LA-UR- 03-2892

Approved for public release; distribution is unlimited.

c.1

Title: An Intelligent, Onboard Signal Processing Payload Concept

Author(s):
Patrick Shriver
Jayashree Harikumar
Scott Briles
Maya Gokhale

LOS ALAMOS NATIONAL LABORATORY
3 9338 01045 2539

Submitted to: SPIE AeroSense Conference 2003
April 21 - 25, 2003
Orlando, FL

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Form 836 (8/00)

An Intelligent, Onboard Signal Processing Payload Concept

Patrick Shriver, Jayashree Harikumar, Scott Briles, and Maya Gokhale

Los Alamos National Laboratory, NIS-3 Space Data Systems Group, Los Alamos, NM

ABSTRACT

Our approach to onboard processing will enable a quicker return and improved quality of processed data from small, remote-sensing satellites. We describe an intelligent payload concept which processes RF lightning signal data onboard the spacecraft in a power-aware manner. Presently, onboard processing is severely curtailed due to the conventional management of limited resources and *power-unaware* payload designs. Delays of days to weeks are commonly experienced before raw data is received, processed into a human-usable format, and finally transmitted to the end-user. We enable this resource-critical technology of onboard processing through the concept of Algorithm Power Modulation (APM). APM is a decision process used to execute a specific software algorithm, from a suite of possible algorithms, to make the best use of the available power. The suite of software algorithms chosen for our application is intended to reduce the probability of false alarms through post-processing. Each algorithm however also has a cost in energy usage. A heuristic decision tree procedure is used which selects an algorithm based on the available power, time allocated, algorithm priority, and algorithm performance. We demonstrate our approach to power-aware onboard processing through a preliminary software simulation.

Keywords: power-aware, onboard processing

1. INTRODUCTION

Our objective is to maximize capability and performance of remote-sensing satellites through smarter use of processing resources. Technological advances in instrument designs provide increases in resolution at a cost of transmission power. A byproduct this of increased resolution is the generation of large volumes of data. The more data generated, the more power required to transmit the data. Large power requirements in turn increase the spacecraft financial costs and "design-to-launch" time. Onboard compression and feature extraction can be used to reduce the data rates required for transmission, however, traditional power management strategies make significant onboard processing impractical. Conventional spacecraft architectures only collect raw data and leave processing tasks to be performed on the ground. By enabling onboard processing, it is hoped to quickly provide accurately processed, field-usable data to the end-user. However, in this resource critical problem, a smarter distribution of power resources is essential. We are developing a power management design paradigm for intelligent, next-generation processing payload systems through the unique concept of Algorithm Power Modulation (APM). With APM, a spacecraft payload processor can operate at multiple levels of power consumption in a power-aware manner while remaining sensitive to the mission tasks. Through judicious selection of the processing algorithms, data can still be processed onboard at some level of accuracy even during periods of low power availability. In this paper, we describe a power-aware processing payload concept that utilizes APM to process simulated orbital lightning signal data. This application is focused on improving detection capability through accurate, post-trigger processing.

Based on our payload concept, we have developed a simplified software simulation to demonstrate APM. The simulation consists of three main modules: battery model, event rate model, and the APM decision scheduler. These modules were incorporated and run through a 36 minute eclipse orbit test case. APM is compared to a more non-intelligent approach that just executes one specific algorithm. Our preliminary results show that APM drains more battery capacity than the benchmark, one-algorithm cases, but APM balances the power used with parameter estimation accuracy and detector performance while minimizing the number of lost events.

2. BACKGROUND AND MOTIVATION

We are currently seeing spacecraft system-related complications in implementing cutting-edge, advanced remote-sensing instrument payloads. Continual improvements in instrument technologies are providing ever-increasing resolution accuracies that result in growing demands of required transmitter power. The typical remote-sensing satellite has only a finite number of ground contacts with a limited duration over each contact. Hence, the high volume of raw data generated by these advanced instruments translates directly into high data rates during downlink opportunities. For example, one such advanced instrument class is the Hyperspectral Imagery (HSI) sensors. HSI instruments collect both spatial and spectral information generating hundreds of megabytes to gigabytes of data within a few minutes.[1] These massive amounts of data push the limits of existing system technology. As defined by the communication link equation, there is a direct cost relationship in required transmitter power for a given data rate.[2] A ten-fold increase in the data rate requires a subsequent ten-fold increase in the power for transmission.

From the systems design perspective, power is a significant factor in the physical size and resultant cost of the mission. Spacecraft within the orbit of Mars typically operate on photovoltaic power in combination with rechargeable, secondary batteries*.[3] Available power to spacecraft systems is therefore constrained by the amount of electricity generated by photovoltaics and the capacity that can be stored in secondary batteries. Increased solar panel surface area provides increased power, but also contributes to increased physical size. Likewise, more batteries provide additional power at a cost of increased mass. As a general rule, the larger the spacecraft, the higher the financial costs and the longer the "design-to-launch" time. Although a difficult metric to measure, due to a wide variety of mission and performance requirements, an average spacecraft cost for Low Earth Orbit (LEO) missions is estimated at $1,000 per kilogram.[4] In remote-sensing applications, there are two possible approaches to onboard processing techniques that will reduce the amount of data rates required for downlink: data compression and feature extraction.

