided in the ENERGY_LABEL_SPECT extension. These count rate spectrograms provide a way to examine our data at high time resolution over the full energy range of the instrument. These kinds of spectrograms have proven useful for analyzing high time resolution plasma in situ measurements. Having a high time resolution product is critical for identifying plasma boundaries and shocks. In Figure 14-8 we show examples of what coincidence and an error coincidence spectrogram arrays look like when opened in FV.

14.5.4 Errors

In the Level 2 (calibrated) data files an error value for every measurement is given in the extension labeled ERROR_BARS. We also provide spectrogram arrays for each signal type for the errors in the extensions labeled X_ERROR_BARS_SPECT_HZ where X is PCEM, SCEM, or COIN. The errors provided are errors for the rates, and an error for the sample time, and data compression when compression occurs. The raw rates (count counts per sample (C)) are converted to Hz using the 0.390 sec sample/accumulation time (t) (Equation 1). The error squared is given in Equations 2 and 3, and the fractional error squared is shown Equation 4. Taking the square root the resulting fractional error is given by Equation 5. The final error given in the data files is shown in Equation 6.

$$R_c = \frac{C}{t} \quad \text{(Equation 1)}$$

$$(\Delta R_c)^2 = \left(\frac{\partial R_c}{\partial t}\right)^2 (\Delta t)^2 + \left(\frac{\partial R_c}{\partial C}\right)^2 (\Delta C)^2 \quad \text{(Equation 2)}$$

$$(\Delta R_c)^2 = \left(\frac{-C}{t^2}\right)^2 (\Delta t)^2 + \left(\frac{1}{t}\right)^2 (\Delta C)^2 \quad \text{(Equation 3)}$$

$$\frac{(\Delta R_c)^2}{R_c^2} = \frac{(\Delta t)^2}{t^2} + \frac{(\Delta C)^2}{C^2} \quad \text{(Equation 4)}$$

$$\frac{\Delta R_c}{R_c} = \sqrt{\frac{(\Delta t)^2}{t^2} + \frac{(\Delta C)^2}{C^2}} \quad \text{(Equation 5)}$$

$$\Delta R_c = R_c \sqrt{\frac{(\Delta t)^2}{t^2} + \frac{(\Delta C)^2}{C^2}} \quad \text{(Equation 6)}$$

If the count rates are not compressed (C_{un}) then $C = C_{un}$ and $\Delta C = \sqrt{C_{un}}$. However, if the counts are compressed (C_{comp}) then $C = 16C_{comp} + 7.5$ and $\Delta C = \sqrt{(16C_{comp} + 7.5)}$. In the ERROR table there is a column indicating if a background has been removed. The background is described in Section 14.5.10 and described in detail in the calibration document.

14.5.5 Quality Flags

Flags assessing the quality of the data are based on operational housekeeping alarms, but in the future additional ones may need to be added which are based on orbit and attitude, and additional calculations. All current flags are stored in a table extension.

14.5.6 Thruster Firings

As mentioned earlier the calibrated (Level 2) code reorganizes the thruster data into a table where each column refers to a given thruster name and each row is the start time of the thruster firing. The numbers under each thruster column indicate the duration of the thruster firings. Each thruster column has a title that looks like GC1_A2_FIRE where GC1 indicates it originated in a GN&C packet, and FIRE indicates thruster firing. The thruster names are A1, A2, B1, B2, B3, C1, C2, C3, C4, D1, D2, D3, D4, F1, and F2. The value for each thruster firing corresponds to the duration of the thruster firing (0=0msec,1=5msec,2=20msec,3=40msec). In the Level 1 data each row is a major frame, and the columns are minor frames where each minor frame is 40msecs. Thus, there are 25 columns with numbers between 0 and 24. In the Level 2 (calibrated) data we calculate the start time of the firings for a given minor frame which means we have already taken the major frame start time and added in the time to the start of the minor frame where the firing occurred ($\text{Start_time} = \text{major_frame_start_time} + 0.040 * (\text{minor_frame_number} + 1)$). The implication of this is that one row in the raw file may result in several rows in the calibrated file if there are multiple thruster firings.

14.5.7 SPICE Orbit and Attitude Calculations

Our orbit and attitude calculations are contained in the SPICE_ORBIT_ATTITUDE_CALC extension. We calculate times for each SWAP measurement in the REAL_TIME extension. The MET for the packet is listed along with the UTC, and ET for the start and stop time of each measurement. There are two start times and two stop times one since each packet stores two measurements one in the first half second (labeled with a 0) and one in the second half second (labeled with a 1). In the tail of the Jupiter encounter the spacecraft is spinning so we have included the angle in the Xsc-Zsc plane between the Zsc axis and Jupiter's spin axis (north). This angle is 0 deg (90 deg) when Zsc (Xsc) is aligned with the North end of Jupiter's spin axis. These angles are named ANGLE_JSP_XZ in the files and calculations were done for the start, middle, and stop for each observation, since the spacecraft rotates quickly (5RPM). All other parameters are calculated at the middle of each observation. We also calculated the angle between the Ysc and the Sun, Jupiter, and Earth. The label for the angle between Ysc and the Sun for the 1st measurement is called Y_SUN_ANG0_MIDDLE. The distances from the spacecraft to Earth, Jupiter and the Sun are calculated (i.e., SUN_SC_0_MIDDLE). We calculate the angle to the Sun is in the X-Y plane (phi), and the latitude angle from the X-Y plane (theta). Positive phi values are toward the +Xsc axis and negative phi angles are towards the -Xsc axis. Negative theta values are towards the top of the instrument since the -Zsc axis is at the top of the instrument. Note this is the opposite convention used in the calibration chamber. However, the phi angle is analogous to the roll angle in calibration (see calibration document). We also calculate the position and velocity of the spacecraft in IAU Jupiter Cartesian coordinates. The naming convention is such that the X component of position in IAU Jupiter for the 1st half

second measurement is labeled as SC_IAU_JUP_X_0. Likewise the X component of the velocity is SC_IAU_JUP_VX_0. In addition to IAU Jupiter coordinates we calculate the spacecraft position in J2000 Jupiter coordinates the X component for the 1st measurement is labeled as SC_J2000_JUP_X_0. Column name descriptions are given in the header for the SPICE extension as well as the names of the SPICE kernel files used to perform the calculations.

As of January, 2014, these SWAP instrument attitude calculations are only made for the Jupiter reference frame. There is a task in its final stages to formalize similar calculations for the SWAP instrument attitude in heliocentric, heliographic and inertial reference frames for solar wind observations, as well as future plans for a Pluto-based reference frame for Pluto Encounter observations. These products will be added to SWAP PDS data sets with future deliveries. In the meantime, the PDS user should note that a table containing spacecraft mission trajectory and PDS label keywords describing the attitude of the spacecraft, both in the J2000 inertial reference frame, are provided.

14.5.8 Summary and Histogram Files

Both the summary (Figure 14-9 top) and histogram (Figure 14-9 bottom) files also have the primary header, and housekeeping, quality, thruster and SPICE orbit attitude extensions the same as in the real-time files. In summary files, the Primary Data Unit (PDU; 0th extension) is empty, and extension 1 contains the summary data table converted to engineering units. In histogram files, the histogram data are stored as images in the 0th and 1st extensions. The 0th extension contains the number of times data were added to each bin, and the 1st extension contains the histogram count rates.

