

AD-A278 406



Data Processing Algorithms for On-Board Satellite Event Analysis

30 January 1994

Prepared by

R. KOGA, S. D. PINKERTON, and N. KATZ
Space and Environment Technology Center
Technology Operations

Prepared for

SPACE AND MISSILE SYSTEMS CENTER
AIR FORCE MATERIEL COMMAND
2430 E. El Segundo Boulevard
Los Angeles Air Force Base, CA 90245



Engineering and Technology Group

94-11855



94 4 19 077

APPROVED FOR PUBLIC RELEASE;
DISTRIBUTION UNLIMITED

This report was submitted by The Aerospace Corporation, El Segundo, CA 90245-4691, under Contract No. F04701-93-C-0094 with the Space and Missile Systems Center, 2430 E. El Segundo Blvd., Los Angeles Air Force Base, CA 90245. It was reviewed and approved for The Aerospace Corporation by A. B. Christensen, Principal Director, Space and Environment Technology Center. Major Leslie Belsma was the project officer for the Mission-Oriented Investigation and Experimentation (MOIE) program.

This report has been reviewed by the Public Affairs Office (PAS) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nationals.

This technical report has been reviewed and is approved for publication. Publication of this report does not constitute Air Force approval of the report's findings or conclusions. It is published only for the exchange and stimulation of ideas.



Leslie Belsma, Major, USAF
Project Officer



Wm. Kyle Sneddon, Captain, USAF
Deputy, Industrial & International Division

REPORT DOCUMENTATION PAGEForm Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE 30 Jan 1994	3. REPORT TYPE AND DATES COVERED	
4. TITLE AND SUBTITLE Data Processing Algorithms for On-Board Satellite Event Analysis			5. FUNDING NUMBERS F04701-93-C-0094	
6. AUTHOR(S) R. Koga, S. D. Pinkerton, and N. Katz				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The Aerospace Corporation Technology Operations El Segundo, CA 90245-4691			8. PERFORMING ORGANIZATION REPORT NUMBER TR-0091(6940-05)-8	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Space and Missile Systems Center Air Force Materiel Command 2430 E. El Segundo Boulevard Los Angeles Air Force Base, CA 90245			10. SPONSORING/MONITORING AGENCY REPORT NUMBER SMC-TR-94-16	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Event rates encountered by satellite-borne sensors typically exceed the available telemetry bandwidth, necessitating on-board data analysis. One method of reducing output telemetry requirements, employed by data processing units (DPUs) on several recent and forthcoming satellites (VIKING, CRRES, CAMMICE) is to partition the underlying event parameter spaces into disjoint classes, and to downlink statistics based on these classes rather than individual events. The DPUs utilize a hardware implemented binary search algorithm to classify each valid event into one of several mass and mass per charge-state (M/Q) groups, using look-up tables written under software control. Mass and M/Q group boundaries are defined via polynomial curves in the underlying parameter spaces. The boundary curves were obtained using a combination of global and local least-squares approximation algorithms applied to pre-flight sensor calibration data.				
14. SUBJECT TERMS Data processing unit, On-board data analysis, Least-squares approximation, Satellite-borne sensors			15. NUMBER OF PAGES 10	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT	

CONTENTS

Abstract	3
I. Introduction	3
II. DPU Hardware	3
A. Overview of DPUA	3
B. Successive Approximation Routine	4
III. On-Board Data Processing Schemes	5
A. M vs. M/Q Matrix Scalers	5
B. Rate Scalers	5
IV. Mass and M/Q Group Determination	6
A. Overview	6
B. Curve-Fitting Software	6
C. Mathematical Foundations	7
V. Table Writing Software	8
A. Mass Group Table Writing	8
VI. Summary	9
VII. Conclusion	9
Acknowledgments	9
References	9

FIGURES

1. DPUA Functional Block Diagram	4
2. SAR Mass Group Classification	5
3. Mass Group Boundary Curves Calculated by FIT	6
4. Mass Group Classification Schemes.....	8

Abstract

Event rates encountered by satellite-borne sensors typically exceed the available telemetry bandwidth, necessitating on-board data analysis. One method of reducing output telemetry requirements, employed by data processing units (DPUs) on several recent and forthcoming satellites (VIKING, CRRES, CAMMICE) is to partition the underlying event parameter spaces into disjoint classes, and to downlink statistics based on these classes rather than individual events. The DPUs utilize a hardware implemented binary search algorithm to classify each valid event into one of several mass and mass per charge-state (M/Q) groups, using look-up tables written under software control. Mass and M/Q group boundaries are defined via polynomial curves in the underlying parameter spaces. The boundary curves were obtained using a combination of global and local least-squares approximation algorithms applied to pre-flight sensor calibration data.

I. Introduction

Event rates encountered by satellite-borne sensors typically exceed the available telemetry bandwidth, necessitating on-board data analysis and compression. One method of reducing telemetry requirements is to compute summary statistics on-board and to downlink these statistics rather than individual events.

