Chapter 7

EDS Instrument Description

7.1 Introduction

This manual is intended to help *X-ray timing explorer* (XTE) observers use the *experiment data system* (EDS) in a proper and optimal fashion. The EDS is an on-board microprocessor-based electronics package that controls data acquisition and formatting for the *proportional counter array* (PCA) and the *all-sky monitor* (ASM).

The EDS can run any of 10 'modes' (programs), 7 for PCA data and 3 for ASM data, each of which is further defined with a set of parameters (e.g. energy bin boundaries, time bin widths, type of event). Examples are 'binned data' and 'FFT'. A designation of a particular mode together with a set of parameters is known as a *configuration*. The number of possible configurations consists of all combinations of modes and parameters and is therefore very large. A set of ~650 configurations are permanently recorded in the EDS and are available for selection by the observer. In addition, others are available for uplink to the EDS. Observers may suggest additional configurations.

The EDS consists of 8 parallel processing systems known as *event analyzers* (EAs) which can carry out different analyses in parallel on the incoming data streams. A configuration specifies completely the setup and operation of a single EA. The setup will be fixed during an *EA run.* An observation or a single pointing of the spacecraft will usually require only one EA setup for each of the available EAs. However, an observer might choose to change modes (i.e. configuration) of one or more EAs during an observation. In this chapter, an 'observation' usually alludes to a fixed EA setup. Occasionally we use the more precise 'EA run'.

This chapter covers the (1) basic theory of operation, and (2) the modes of data processing provided by EDS flight software. An adequate understanding of the operation of the system (for PCA observations) may be obtained by reading through “Time Signals and Counters” (Section 4.4 of this chapter). This should be sufficient preparation for Section 5 (“PCA Data Modes”) and for the chapter on “EDS Configurations” (Appendix 1). The EDS operation for the ASM is described in Section 6, and the ASM data modes are described in Section 7.

The configurations are described in another chapter (Appendix 1).

7.2 Experiment Data System (EDS)

The large effective area and excellent time resolution capability of the PCA creates a great strain on telemetry resources. If the data were to be telemetered simply as 48 bits (time, detector, and energy information) per individual PCA event, the nominal sustained PCA telemetry rate limit of ~20 kb/s would be exceeded for average total count rates above ~420 c/s. This limit is equivalent to an X-ray source intensity of only ~50 mCrab and makes no allowance for non-source background. The EDS thus serves to compress the PCA data by means of time and energy binning, folding to study periodic signals, etc. The EDS also converts ASM event data into a low-bandwidth form.

7.2.1 Interfaces

Figure 1 is a diagram of the flow of data into and out of the EDS. The EDS consists of two nearly identical halves, EDS-A and EDS-B; each contains 4 EA’s. Detector events in digital form are piped to the EAs from the five *proportional counter units* (PCUs) of the PCA and from the three *scanning shadow cameras* (SSCs) of the ASM. EDS-B normally controls the rotation of the ASM and handles housekeeping data from the SSCs and *drive assembly*. In case of failure, EDS-A can control the rotation drive, but can not monitor housekeeping data. Burst trigger signals can be sent to or from the HEXTE from EDS-A (not shown).

FIGURE 7.1. EDS Subsystems and Interfaces

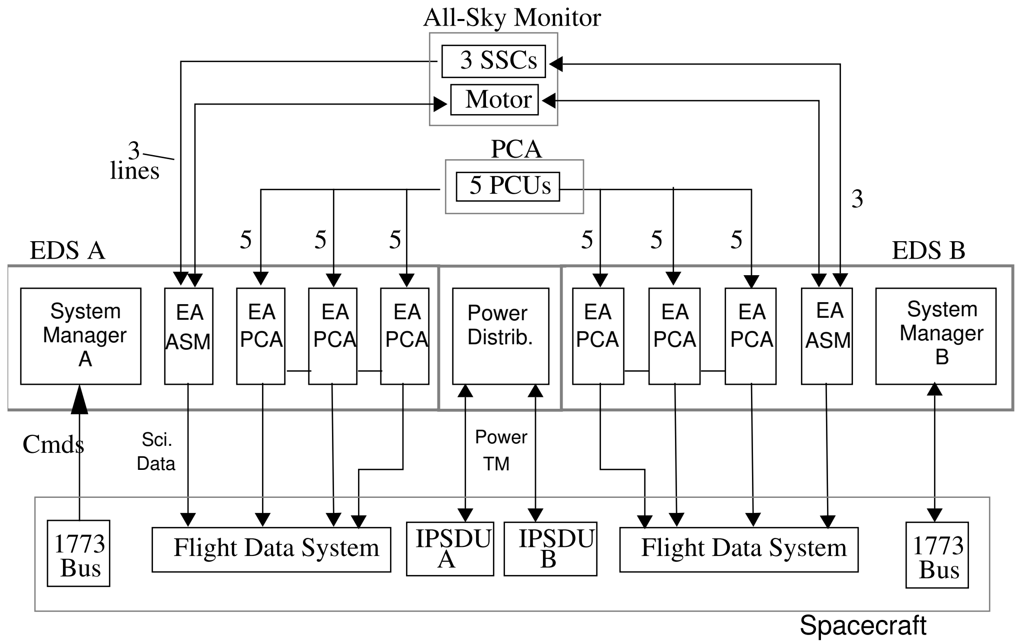

Power is received from the S/C *instrument power switching and distribution units* (IPSDUs) together with clock signals, while housekeeping data flow from the EDS to the IPSDUs. Either IPSDU can power the entire EDS. A fiber optic link (1773 bus) to the spacecraft *flight data system* (FDS) provides discrete and serial commands. A set of eight RS422 serial links, one for each EA, provides high bandwidth paths (1 MHz each) over which science data is carried to the flight data system.

7.2.2 EDS Functions

The EDS allows the observer to recover the most desired information by providing multiple data-processing options and opportunities on board the XTE spacecraft. While the EDS can send PCA data on an event-by-event basis, it can also bin events with regard to time, photon energy, detector number, and event type (e.g., measurement chain). Pulse profile averaging (pulsar folding), and Fourier power spectrum averaging can be performed. There are supplementary options to record data at very high time resolution, and it is possible to identify X-ray burst-like episodes and send detailed data from them.

The eight EAs operate independently of one another. Six of the EAs are configured (in hardware) to service the PCA while two support the ASM. These are known as PCA EAs and ASM EAs respectively. Each PCA EA receives all PCA data. This permits several different data processing routines to operate on the data stream independently and simultaneously. The data system allows the outputs of the several EAs to be independently sent to the FDS as long as the total data rate remains within the limits of the overall system.

Two of the six PCA EAs will normally be dedicated to collecting *standard data*. The standard data produce a uniform database of time series measurements and X-ray spectra for all PCA observations. The remaining four PCA EAs will be configured to meet the needs of each particular observation, and thereby produce *selectable data*.

The observer should expect to receive the standard data obtained during the observation and should use this manual to understand the standard data contents and the options for choosing parallel configurations for the selectable data. The ASM data will not be provided to the observer but will be placed in the public domain in near real-time.

7.2.3 Telemetry Rates to FDS Buffer

The EDS can send data to the FDS at various rates, up to an aggregate total for all 8 EA's of ~500 kbps. (The FDS will accept 65 packets @ ≤1 kBy per second.) The rate from an individual EA is limited to ~300 kbps. In practice, the telemetry rate will generally be limited by other considerations, namely the ability of the spacecraft to store and telemeter the data, and the capacity of the ground-based systems to process and archive the data. Nominal sustained rates (or daily averages) will be of order 20 kb/s for the PCA, 5 kb/s for HEXTE and 3 kbs for the ASM. Much higher rates are possible for limited durations of time. It is the responsibility of the observer to configure EDS data modes such that the telemetry rates lie within the levels specified by the *science operations center* (SOC); see Chapter 8 on simulating data for the tools to use for estimating telemetry rates.

7.3 EDS Architecture

The EDS consists of three major types of subsystems. The types and total numbers of these subsystems in the entire EDS are (see Fig. 1):

- 6 PCA Event Analyzers,
- 2 ASM Event Analyzers,
- 2 System Managers, and
- 1 Power Distribution Board.

In the physical EDS, each of these systems is contained on a single multilayer printed circuit board. The EAs and system managers are grouped into two halves, EDS-A and EDS-B, as shown. This provides some degree of failure tolerance.

7.3.1 Event Analyzer

As noted above, the six identical PCA EAs operate on the data from the PCA, while two identical ASM EAs process the data from the ASM. Each PCA EA receives data from all 5 of the PCUs and each ASM EA receives data from all 3 SSCs.

Each EA contains (1) dedicated circuitry to select and classify event data, (2) a *digital signal processor* (DSP; TI 320C25) to operate on the selected data (e.g., to accumulate histograms), (3) a microprocessor (Harris 80C286) to carry out control functions, calculations, and transfers of data to the S/C, and finally (4) memory for storage of data and instructions. The PCA and ASM EA's are quite similar to each other but are not identical. All science data to be telemetered are sent to the S/C FDS directly (and independently) from the individual event analyzers.

If the interface limit of 65 packets per second is exceeded, telemetry packets will backlog in the EAs. In this situation, an EA generating less than 8 packets/s is guaranteed to transmit all its data. For EAs generating more than 8 packets/s, *complete partitions* (which typically consist of groups of packets) will be discarded if they cannot be telemetered before the EA back-end memory is needed for the next partition. This avoids the situation where every partition could be missing packets.