14.5.9 Calibration

Analysis of ground calibration data provides critical information used to process the SWAP data, and is consequently a crucial input to our software; therefore, the use of calibration information in the pipeline is described in a section of the calibration document. The SWAP lab calibration consists of an effective area; an angular response function for the ESA (function of α and ϕ); an energy response curve for the RPA that depends on the RPA and deflector voltages; a function representing how the RPA changes the energy of ions prior to enter the ESA; an energy response function for the ESA; the functions for how the width and center of the ESA passband varies with α ; and a function for how the conversion factor for converting ESA voltage to energy depends on ϕ . Our pipeline incorporates ESA and RPA response curves by precalculating the energy passband for each RPA and ESA voltage pair stored in the onboard tables. This information is used as a lookup table in the pipeline code so that each RPA and ESA voltage step can be assigned an energy per charge value. The code determines which set of tables to use by examining the time range since we have had different sets of RPA and ESA tables loaded at different times in the mission. To determine which of the 4 RPA tables to use the code compares the sweep and plan number in the data to the sweep and plan number in the precalculated tables. The results of the pipeline energy determination are stored in the energy axis label extension and can be used directly to label the y-axis of the spectrogram.

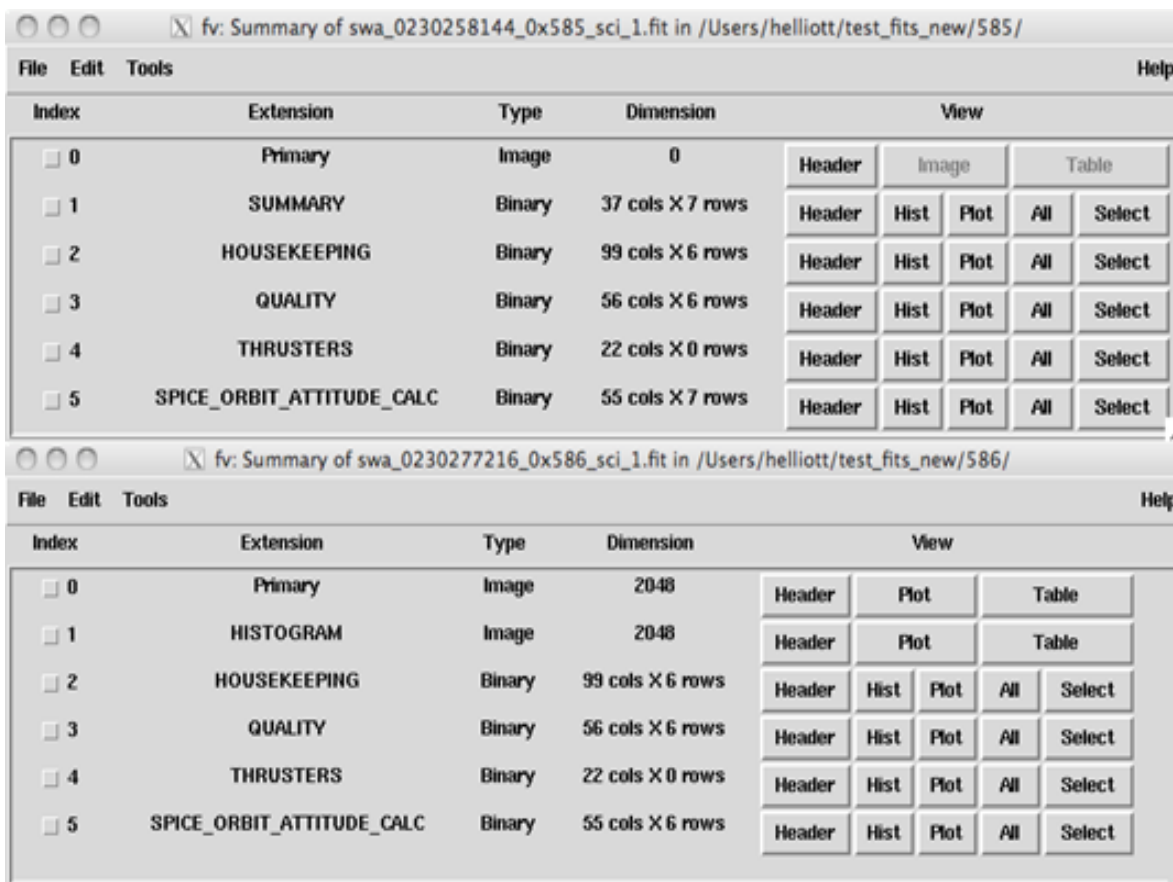


Figure 14-9: Examples of a summary (top) and histogram (bottom) files viewed in FV.

14.5.10 Background Subtraction

There is a background signal (count rate) in our data when the RPA is on. This background decreases as the distance from the Sun increases and will most likely not be a problem for the Pluto encounter. We can operate using only the ESA and have done so for the Jupiter encounter (plan 5). For the solar wind measurements inbound we needed to use the RPA (plan 0) to shield the instrument from high solar wind fluxes. Background subtraction information is processed in a way similar to the energy bin information. A given background subtraction, provided in a calibration file, is only valid for a given time range; therefore, a list of background files with valid application times for each are read in and used to select and read the applicable background, and then the background is subtracted. The background files are stored in the calibration directory and the names of the files used are stored in the list file. The instrument detector gain will evolve with time and this information will be incorporated in fashion similar to the background subtraction. Gain calibrations occur during annual checkouts. The background subtraction is described in more detail in the calibration document.




Figure 14-10: Timelines for SWAP plan numbers, data rates, and spacecraft maneuvers (top panel). Plan 0 is blue plan 5 is red. Two minutes of data set measurements every two hours is in orange where one set consists of a two coarse-fine sweeps. Purple is one set of measurements every 5 minutes. Green is 3-axis stable and black is spin mode where the spinning is about Ysc. The 2nd panel is the spacecraft-Jupiter distance in Rj. In the 3rd panel is the spacecraft-Sun distance in AU, and in the bottom panel is the Ysc-Sun angle in degrees.

14.6 Operations

14.6.1 Commissioning

Plan 3 was used to take the first complete sweeps of the solar wind in sequence SWAP-009-6 during commissioning.

14.6.2 Jupiter

The voltage tables were updated before the Jupiter mission phase to improve the voltage sweeps for the Jupiter encounter. There were only two plans used in the Jupiter phase: plan 0 and plan 5. Plan 0 is designed for solar wind and plan 5 is designed for magnetosphere. The voltage settings in plan 0 protect the instrument from high solar wind proton fluxes. Early on the data rate was two coarse-fine scans taken approximately every two hours; the exact timing was adjusted so as not to interfere with other instruments. At 2/21/2007 17:27:31 we kept the plan number set at 0 and changed the sampling rate to 1 coarse and 1 fine scan every 5 minutes. Then at 02/25/2007

19:22:43 we switched to plan 5, which consists of two coarse scans in 64 seconds every 5 minutes. In plan 5 the RPA is off and only the ESA voltage is swept. If the instrument encountered high proton fluxes in the solar wind with the RPA off, the energy passband is large enough that high count rates could cause the instrument to shut down via autonomy rules. Therefore, we took some plan 0 measurement in the magnetosphere to ensure that some measurements would be obtained even if the spacecraft did encounter solar wind. We show timelines for operations in Figure 14-10 along with the distance to Jupiter and the Sun. And in the bottom panel we show the angle between Y_{sc} and the Sun. In spin mode the spacecraft spins about Y_{sc} . The spacecraft entered the magnetosphere on day 056(02/25) in 2007 and the first time it exited in the tail was approximately on day 132 (05/12) of 2007. There were many tail crossing and there appears to be a boundary layer inside the sheath; therefore, definitive boundary crossing times are difficult to determine. When SWAP turned off on day 168 (06/17) of 2007 for hibernation, *New Horizons* was once again in the magnetosphere.

14.6.3 Pluto Cruise

Per the original mission plan for cruise operations, the SWAP instrument was off for the first 460+ days of Pluto Cruise. After that the operations were sporadic (just a few days in 2009) and mostly Science, alternating with Channel Electron Multiplier gain tests during Annual CheckOuts. After extensive testing in early 2012, in July of that year the project approved daily science operations for the SWAP and PEPSSI instruments throughout the rest of the cruise to Pluto.