For example, the microprocessor-based data processing units (DPUs) that support the Magnetospheric Ion Composition Spectrometer (MICS) aboard both the Swedish satellite VIKING and the Combined Release and Radiation Effects Satellite (CRRES) utilize a hardware implemented binary search strategy to classify each valid event into one of several "mass groups" based on energy (E) and time-of-flight (T) values generated by the sensor. Each event is similarly assigned to a mass per charge-state group based on the time-of-flight and energy per charge-state (E/Q) of the incident particle. The assigned mass (M) and mass per charge-state (M/Q) groups are then used to index a matrix of hardware accumulators (called "scalers") in such a way that each event is counted by exactly one M vs. M/Q "matrix" scaler. By periodically downlinking the complete set of M vs. M/Q matrix scalers, a coarse determination of event populations can be obtained through the limited bandwidth telemetry channel.

Finer temporal and parametric resolution is achieved through the use of individual "rate scalers" which may be programmed to accept only those events having particular combinations of M and M/Q values. The additional inclusion in the telemetry output of a small number of "direct events" (raw energy and time-of-flight values) allows resolution at the scale of a single event. The combined use of M vs. M/Q matrix scalers, rate scalers, and sample direct events thereby provides coverage of event parameter space at a range of resolutions, from coarse to fine.

The following sections describe in detail the data analysis and reduction schemes employed by the data processing unit (DPUA) supporting the MICS sensor aboard the CRRES satellite. Similar techniques are utilized by the DPUs associated with the VIKING MICS instrument [1] and the Low Energy Magnetospheric Ion Composition Sensor aboard CRRES. Data from the MICS sensor to be flown on the forthcoming CAMMICE satellite will also be analyzed via M vs. M/Q reduction.

Classification of E-T events into mass and M/Q groups is performed automatically by a hardware-implemented binary search algorithm utilizing look-up tables written by the DPU software. The tables contain discrete approximations to polynomial curves which separate adjacent groups. The curves were derived from extensive pre-flight sensor calibration data using a novel curve-fitting and extrapolation algorithm combining local and global least-squares approximations. Both the on-board table writing software and the techniques employed to derive the boundary curves are described in detail below.

II. DPU Hardware

The data processing unit (DPUA) that interfaces with the CRRES MICS instrument performs the following functions: (1) Accepts, decodes, and executes ground commands uplinked to the DPU by way of the spacecraft; (2) Distributes power to the sensor and controls the high voltage levels of the MICS electrostatic analyzer and post-accelerator; (3) Controls the experiment by configuring the sensor in particular data collection modes through the use of "MICS instructions"; and (4) Reads, processes, and outputs to telemetry data from the MICS instrument. The following sections are concerned primarily with the last of these functions; additional information on the remaining functions may be found in [2].

A. Overview of DPUA

As shown in Fig. 1, DPUA consists of a computer, peripheral sections such as various I/O buffers and the MICS Event Processor, a power supply, and an analog signal handler. The computer consists of a microprocessor (SANDIA 1802), control logic, and memory (12K bytes of ROM and 16K bytes of RAM). The radiation-hardened 1802 microprocessor operates with a 1.6 MHz clock. The DPUA software is initially stored in ROM and is copied to RAM shortly after the DPU is turned on. Power to the ROMs is then shut off, and thereafter the software executes exclusively in RAM.

One important group of peripheral sections functions as the MICS Event Processor (see the area inside the dashed lines in Fig. 1) which analyzes MICS sensor data semi-autonomously after initial preparation (table writing by the computer) has been performed. The MICS sensor generates energy (E) and time-of-flight (T) values for each event. In the Event Processor the E and T values, along with an energy per charge-state

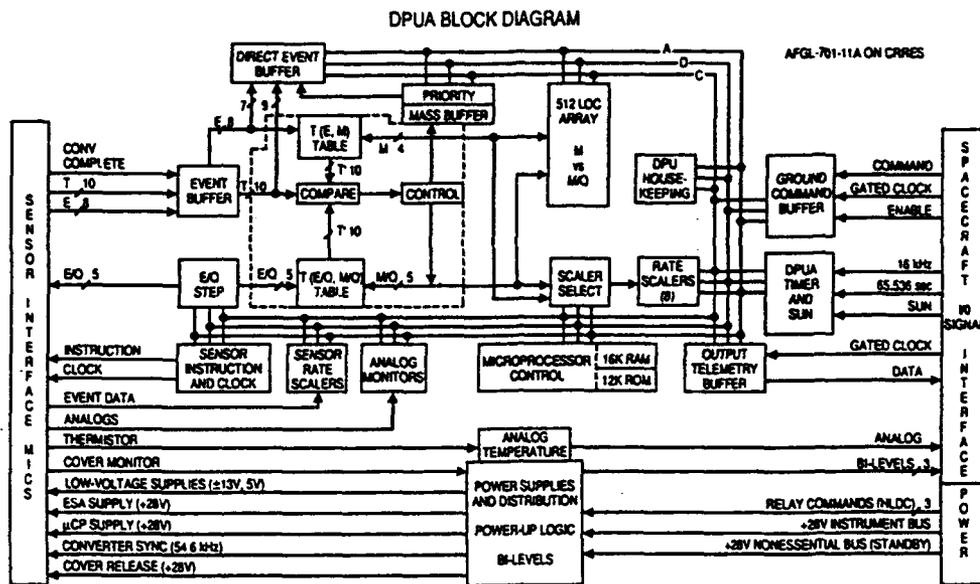


Fig. 1. DPUA Functional Block Diagram.

(E/Q) value, are used to classify the event into one of several mass (M) and mass per charge-state (M/Q) groups. This classification is performed by the hardware-based Successive Approximation Routine, which utilizes look-up tables written by the DPU under software control.