7.3.2 System Manager

The two system managers provide command and power services to the PCA and ASM event analyzers, communicate with the spacecraft via the 1773 data bus, and receive 28 VDC power and clock signals from the IPSDUs. The power is received via the power distribution board. Each system manager includes ground-preprogrammed read-only memory (ROM) in which is resident the flight software and preprogrammed EA configurations. The system manager also includes single-event-upset resistant read-write memory (Protected RAM = PRAM) which is used to hold additional uploaded flight software, additional EA configurations, and data indicating the proper configuration of the moment for the EAs. The PRAM can be loaded as needed during flight operations.

7.3.3 EA Front-End Back-End Concept

The architecture of an event analyzer (either for the PCA or the ASM version) can be described in terms of a *front end* and *back end* (Fig. 2). The Front End interfaces to the science instrument and is responsible for data collection and reduction. It operates as a slave to the Back End. The Back End is responsible for the overall management of the EA.

FIGURE 7.2. EDS: Front End – Back End functions

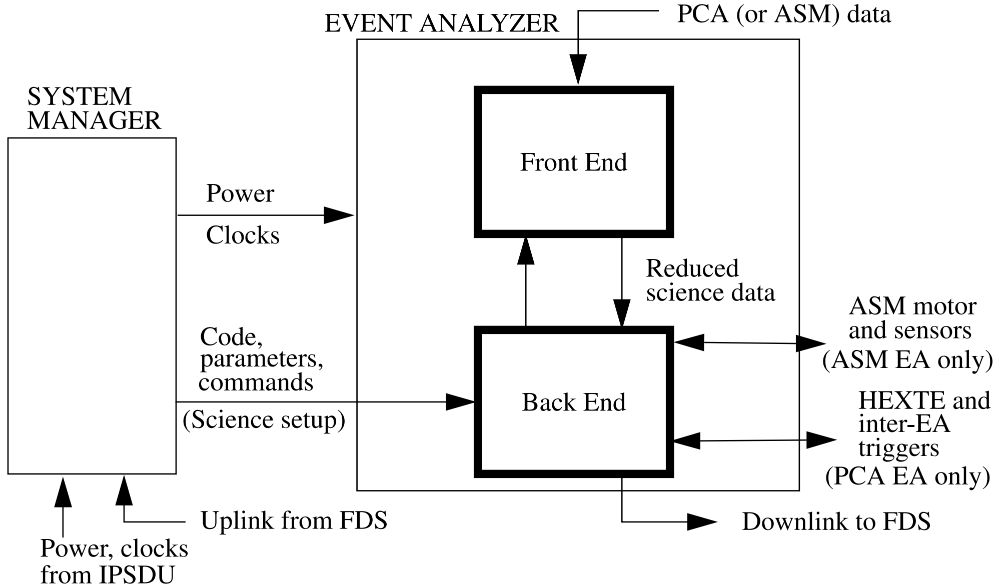

7.4 PCA Event-Analyzer Operations

EA operation is described in this section in sufficient detail to facilitate understanding of the individual data modes to be discussed below. An overview precedes a more detailed discussion. For purposes of illustration we usually assume that a PCA EA is being used in a binning mode.

7.4.1 PCA Event Analyzer Architecture

The internal architecture of the front and back ends of a PCA EA is shown in Fig. 3. The Front End comprises holding registers, event processing hardware (including timekeeping hardware and a "Giant Or", see below), typer and energy-recoder lookup tables, a FIFO (*first in first out*) register, the DSP, and a data buffer. The Back End comprises a microprocessor, memory, and support hardware.

FIGURE 7.3. EA Architecture

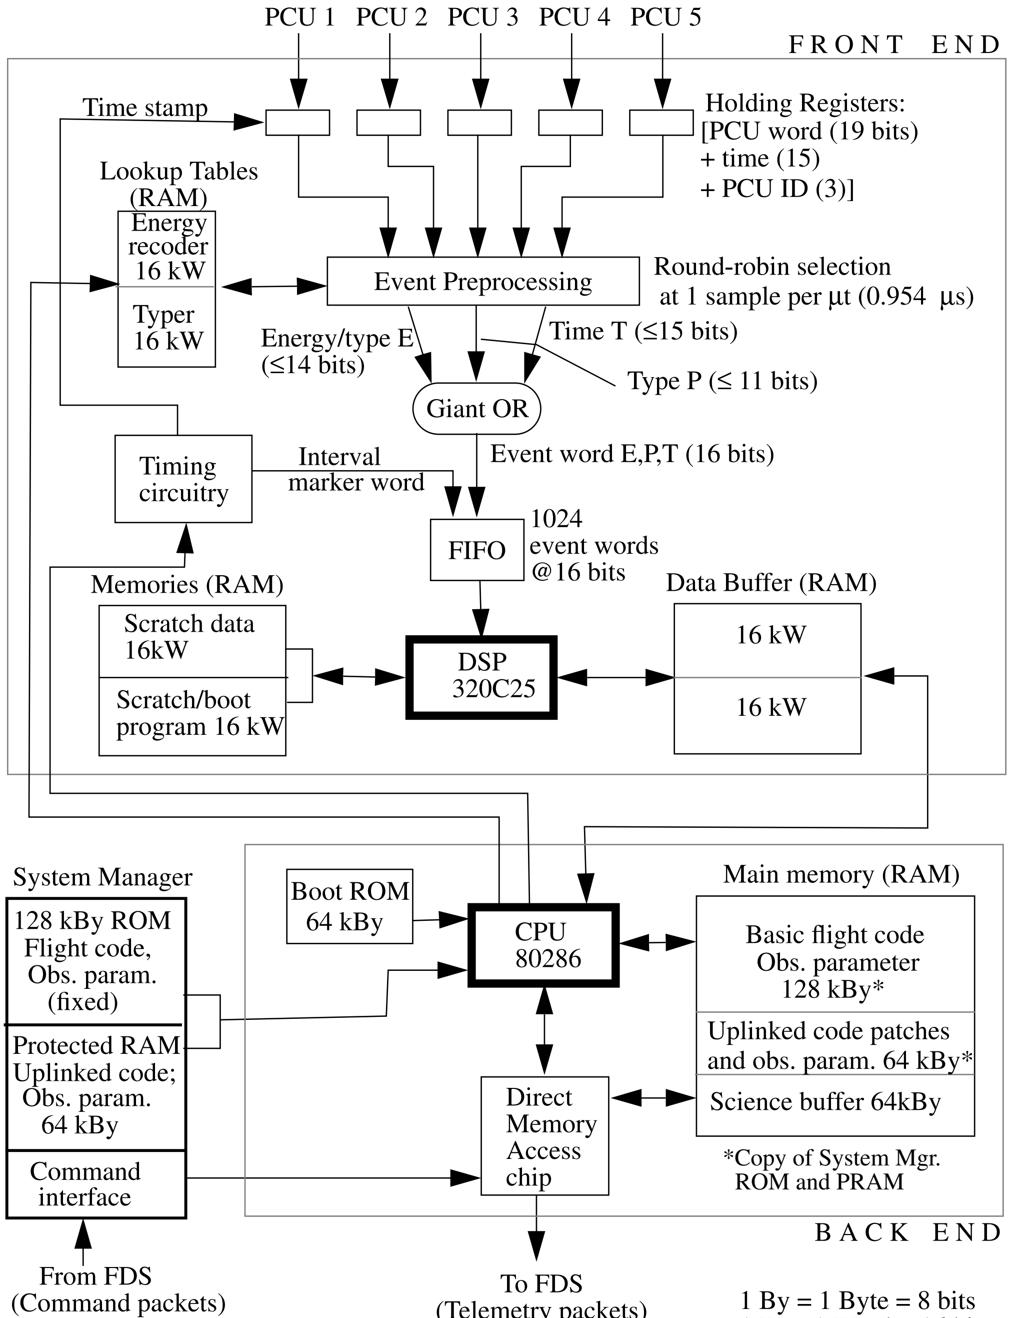

7.4.2 Overview of EA Operation.

Prior to the beginning of an observation, the desired configuration (mode and parameters) for each EA is specified as an address of the ROM or PRAM location where the configuration is located. If the configuration is not already in the ROM or PRAM, it is uploaded to the PRAM.

At the appropriate time, the system manager is commanded to send a reset command to the EA. The EA then reinitializes its own memory and hardware and then copies the contents of the system-manager ROM and PRAM into the EA back-end memory. This makes available to the CPU (in the Back End) the code, parameters, and commands required for the observation.

The Back End then prepares for the observation by loading *lookup tables* (LUT) and special registers in the Front End with values needed for recoding and classifying events, and it sets up the *timing circuitry*. It also provides software with appropriate parameters to the DSP. It then starts the accumulation of data at the beginning of the next even 16-s time boundary (lowest 4 bits of the spacecraft elapsed time in seconds are zero).

During the observation, as events arrive at the Front End, they are time stamped. Then, in the PCA EA only, each event is typed and recoded. For example, the 256 channel (8-bit) pulse height information may be reduced to 8 channels (3 bits). The recoded events (including time bits) are then sent to the DSP which builds sets of reduced data in the *data buffer* (DB). The structure of these sets, which may consist of histograms (binned data), lists of words describing individual events, Fourier power density spectra, etc. is described in the sections on the various EA modes. The DB in which the DSP collects the reduced data is typically divided into two sections of equal size.

The Back End defines time *Intervals*, which are of equal duration for the standard binning mode. During each of these Intervals, the DSP steadily deposits data into the active half of the data buffer (*half data buffer*, HDB). When the DSP has filled the HDB, the Front End starts to load the other half of the data buffer, and the Back End moves the just-completed data into its own memory and zeroes the locations of the just-read HDB. The data set just read from the HDB is known as a *partition.* Over the course of an observation, the DSP loads data alternately into the two halves of the DB, completing a partition with the passage of each Interval. This process allows data to be acquired and binned by the DSP without interruption.