14.7 Observation Examples

In this section we show several examples of SWAP data. Figure 14-11 shows the solar wind measurements on approach to Jupiter, and all of the Jupiter magnetosphere and sheath observations. The format is that of a color-coded spectrogram of the background-subtracted coincidence count rates in Hz of solar wind ions as a function energy per charge (E/q) as measured by the Solar Wind Around Pluto (SWAP) instrument on *New Horizons* at ~ 4.9 AU from the Sun (~ 0.4 AU upstream from Jupiter). We used plan 0 for these measurements since this plan helps protect the instrument from high fluxes by reducing the overall energy passband. In plan 0 the RPA is on at energies below 2000 eV/q. The RPA creates a background and this background has been removed (see calibration document). The lower trace shows the solar wind protons, while the upper trace shows the alpha particles (He^{++}), with enhanced sensitivity ($\sim 100\times$ larger) above 2 keV/q. Solar wind speed is a function E/q , with 1 keV protons corresponding to typical, ~ 440 km/s solar wind and larger (smaller) E/q representing faster (slower) wind speeds. Interplanetary shocks passed over the *New Horizons* spacecraft at $\sim 18:00$ on Day-of-Year (DOY) 11 and at ~ 1300 on DOY 14 causing the abrupt jumps in solar wind speed; the speed immediately following the latter shock was in excess of 600 km/s. The slowly decreasing speed after the second shock (falling E/q of the proton and alpha beams) is a rarefaction region, which forms as faster solar wind outruns the slower solar wind behind. In all, these SWAP observations show a clear solar wind stream structure.

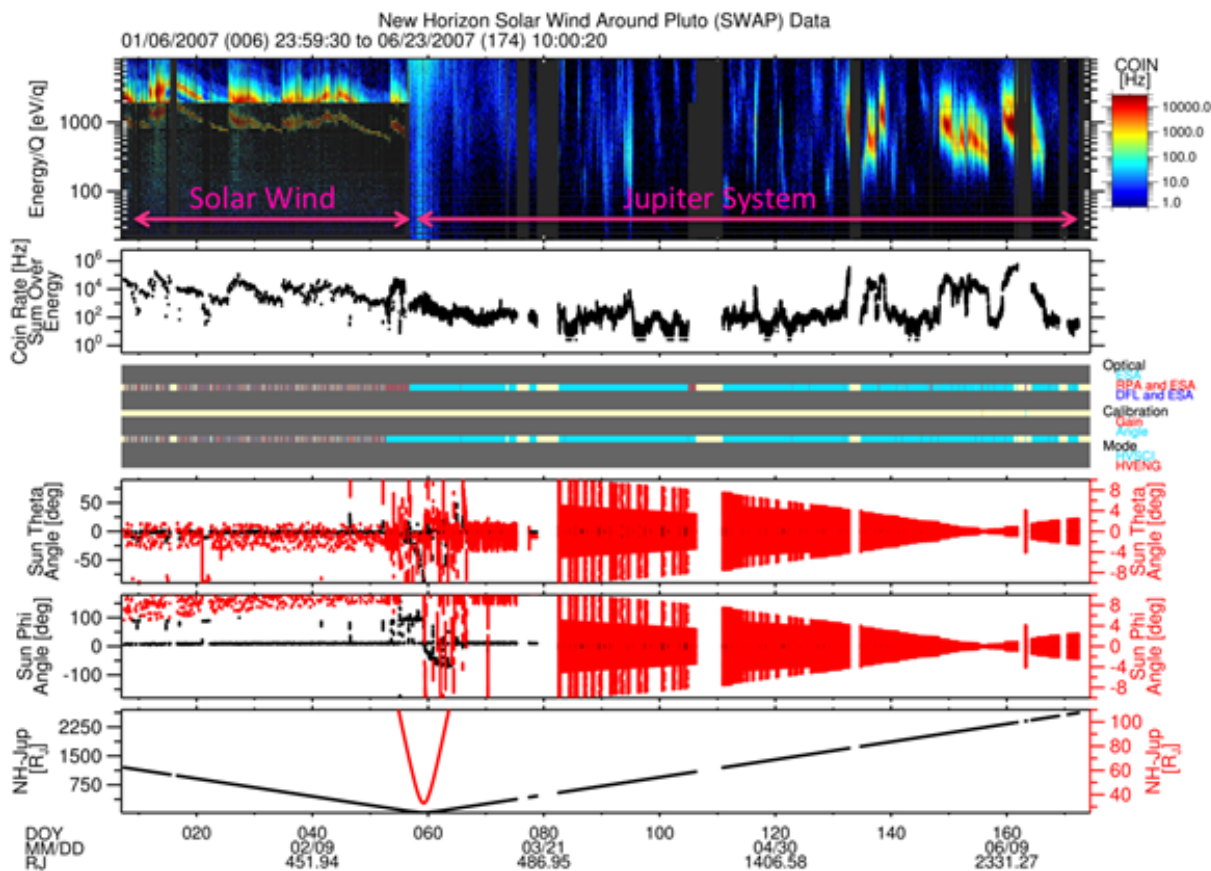


Figure 14-11: Top panel: Coincidence count rate spectrogram of solar wind measurements. The y-axis is the energy per charge. The x-axis is day of year. The color indicates the count rate in Hz. In the top trace are He⁺⁺ ions and in the bottom trace are H⁺ ions. Second panel: Total count rate for a given energy sweep. The dips in the count rates occur when the deflector is operating. Third panel: The top bar indicates if the tophat ElectroStatic Analyzer (ESA), Retarding Potential Analyzer (RPA), or Deflector (DFL) are operating. The middle bar indicates when either a Channel Electron Multiplier (CEM) gain test is occurring or if an angle deflector test is occurring. In the 4th and 5th panels is the angle of the Sun in a spacecraft coordinate system where theta is a latitude form the spacecraft X-Y plane and phi is the longitude angle in the X-Y plane. The bottom panel indicates the spacecraft-Jupiter distance in units of Jupiter radius. Note that data gaps in the top panel are shown as charcoal grey. Since the instrument was designed to measure low fluxes near Pluto, we applied a voltage on the RPA inside the ESA energy passband to reduce the overall instrument energy passband in order to measure the high solar wind flux upstream of Jupiter. The highest RPA voltage we could apply was 2000 Volts. Protons typically have energies below 2000 eV so this method was quite successful at reducing the detected rates to protect the instrument detectors. A background signal that occurs when the RPA operates has been removed. Since the alpha particles have a much lower flux and have about twice proton energy per charge, the instrument rates were not too high above 2000 eV where we only operated the ESA. Inside the Jupiter system most of the time we only used the ESA since we anticipated lower fluxes inside the Jupiter system. Inside the Jupiter system, we have colored all the RPA operation intervals grey since the RPA significantly reduces the count rates and the rates in the Jupiter's magnetosphere are much lower than in the solar wind. The gain test and deflector angle test were also removed since these were calibration activities.

Most of the magnetospheric observations were taken using plan 5 since the plan 5 energy passband is wider and the flux rates are lower in the magnetosphere than in the solar wind (Figure 14-11). The only time the plan 5 data has any background is when there is penetrating radiation, and we do not remove background due to penetrating radiation. The penetrating radiation occurs close to Jupiter and usually is greatest in the secondary and primary signals. When it occurs, the count rates are usually elevated at all energy steps; therefore, in spectrograms the penetrating radiation usually shows up as vertical stripes. The penetrating radiation is significantly reduced in the coincidence signal. The plan 5 data consists of one measurement set every 5th SWAP minute, and one set consists of 2 coarse scans performed in 64 seconds.