The assigned M and M/Q values for an event are used to address one scaler in a 512-element "matrix" of scalers (M vs M/Q in Fig. 1). DPUA also contains eight "rate scalers" in which events with special combinations of M and M/Q are accumulated. The Rate Scaler Select Table determines which combinations of M and M/Q are accumulated by the individual rate scalers. Sample "direct events" (E and associated T values) are collected from the sensor and stored in the event buffer. Direct events are selected according to a priority scheme determined using DPUA's Event Priority Table. The contents of the various scalers and the direct event buffer are read-out periodically by the computer, processed (e.g., compressed), and placed in the telemetry stream.

B. Successive Approximation Routine

For each valid event, the MICS sensor outputs an 8-bit energy (E) value and a 10-bit time-of-flight (T) value. These values, together with a 5-bit quantization (E/Q) of the high voltage level of the MICS electrostatic analyzer, are used by DPUA's MICS Event Processor to classify the event into one of up to 16 mass groups (M) and one of up to 32 mass per charge-state groups (M/Q). Mass and M/Q group classification is performed directly by the DPU hardware using look-up tables written under software control. Each mass group is defined as the region of E vs. T parameter space lying between two predefined boundary polynomial curves of the form $T = P(E)$, as discussed further in Section IV. Soon after the DPU is powered-up in space, the DPU software calculates discrete approximations to these curves, based on stored coefficients,

and writes the calculated values to an area of DPU memory known as the Mass Group Table.

The Mass Group Table, or MGT, is a two-dimensional array indexed by mass group M (4 bits), and energy E (8 bits). For each mass group M, the 256 table values $MGT(E=0, M) \dots MGT(E=255, M)$ give the discrete approximation to the upper boundary polynomial for the mass group - that is, $MGT(E, M)$ is the nominal time-of-flight value (10 bits) for a particle of energy E belonging to mass group M. The DPU's hardware-based Successive Approximation Routine (SAR) uses the precomputed Mass Group Table values to classify each incoming event (specified by E and T) via a binary search for the mass group M satisfying the inequality $MGT(E, M) \leq T < MGT(E, M+1)$. In particular, the SAR first compares T to the table value $MGT(E, M=8)$, i.e., to the table value corresponding to E for mass group number 8. Depending on whether $T < MGT(E, M=8)$ or $T \geq MGT(E, M=8)$, the SAR then compares T to $MGT(E, M=4)$ or $MGT(E, M=12)$, respectively. The first comparison determines in which half of "mass group space" the event belongs; subsequent comparisons (to $MGT(E, M=4)$ or $MGT(E, M=12)$, for example) determine the appropriate quadrant, octant, and sextant. Thus, in at most four comparisons, the unique mass group M is found which satisfies $MGT(E, M) \leq T < MGT(E, M+1)$; the SAR then assigns the event being classified to mass group M, as illustrated in Fig. 2.

While this classification algorithm imposes slightly greater hardware processing demands than would direct table look-up (storing the mass group directly as a function of E and T), the reduction in memory requirements is quite dramatic: $(16 \text{ groups}) \cdot (256 \text{ E-values/group}) \cdot (10 \text{ bits/T-value}) = 40,960$ bits for the SAR-based scheme, versus $(256 \text{ E-values}) \cdot (1024 \text{ T-values}) \cdot (4 \text{ bits/group}) = 1,048,576$ bits for direct table look-up. Thus the slight increase in hardware complexity is more than offset by the 96% decrease in the amount of

memory required (it should be noted that the throughput of the SAR exceeds the event rate of the MICS instrument, hence no bottleneck is introduced by this increase in processing).

The M/Q Group Table is organized in a similar fashion. The values in this table represent discrete approximations to polynomial curves of the form $T = P(E/Q)$, where E/Q is a quantization of the energy per charge-state of the event. Up to 32 M/Q groups can be defined, so that the table requires a total of $32 \cdot 32 \cdot 10 = 10,240$ bits, as opposed to $32 \cdot 1024 \cdot 5 = 163,840$ bits for a direct table look-up approach – the reduction in storage requirements is thus approximately 94%.

III. On-Board Data Processing Schemes

The principal sensor-related data entities processed by the DPU are the R0 – R7 rate scalers and the M vs. M/Q matrix scalers. The contents of scalers (RAM accumulators) are defined via software-configurable look-up tables residing in DPU memory. The DPU software periodically reads and resets the scalers, compresses the values read from up to 19 significant bits to 8 bits, and then places the compressed data in the telemetry stream. The matrix and rate scalers are described in turn below; a detailed description of the scaler compression scheme utilized by DPUA may be found in [2].

A. M vs. M/Q Matrix Scalars

As discussed above, each valid sensor event is assigned to one of up to 16 mass (M) and one of up to 32 mass per charge-state (M/Q) groups by DPUA's MICS Event Processor, based on the contents of the Mass Group and M/Q Group Tables. The DPUA hardware maintains a $16 \times 32 = 512$ -element matrix of scalars (RAM accumulators), indexed by M and M/Q, respectively. A given M vs. M/Q matrix scaler counts only those events assigned to the corresponding mass and mass per charge-state groups. The scalars are 19 bits each, allowing up to $2^{19}-1$ events to be accumulated without overflowing (in practice, these scalars never overflow). Prior to telemetry output, all scalars are compressed to 8 bits.