Finally, the Back End converts the partition of data to *data packets* and forwards them with engineering data to the spacecraft FDS.

The next few sections of this chapter present the PCA EA operation in some more detail. We suggest that the reader review at least the next two sections, “PCA-EDS Interface” and “Time Signals and Counters” before going on to “Data Modes”.

7.4.3 PCA-EDS Interface

An X-ray photon produces an analog pulse in a PCU. Electronics located at the detector amplify this pulse and convert it to a digital pulse height. The event is also labeled with measurement chain identification and with flags giving the status of other measurement chains in the same detector for a total of 19 bits:

- 8 bits of pulse height (256 energy channels)
- 6 bits for xenon LLDs (each bit represents the state of the *lower-level discriminator* of one of the six principal xenon measurement chains: L1, R1, L2, R2, L3, R3
- VLE flag (*very large event*)
- Alpha flag (event detected in ^{241}Am calibration source layer)
- Propane flag (event from propane chamber)
- 2 bits for veto flag (veto-layer event; 3 pulse-height channels).

The PCU electronics transfers the data for each event directly to the EDS (via a 4 MHz serial link) where it is distributed to each of the 6 PCA EAs. This transfer proceeds asynchronously with sequential events (from one PCU) separated by at least 7 µs (start bit to start bit). Any event from a given PCU that occurs within ~10 µs of a previous event from that PCU is not transferred to the EDS by the PCU. This allows for detector recovery and insures that the EDS electronics has time to read every event from the PCA.

7.4.4 Time Signals and Counters

The time reference for the EDS is the spacecraft clock which outputs clock ticks on three lines with different intervals, each a binary subdivision of 1.0000 s.

- 2^{-20} s = 0.953674316406 µs = 1.0 *microtick* (µt)
- 2^{-5} s = 2^{15} µt = 31.2500 ms = 1.0 *minor tick*
- 1.00000000 s = 2^{20} µt = 1.0 *major tick*

The underlying clock is stable to about one part in 10^9. (See elsewhere in this manual.) The spacecraft also provides a count of the major ticks, which is known as *mission elapsed time* (MET). The logic of the EDS described herein makes use, almost exclusively, of the microticks and the three EA counters driven by them. All EDS data are time stamped relative to the MET reference provided by the S/C.

The microticks are counted in a 32-bit counter called an *Interval-tick generator* which can output pulses (*Interval ticks*) separated by a programmable *Interval duration* ranging from 2 to 2^{32} µt (1.9 µs to 4096 s). In practice the Interval duration will be greater than ~5 ms. The Interval duration can be programmed to change from one Interval to the next; the Back End accomplishes this by writing the desired time in microticks into an *Interval duration register* during the previous Interval. The Interval ticks are used to mark the times that the data buffer should be read out, to define the pulse period in the pulse folding mode, and as a reference for the time stamping of events. We distinguish this Interval with an initial capital. The spacecraft times of the Interval ticks can be derived from the telemetered data.

The first Interval of an observation always begins synchronously with a major tick which corresponds to a MET (mission elapsed time) value that is an integer multiple of 16 (seconds). This ensures synchronization of data collection between EAs and also between an EA and HEXTE, *if* the interval duration is less than 16 s and also a factor of 2^{-n} less than 16 s (n is an integer).

The microticks are counted in another 32-bit counter called a *bin-tick generator* which can output pulses (*bin ticks*) separated by a fixed programmable time ranging from 0.954 µs to 4096 s, in multiples of microticks. This time can not vary during an EA run. The bin ticks serve to define time bins into which data will be accumulated (in binned mode). The bin tick generator is reset at each Interval tick, so the first time bin of an Interval will commence synchronously with the beginning of an Interval. The time between bin ticks need not be a submultiple of the Interval though it is in most configurations. (Pulsar fold mode is the exception.)

Although the bin tick generator may be set to produce a tick at any multiple of 1 µt, in practice powers of two are generally preferred. This ensures that time bin boundaries between different EAs remain synchronous with each other and with HEXTE.

Finally, a *timebin counter* (15 bits) counts the bin ticks just defined. It too is reset with each Interval tick. This counter thus registers the number of bin ticks (or time bins) since the previous Interval tick. These 15 bits are used to time stamp individual events as they arrive from the PCUs. The time resolution of the time stamp is the time between adjacent bin-ticks, which is programmable as just described. A particular configuration may use fewer bits to conserve telemetry by limiting the count to, say, 16, so only 4 bits are used. This is accomplished by choosing a large interval between bin ticks (low time resolution) and/or by choosing frequent resets (with a small Interval). The entire 15 bits will constitute the time stamp, but the unused bits will be zeroes.

There is a *time bin mask register* that can mask out sets of 4-bits in the time stamp. This is also known as the ‘time-bin nibble mask’. If the value of the nibble mask is set to 1, the lowest 4 bits of the the time tag will be masked out (will read 0). If the value of the nibble mask is set to 2, the second 4 bits will be masked. If the value is set to 8 , the top 4 bits will be masked. The nibble mask is currently used only to reduce the number of time bits handled by a given EA when 1–µt time tagging is desired in the ‘*transparent-3 event mode*’. In this case, three EA’s are used together to make the EDS completely transparent to the PCU words (including PCU ID and time stamps). Each EA will have a different mask applied. Between all three EA’s, all the time bits are recovered.

7.4.5 Holding Registers and Time Stamp

Upon arrival at an EA, the 19-bit event word is latched in a register (one for each PCU). Here 3 bits of PCU ID and the 15-bit time stamp (i.e., the contents of the timebin counter) are immediately (within 1 µt) appended to the 19-bit event word received from the PCU. The time stamp specifies the bin that will be incremented by the DSP in binning modes (see below). In event modes, the time-stamp bits are eventually passed on to telemetry.

The 15-bit time stamp allows one to count, for example, 2^{15} time bins, e.g. 1-µt (2^{-20} s) time bins if the Interval duration is set to 2^{-5} s. The preprogrammed configuration specifies the number of bits actually used and the time bin width. The bin boundaries are known to 1 µt for any choice of bin widths.

The time stamp is applied about ~18 µs after the arrival of an X-ray photon at the PCU. (See PCA chapter.) This delay is due to processing time in the PCU electronics and to the 5-µs transfer time (19 bits at 4 MHz). The delay has a random +/– 250 ns component, so if 5 photons arrive at the PCA simultaneously, they will arrive at an EA with a spread of times. The jitter in the application of the time-stamp is 1 – 2 ns. If two different EAs are configured to examine the same events, an occasional individual event can acquire different time stamps in the two EAs because of this latter jitter.

The preprocessing hardware cycles through and reads the five registers at a rate of 1 µt per event. Since each PCU is inhibited from sending a second event for at least 7 µs, this scheme assures that no events will be lost in this stage of processing, even in the extreme case when five events, one from each PCU, all arrive at the EDS within the same microsecond.

Very closely-spaced events may appear in different orders in two EAs, even if they acquire the same time stamp. Thus if two EAs in event mode are telemetering different information about a single event (as in transparent event mode), one should place the PCU ID in the event word in each participating EA to insure that the several parts will be properly recombined. This labeling is sufficient because no one PCU can telemeter two events with less than 7 µs separation.

7.4.6 Event Classification (Lookup Tables)

As noted earlier, front-end lookup tables had been loaded by the Back End as part of the setup for the observation. These are the *typer* and *energy recoder* tables which are used to select and recode the events. Portions of the digital word from the holding register are used as addresses within the tables, and the table content at each address contains the desired remapped value of the event type or pulse height. This content is the *output* of the table lookup operation. The outputs of the two tables together with the time stamp constitute the information upon which the DSP will operate as it lists or bins the events. Considerable compression originates here. The lookup cycle takes less than 1 µt; the lookup process thus keeps up with the previous processing; again no events will be lost.

7.4.6.1 Typer table lookup

The typer table contains 2^{14} words at 16 bits each. Fourteen bits describing the event are used as an *input address* into the table. They are:

- 3 bits of PCU ID (counter number, 1 to 5)
- 6 bits for xenon LLDs
- VLE Flag
- Alpha Flag
- Propane Flag
- 2-bit Veto Flag

The PCU ID bits allow one to classify events by PCU or to select/reject events from particular PCUs.

The 16 bits of content at each address are allocated as follows:

- 1-bit Reject code (MSB = 1)
- 3 bits of ‘alternate recoder selection’ (to be fed to energy recoder table)
- 11 bits of event ‘type’
- 1-bit time-disable code

The most significant bit of the typer table output indicates an event that should be rejected if MSB = 1; this event is not processed further. The most common use of the typer is to reject all but good xenon X-ray events. Those with flags indicating penetrating particles (detection in more than one xenon measurement chain), veto layer events, very large events, propane events, or calibration events would be rejected.

The typer feeds 3 bits of its output (*alternate recoder selection*) to the energy recoder table. The utility of this is discussed in Section 4.6.2 “Energy-recoder table lookup” below.

The event type can mimic the 11 input bits of flag and LLD information, or it can be reduced to only a few (or no) bits of information. For example one might wish to select and output labels for three types of events: good xenon events, penetrating charged particles (two or more LLD flags), and calibration events. In this example, the output ‘type’would only require the use of 2 bits. The input PCU ID may be combined with other flags to conserve telemetry. For example, for “good” xenon events the 6 LLDs per detector could be combined with the PCU ID to produce an LLD number of 0 to 29, which requires 5 bits. Alternately the typer (or the energy recoder) can output the PCU ID verbatim (3 bits) with somewhat less efficiency.