In Figure 14-12 we show an example of solar wind measurements taken during the 2013 Annual Checkout (ACO). The data format is similar to the previous figure except now the Sun-spacecraft-Earth angle is in the bottom panel. Dips in the count rate occur when the SWAP instrument deflector (DFL) is turned on. The deflector intervals are denoted by dark blue in the top bar of the 3rd panel. The main species in the solar wind are protons (H⁺) and alpha particles (He⁺⁺). In the spectrogram (top panel), the peak with the highest count rate is the solar wind proton (H⁺) peak, and the alpha particle peak (He⁺⁺) has energy per charge twice that of the proton peak. Therefore, the 2nd brightest peak in yellow and green is the alpha particle (He⁺⁺) peak. The blue band above alphas is the interstellar proton pickup ions. The cutoff for these ions occurs at 4 times the energy per charge of the proton peak.

To help users know when SWAP was collecting observations, in Figure 14-13 we show an overview of all the SWAP science measurements for the entire *New Horizons* mission to date. Note that starting in 2012 the coverage increases since we began taking measurements during hibernation. In the heliospheric measurements there are times when the solar wind is outside the SWAP FOV. Since the solar wind is quite radial, the Sun location can be used to find the approximate solar wind direction. The Sun direction is given in spacecraft coordinates in the SPICE extension. The Sun location is given as a latitude (theta) and longitude (phi). The names of the parameters THETA_SUN_SC_0, THETA_SUN_SC_1, PHI_SUN_SC_0, PHI_SUN_SC_1 where the "0" and "1" refer to the first and second measurement in a given packet. There for each count rate measurement in the REAL_TIME extension there is a corresponding set of theta and phi angle values in the SPICE extension. The SWAP FOV spans from -10 to +10 degrees in theta and from -138 to +138 degrees in phi.




Figure 14-12: Top panel: Coincidence count rate spectrogram of solar wind measurements. The y-axis is the energy per charge. The x-axis is day of year. The color indicates the count rate in Hz. In the top trace are He⁺⁺ ions and in the bottom trace are H⁺ ions. Second panel: Total count rate for a given energy sweep. The dips in the count rates occur when the deflector is operating. Third panel: The top bar indicates if the tophat ElectroStatic Analyzer (ESA), Retarding Potential Analyzer (RPA), or Deflector (DFL) are operating. The middle bar indicates when either a Channel Electron Multiplier (CEM) gain test is occurring or if an angle deflector test is occurring. In the fourth and fifth panels is the angle of the Sun in a spacecraft coordinate system where theta is a latitude form the spacecraft X-Y plane and phi is the longitude angle in the X-Y plane. The bottom panel indicates the Sun-Spacecraft-Earth angle.

14.8 Updates to the Level 2 Calibrated SWAP Data Pipeline

In 2007 we added the center energy for given RPA and ESA voltages to the REAL_TIME data extension. The column names are ENERGY_0 and ENERGY_1 in eV and correspond to the 1st and 2nd measurement in a given packet (row). We corrected the background subtraction. We corrected a rounding error in the time used to calculate the spin angles in the SPICE extension and fixed a small offset in the times for the spectrogram in the TIME_LABEL_SPECT extension.




Figure 14-13: Examples of yearly coincidence spectrograms of SWAP science operations. On the right we indicate the year and range of distances for each year.

In 2012 and 2013 the SWAP Level 2 calibrated data pipeline was rewritten in IDL. The original code was written in C and had some memory overwrite issues that proved difficult to solve. The code was rewritten into IDL since it is easier to avoid memory overwrite issues in IDL. The old files were compared against the new files by examining differences between the old and new values for all the quantities in all the extensions. All the differences were flagged and each difference was addressed. When a problem was found with the new files, the problems were resolved. We found many issues with old files and have made many corrections. We documented the differences that remain which are corrections. The some small errors in the time calculations were fixed and this also affected some of the angle calculations. Many of the comment fields were updated. A few typos in parameter names were fixed. We provided a PowerPoint(tm) presentation to the SOC with a summary of the changes and corrections (SWAP_level2_C_IDL_10152012.pdf; N.B. this is a project-internal presentation and it is not delivered with SWAP PDS data sets).

14.9 Gain and Deflector Angle Sweep Tests

After October 28, 2008 the Channel Electron Multiplier (CEM) detector gain tests are performed using plan 2 and the only plan 2 data after that date occurs during a gain test. Table 14-2 shows the start and stop times for the gain tests and the angle test prior to this date.

Table 14-2: Start and stop times for gain tests and angle test prior to 2008-10-28.

Type of Test	Start Date and Time	Stop Date and Time
CEM Gain Test	2007-06-04T17:12:34.41	2007-06-04T18:54:58.41
Deflector Angle Sweep Test	2007-06-12T05:30:42.51	2007-06-12T10:23:18.41
CEM Gain Test	2008-09-28T10:06:58.62	2008-09-28T11:36:34.62
CEM Gain Test	2008-10-27T04:07:30.63	2008-10-27T05:38:09.63

14.10 SWAP Science Goals

These Level 2 (calibrated) data products described above will allow us to meet two key science goals as outlined in Section 6.3.1.1 of the SWAP Specification Document (Document No. 05310-03-SWAPSPEC-01). Below we quote this section.

The Mission Science Requirements document specifies that SWAP should make the following measurements.

- *Measure solar wind standoff to ~ 3000 km.*
- *Determine nature of solar wind interaction at Pluto. Distinguish between magnetic, cometary, & ionospheric interactions.*

14.10.1 Dataflow Block Diagram

Figure 14-14 provides a summary of the inputs (in gold and violet) and steps (in blue) in the calibration pipeline which convert the raw data to scientific units.




Figure 14-14: SWAP Level 3 Pipeline Diagram

14.10.2 Extra FITS Extensions and Their Definitions

We have binary tables for both real-time, and summary data. Likewise, we have binary table extensions for housekeeping data, flags, errors, and thruster packet spacecraft data. We have image extensions for both count rate spectrograms and error spectrograms image. For histogram data we have two image extensions: one for counts and one for accumulations (number of samples).

14.10.3 Scientific Units

14.10.4 Additional FITS and PDS Keywords Added

The Level 2 calibrated code has rate spectrogram in Hz, and all the engineering data converted to engineering units.

14.10.5 *Hardware/OS Development Platform*

Intel, Linux

14.10.6 *Language(s) Used*

IDL

14.10.7 *Third Party Libraries Required*

SPICE and FITS.

14.10.8 *Memory Required***14.10.9 *Temporary File System Space Needed***

None

14.10.10 *Predicted Size of Output File(s)***14.10.11 *Predicted Execution time***

On the order of seconds

14.10.12 *Contact/Support Person(s)*

PI: Dave McComas

Lead Engineer and Project Manager: Scott Weidner

Onboard Software and Commanding: John Hanley

Pipeline and Science Operations: Heather Elliott

Sequencing: Helen Hart

14.10.13 *Maintenance Schedule (Code/Data Updates, Documentation)***14.11 SWAP Publications Examples**

Ebert, R. W., D. J. McComas, B. Rodriguez, P. Valek, and S. Weidner, A composition analysis tool for the Solar Wind Around Pluto (SWAP) instrument on *New Horizons*, *Space Sci. Rev.*, doi 10.1007/s11214-010-9683-6, 156, 1-12, 2010.

Ebert, R. W., D. J. McComas, F. Bagenal, and H. A. Elliott, Location, structure, and motion of Jupiter's dusk magnetospheric boundary from ~1625 to 2550 RJ, *J. Geophys. Res.*, 115, A12223, doi: 10.1029/2010JA015938, 2010.

McComas, D. J. and F. Bagenal, Jupiter: A fundamentally different magnetospheric interaction with the solar wind, *Geophys. Res. Lett.*, 34, L20106, doi:10.1029/2007GL031078, 2007.

McComas, D.J., F. Allegrini, F. Bagenal, F. Crary, R. W. Ebert, H. Elliott, A. Stern, and P. Valek, Diverse Plasma Populations and Structures in Jupiter's Magnetotail, *Science*, 318, doi: 10.1126/science.1147393, 217-220, 2007.