Notice that the mapping from energy and time-of-flight values (sensor outputs) to M and M/Q groups (Event Processor outputs) effectively reduces the dimensionality of event parameter space from $8 \cdot 1024$ to $16 \cdot 32$, or approximately 94%. By periodically downlinking the entire complement of M vs. M/Q matrix scalars, a coarse determination of event populations can be obtained through the limited bandwidth telemetry channel.

B. Rate Scalars

The outputs of the Event Processor (M and M/Q) can also be stored in eight rate scalars that are output to telemetry much more frequently than the M vs M/Q matrix scalars. Each rate scaler (R0 – R7) may be "programmed" to accept only those events with particular combinations of M and M/Q. This scheme allows multiple values of M and M/Q to be grouped together and counted by a single scaler.

The assignment of (M,M/Q) combinations to particular scalars is accomplished using the software-configurable Rate Scaler Select Table. There are eight distinct "pages" in the Rate Scaler Select Table, each of which defines a different assignment rule for the eight rate scalars. Only one page at a

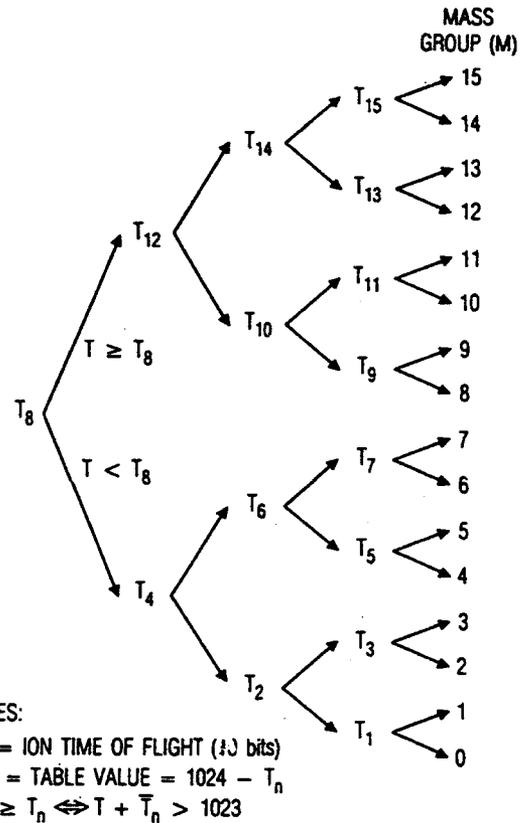


Fig. 2. SAR Mass Group Classification: Incoming events specified by energy (E) and time-of-flight (T) are classified in one of sixteen mass groups (M) by successive comparisons of T with the nominal time-of-flight values stored in the DPU's Mass Group Table. Shown is a schematic representation of table values for a particular value of E.

time is active, i.e., defines the contents of the rate scalars (the currently active page may be changed by uplinking the appropriate ground command to the DPU). The rate scalars have a capacity of $2^{16}-1$ (16 bits) per scaler and are compressed to 8 bits prior to telemetry.

Events arriving at the Event Processor are automatically routed to the appropriate rate scaler based on the definitions of R0 – R7 contained in the currently active page of the Rate Scaler Select Table. Each of the eight pages can be viewed as a matrix indexed by M (4 bits) and M/Q (5 bits). Each location within this 32×16 matrix contains a 3-bit value in the range 0 to 7 specifying the rate scaler (one of R0 – R7) in which events with the corresponding M and M/Q values are to be accumulated.

In this way, regions of particular interest in M vs. M/Q space can be isolated and monitored with greater parametric and temporal resolution than is possible using the matrix scalars (the accumulation period for the DPUA rate scalars is $1/256$ th that of the matrix scalars).

IV. Mass and M/Q Group Determination

As discussed above, DPUA's MICS Event Processor classifies each sensor event into the mass (M) and mass per charge-state (M/Q) groups determined by values contained in the Mass and M/Q Group Tables, respectively. The contents of the Mass Group Table (for example) represent discrete approximations to polynomial curves of the form $T = P(E)$, where T and E are nominal time-of-flight and energy values. These curves divide the underlying E vs. T parameter space into disjoint regions, each of which constitutes a single mass group.

In this section the genesis of the mass group boundary polynomials is described; similar comments apply to the M/Q group boundary polynomials. An overview of the techniques used to obtain the polynomial coefficients is presented first, followed by a brief description of FIT, which is an integrated software package implementing the procedures outlined in the overview. Finally, the mathematical foundations of the curve-fitting algorithms are reviewed.

A. Overview

The boundary curves separating adjacent mass groups are derived from polynomial least-squares fits to calibration data obtained from the MICS instrument prior to launch. At first glance it might appear that one need only calculate a single approximating curve for each ion species and then translate these curves slightly so that the curves separate, rather than fit, the data corresponding to distinct ion species. However, the curves so obtained tend to intersect for low energy values and are therefore unsuitable for use with an SAR-based classification scheme. In particular, for each pair of E and T values, the SAR requires that there be a unique mass group M satisfying the inequality, $MGT(E,M) \leq T < MGT(E,M+1)$, which is equivalent to the requirement that the mass group boundary curves be non-intersecting. In order to satisfy this constraint, a model-based approach was employed to determine the polynomials.