Since the typer can only output 11 bits of type information to the Giant OR and since the placement of these bits within the Giant OR word is restricted, it is sometimes desirable for the energy recoder to feed the PCU ID to the Giant OR. (The PCU ID is input to both tables.)

The use of a small number of event types provides significant data compression and leaves room in the final event word for more pulse height and timing information. For example, to maximize time resolution, one could choose to simply sum the events from all xenon measurement chains together and to reject all particle and calibration events. In this case, there would only be one type of selected event, and the reject/select code alone would suffice as the output. All other bits of each word in the table, and hence of the output, would be zero.

The time disable bit in the output inhibits the time-tag bits in the next step of the processing. It is used in configurations where it is desired to use one EA to bin two different event types with two different time scales. This is currently used only in the Standard 1 configuration wherein a single EA yields 1/8 second integrations of X-ray events and also 128-s integrations of calibration-source spectra. Here, the time stamp resolution is 1/8 s, so the high-rate X-ray events are binned with this resolution, while the time to accumulate a partition of data is set to 128 s (the Interval duration). For the calibration events, the time tag is disabled and the 128-s Interval duration becomes the binning time.

7.4.6.2 Energy-recoder table lookup

The energy recoder table also contains 2^{14} words each of 16 bits. The input address is a 14-bit word containing:

- 8 pulse height bits (256 channels)
- 3 PCU ID bits (5 detectors)
- 3 alternate recoder selection bits (8 pulse-height mappings).

The recoded information (the output) is less than or equal to the 14 bits of the input data. Again the MSB (bit 15) is reserved for a reject code; rejected events are not processed further. The energy/type bits may be located at any of the other 15 locations. These can include any combination of the input bits, including the PCU ID and the alternate recoder bits. The content of the table at each address thus consists of:

- 1-bit rejection code (MSB)
- 14 bits of recoded energy
- 1 bit not used (set to zero), can be at any of position except MSB.

The recoded energy consists typically of the recoded pulse height, e.g., a 3-bit pulse-height (8 broad channels), but may also contain PCU ID and alternate recoder selection bits. The pulse-height remapping potential is quite general. Thus, for example, one can recode pulse heights so that the pulse-height bin width increases with pulse height.

Since the PCU ID bits are part of the energy recoder address, the table content can reflect pulse-height calibrations that differ from detector to detector. For example, the recoded pulse height can be adjusted channel by channel to approximate a common 8-bit energy scale. This is called *gain correction.* It is useful for configurations which bin the events from all five PCUs together. The gain correction values are set by the PCA group. By default, gain corrections are applied to all configurations except for the Standard and transparent configurations. Gain correction may be disabled for one or both of the EDS sides (A,B).

The alternate recoder selection bits (from the typer) are used as part of the address to distinguish different types of events (in one PCU) for which one may have different energy calibrations. It is used, for example, to provide different energy mappings for events from the xenon and propane chambers of the PCUs (see PCU description) in Standard mode 2.

7.4.7 Giant OR operation

The outputs of the typer and energy-recoder lookup tables and the time stamp are combined into a single 16-bit word by the *Giant OR* operator (see Figure 4). Fifteen bits of the 16-bit field are allocated to four subfields which receive the four inputs respectively. The typer contributes up to 11 bits of event type to the Giant OR within positions 3 – 13 (Fig. 4a). Similarly, the energy recoder contributes up to 14 bits of energy/type information within positions 0–14. The time stamp contains up to 15 bits which are in positions 0 – 14. The Interval constant contains up to 2 bits which are in positions 13–14. At the output of the Giant OR, a value of “1” for the most significant bit (MSB = bit 15) is asserted to indicate that this is a PCU event (and not an Interval marker word, see below).

FIGURE 7.4. Giant OR

(a) Available locations for science bits

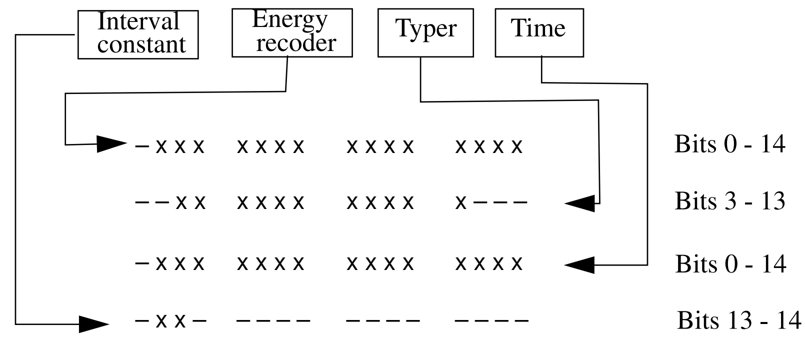

(b) Example - 256 pulse height resolution

0 0 E E   E E E E   E E 0 0   0 0 0 0   Energy recoder: 256 pulse height chan
0 0 0 0   0 0 0 0   0 0 P P   0 0 0 0   Typer: 4 types (4 layers; 3 Xe and 1 P
0 0 0 0   0 0 0 0   0 0 0 0   T T T T   Time: 16 time bins (per Interval)
F 0 0 0   0 0 0 0   0 0 0 0   0 0 0 0   Real event marker; set to one; F = 1
0 N 0 0   0 0 0 0   0 0 0 0   0 0 0 0   Interval constant: Indicates half of data buffer to be loaded

F N E E   E E E E   E E P P   T T T T   Result of OR operation.

(c) Example - 4 pulse height channels; high time resolution

0 0 E E   0 0 0 0   0 0 0 0   0 0 0 0   Energy recoder: 256 pulse height channels
0 0 0 0   P P P P   0 0 0 0   0 0 0 0   Typer: 15 types (3 layers; 5 PCU)
0 0 0 0   0 0 0 0   T T T T   T T T T   Time: 256 time bins (per Interval)
F 0 0 0   0 0 0 0   0 0 0 0   0 0 0 0   Real event marker; set to one; F = 1
0 N 0 0   0 0 0 0   0 0 0 0   0 0 0 0   Interval constant: Indicates half of data buffer to be loaded

F N E E   P P PT   T T T T   T T T T   Result of OR operation.

* Interval Marker Words (F = 0) do not pass through the Giant OR. They are inserted into the Giant Or output stream in proper time-ordered sequence. They have the form: bit 15 = 0, bits 3-14 = 1, bit 2 = Interval-interrupt Enable F bit 1 = Spillage Flag; bit 0 = 0.

This design assumes that the values for the four inputs are arranged (by proper loading of the lookup tables) so that each quantity exclusively uses only a portion of the subfield available to it (Fig. 4b,c); these portions are chosen to not overlap. The effective number of bits assigned to time, type, and pulse height are assigned according to the requirements of the observation. Time-stamp bits are always arranged to the least significant bit side, while the energy recoder bits are usually placed in the most significant side with the typer bits in the middle. The locations within the input words that are outside the restricted fields are filled with zeros. This allows the inputs to be combined into a single word with an inclusive OR operation.

Since the time-stamp bits indicate the time since the previous Interval tick in units of bin ticks, many of the higher order bits will always be zeroes. For example, if one is binning data with 1-s timebins with Interval markers every 64 s, the time stamp will have values from 0 to 63 (6 bits).

In the example of Fig. 4b, a large number (8) of pulse-height bits are output from the energy recoder thereby restricting the number of bits allocated to type and time. In Fig. 4c, a small number of pulse-height bits makes room for more information about type and also for improved time resolution (for the same Interval duration).

There is a fourth input to the Giant OR, an operational code called the *Interval constant*. It consists of 1 or 2 bits depending upon the data mode. It appears in bits 14 (or 13–14) in the output and is obtained from a register previously loaded by the Back End. Typically its count changes by unity with each new Interval. It is used by the DSP as part of an address to the DB (see Section 4.9 just below).

7.4.8 Interval Markers

At the occurrence of an Interval tick, an *Interval marker* is placed in proper time-order among the time-stamped events at the output of the Giant OR. This marker is indicated by a value of “0” in the MSB. The Interval marker event also contains status flags, specifically *spillage* and *interval-interrupt enable* flags (see next two sections, 4.9 and 4.10).

7.4.9 Accumulation of a Partition; DSP Operation

The 16-bit recoded events from the Giant OR and the Interval markers are fed into a 1024-word-deep *first-in-first-out* (FIFO) register to absorb brief bursts of high counting rate. The DSP reads the event words and Interval markers from the FIFO and modifies the *data buffer* (DB) in a manner dependent upon the data mode in use. The data modes are described in Section 5 below “EDS PCA Data Modes”.

In binned mode, for example, 14 bits (0–13) are allocated to energy/type/time. These bits of the event word are the address of the energy/type/time bin within the DB that should be incremented. The task of the DSP for an individual event is simply to increment that bin. This is straightforward and rapid.

The DB consists of two half data buffers (HDB), each of 16 kWords @16 bits (Fig. 3). For most modes, the accumulation takes place alternately in the two halves of the buffer. The data will accumulate in a given HDB until the next Interval marker is received by the DSP.

The switch to the next HDB requires no additional calculation by the DSP; it simply uses the interval constant (bit 14 in binned mode; see above) in the event word as part of the address to the DB, so the

DSP automatically operates on the appropriate HDB. The DSP immediately begins to accumulate data in the other HDB while the CPU processes the data of the just-completed partition. The DSP thus continues in its tight loop without interruption through the entire observation.