- McComas, D. J., F. Allegrini, F. Bagenal, P. Casey, P. Delamere, D. Demkee, G. Dunn, H. Elliott, J. Hanley, K. Johnson, J. Langle, G. Miller, S. Pope, M. Reno, B. Rodriguez, N. Schwadron, P. Valek, and S. Weidner, The Solar Wind Around Pluto (SWAP) instrument aboard *New Horizons*, *Space Sci. Rev.*, doi: 10.1007/s11214-007-9205-3, 140, 261-313, 2008.
- McComas, D. J. and F. Bagenal, Re: Jupiter: A fundamentally different magnetospheric interaction with the solar wind, Response to Comment, *Geophys. Res. Lett.*, 35, L10103, doi:10.1029/2008GL034351, 2008.
- McComas, D. J., H. A. Elliott, and N. A. Schwadron, Pickup hydrogen distributions in the solar wind at ~11 AU: Do we understand pickup ions in the outer heliosphere?, *J. Geophys. Res.*, 115, A03102, doi: 10.1029/2009JA014604, 2010.
- Nicolaou, G., D. J. McComas, F. Bagenal, and H.A. Elliott, Properties of plasma ions in the distant Jovian magnetosheath using Solar Wind Around Pluto data on *New Horizons*, *J. Geophys. Res.*, 119, 3463-3479, doi: 10.1002/2013JA019665, 2014.
- Randol, B. M., R. W. Ebert, F. Allegrini, D. J. McComas, and N. A. Schwadron, Reflections of ions in electrostatic analyzers: A case study with *New Horizons*/Solar Wind Around Pluto, *Rev. Sci. Instrum.*, 81, 114501, 2010.
- Randol, B. M. and D. J. McComas, Density correlations between solar wind and pick-up ions, Pickup Ions throughout the Heliosphere and Beyond: Proceedings of the 9th Annual International Astrophysics Conference, AIP CP1302, pp. 64-69, 2010.
- Randol, B. M., H. A. Elliott, J. T. Gosling, D. J. McComas, and N. A. Schwadron, Observations of isotropic interstellar pick-up ions at 11 and 17 AU from *New Horizons*, *Astrophys. J.*, 755:75, doi:10.1088/0004-637X/755/1/75, 2012.
- Randol, B. M., D. J. McComas, and N. A. Schwadron, Interstellar pick-up ions observed between 11 and 22 AU by *New Horizons*, *Astrophys. J.*, 768:120, doi:10.1088/0004-637X/768/2/120, 2013.
- Young, L. A., S. A. Stern, H. A. Weaver, F. Bagenal, R. P. Binzel, B. Buratti, A. F. Cheng, D. Cruikshank, G. R. Gladstone, W. M. Grundy, D. P. Hinson, M. Horanyi, D. E. Jennings, I. R. Linscott, D. J. McComas, W. B. McKinnon, R. McNutt, J. M. Moore, S. Murchie, C. B. Olkin, C. C. Porco, H. Reitsema, D. C. Reuter, J. R. Spencer, D. C. Slater, D. Strobel, M. E. Summers, and G. L. Tyler, *New Horizons*: Anticipated scientific investigations at the Pluto system, *Space Sci. Rev.*, doi: 10.1007/s11214-008-9462-9, 140, 93-127, 2008.

14.12 Packet Description

This section gives the parameters found in SWAP science packets. The names are as given in the APL spreadsheets used to view the data in real-time and playback modes. The Level 1 names are shortened versions of these names.

REAL_TIME

SWAP_RT.CC_PID_VER	Primary Header - Version Number
SWAP_RT.CC_PID_TYPE	Primary Header - Packet Type
SWAP_RT.CC_PID_SEC_FLAG	Primary Header - Secondary Header Flag
SWAP_RT.CC_PID_APID	Primary Header - Application ID; type of data i.e. real-time, housekeeping, etc

SWAP_RT.CC_PSC_GRP_FLAG	Primary Header - Grouping Flag; set to a constant value, not used
SWAP_RT.CC_PSC_SSC	Primary Header - Sequence Count
SWAP_RT.CC_PKT_DATA_LN	Primary Header - Packet Length; this is used to
SWAP_RT.CC_MET_TIME	Secondary Header - Time
SWAP_RT.SEC64_ST	A bit indicating the beginning of a 64-second cycle
SWAP_RT.PLAN_ID	Plan ID
SWAP_RT.SWEEP_ID	Sweep ID
SWAP_RT.ANGLE	Commanded angle for deflection compensation
SWAP_RT.RPA_LVL0	RPA level during first half-second
SWAP_RT.DFL_LVL0	DFL level during first half-second
SWAP_RT.ESA_LVL0	ESA level during first half-second
SWAP_RT.RPA_LVL1	RPA level during first half-second
SWAP_RT.DFL_LVL1	DFL level during first half-second
SWAP_RT.ESA_LVL1	ESA level during first half-second
SWAP_RT.MODE	Enumerated type representing each of the modes
SWAP_RT.PCEM_RNG_ST0	PCEM counts are compressed or raw during first half-second; this is a flag that determines if the data in SWAP_RT.PCEM_CNT0 is compressed or not. If this flag is set to zero the data are not compressed. If the flag is set to 1 then the data in SWAP_RT.PCEM_CNT0 is 1/16 of the actual count rate.
SWAP_RT.SCEM_RNG_ST0	SCEM counts are compressed or raw during first half-second
SWAP_RT.COIN_RNG_ST0	Coincidence counts are compressed or raw during first half-second
SWAP_RT.PCEM_RNG_ST1	PCEM counts are compressed or raw during second half-second
SWAP_RT.SCEM_RNG_ST1	SCEM counts are compressed or raw during second half-second
SWAP_RT.COIN_RNG_ST1	Coincidence counts are compressed or raw during second half-second
SWAP_RT.PCEM_CNT0	PCEM count value during first half-second
SWAP_RT.SCEM_CNT0	SCEM count value during first half-second
SWAP_RT.COIN_CNT0	Coincidence count value during first half-second
SWAP_RT.PCEM_CNT1	PCEM count value during second half-second
SWAP_RT.SCEM_CNT1	SCEM count value during second half-second
SWAP_RT.COIN_CNT1	Coincidence count value during second half-second
SWAP_RT.CHKSUM	XOR checksum; this value is based on all the other real-time quantities listed above. It is used to determine if a bit was flipped in one of the real-time quantities. If check sum is not what is predicted based on the other real-time quantities then the data need to be thrown away.

SUMMARY

SWAP_SM.CC_PID_VER	Primary Header - Version Number
SWAP_SM.CC_PID_TYPE	Primary Header - Packet Type
SWAP_SM.CC_PID_SEC_FLAG	Primary Header - Secondary Header Flag
SWAP_SM.CC_PID_APID	Primary Header - Application ID
SWAP_SM.CC_PSC_GRP_FLAG	Primary Header - Grouping Flag
SWAP_SM.CC_PSC_SSC	Primary Header - Sequence Count
SWAP_SM.CC_PKT_DATA_LN	Primary Header - Packet Length
SWAP_SM.CC_MET_TIME	Secondary Header - Time
SWAP_SM.BEG_TIME	Time stamp at the beginning of this summary period
SWAP_SM.END_TIME	The time stamp at the end of this summary period
SWAP_SM.N64_CNT	Number of 64-second sample sets used
SWAP_SM.NDFL	DFL
SWAP_SM.ANGLE_SUM	Angle
SWAP_SM.DENSITY_SUM	An estimate of the pseudo density
SWAP_SM.VELOCITY_SUM	An estimate of the pseudo speed using energy
SWAP_SM.TEMP_SUM	Estimate of the pseudo temperature
SWAP_SM.ANGLE_SSQ_HI	The variance of the angles
SWAP_SM.ANGLE_SSQ_LO	The variance of the angles
SWAP_SM.DENSITY_SSQ_HI	The variance of the pseudo n-values
SWAP_SM.DENSITY_SSQ_LO	The variance of the pseudo n-values
SWAP_SM.VELOCITY_SSQ_HI	The variance of the pseudo V-values
SWAP_SM.VELOCITY_SSQ_LO	The variance of the pseudo V-values
SWAP_SM.TEMP_SSQ_HI	The variance of the pseudo T-values
SWAP_SM.TEMP_SSQ_LO	The variance of the pseudo T-values
SWAP_SM.ANGLE_MIN	Minimum Angle
SWAP_SM.DENSITY_MIN	Minimum pseudo density
SWAP_SM.VELOCITY_MIN	Minimum pseudo speed
SWAP_SM.TEMP_MIN	Minimum pseudo Temperature
SWAP_SM.ANGLE_MAX	Maximum Angle
SWAP_SM.DENSITY_MAX	Maximum pseudo density
SWAP_SM.VELOCITY_MAX	Maximum pseudo speed
SWAP_SM.TEMP_MAX	Maximum pseudo temperature
SWAP_SM.CHKSUM	XOR checksum