Assuming mass, energy, and time-of-flight to be related by a multivariate polynomial of the form $M = Q(E,T)$ (e.g. $Q(E,T) = P_0(E) + T \cdot P_1(E) + T^2 \cdot P_2(E)$, where P_0 , P_1 , and P_2 are polynomials in E), a "global" least-squares approximation was performed over the entire calibration data set, resulting in the determination of the polynomial coefficients relating E and T to the atomic mass M of each ion species.

The relationship summarized by $M = Q(E,T)$ was then used to interpolate data points for "fictitious" values of M by letting E range from 0 to 255 and solving $M = Q(E,T)$ for the corresponding values of T . For example, with $Q(E,T)$ as above, evaluation of P_0 , P_1 , and P_2 at a particular value $E = E'$ yields three values: $E_0 = P_0(E')$, $E_1 = P_1(E')$, and $E_2 = P_2(E')$. Substituting these values into the equation for M then yields: $M = E_0 + E_1 \cdot T + E_2 \cdot T^2$, or equivalently, $0 = (E_0 - M) + E_1 \cdot T + E_2 \cdot T^2$, where the E_k 's are constants. After setting M to some "fictitious" value (e.g., $M = 1.5$), the quadratic equation $0 = (E_0 - M) + E_1 \cdot T + E_2 \cdot T^2$ can be solved for T .

By fixing M and letting E vary from 0 to 255, a data set of "simulated" $E-T$ pairs satisfying $M = Q(E,T)$ was obtained, as described above. A polynomial least-squares fit to this data set was then performed, producing a curve $T = P(E) = a_0 + a_1 \cdot E + a_2 \cdot E^2 + a_3 \cdot E^3$, describing the nominal relationship between E

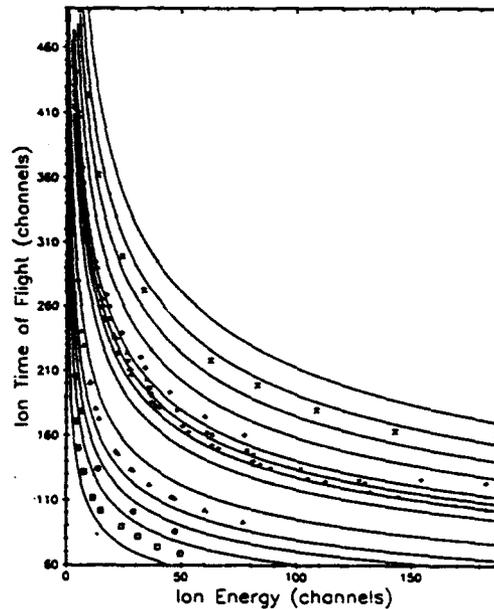


Fig. 3. Mass Group Boundary Curves Calculated by FIT.

and T for the "fictitious" mass value M . (In practice the logarithms of E and T were utilized.) For example, with $M = 1.5$, a curve separating the data for $M = 1.0$ (hydrogen) from the data for $M = 2.0$ (deuterium) was obtained. This technique was repeated for various "fictitious" mass values, producing a set of non-intersecting mass group boundary curves suitable for use with the SAR-based classification algorithm, as shown in Fig. 3.

It should be noted that the "data points" utilized in the curve fitting algorithms were not simple $E-T$ pairs as implied above. Rather, to reduce the influence of noise in the raw $E-T$ data obtained from the sensor, and to limit the number of data points to a computationally tractable range, each data point was obtained as an approximation to the mean of a distribution of individual $E-T$ measurements.

B. Curve-Fitting Software

The several steps outlined in the above overview have been codified into an integrated software package known as "FIT." The FIT package consists of several dozen individual routines written in VAX FORTRAN. Included are routines for performing least-squares approximations, plotting, error calculation, and generating the look-up tables required by the DPU's SAR (this last function is used to verify the performance of the on-board table writing software). The individual routines may be invoked either by way of a VAX/VMS DCL command file, or via a menu-driven user-interface. A number of options exist for customizing the general-purpose routines in the FIT package to particular applications, however the primary utility of these routines is in the generation of the coefficients of the boundary polynomials upon which the mass and M/Q group look-up tables are based.

C. Mathematical Foundations

The least-squares approximation algorithms utilized in the FIT program will be briefly reviewed here (see [3] for a more extensive development).

Suppose that $(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)$ are independent measurements (for example, x = energy and y = time-of-flight), and that the distribution is approximated by the polynomial $y = P(x) = c_1 + c_2x + \dots + c_n x^{n-1}$. One measure of the total error introduced in the approximation P is given by the sum,

$$E(c_1, c_2, \dots, c_n) = \chi^2 = \sum_{k=1}^N (y_k - P(x_k))^2. \quad (1)$$

In the method of least-squares, a polynomial P which minimizes the error E is sought. From elementary calculus, the coefficients of such a polynomial must satisfy $\partial E / \partial c_i = 0$ for $i = 1, \dots, n$. Differentiation of both sides of Eq. 1 yields,

$$\begin{aligned} \partial E / \partial c_i &= -2 \sum_{k=1}^N (y_k - P(x_k)) \cdot \partial P(x_k) / \partial c_i \\ &= -2 \sum_{k=1}^N (y_k - P(x_k)) \cdot x_k^{i-1}. \end{aligned} \quad (2)$$