The DSP processes the events that emerge from the FIFO at rates which depend on the mode (e.g. binned, FFT, etc.). Overflow of the FIFO is indicated by a *spillage flag* (in the next Interval marker word). Upon its occurrence, the hardware ceases writing into the FIFO until the beginning of the following Interval.

Events are lost when under the following circumstances: (1) the FIFO overflows (this is mode dependent, but generally unlikely), (2) in binned modes, the allocated 4, 8, or 16 bits per word in the DB overflow, (3) high rates in event modes fill the DB before readouts, (4) the FFT mode is in a processing interval, and (5) the EDS–FDS interface TM rate is saturated (500 kb/s).

7.4.10 Compression and Transmission

When the DSP reads the Interval marker from the FIFO with the *interval-interrupt enable flag* set, an interrupt signal is sent to the Back End which then performs mode-dependent functions. In binned mode, for example, it reads the just completed partition of data into its own memory and zeroes the memory locations just read. The back-end CPU can then operate on the data, e.g., to search for a flare or burst.

A *partition map*, unique to each configuration, is used to specify which bits of each 16-bit word are to be telemetered. The map can specify the lowest 4, 8, or 16 bits of any given word on a word-by-word basis. Thus a configuration can specify different word sizes for different energy channels, e.g., fewer bits at higher X-ray energies. The partition map can also indicate that entire words are not to be telemetered.

If the count in a DB word exceeds the maximum value that can be telemetered according to the partition map (i.e., 2^4 - 1, 2^8 - 1, or 2^{16} - 1), the value appears to have ‘rolled over’ in the telemetered data. No notice of such rolled over values is given to the observer.

The data are then formed into data packets and transferred to the spacecraft FDS.

7.4.11 Back-end (CPU) Operations

There are ~650 configurations in the system manager ROM, and others will be uploaded to system-manager PRAM. The PRAM will also contain the desired configuration for the upcoming EA run.

Upon an EA reset, the CPU memory is reinitialized. The CPU then copies the contents of both the ROM and PRAM into its own memory. It then sends to the spacecraft an *observation parameter dump* for transfer to the ground. This consists of all parameters that specify the configuration, including those for the two lookup tables and the partitiom map that will be used in the upcoming EA run. The CPU then loads the lookup tables and also the code needed by the DSP.

The Back End (with its CPU) controls the Front End. It starts and stops observations. During the observation, it controls interval generation (interval constant and interval duration) by loading the appropriate front-end registers during the preceding interval, it retrieves and compresses the data in the DB, and it creates data packets and transmits them to the FDS. It also can carry out simple analyses of the data it has retrieved from the DB. The CPU prefixes the data with a header (type of data, start time of integration, time interval, etc.).

During data acquisition, an interrupt to the CPU (see above) informs the CPU that it can (1) calculate and load into front-end registers the Interval duration and interval constant for the interval following that just started, (2) read out the partition from the DB or HDB, and (3) begin to process the just-read partition of data.

The CPU places a *partition number* in the header of the partition, as noted. This is referenced to the time of the beginning of the observation (or EA run). Since the time between partitions is well known (i.e., it was specified in the configuration of the observation), the time of the first data in a partition may be calculated. In modes where a single partition contains multiple Interval markers embedded in a stream of events (e.g. event encoded mode), the ground analysis software can count Interval markers within the partition. Since the Interval duration is known, absolute times (MET) can be calculated.

At the end of an EA run, the CPU checks the tables (typer, energy recoder, and partition map) for single-event upsets. The number and locations of errors are telemetered.

The Back End also receives and processes commands uplinked from the ground.

7.4.12 System Managers

7.4.12.1 EA memory repositories

The system manager ROM and Protected RAM (PRAM) is made available to all of the event analyzers in that half of the EDS. Upon reset (or power up) the EAs read first the ROM image, which contains the basic flight software code, and then the PRAM image. The EA can never write to the PRAM and it is, therefore, immune to corruption from software faults; it is also single-event-upset (SEU) immune.

7.4.12.2 Reboots

If an EA detects any error, it immediately does a reboot which initializes its memory and it again copies the ROM and PRAM images from the system manager. The reboot takes about 20 s.

As protection against buildup of undetected single-event upsets, each EA will be rebooted, by command, approximately once per orbit.

7.5 EDS PCA Data Modes

The individual event analyzers for the PCA can be operated in 7 modes, namely:

- Binned Data: Histograms according to pulse height, time, and event type; useful for time series and spectra with selectable resolutions.
- Pulsar Fold: Binned mode with pulse phase substituted for time; for pulse phase spectroscopy of short-period pulsars.
- Delta Binned: Histogram of time differences between events (one type of event only)
- Event Encoded: Pulse-height, time, and event type for each event transmitted as a 15-bit word. For rapid timing on less bright sources. There are two submodes:
    - Continuous: Data buffer halves read out alternately; continuous accumulation of data for modest rates.
    - Snapshot: Entire data buffer loaded and then read out; data accumulation interrupted during readout (i.e., some data lost).
- Burst Catcher:Two EA's used, one to search data and generate a trigger ('trigger EA') and one to capture data at high time resolution ('catcher EA'). Useful for brief unpredictable X-ray burst-like phenomena. There are three trigger modes, and two catcher submodes:
    - Triggers:Selectable trigger criteria (absolute rate, delta rise, and hardness)
    - Binned:Binned data with 25%–50% of data before trigger.
    - Event:Event mode with 32 kWord buffer used for capture. All burst-catcher ROM configurations use this submode.
- Single-Bit Code: String of ones and zeros representing events and clock ticks, respectively. For fast timing without danger of bin overflow. Highly efficient when time resolution is comparable to interval between events.
- Fast Fourier Transform (FFT): Accumulates simultaneously two channels of 256 time bins (e.g., two energy bands). Averages power density spectra and cross spectrum. For low telemetry rate surveillance of rapid time variability.

7.5.1 Binned Mode

The *binned mode* builds histograms consisting of counts of events with the same time-bin number, recoded pulse height and event type. One configuration could include separate histograms for "good" events, background events, propane layer events, and calibration events. The time bin duration is a selectable value in the range 1 μt (0.954 μs) to 2^{32} μt (4096 s). [1.0 microtick = 1.0 μt = 1.000×2^{-20} s]

The DSP executes this mode as follows. It uses the least significant 14 bits (0–13) of the 16-bit event word from the FIFO as an address within the *half data buffer* (HDB) currently accumulating data. Bit 14 (interval constant) defines which HDB is being addressed. The HDB contains 2^{14} words (16 kW), i.e. it can hold 2^{14} time/energy/type bins. The value stored at that address is then incremented by unity. The energy/type/time bins thus fill sequentially (in time). The data values “roll-over” to zero when incremented beyond their maximum (4, 8, or 16 bits according to the configuration). No flag is set to indicate a rollover.

The DSP also interprets an Interval marker as an address. Since the MSB is zero and since, in this mode, all the Interval marker words are identical, the address is a single location in the scratch RAM of the DSP, i.e. not in the DB. Thus this non-event word is conveniently trashed with no additional logic required of the DSP. However, when this word arrives at the DSP, the combination of MSB = 0 and the *interval-interrupt enable* bit asserted results in an interrupt to the Back End. This causes the Back End to read out the just-completed partition of data from the appropriate HDB, and to reinitialize that HDB with zeroes.

Configurations for the binned mode can be constructed that emphasize time resolution or, alternatively, pulse-height resolution. Yet another choice is to emphasize the ‘type’ of event. A wide range of intermediate choices are also available.

The selectable parameters for the *binned mode* are:

- Energy/type selections
- PCU selections (Any of 1,2,3,4,5)
- Time bin duration (0.954 μs to 4096 s) (Interval durations and hence timebins exceeding 16s can lead to loss of synchronization of data transfers in different EAs; see Section 4.4 “Time Signals and Counters”.)
- Word size per bin (4, 8, or 16 bits); can vary with X-ray energy

7.5.2 Pulsar Fold Mode

The *pulsar fold mode* is essentially the same as binned data mode, except that the time bins are replaced with phase bins. The data can be binned according to pulse height and event type as well as phase, which allows one to do ‘pulse-phase spectroscopy’.

This mode supports pulse periods in the range ~3 ms to 4096 seconds. Observers desiring lesser periods may fold on a multiple of the desired period, preferably at least 5 ms. (Observers wishing to use fold periods less than 5 ms are advised to consult with the operations center to obtain the latest performance characteristics of the EDS.) The pulse period may be specified with a precision of 2^{-16} μt (~1.5 $\times 10^{-11}$ s) to prevent a buildup of phase error during the accumulation of data.

In this mode, the Interval ticks coincide, to an accuracy of 1 μt, with the times of zero pulse phase. The time between Interval ticks, i.e., the Interval duration, is an integer number of microticks (0.954 μs); this number of microticks is modulated up and down by unity as needed from pulse cycle to pulse cycle to approximate the desired period and to keep the Interval ticks close to zero phase. If the observer accurately specifies the pulse period with the full available precision, the pulse folding could proceed for ~30,000 periods with a phase error that is less than 1 μt . We expect that, in practice, pulse profiles will usually be averaged over shorter times.

The duration of a phase bin is defined by the bin-tick duration (also in integer microticks). The first phase bin begins at the Interval tick. The last phase bin will be shorter than the other phase bins unless the bin width happens to be commensurate with the pulse period. It is the length of the last phase bin that is varied by the aforementioned modulation of the pulse period. Thus its average width will not be an integer number of microticks.

The number of bins to be accumulated is given by the product of the numbers of phase bins, pulse-height bins, and event types. Each of these numbers is selectable, but the total is limited to a maximum of 16384 by the size of the HDB. Again, the histogram values may be telemetered as 4, 8 or 16 bit words.