HISTOGRAM Type 1

SWAP_H0.CC_PKT_DATA_LN	Primary Header - Packet Length
SWAP_H0.CC_MET_TIME	Secondary Header - Time
SWAP_H0.BEG_TIME	Time stamp at the beginning of this histogram period
SWAP_H0.END_TIME	The time stamp at the end of this histogram period
SWAP_H0.SMPLS_CNT	Number of 64-second samples used
SWAP_H0.PLAN_ID	The Plan ID used for the current science sweeping mode
SWAP_H0.TABLE_ID	Table ID within the Plan ID that is being used
SWAP_H0.SEQNUM_CNT	Sequence number starting at 0 for header
SWAP_H0.DATA	Histogram data
SWAP_H0.CHKSUM	XOR checksum

HISTOGRAM Type 2

SWAP_H1.CC_PID_VER	Primary Header - Version Number
SWAP_H1.CC_PID_TYPE	Primary Header - Packet Type
SWAP_H1.CC_PID_SEC_FLAG	Primary Header - Secondary Header Flag
SWAP_H1.CC_PID_APID	Primary Header - Application ID
SWAP_H1.CC_PSC_GRP_FLAG	Primary Header - Grouping Flag
SWAP_H1.CC_PSC_SSC	Primary Header - Sequence Count
SWAP_H1.CC_PKT_DATA_LN	Primary Header - Packet Length
SWAP_H1.CC_MET_TIME	Secondary Header - Time
SWAP_H1.SEQNUM_CNT	Sequence number
SWAP_H1.DATA	Histogram data
SWAP_H1.CHKSUM	XOR checksum starting with beginning of CCSDS header

HOUSEKEEPING

SWAP_HK.CC_PID_VER	Primary Header - Version Number
SWAP_HK.CC_PID_TYPE	Primary Header - Packet Type
SWAP_HK.CC_PID_SEC_FLAG	Primary Header - Secondary Header Flag
SWAP_HK.CC_PID_APID	Primary Header - Application ID
SWAP_HK.CC_PSC_GRP_FLAG	Primary Header - Grouping Flag
SWAP_HK.CC_PSC_SSC	Primary Header - Sequence Count
SWAP_HK.CC_PKT_DATA_LN	Primary Header - Packet Length
SWAP_HK.CC_MET_TIME	Secondary Header - Time
SWAP_HK.CMD_EXE_CNT	Cumulative mod-256 count of successfully executed commands
SWAP_HK.CMD_REJ_CNT	Cumulative mod-256 count of rejected commands
SWAP_HK.LUT_CHOICE	Which LUT is in use
SWAP_HK.PCEM_SAFE	PCEM was safed due to CEM interrupts

SWAP_HK.SCEM_SAFE	SCEM was safed due to CEM interrupts
SWAP_HK.WDT_ST	SWAP has rebooted due to a watchdog expiration.
SWAP_HK.RCV_SAFE_ST	SWAP has received safe command from S/C
SWAP_HK.SAFE_ST	SWAP has safed itself.
SWAP_HK.PCEM_RATE_ST	Count rate threshold for PCEM counter has been exceeded
SWAP_HK.SCEM_RATE_ST	Count rate threshold for the SCEM counter has been exceeded
SWAP_HK.PCEM_CURR_ST	Current threshold for the PCEM has been exceeded
SWAP_HK.SCEM_CURR_ST	Current threshold for the SCEM has been exceeded
SWAP_HK.PCEM_VOLT_ST	Voltage tolerance for the PCEM has been exceeded.
SWAP_HK.SCEM_VOLT_ST	Voltage tolerance for the SCEM has been exceeded.
SWAP_HK.LVPS_VOLT_ST	Voltage tolerance for +5 V or -5V supply has been exceeded.
SWAP_HK.LVPS_CURR_ST	Current tolerance for +5 V or -5V supply has been exceeded.
SWAP_HK.OVR_TEMP_ST	Upper temperature limit exceeded.
SWAP_HK.UND_TEMP_ST	Lower temperature limit exceeded.
SWAP_HK.MODE	Enumerated type representing each of the modes
SWAP_HK.MEMDP_ST	MEMDUMP State
SWAP_HK.SENSOR_TEMP	Temperature of sensor detector. AD mux = 0x10
SWAP_HK.HVSUPP_TEMP	Temperature of HVPS. AD mux = 0x11
SWAP_HK.CNTRLR_TEMP	Temperature of controller. AD mux = 0x12
SWAP_HK.PCEM_VOLT	Voltage monitor of PCEM HVPS. AD mux = 0x02
SWAP_HK.SCEM_VOLT	Voltage monitor of SCEM HVPS. AD mux = 0x03
SWAP_HK.PCEM_CURR	Strip current monitor of PCEM HVPS. AD mux = 0x04
SWAP_HK.SCEM_CURR	Strip current monitor of SCEM HVPS. AD mux = 0x05
SWAP_HK.P5_VOLT	Voltage monitor of +5V power supply. AD mux = 0x0c
SWAP_HK.N5_VOLT	Voltage monitor of -5V power supply. AD mux = 0x0d
SWAP_HK.P5_CURR	Current monitor of +5V power supply. AD mux = 0x0e
SWAP_HK.N5_CURR	Current monitor of -5V power supply. AD mux = 0x0f
SWAP_HK.SWAP_REV	Revision number for the SWAP software
SWAP_HK.LAST_OPCODE	Opcode of last executed command
SWAP_HK.PHD_LLD_LVL	DAC Level of PHD LLD
SWAP_HK.MEMLD_ST	MEMLOAD state
SWAP_HK.OPT1_ST	State of primary optics
SWAP_HK.OPT2_ST	State of backup optics
SWAP_HK.PCEM_ST	State of Primary Channel Electron Multiplier disable/enable
SWAP_HK.SCEM_ST	State of Secondary Channel Electron disable/enable
SWAP_HK.SPARE1	SPARE
SWAP_HK.PCEM_CNT_ST	The PCEM count rate was tripped but handled by SWAPFW
SWAP_HK.SCEM_CNT_ST	The SCEM count rate was tripped but handled by SWAPFW