Hence, $\partial E / \partial c_i = 0$ if and only if $\sum_{k=1}^N (y_k - P(x_k)) x_k^{i-1} = 0$,

or equivalently,

$$\sum_{k=1}^N y_k x_k^{i-1} = \sum_{k=1}^N P(x_k) \cdot x_k^{i-1}, \quad (3)$$

which can also be written,

$$\begin{aligned} \sum_{k=1}^N y_k x_k^{i-1} &= \\ c_1 \sum_{k=1}^N x_k^{i-1} + c_2 \sum_{k=1}^N x_k^i + \dots + c_n \sum_{k=1}^N x_k^{n+i-2}. \end{aligned} \quad (4)$$

The final equation above (Eq. 4), taken for each $i = 1, \dots, n$, defines a system of n , n^{th} order, non-homogeneous linear equations with respect to c_1, \dots, c_n :

$$\begin{aligned} \sum y_k &= c_1 \sum 1 + c_2 \sum x_k + \dots + c_n \sum x_k^{n-1} \\ \sum y_k x_k &= c_1 \sum x_k + c_2 \sum x_k^2 + \dots + c_n \sum x_k^n \end{aligned}$$

$$\sum y_k x_k^{n-1} = c_1 \sum x_k^{n-1} + c_2 \sum x_k^n + \dots + c_n \sum x_k^{2n-2}.$$

Now define the column vectors:

$$Y = \left(\sum_{k=1}^N y_k, \sum_{k=1}^N y_k x_k, \dots, \sum_{k=1}^N y_k x_k^{n-1} \right)^T$$

and $C = (c_1, c_2, \dots, c_n)^T$,

and let X be the $(n \times n)$ matrix whose j^{th} row is the vector,

$$\left(\sum_{k=1}^N x_k^{j-1}, \sum_{k=1}^N x_k^j, \dots, \sum_{k=1}^N x_k^{n+j-2} \right).$$

The system of equations given above can then be written compactly as $Y = XC$, and (provided that X is invertible) the solution vector is simply $C = X^{-1}Y$. (In practice, it is more economical to calculate C via the equation $C = (Z^T Z)^{-1} Z^T Y$ where Z is the $(N \times n)$ matrix whose j^{th} column is the vector $(x_1^{j-1}, x_2^{j-1}, \dots, x_N^{j-1})$, and $V = (y_1, y_2, \dots, y_N)$. As is easily verified, $Z^T Z = X$ and $Z^T V = Y$.)

The situation for multivariate polynomials is entirely analogous since it is with respect to the coefficients of the polynomial that the error must be minimized. For example, if $(x_1, y_1, z_1), (x_2, y_2, z_2), \dots, (x_N, y_N, z_N)$ is a set of measurements approximated by a multivariate polynomial of the form: $z = P(x, y) = (c_{11} + c_{21}x + \dots + c_{n1}x^{n-1}) + (c_{12} + c_{22}x + \dots + c_{n2}x^{n-1})y + \dots + (c_{1m} + c_{2m}x + \dots + c_{nm}x^{n-1})y^{m-1}$, then the error in the approximation is given by,

$$E = \chi^2 = \sum_{k=1}^N (z_k - P(x_k, y_k))^2. \quad (5)$$

Hence,

$$\begin{aligned} \partial E / \partial c_{ij} &= -2 \sum_{k=1}^N (z_k - P(x_k, y_k)) \cdot \partial P(x_k, y_k) / \partial c_{ij} \\ &= -2 \sum_{k=1}^N (z_k - P(x_k, y_k)) \cdot x_k^{i-1} y_k^{j-1}, \end{aligned} \quad (6)$$

and $\partial E / \partial c_{ij} = 0$ if and only if $\sum_{k=1}^N (z_k - P(x_k, y_k)) x_k^{i-1} y_k^{j-1} = 0$, or equivalently,

$$\sum_{k=1}^N z_k x_k^{i-1} y_k^{j-1} = \sum_{k=1}^N P(x_k) \cdot x_k^{i-1} y_k^{j-1}. \quad (7)$$

As in the single variable case, this equation leads to a system of $n \cdot m$ non-homogeneous linear equations, the solution of which provides the coefficients of the multivariate polynomial for which the approximation error E is minimal.

V. Table Writing Software

The look-up tables required by DPUA's MICS Event Processor are written by the SYST table calculation routine. SYST writes three distinct tables: The Mass Group Table used by DPUA's Event Processor to classify events into mass (M) groups; the Mass per Charge-State (M/Q) Group Table used by the Event Processor to classify events into M/Q groups; and the Rate Scaler Select Table used to define the contents of the R0 - R7 rate scalers.

The procedure followed by SYST in writing the Mass Group Table is described below (the M/Q Group Table is written in an analogous fashion). The role played by SYST in constructing the Rate Scaler Select Table is minimal, and will not be described here.

A. Mass Group Table Writing

As described in Section IIB, the entries in the Mass Group Table represent overlaid polynomial curves in E vs. T parameter space, arranged such that adjacent curves define the upper and lower bounds of a single mass group. Table values are calculated using coefficients stored in DPU memory. The default coefficients are derived from polynomial fits to sensor calibration data as described in Section IV; new sets of coefficients can also be uplinked to DPUA using ground commands.

The Mass Group Table is used by the hardware-based Successive Approximation Routine of the Event Processor to classify each event into one of up to 16 mass groups, based on the 8-bit energy (E) and 10-bit time-of-flight (T) values obtained from the sensor. The table is divided into 16 sections, one per mass group, each of which contains 256 10-bit entries.