In pulsar fold mode, the Interval markers occur at the end of each pulsar period and do not initiate a readout, until the end of the accumulation period after *N* Intervals. This does not require any change in the DSP code from binned mode.

The selectable parameters for the *pulsar fold mode* are:

- Energy/type selections
- PCU selections (Any of 1,2,3,4,5)
- Pulsar period (5 ms to 4096 s)
- Number of phase bins (2 to 16,384)
- Word size per bin (4, 8, or 16 bits)
- Data accumulation time (no. of pulse periods)

7.5.3 Delta Binned Mode

The EDS is also capable of running a variation of binned mode, called *delta binned mode*, in which a histogram of the time between events is generated for a single selectable energy range and a single selectable type.

The desired events are selected by the Front End; all other events are rejected. The time resolution is set by the period of the bin-tick generator, just as for all other modes. Upon receipt of an event, the DSP subtracts the time of the previous event from the time of the current event. The difference is used as an address to the data buffer, and that location is incremented by unity. There are 2^{14} = 16,384 bins in the HDB; this is the maximum number of time-difference bins. For most observations, fewer than 1000 need be telemetered. Thus, if the time resolution is 1 µt, the maximum tabulated inter-event time would be 16.384 ms, or ~ 1ms if only 1000 bins are used. If the time resolution is 16 µt, the maximum would be ~16 ms for 1024 bins, and 0.26 s for the full memory. The latter is the average time between events for 2–10 keV background from a single PCU. Thus the practical time resolutions range from 1 to ~16 µt.

If an inter-event time exceeds 16,384 times the time resolution, the time difference rolls over, and the bin that is incremented corresponds to the inter-event time modulo 16,384 times the time resolution .

The accumulation and readouts occur just as in normal binned mode. The histogram values may be telemetered as 4, 8 or 16 bit words, and the Interval marker signals the switching of the two HDBs. The bins roll over if filled; no flag is set.

One possible configuration of this mode is to provide 1 s integrations of histograms containing 1024 bins with 1 µt time differences. Time differences of 1 µt to ~1 ms are measured. With this, one can explore the dead time of an individual PCU.

The selectable parameters for the *delta time bin mode* are:

- Energy/type selection (one energy/type only)
- PCU selections (1 group consisting of any of 1,2,3,4,5)
- Time resolution (0.954 µs to 4096 s; ~16 µs is the practical maximum).
- No. of bins to telemeter; max is 16,384.
- Energy/type/time selection
- Word size per bin (4, 8, or 16 bits); can vary with delay time
- Readout time (Interval duration)

7.5.4 Event Encoded Mode

The *event encoded* (EE) mode provides data in which each event is individually described. The DSP stores the 16-bit recoded event word directly and sequentially into the active HDB (16 kWords) or the full DB (32 kWords). Thus, the time of each event is obtained together with some combination of detector ID, pulse height, and event type according to the details of the configuration. All of this information can occupy no more than 15 bits of the recoded 16-bit event word. The most significant bit (MSB) is used to distinguish an event from an Interval marker.

The Interval markers are used for two purposes in this mode. First, the markers are recorded in the buffer at their proper sequential position to serve as time references. (These are labeled with MSB = 0.) Thus the time stamp of any event in the buffer is a count of bin ticks since the previous Interval marker word in the stored sequence of events. For example, if Interval markers occur every 2^{-7} s = 7.8125 ms or, equivalently $2^{13} = 8192$ µt, a time stamp of 13 bits would suffice to subdivide the Interval into 1 µt resolution elements: $2^{-13} \times 7.812$ ms = 1 µt = 0.954 µs.

Secondly, the Interval markers are used to trigger the readout of the HDB by the BackEnd at the end of an *accumulation interval*. This occurs when the *N*th Interval marker is counted. The number ***N*** is a parameter of the event encoded mode configuration. The accumulation interval can be set to values between 1/8 s to 4096 s. (These are the practical limits; the technically allowed range is much greater.) After the readout, the Back End loads the just-read buffer with *end-of-data marks*. Thus, at the subsequent readout, the Back End will retrieve data from only the locations containing real events.

The data must be sent in 16-bit words for proper operation of the mode. Use of the partition map to reduce the bits per word eliminates the higher-order bits which contain the indicator that the word is a true event or an interval marker (MSB).

There are two basic methods for reading out the data, *continuous mode* and *snapshot mode*. In the former, the two halves of the data buffer are sequentially filled as in the binned mode. In the latter, the entire buffer is used for each accumulation. These are discussed below.

This mode can consume large amounts of telemetry for bright sources. For example, the Crab nebula at 10^4 cts/s and 16 bits per event yields a telemetry rate of 160,000 kb/s, or 8 times the nominal sustained PCA telemetry rate quoted earlier. Such an observation could be carried out only for a short duration (~40 minutes, once per day).

The selectable parameters for the *event encoded mode* are:

- Energy/type selections
- PCU selections (Any of 1,2,3,4,5)
- Time resolution (0.954 µs – 4096 s)
- Time-marker interval (Interval marker; > 4 ms)
- Accumulation interval (1/8 s – 4096 s; practical limits)
- Continuous or snapshot modes (see below)

7.5.4.1 Continuous EE mode

In the *continuous EE mode*, the two halves of the data buffer are alternately read out, as in the binned mode. Each readout occurs at the end of an accumulation interval (*N* Interval markers).

The HDB can store up to $2^{14} - 1 = 16383$ events and time markers @16 bits (16 kWords). If the rate of events is sufficiently low and/or the frequency of readouts is sufficiently high, the accumulation is ‘continuous’; no events are lost. On the other hand, if the event rate is too high, the HDBs will fill before the accumulation interval ends. In this case, one obtains regularly spaced samples of 16 k events with temporal gaps between them. The lost events are counted and the count is telemetered. In this high-rate situation, the continuous mode becomes discontinuous! Generally, one would prefer to make the accumulation interval sufficiently short so that the events are not likely to fill the HDB.

For our example of the Crab nebula at 10 kct/s, an accumulation interval of 1 s would clear the 16-k HDB when it is about 2/3 filled. But if the counting rate should exceed 16 kct/s, gaps would appear in the data at intervals of 1 s.

As noted above, the Back End telemeters only recorded events (and not the unfilled part of the buffer). Thus the telemetry rate will be directly tied to the actual event rate plus overhead. If the source suddenly becomes very strong, an unexpectedly high telemetry rate would result. The data produced could exceed the capacity of the spacecraft solid-state recorder.

A different strategy would be to limit the telemetry by specifying a rather long accumulation interval. In this case, the event rate would fill the half data buffer sometime before the readout occurred, and events would be lost. This is the ‘discontinuous’ high-rate situation described above. In this case, the telemetry rate is limited by the size of the data buffer and the number of times it is read out per second.

Most observational needs should be met with the configurations in ROM, each of which is based on an accumulation interval of 1 s or 8 s.

7.5.4.2 Snapshot mode

If the event rate is so high that one cannot telemeter all events, the snapshot mode is an alternative to the ‘continuous’ mode. The snapshot mode can yield a longer train of contiguous data (32 k events and time markers) than the 16 k events provided by the continuous mode.

In this case, the entire (both halves) 32 kWord DB is loaded sequentially until the accumulation interval ends. Usually, the accumulation interval will be set sufficiently long to allow the DB to fill completely with 32 k events and time markers. In either case, at the end of the accumulation interval, the Back End reads out the data buffer and loads it with end-of-data marks. Accumulation begins again after a *dummy interval* sufficient to allow these operations (~70 ms). The temporal gap could be significantly longer than the dummy interval if the DB is filled well before the end of the accumulation interval.

7.5.4.3 Event-encoded configurations

The following three EE configurations, *mixed, time,* and *transparent*, illustrate the capability of the EE mode. Such configurations may be used in either continuous or snapshot modes. For more information see the Configurations chapter.

- In the *mixed EE configuration*s, an event is specified by a mixture of time and type/energy bits. For example, if 8 time bits are used, up to 7 bits may be selected in any combination from the bits available to the typer or energy recoder. If time markers are specified to occur every 2^{-7} s = 8 ms, the time bits are sufficient to specify the time of the event with an accuracy of 2^{-15} s (~31 µs). This configuration is useful for the less frequent high-energy events in high-rate sources. It will typically have a low-energy cutoff, and events below that energy will be telemetered in binned mode.
- In the *timed EE configuration*, an event is specified by time bits only. The 15 time bits are sufficient to specify the time of the event with an accuracy of 2^{-20} s (1 µt). In this case. the Interval markers would be programmed to occur at 2^{-5} s (31.250 ms).
- In the *transparent EE configuration,*, the EDS telemeters *all* information about every selected event. This includes all information received from the PCU, the detector ID and all the time bits. The time bits are sufficient to specify the time of the event with an accuracy of 2^{-20} s (1 µt = 0.9537 µs). This format uses the resources of three EAs, each of which, for example, processes only a portion of the time stamp. This mode is useful for system diagnostics.

7.5.5 Burst Trigger and Catcher Modes

7.5.5.1 Trigger and catcher event analyzers

The objective of *burst catcher mode* is to capture a brief segment of data during an X-ray burst or similar phenomenon with a high time resolution that could not be sustained continuously by the telemetry system. This mode normally operates with at least one *trigger EA* and at least one *catcher EA*. The trigger EA detects X-ray burst-like phenomena and sets a *burst flag* upon each such occurrence. In the absence of a burst, the catcher EA continuously writes data into the DB with old data being overwritten; the data is not telemetered to the ground. When the catcher EA receives a burst flag signal, its Back End transfers the data currently in the DB into its own memory. The timing is such that both pre and post '*burst-onset*' data can be captured. These data may then be sent to the ground.