SWAP_HK.PCEM_CURRTHR	Current level when PCEM safety algos are tripped
SWAP_HK.SCEM_CURRTHR	Current level when SCEM safety algos are tripped
SWAP_HK.PCEM_LVL	PCEM DAC level
SWAP_HK.SCEM_LVL	SCEM DAC level
SWAP_HK.AGND_VOLT	Voltage monitor of Analog ground. A/D mux = 0x00
SWAP_HK.CEM_CURR	Current level when SCEM safety interrupt is tripped
SWAP_HK.ESA1_VOLT	Voltage monitor of ESA HVPS. AD mux = 0x06
SWAP_HK.ESA2_VOLT	Voltage monitor of ESA HVPS. AD mux = 0x07
SWAP_HK.DFL1_VOLT	Voltage monitor of DFL HVPS. AD mux = 0x08
SWAP_HK.DFL2_VOLT	Voltage monitor of DFL HVPS. AD mux = 0x09
SWAP_HK.RPA1_VOLT	Voltage monitor of RPA HVPS. AD mux = 0x0a
SWAP_HK.RPA2_VOLT	Voltage monitor of RPA HVPS. AD mux = 0x0b
SWAP_HK.P2_5_VOLT	Voltage monitor of +2.5V reference. AD mux = 0x13
SWAP_HK.PHD_LLD_VOLT	Voltage monitor of PHD_LLD. AD mux = 0x14
SWAP_HK.PCEM_RATELIM	The count value at which the primary CEM safety limit is set.
SWAP_HK.SCEM_RATELIM	The count value at which the secondary CEM safety limit is set.
SWAP_HK.STIM_ENA	State of whether the stimulator pulsers are enabled or disabled.
SWAP_HK.PPS_SEL_ST	State of which side of the IEM interface is being used.
SWAP_HK.PPS_DET_ST	1 PPS detected state
SWAP_HK.CEM_INT_LIM	Current limit at which the CEM triggers an interrupt.
SWAP_HK.CMD_ECHO_ST	Whether command echo is enabled or disabled
SWAP_HK.HV_PGSAFE_ST	State of safe / arm plug
SWAP_HK.HV_PGENA_ST	State of high-voltage disable/enable plug
SWAP_HK.HV_ARM_ST	State of high-voltage software disable/enable
SWAP_HK.CEM_INT_DIP	Counts to dip the CEM supplies
SWAP_HK.PLAN_ID	The Plan ID used for the current science sweeping mode if any.
SWAP_HK.SWEEP_ID	Current Sweep table ID
SWAP_HK.ANGLE	The commanded angle for deflection compensation
SWAP_HK.PCEM_VLT1_ST	PCEM voltage is out of tolerance for only one 0.5 second sample
SWAP_HK.PCEM_CUR1_ST	PCEM current is out of tolerance for only one 0.5 second sample
SWAP_HK.SCEM_VLT1_ST	SCEM voltage is out of tolerance for only one 0.5 second sample
SWAP_HK.SCEM_CUR1_ST	SCEM current is out of tolerance for only one 0.5 second sample
SWAP_HK.PCEM_INT_ST	The CEM current interrupt was tripped, but handled by SWAPFW
SWAP_HK.SCEM_INT_ST	The CEM current interrupt was tripped, but handled by SWAPFW

SWAP_HK.EEP2_RDY	EEPROM 2 is ready to be written
SWAP_HK.EEP1_RDY	EEPROM 1 is ready to be written
SWAP_HK.FPGA_TYPE	Type number of the FPGA
SWAP_HK.FPGA_REV	Revision number of the FPGA
SWAP_HK.SM_TLM	How often the science summary packet is output.
SWAP_HK.HX_TLM	How often the histogram telemetry packet is output.
SWAP_HK.RT_TLM	How often all of the 64-second real-time packets are output
SWAP_HK.HK_TLM	How often the housekeeping packet is output.
SWAP_HK.FPGA_PUP_ST	A status of the power on check of the FPGA initialization check
SWAP_HK.EEPL2_CKS_ST	Status of the power on check of the EEP_L2 checksum
SWAP_HK.EEPL1_CKS_ST	Status of the power on check of the EEP_L1 checksum
SWAP_HK.RAM_D_ST	Status of the power on check of the RAM_D memory test
SWAP_HK.EEPC2_CKS_ST	Status of the power on check of the EEP_C2 checksum
SWAP_HK.EEPC1_CKS_ST	Status of the power on check of the EEP_C1 checksum
SWAP_HK.RAM_C_ST	Status of the power on check of the RAM_C memory test
SWAP_HK.PROM_CKS_ST	Status of the power on check of the PROM checksum
SWAP_HK.CHKSUM	XOR checksum

15. ACRONYMS AND ABBREVIATIONS

ACO	Annual CheckOut
AD	Analog-to-Digital (converter)
ADU	Analog-to-Digital Units; same as DN
ADC	Analog-to-Digital Converter; converts electrical quantity, typically Volts, to DN or ADU
ADDR	Address
AGC	Automatic Gain Control
AGND	Analog GrouND
AIP	American Institute of Physics
a.k.a.	Also Known As
ANA	ANAlog
ApID	Application (Process) Identifier; 12-bit integer for telemetry packets, usually in hexadecimal
APL	Alternate form of JHUAPL
APP	APPLication
ASCII	American Standard Code for Information Exchange
ASE	Associate Support Electronics
ATM	ATMospheric (science theme team)
AU	Astronomical Unit (distance)
BIL	Band Interleaved by Line
bit	Binary Digit; 0 or 1 in base two
BTI	Bad Time Interval (PEPSSI)
C&DH	see CDH
CAL	CALibration; CALibrated data or signal
CALIB	CALibration; CALibrated data or signal
CCD	Charge-Coupled Device (Instrument) or Critical Command Decoder (RF/Telecomm.)
CCSDS	Consultative Committee for Space Data Systems; telemetry packet format specification
CDH	Command and Data Handling; S/C or instrument; typically redundant, so also CDH1, CDH-2, ...
CDS	Correlated Double Sample
CEM	Channel Electron Multiplier
CH	CHannel
CKS	CheckSum
CKSM	CheckSuM
CM	Configuration Management (do not confuse with cm => centimeter)
CNO	Carbon, Nitrogen, Oxygen (ions)
CNT	CouNT
CODMAC	Committee on Data Management and Computation; defines data levels within NASA
COMP	COMPosition (science theme team)
CPU	Central Processing Unit

CSOC	Christy Science Operations Center at MOPS JHU/APL in Laurel, MD
CTI	Charge Transfer Inefficiency
CW	Clock-Wise
CU	Compression Unit
DAC	Digital-to-Analog Converter
DDL	Double Delay Line
DEC	DEClination
DES	DEScriptor
DET	DETelector
DFL	DeFLector
DIAG	DIAGnostic
DIR	DIRectory
DMAP	Data Management and Archive Plan
DMP	DuMP
DN	Data Number; raw value measured by Analog-to-Digital Converter; typically an integer
DOI	Digital Object Identifier
DOY	Day Of Year (1-366)
DSN	Deep Space Network
DT	Delta (change in, accumulation of) Time
DU	Data Unit (of a FITS file)
EDU	Extension Data Unit (of a FITS file); EDUs follow PDU; EDU numbering starts at 1
EHDU	Extension Header Data Unit – FITS header of Extension Data Unit
ELEC	ELECTronics
EME	Earth Mean Equator
ENA	ENABle; ENABled
EPS	Energetic Particle Sensor
EPU	Event Processor Unit
ERR	ERRor
ESA	ElectroStatic Analyzer; part of SWAP instrument
ET	Ephemeris Time, s path J2000 epoch
ETIME	Effective exposure TIME
EXE	EXEcutable (file; program)
FF	Flat Field
FIR	Finite Impulse Response (filter)
FITS	Flexible Image Transport System; data storage format
FLG	FLaG
FOV	Field Of View
FPGA	Field Programmable Gate Array
FPN	Fixed Pattern Noise
FPU	Floating Point Unit
FUV	Far-UltraViolet
FV	FITS Viewing application; also Fv
GB	GigaByte; 1,000,000,000 bytes