Each mass group section of the table is indexed by energy values ranging from $E = 0$ to $E = 255$. For each energy value E, a corresponding nominal time-of-flight value T is calculated by SYST as $T = 2^t$, where $t = P(x) = c_0 + c_1x + c_2x^2 + c_3x^3$ and $x = \text{Log}_2E$, and then the complement of this value relative to 1024 is written to the Mass Group Table. A different set of coefficients is used to calculate the table values for each mass group. The coefficients define non-intersecting curves in E vs. T parameter space (see Fig. 3) which represent the boundaries between successive mass groups (fifteen such curves are therefore necessary to partition the space into sixteen mass groups). In particular, the time-of-flight values in a given mass group section of the table represent the lower boundary of the mass group, hence no values need be written for Mass Group 0.

In the basic classification scheme up to 16 mass groups may be defined via stored coefficients as above. It is important to notice that fewer than 16 mass groups may be allocated, simply by specifying appropriate coefficients. For example, specifying "barrier-down coefficients" for Mass Group 12 ensures that only the lowest 12 mass groups will be used. These coefficients result in zeros being written to the Mass Group Table for the specified group; hence, since for any T, $T + 0 \leq 1023$, no event will ever be placed in Mass Group 12 or any higher numbered group, i.e., only Mass Groups 0 through 11 will be used.

In addition to the basic classification scheme outlined above and illustrated in Fig. 4a, SYST also offers a special classification

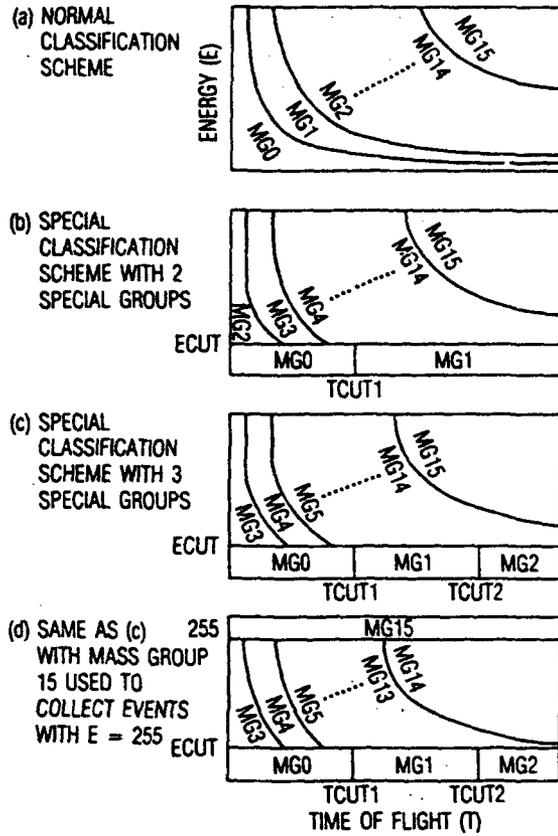


Fig. 4. Mass Group Classification Schemes.

classification scheme in which all events with sufficiently small energy values are placed in "special" mass groups. Either two or three special mass groups may be allocated. With two special mass groups (Fig. 4b), every event with energy less than or equal to the pre-defined constant ECUT is placed in either Mass Group 0 or Mass Group 1, depending on whether the corresponding time-of-flight value is less than or equal to the constant TCUT1 (Mass Group 0), or greater than TCUT1 (Mass Group 1). Events with energy values greater than ECUT are classified as in the basic scheme described above. In particular, if an event with $E > ECUT$ falls into either Mass Group 0 or Mass Group 1 under the basic scheme, then it will also be placed there under the special scheme, thereby corrupting the interpretation of these mass groups as low-energy event collectors. To avoid this problem, the coefficients for the Mass Group 2 section of the table should be changed to "barrier-up coefficients". These coefficients cause a saturated T value to be written to the table locations corresponding to $E > ECUT$ in the Mass Group 2 section of the table, the net result of which is that all events with $E > ECUT$ which would have fallen into Mass Group 0 or Mass Group 1 under the basic classification scheme are instead placed in Mass Group 2.

The situation with three special mass groups is similar (Fig. 4c) except that an additional constant, TCUT2, is used in conjunction with TCUT1 to place events with $E \leq ECUT$ in either Mass Group 0 ($T \leq TCUT1$), Mass Group 1 ($TCUT1 < T \leq TCUT2$), or Mass Group 2 ($T > TCUT2$). As noted above, to maintain the integrity of the 3 special mass groups, the coefficients for both Mass Group 2 and Mass Group 3 should be changed to "barrier-up coefficients".

SYST also offers an option through which all events with energy equal to the maximum value of 255 are placed in Mass Group 15 (shown in Fig. 4d together with the special classification scheme). This option may be used in conjunction with any of the above classification schemes, that is, independent of the number of special mass groups, if any. However, it should be noted that if the selected classification scheme makes use of all 16 mass groups (including Mass Group 15), then Mass Group 15 will collect lower energy events in addition to events with $E = 255$. Thus this option is only used when "barrier-down coefficients" are simultaneously being employed to deflect lower-energy events away from Mass Group 15.