7.5.5.2 Burst Trigger Mode

The trigger EA runs a variation of binned mode in which the Back End analyzes each partition of data to see if a programmed counting rate or hardness ratio threshold has been exceeded. The trigger EA configuration specifies that one of the following three algorithms is to be used in carrying out this analysis:

- Absolute rate (or level) Trigger: Does the number of counts in any bin from a single energy range exceed the threshold?
- Delta rise (or edge) Trigger: Does the number of counts in any bin from a single energy range exceed the mean of the N bins preceding it by the threshold value?
- Hardness Trigger: Does the ratio of the count in any bin from one energy range to that of the corresponding bin for another energy range exceed a threshold? The calculation ignores any pair of bins for which the denominator is zero.

If these algorithms are not sufficient, the EDS has the capability to receive new trigger algorithms via upload after launch.

The trigger may operate in either *absolute* or *most interestin*g mode, and thereby generate, respectively, either *high-priority burst flag*s or *low-priority burst fla*gs. In absolute mode, the trigger threshold is maintained at the value specified by the configuration. In most interesting mode, the trigger threshold is raised with each burst to the level of that burst. The flag priority determines whether and when the burst is telemetered (see below).

The trigger EA may optionally telemeter its accumulated binned data to the ground. The partition header packets for these data are always telemetered; these indicate whether or not the trigger threshold was exceeded. Even if the partition data is not telemetered, each partition header requires one packet of telemetry. At a readout time of 32 ms, one uses 1/2 the allowable packet count for the EDS-FDS interface. A readout time of at least 250 ms is a more conservative and practical value.

The trigger EA normally will bin with relatively coarse time resolution to obtain sufficient statistics per time bin and to minimize the number of bins the Back End must examine during each Interval. Since the trigger EA analyzes a partition of data after the partition is transferred to the back-end memory, one may choose a relatively short readout time to minimize the time delay between the occurrence of a burst and the generation of a burst flag.

The selectable parameters for burst trigger mode (*trigger EA)* are:

- Energy/type selections
- PCU selections (Any of 1,2,3,4,5)
- Time bin duration (0.954 µs to 4096 s)
- Word size per bin (4, 8, or 16 bits)
- Trigger algorithm
- Absolute threshold or most-interesting mode
- Telemeter header data only
- Readout time (32 ms to 4096 s; 32 ms could saturate the EDS–FDS interface)

- Send trigger to HEXTE

7.5.5.3 HEXTE Burst Flags

A burst flag may also be received from HEXTE, and this is treated as a high-priority burst flag (see below). Catcher EAs have the option of ignoring the HEXTE burst flag. EDS high-priority burst flags may also be sent to HEXTE.

7.5.5.4 Catcher EA Priorities and Modes

When a burst flag is received by an EA running the burst catcher mode, data from both before and after the time of the trigger is captured. If the flag is *high priority*, the captured burst is immediately sent to the S/C to be telemetered. If the flag is *low priority*, the burst is stored in the back-end memory, and it will be overwritten by subsequent bursts. At the end of an observation (EA run), the last low priority burst detected is sent to the FDS. This burst is the one that most exceeded the original trigger threshold. If additional EAs and/or HEXTE are generating triggers, a high-priority burst from one of them could overwrite a low-priority burst awaiting the end of the EA run.

The DSP program code supports two catcher modes: a *binned catcher mode* or an *event catcher mode*. These are fundamentally similar to the binned and event encoded modes described above, except for the extra provisions to sense the trigger and capture the burst data.

The event catcher mode should be useful for short events (< few seconds) at Crab-like intensities. This is limited by the 32,767 events that can be stored in the data buffer. The normal (non-burst-catcher) modes should prove adequate in most cases of longer bright events, e.g. classic X-ray bursts. Binned catcher mode should prove useful for sustained events. All burst-catcher configurations in the ROM utilize event-catcher mode. Binned catcher modes may be requested by the observer; many such modes should be in PRAM prior to launch.

7.5.5.5 Binned burst-catcher mode

In binned catcher mode, the Interval ticks of the catcher EA must be synchronized to the Interval ticks of the trigger EA for the mode to operate as designed, and for the description below to be valid. This condition ensures that 25% to 50% of the ‘captured’ data precedes the ‘burst onset’ time, i. e. the time of the rate or spectral change that caused the trigger flag to be set.

In this mode, the entire DB is logically divided into four equal parts rather than the usual two. The four quarters fill cyclically. The switch from one quarter to another occurs at Interval ticks via changes of the Interval constant (bits 13 - 14 in the event word). As each quarter is about to be filled, the Back End loads it with zeros.

If the trigger threshold is exceeded during Interval N, the trigger Back End will discover this fact during Interval N+1, and the burst flag will be set. During Interval N+2, the flag causes the catcher Back

End to read into its own memory the data for Interval N–1 (the oldest data in the DB). At the same time, the catcher Front End is accumulating N+2 data in the appropriate quarter buffer. During accumulation of N+3 data, the N data are read out. Similarly the post-onset data N+1 and N+2 are read out while accumulation continues. The readouts then cease. One thus obtains a train of data consisting of four quarters with the burst onset in the 2nd of the four quarters.

If the high-priority trigger flag continues to be set because the trigger criterion continues to be satisfied, the readout of the quarter buffers continues until the two quarters following the last trigger are read out. This continuous stream of high-time resolution data could well saturate the telemetry interface or the FDS solid-state recorder.

Four quarter buffers of data yield a total of 2^{15}= 32,768 bins for accumulation according to event type, pulse height and time.

The Interval-tick duration for an X-ray burst might typically be set to 1 or 4 seconds, giving a data train of 4 or 16 s. As usual, one will specify the word size as 4, 8, or 16 bits to minimize telemetry while avoiding overflows.

The selectable parameters for *binned burst-catcher mode* (*catcher EA*) are:

- Energy/type selections
- PCU selections (Any of 1,2,3,4,5)
- Time bin duration (0.954 µs to 4096 s)
- Word size per bin (4, 8, or 16 bits); can vary with X-ray energy
- Trigger selection (any of high priority, low priority, HEXTE)
- Interval duration.

7.5.5.6 Event burst-catcher mode

In event burst-catcher mode, the catcher EA operates similarly to the event encoded mode (including mixed, timed, and transparent formats) described above. It makes use of the entire data buffer; it is capable of capturing $2^{15} - 1$ = 32,767 contiguous events and/or Interval markers per burst, each consisting of 16 bits (15 bits of science). The entire data buffer is used as a cyclic buffer with the oldest data being continuously overwritten until a trigger is received. A floating end-of-data mark indicates the boundary between the oldest and newest events. The Interval markers interspersed in the data as time references need not be synchronized with the Interval ticks of the trigger EA.

There is a configurable *capture delay* (an integer number of time-marker intervals) measured from the reception of the trigger flag which allows the capture of post-onset data. At the end of the delay, the accumulation of data in the DB is stopped. The entire 32 kWords of the catcher DB is then read out. After a transfer interval, the accumulation resumes.

In this mode, an unexpectedly bright burst could fill the buffer and begin to overwrite itself before the trigger (with delay) stops the accumulation. If the trigger flag remains on, the CPU will continue to grab and process contiguous groups of 32 k events with gaps for the readouts. This also could saturate the telemetry system.

The selectable parameters for *event burst-catcher mode* (*catcher EA*) are:

- Energy/type selections
- PCU selections (Any of 1,2,3,4,5)
- Time resolution (0.954 μs – 4096 s)
- Word size per bin (4, 8, or 16 bits)
- Time marker intervals (> 4 ms)
- Capture delay (no. of time-marker intervals)
- Trigger selection (any of high priority, low priority, HEXTE)

7.5.6 Single-Bit Mode

In *single-bit mode*, the reduced data consist of a string of zeroes and ones; a ‘zero’ denotes the occurrence of a regularly-spaced time-bin boundary while a ‘one’ denotes the occurrence of a detector event. This mode provides excellent time resolution with minimal telemetry without the danger of overflow of time bins. However, it generates telemetry data at a source intensity dependent rate. The time bin size is selectable between 2^{-20} s (1 μt) and 4,096 s. The single-bit mode can only be carried out on a single choice of event/type, including a choice of a set of PCU detectors.

This mode is optimal when the bin size is chosen to be comparable to the reciprocal of the event rate. For example, 1 ms time bins are optimum for an event rate of 1000 c/s. If the rate is substantially lower, an excessive number of boundaries (with no events) are telemetered, and event mode might be more efficient. At a substantially higher rate than optimal, a large number of event bits per boundary could be more efficiently telemetered in binned mode.

The DSP uses 16-bit arithmetic, so it can accumulate up to $2^{16} - 16$ ‘ones’ and ‘zeros’ in a given partition. i.e. 65520 events and time boundaries. These bits are loaded sequentially into the active half of the DB; the first 16 events/boundaries will occupy the first 16-bit word and so forth. The 65520 events/boundaries occupy only about 1/4 of the 16 kW in the HDB.

Upon the occurrence of the Interval tick, the accumulation begins in the other HDB. If the Interval tick occurs before all 65520 locations are filled, no data are lost. If, on the other hand, the maximum number is reached before the Interval tick occurs, the mode will keep a count of the lost events.

The SOC will convert these data to a standard binned-mode format before delivery to observers.