GC	Guidance and Control
GGI	Geology and Geophysics Investigation (science theme team)
GN&C	Guidance, Navigation, and Control
GNU	Gnu's Not Unix
GRP	GRouP; GRouPing
GSEOS	Ground Support Equipment Operating System
GSFC	Goddard Space Flight Center
HDU	Header Data Unit
HGA	High-Gain Antenna
HK	HouseKeeping
HST	Hubble Space Telescope
HTRC	HeaTeR Current
HV	High-Voltage
HVPS	High-Voltage Power Supply
HVSCI	High-Voltage SCIENCE
HW	HardWare
HX	Histogram
Hz	Hertz (frequency unit); also HZ
i.e.	Id Est (Latin: "that is")
IAU	Internation Astronomical Union
ICD	Interface Control Document
ID	IDentifier
IDL	Interactive Data Language
IEM	Integrated Electronic Module
IF	Intermediate Frequency
IFOV	Incrmental Field-Of-View
IMG	IMaGe
INT	INTerrupt
IR	InfraRed
IUE	International Ultraviolet Explorer
IPM	Inter-Planetary Medium
ITF	Instrument Transfer Frame
JHK	1220 (J), 1630 (H), 2190 (K) nm (photometric system)
JHU	Johns Hopkins University
JHUAPL	Johns Hopkins University/Applied Physics Laboratory; also JHU/APL
JPEG	Joint Photographic Experts Group
JPL	Jet Propulsion Laboratory
JSP	Jupiter SPin
JUP	JUPiter
kB	kilobyte; 1,000 bytes
KBO	Kuiper Belt Object
KiB	kibibyte; 1,024 = 2 ¹⁰ bytes
LASP	Laboratory for Atmospheric and Space Physics, University of Colorado Boulder
LBL	Label (file name suffix)

LCP	Left Circular Polarization
LEI	LEIsa
LEISA	Linear Etalon Imaging Spectral Array
LEN	LENgth
LIM	LIMit
LLD	Low Level Discriminator
LO	LOw
LORRI	LOng Range Reconnaissance Imager
LPS	LEISA Processing System
LSB	Least-Significant Bit/Byte first; ordering of bits or bytes; also lsb, LSByte-first, ...
LUT	LookUp Table
LVDS	Low Voltage Difference Signal
LVL	LeVeL
LVPS	Low-Voltage Power Supply
MAC	MACro
MAX	MAXimum
MB	MegaByte; 1,000,000 bytes
Mbit	MegaBit; 1,000,000 bits
MCP	MicroChannel Plate; detector
MDM	Master Data Manager
MEM	MEMory
MEMDP	MEMory DumP
MEMLD	MEMory Load
MET	Mission Event (or Elapsed) Time; 10-digit, zero-padded, decimal representation of S/C seconds; may have decimal point and fractional suffix
MeV	Mega-electron Volt
MiB	MebiByte; $1,048,576 = 2^{20}$ bytes
MIDL	Mission Independent Data Layer
MIKE	Nickname for UltraViolet data processing pipeline
MIN	MINimum
MOC	Mission Operations Center
MOD	MODulo; MODulus
MON	MONitor
MOPS	Mission OPerationS
MSB	Most-Significant Bit/Byte first; ordering of bits or bytes; also msb, MSByte-first, ...
MVIC	Multispectral Visible Imaging Camera
N.B.	Nota Bene (Latin: "note well" i.e. pay particular attention to what follows)
NaN	Not a Number
NASW	prefix for contract designators
NEG	NEGative (voltage)
NIR	Near-InfraRed
NH	<i>New Horizons</i> ; may refer to spacecraft, mission or project
NUM	NUMber
NYI	Not Yet Implemented

OAP	Off-Axis Paraboloidal (mirror)
OPT	OPTics
OS	Operating System
OVR	OVeR (limit)
PAN	PANcrhomatic
PARM	PARaMeter
PC	Personal Computer
PCEM	Primary Channel Electron Multiplier
PCI	Peripheral Component Interconnect (backplane)
PDF	Portable Document Format
PDS	Planetary Data System; NASA archive organization for planetary data
PDU	Primary Data Unit (of FITS file); compare to EDU; often interpreted as (E)DU 0
PERSI	Pluto Exploration Remote Sensing Investigation; sensor suite: MVIC; LEISA; ALICE.
PEPSSI	Pluto Energetic Particles Spectrometer Science Investigation
PGENA	PluG ENAbLe
PH	Primary Header
PHA	Pulse-Height Analysis
PHD	Pulse-Height Distribution
PHDU	Primary Header Data Unit – FITS header of Primary Data Unit
PI	Principal Investigator
PKT	PacKeT
PNL	PaNeL
POS	POSition
PPS	Pulse-Per-Second
PRN	Pseudo-RaNdom
PROM	Programmable Read-Only Memory
PSF	Point Spread Function
PUP	Power-UP
PV	Positive Voltage
PVDF	PolyVinylidene Flouride
RA	Right Ascension
RAM	Random Access Memory
RCP	Right Circular Polarization
RCS	Reaction Control System
RCV	Receive; Received
RDY	ReaDY
REJ	REJected
REV	REVision
REX	Radio Experiment
RF	Radio Frequency
RI	658 (R) and 806 (I) nm; photometric system
Rj	Radius of Jupiter
RN	Read Noise

ROF	REX Output Frame
R _p	Radius of Pluto
RPA	Retarding Potential Analyzer
RPM	Revolutions Per Minute
RSLT	ReSuLT
RT	Real-Time
RTG	Radioisotope Thermoelectric Generator
R _x	Receive (line; pin-out; channel)
S/C	Spacecraft; also SC, sc (subscript e.g. Y _{sc})
SC	SpaceCraft
SCEM	Secondary-Channel Electron Multiplier
SCLMP	Secondary Calibration LaMP
SCSI	Small Computer System Interface
SDC	(Venetia Burney) Student Dust Counter
SEC	SECOndary
SEL	SELEct
SH	Secondary Header
SIA	Solar Illumination Assembly (see [REUTERETAL2008])
SIG	SIGma (one standard deviation)
SIP	Simple Imaging Polynomial
SM	SuMmary
SMM	State Machine Mode
SMPLS	SaMPLeS
SNR	Signal-to-Noise Ratio
SOC	Science Operations Center
SOCC	Solar OCCultation or Solar OCultation Channel (ALICE instrument)
SPC	SPiCe
SPCK	SPiCe Kernel
SPEC	SPECTrogram
SPECT	SPECTrogram
SPYCE	Python code accessing SPICE toolkit
SQL	Structured Query Language
SRAM	Static Random Access Memory
SRIM	Stopping and Range of Ions in Matter
SSC	Sequence Count
SSD	Solid-State Detector
SSF	SeqGen Sequence Files
SSG	Company that built LORRI structure
SSR	Solid-State Recorder; two on-board NH, so also SSR1, SSR-2, ...
ST	STate
STAT	STATus
STIM	STIMulus; STIMulator
STP	Supplemented Telemetry Packet
STScI	Space Telescope Science Institute

SWAP	Solar Wind Around Pluto
SWAPFW	SWAP FirmWare
SwRI	Southwest Research Institute; office in Boulder, CO, USA
TBD	To Be Determined (i.e. not known now)
TDB	Barycentric Dynamical Time
TDI	Time-Delay Integration
TEXP	EXPosure Time
TF	Thruster Firing
TLM	TeLeMetry
TM	TeLeMetry
TNF	Tracking and Navigation File
TOD	Time Of Day
TOF	Time Of Flight; sometimes used as an analog for along-track distance
TRK	TRAcK
TSOC	Tombaugh Science Operations Center - SOC at SwRI in Boulder, CO
Tx	Transmit (line; pin-out; channel)
TYP	TYPe
UBV	365 (U), 445 (B), 551 (V) nm; photometric system
UBVRI	See UVB and RI
UKIRT	United Kingdom Infrared Telescope
UND	UNDer
USO	Ultra-Stable Oscillator
UTC	Coordinated Universal Time
UV	UltraViolet
VER	VERsion
VLT	VoLTage
VO	VOltage
VX	Velocity in X
WCS	World Coordinate System
WDT	WatchDog Timeout
XOR	eXclusive OR; binary logic operator