VI. Summary

The VIKING MICS DPU, the CRRES LOMICS and MICS DPUs, and the forthcoming CAMMICE DPUs all employ the M vs. M/Q analysis techniques described here. The combined use of M vs. M/Q scalars, rate scalars, and sample direct events allows event parameter space to be monitored at a range of parametric and temporal resolutions within the constraints imposed by the limited-bandwidth telemetry channel of the spacecraft. Scalar compression further conserves telemetry bandwidth, without significant loss of precision.

A hardware-implemented binary search strategy called the Successive Approximation Routine (SAR) is used to classify events into mass and M/Q groups using memory-resident tables written by the DPU software. Use of the SAR-based scheme results in a significant reduction in memory requirements compared with direct table look-up techniques, without creating a processing bottleneck at the DPU.

The tables used by the SAR are user-configurable. A number of options exist whereby events with significantly small or large parameter values may be placed in "special" groups, segregated from the general event population. The tables may be reconfigured in-flight by simply changing the variables controlling operation of the SYST table writing routine.

The discrete polynomial approximations contained within the tables were obtained using a novel combination of local (single group) and global curve-fitting algorithms. The resulting approximations are globally consistent in the fashion required by the SAR-based classification scheme.

VII. Conclusion

Clearly techniques similar to those described here are necessary to reduce the enormous quantity of sensor-generated data and maximize the utilization of the available telemetry bandwidth. Although other on-board analysis and reduction schemes might be considered, the M vs. M/Q classification technique utilized by the VIKING and CRRES DPUs has proven to be quite successful, and reasonably simple to implement. Slight modifications to the existing algorithms are anticipated for the forthcoming CAMMICE DPUs, including the addition of several new options to the SYST table writing routine to allow further specialization of mass and M/Q groups. The curve-fitting routines are also being refined to provide additional tools for evaluating the "goodness" of the resultant classification scheme. In general, the proposed modifications seek to add greater flexibility to the existing algorithms, without altering the essence of these flight-tested techniques.

Acknowledgments

The authors wish to thank W.R. Crain, S.S. Imamoto, D.T. Katsuda, M.T. Marra, and W.J. Wong for their contributions to the design, fabrication, and testing of the DPUs, and also to acknowledge the contributions of D.L. Chenette and J. Jennings, who participated in the development of earlier versions of the software.

References

1. R. Koga, C.G. King, N. Katz, S. Imamoto, J.F. Fennell, D.L. Chenette, and J.B. Blake, "A Data Processing Unit (DPU) for a Satellite-Borne Charge Composition Experiment," *IEEE Transactions on Nuclear Science*, NS-32, 163-167, 1985.
2. R. Koga, S.S. Imamoto, N. Katz, and S.D. Pinkerton, "Data Processing Units for Eight Magnetospheric Particle and Field Sensors." To appear, *Journal of Spacecraft and Rockets*, 1991.
3. P.R. Bevington, *Data Reduction and Error Analysis for the Physical Sciences*, New York: McGraw-Hill, 1969.

TECHNOLOGY OPERATIONS

The Aerospace Corporation functions as an "architect-engineer" for national security programs, specializing in advanced military space systems. The Corporation's Technology Operations supports the effective and timely development and operation of national security systems through scientific research and the application of advanced technology. Vital to the success of the Corporation is the technical staff's wide-ranging expertise and its ability to stay abreast of new technological developments and program support issues associated with rapidly evolving space systems. Contributing capabilities are provided by these individual Technology Centers:

Electronics Technology Center: Microelectronics, solid-state device physics, VLSI reliability, compound semiconductors, radiation hardening, data storage technologies, infrared detector devices and testing; electro-optics, quantum electronics, solid-state lasers, optical propagation and communications; cw and pulsed chemical laser development, optical resonators, beam control, atmospheric propagation, and laser effects and countermeasures; atomic frequency standards, applied laser spectroscopy, laser chemistry, laser optoelectronics, phase conjugation and coherent imaging, solar cell physics, battery electrochemistry, battery testing and evaluation.

Mechanics and Materials Technology Center: Evaluation and characterization of new materials: metals, alloys, ceramics, polymers and their composites, and new forms of carbon; development and analysis of thin films and deposition techniques; nondestructive evaluation, component failure analysis and reliability; fracture mechanics and stress corrosion; development and evaluation of hardened components; analysis and evaluation of materials at cryogenic and elevated temperatures; launch vehicle and reentry fluid mechanics, heat transfer and flight dynamics; chemical and electric propulsion; spacecraft structural mechanics, spacecraft survivability and vulnerability assessment; contamination, thermal and structural control; high temperature thermomechanics, gas kinetics and radiation; lubrication and surface phenomena.

Space and Environment Technology Center: Magnetospheric, auroral and cosmic ray physics, wave-particle interactions, magnetospheric plasma waves; atmospheric and ionospheric physics, density and composition of the upper atmosphere, remote sensing using atmospheric radiation; solar physics, infrared astronomy, infrared signature analysis; effects of solar activity, magnetic storms and nuclear explosions on the earth's atmosphere, ionosphere and magnetosphere; effects of electromagnetic and particulate radiations on space systems; space instrumentation; propellant chemistry, chemical dynamics, environmental chemistry, trace detection; atmospheric chemical reactions, atmospheric optics, light scattering, state-specific chemical reactions and radiative signatures of missile plumes, and sensor out-of-field-of-view rejection.