The selectable parameters in *single bit mode* are:

- Energy/type selection (one only)
- PCU selections (Any of 1,2,3,4,5)
- Time bin duration (0.954 µs to 4096 s)
- Interval duration (Readout interval of HDB; 5 ms to 4096 s)

7.5.7 FFT Mode

In FFT mode, Fourier power density and cross spectra are computed and summed. To accomplish this, the EA first simultaneously generates two time series histograms, each of which consists of counts in 256 sequential time bins. The time series are accumulated and stored using 16-bit words. The events in each time series are those selected according to observer-specified criteria; the two time series will generally differ in their energy ranges. The time during which events are accumulated to form the two time-series is called an *accumulation interval*. Its length is 256 times the time bin size but cannot be less than 4 ms.

Following the accumulation interval, the DSP calculates the mean number of counts per bin for each time series and subtracts the mean from the count in each time bin. The two means are saved. The DSP then computes the power density spectra (PDS, the sum of the squares of the real and complex components of the Fourier transform) of the two mean-subtracted time series and the cross spectrum (CS, the complex product of 2 transforms, $A \times B^*$). The PDS and CS are based upon 256-point decimation-in-time complex FFTs. The individual transforms are calculated using 16-bit integer arithmetic. These computations take place during a *processing interval* of duration 12 ms.

In order to avoid roundoff error due to the use of 16-bit integer arithmetic, a number of scaling operations are performed in FFT mode. First, the two time series are each scaled by a configurable scaling factor. Second, after each of the eight stages of calculation of the FFTs, the intermediate data values are divided by 2. Third, prior to summing, the two PDS and the CS are scaled by a configurable value. Details are in the “FFT Mode Module Specification” (not in this package).

The collection of events and computation of the PDS and CS is repeated a selectable number of times, e.g. a few hundred. As each PDS/CS is completed, the results are summed to the previous total. For example, the second set of spectra (2 PDS and 1 CS) is added, point by point, to the first set. The third set is added to the previous total, and so forth until the configured number of spectra have been summed. The summed spectra constitute, in effect, averages of each spectrum. The PDS and CS are accumulated into 32 bit numbers. The sum of all the accumulation and processing intervals prior to telemetry of the summations is called the *product interval*.

During a *transfer interval*, which immediately follows each product interval, the three summed spectra and the means of the individual data sets are copied into back-end memory, packed and sent to the

FDS as telemetry packets. After the data have been copied, the data buffer is reinitialized with zeroes and the process begins anew.

The FFT mode has a selectable time bin size between 2^{-20} s (0.954 µs) and 4096 s.

Dead time is inherent in the operation of FFT Mode. Events received during Processing or Transfer Intervals are not processed. The percentage of live time (i.e., percentage of time accumulating photons into the time series) of the FFT mode is a function of the time bin size, as follows:

$$Efficiency = \frac{256\Delta t_{bin} N_{PDS}}{N_{PDS}(I_{proc} + I_{accum}) + I_{trans}}$$

where N_{PDS} is the number of summations in the readout, Δt_{bin} is the time bin width, and the three intervals I are the processing, accumulation and transfer intervals. The processing interval is 12 ms, the accumulation interval is $I_{proc} = 256\ \Delta t_{bin}$, and the transfer interval ~70 ms. The efficiency ranges from ~2% at 1 µt bins to ~25% at 16 µt bins and ~70% at 100 µt, where we have taken the transfer interval to be negligible compared compared to the overall product time (sum of accumulation and processing intervals).

The selectable parameters for the *FFT mode* are as follows:

- energy/type selection #1
- energy/type selection #2
- PCU selections (Any of 1,2,3,4,5) #1
- PCU selections (Any of 1,2,3,4,5) #2
- time-bin duration (0.954 µs to 4096 s)
- number of PDS and CS to sum together
- time-series and spectra scaling parameters.

7.6 EDS-ASM Operations

7.6.1 ASM-EDS Interfaces

Each X-ray photon which is detected in a position-sensing chamber of a *scanning shadow camera* (SSC) generates a signal which is amplified by two analog electronic measurement chains, one at each end of a resistive anode. The two pulse heights are converted to digital 12-bit numbers and transferred on a serial line (one for each SSC) to the ASM Event Interface. Background event data are also transferred on an event-by-event basis to the EDS over the same set of serial lines. Each data word received by the EDS comprises a 4-bit event code and two 12-bit pulse heights. The event code contains the anode number for “good” X-ray events or a code indicating one of *four types of background* *events* (multi-anode event, position-sensing anode and veto layer coincidence, veto layer event, A and/or B pulse height exceeds the upper level discriminator).

All events transmitted from the ASM over the event interface are received by each of the two ASM Event Analyzers.

A second interface, the control interface, consists of a shared bi-directional serial data bus. With this interface, the EDS can obtain ASM engineering and housekeeping information, and control the SSC signal processing and the ASM rotation drive. Only the ASM EA on EDS Side B can utilize the command and housekeeping interface; the ASM EA on EDS Side A can send rotation commands but cannot monitor housekeeping.

7.6.2 ASM EA Functions

The data processing in the ASM EAs operates quite similarly to the PCA EAs except that there is no typing and recoding of events in hardware. The front end of an ASM EA receives event data from the three SSCs and places this information (4-bit event code and 2 12-bit pulse heights) together with a 2-bit SSC ID and a 10-bit time tag into a series of three 16-bit words which are transferred into the FIFO register. Interval markers are also inserted into the FIFO at times which depend on the details of the EA configuration. The DSP therefore reads one-word interval markers as usual and three-word event descriptions. It extracts the event code, the SSC ID, the A and B pulse heights, and the timestamp from each event description in order to perform further processing according to the active configuration.

The DSP selects events according to preset criteria, computes event energies (i.e., A+B) and anode positions (i.e., N [A/(A+B)] where N is the number of position bins) as required, computes bin numbers from the SSC and anode IDs, event energies, and event positions, and builds histograms or event lists in the data buffer (DB). This activity results in the creation of a partition of data for each Interval as described above for the PCA event analyzers.

The Back End of an ASM EA processes the partitions in the same manner as described above for the PCA EAs. In addition, it controls and monitors the rotation drive assembly and the acquisition of housekeeping and engineering data from the three SSCs.

7.7 ASM Data Processing Modes:

7.7.1 Position Histogram

This is the primary operating mode of the ASM and serves to monitor source intensities and to search for unpredicted outbursts from sources. The EA running this mode controls the rotations of the ASM *drive assembly*, acquires engineering housekeeping data from the three SSCs and the drive assembly, and processes the event data from the three SSCs to produce and telemeter position histogram data. The ASM EA on Side B of the EDS can run this mode with its full functionality, whereas that on Side A does not monitor housekeeping.

Position histogram mode performs “dwell sequences”, which are planned in advance on the ground and uploaded to the ASM EA via the spacecraft and an EDS system manager PRAM load. A *dwell sequence* is basically a series of “dwells”, each of which is followed by a rotation of the drive assembly. A dwell is an interval of time, typically between 80 and 100 seconds, during which the drive assembly does not rotate so that the SSC’s maintain a fixed orientation relative to the sky. Events are collected and processed into position histograms during dwells. After each dwell, the position histogram data is transferred to the spacecraft to be sent to the ground.

Each set of dwell sequence parameters contains two ASM drive-assembly slew angles; these slews are performed prior to the first dwell in order to begin the dwell sequence with the SSCs in the optimum orientation to avoid viewing the Earth during the subsequent ~100-s dwells. In addition to the two slew parameters, a dwell sequence is specified by a number of dwells, a dwell duration, and an angle to rotate after each dwell (one constant value per dwell sequence).

Position histograms consist of counts binned as a function of position for each of the 8 resistive anodes in each of the three detectors, and for each of, nominally, three energy channels. The number of position bins, the number of energy channels, and the energy channel boundaries are parameters for position histogram mode.

Position histogram mode also collects and telemeters engineering housekeeping data from the drive assembly and the 3 SSCs. This is done once each dwell in parallel with the collection of event data.

7.7.2 Multiple Time Series

This mode produces three types of binned data products, namely (1) *time-series histograms*, (2) *background event histograms*, and (3) *pulse-height histograms*. These data products are produced continuously on an Interval-by-Interval basis without regard to the status of the drive assembly, i.e., regardless of whether the SSCs are stationary or rotating.

Time-series histograms consist of counts of “good X-rays” per time bin for each detector for each of a few (1 – 5) energy bands. Each time bin can be as short as 0.001 seconds but will typically be ~0.1 seconds so that the resulting telemetry rate is modest. These data may be useful for studying the pulsations of a few bright X-ray pulsars or for searching for pulsations from newly-discovered bright transient sources.

Background-event histograms consist of the counts of each of the four types of background event (see Section 6.1 “ASM - EDS Interfaces”) in one-second time bins for each of the three detectors.

Pulse-height histograms consist of pulse-height spectra of “good X-ray” events for each of the eight resistive anodes in each of the three detectors. The pulse-height spectra will consist of 32 or 64 channels and are accumulated over time intervals of ~100 seconds. These data will primarily be used to monitor the location of the 5.9 keV spectral line from the weak ^{55}Fe sources in the instrument and to thereby determine the energy-to-pulse-height conversion gain of each of the detectors. Each spectrum is expected to contain roughly 32 pulse-height bins.

7.7.3 Event by Event

In this mode, the DSP stores three-word event descriptions in sequence in the data buffer. The event descriptions are telemetered after the data collection interval has ended. The mode may be configured so that either *all* types of events are selected and stored or only *good* X-ray events are selected and stored. Only data from one of the three SSCs is selected. Therefore, the selection of an SSC ID and choice of ‘good’ or ‘all events’ are options for this mode. This mode will see restricted use since it can result in telemetry rates that are very high. It will be used for diagnostics and calibration.