The first, and more traditional, approach is to compress the science data as much as possible. There are two primary costs associated with this approach, which include the loss of information as a byproduct of the compression algorithm and the computational effort required to perform the compression. The objective in this case is to maximize the compression ratio that also minimizes or provides an acceptable loss of information in the data. One such method describes a HSI compression algorithm that enables a 100-fold increase in instrument coverage *without* an increase in required transmission power but has a corresponding 100-fold increase in computational cost.[5] Although this first approach does promise to decrease the data size for downlink, it does not address the operational problem of data latency. The data must still be uncompressed and processed on the ground after transmission. Delays of days to weeks can be experienced before the data is collected, transmitted, processed into a human-usable format, and given to the end-user.

A second approach is to extract pertinent features of interest from the data and downlink only the selected portions of the data†. As a well-known spacecraft design example, this issue is discussed in relation to the conceptual forest fire detection spacecraft known as *FireSat*.[3] Observational data gathered over the Earth's polar regions is obviously unnecessary. Storing the useless polar cap data uses up valuable onboard storage space and contributes to latency. A person on the ground must sift through the entire data set and filter out the polar cap data. Processing the data onboard and selecting only that which is of interest can significantly reduce data rates and ground operational costs, as well as data latency.[4]

An important, innovative outcome of feature extraction is onboard science analysis. The onboard analysis can extract and process useful information from the observational data. This information can then be used to improve detection performance or reconfigure a detector based on the incoming data. The algorithm performance characteristics are important considerations in the selection of algorithm to use. In one study of three possible onboard HSI detection algorithms, the algorithm performance is evaluated in regard to false alarm detections versus the image size and computational effort.[1]

*Beyond the orbit of Mars, existing solar panel efficiencies are not practical for use. The required surface areas to operate beyond the orbit of Mars are too large. Instead, these spacecraft rely on Radio-Isotope Thermal Generators (RTGs) for power. Spacecraft using RTGs are not considered for the concepts presented in this paper.

†Note that these two approaches are not necessarily mutually exclusive. One can imagine a scenario in which relevant features are selected *and* compressed to significantly reduce the required transmission power.

However, during operations, most spacecraft do not utilize significant onboard processing and typically store the raw data for downlink, leaving difficult trade decisions in terms of data collection and transmission.[6] Both positive and false-positive detections are saved until the data can be downlinked. This can be wasteful of the onboard data storage resources and limit detection capability performance. In part, this paradigm has been due to concerns within the scientific community of maintaining control over all aspects of the data processing. Although, if the processing methodology is known and well-behaved, or there exists a method to reconstruct events based on a few extracted parameters, the potential advantages of reduced costs, decreased latency, and improved detection performance from onboard processing should be seriously considered.

This operational paradigm has also been due to a traditional approach of power management. Electrical power is handled onboard using static, non-intelligent designs. The conventional strategy is a basic binary approach: a subsystem is either on or off. Simplified, onboard logic protects against power system faults. During severe power-constrained or unexpected conditions, such as eclipse or partial solar panel failure, non-critical subsystems may be turned off altogether. Under such circumstances, payload subsystems are typically the first to be turned off. Although essential to the mission objective, the payload is not necessary for spacecraft survival. This operational paradigm can result in a loss of vital data at a crucial moment.[7]

Although intelligent power management has been a cutting-edge research topic in recent years for ground-based, mobile computing systems, it still remains largely unexplored for spacecraft applications. Existing satellite power management strategies are not flexible enough to take full advantage of these emerging processing technologies. Advances in electronics technology have produced multiple modes of processing operation. A processor can operate in different modes at varying levels of power consumption. For example, the IBM PPC750 266MHz processor has 5 typical modes of operation that range from 30mW to 5.7W.[8] Most notably, multiple operational modes are commonplace in ordinary laptop computing to conserve battery power while disconnected from an external power outlet. A recent area of power research, sponsored by the Defense Advanced Projects Agency (DARPA) Power Aware Computing/Communications (PAC/C) program, is focused on developing intelligent, power-aware systems.[9] This research is based on the premise that systems aware of their own power usage can make better use of the available power resources. A system aware of its power consumption can dynamically adjust workload tasks to change the operational mode. Thus, with power-aware techniques, it is conceivable for a processing payload to process data onboard, at some level of accuracy, even during periods of low power availability[‡].

Our research focuses on developing and demonstrating an onboard processing payload concept that executes signal processing tasks *in-situ* based on the available resources while remaining sensitive to the mission objectives. This is a needed paradigm shift in conventional power management methods to enable advanced, intelligent processing payload systems. We have focused in the area of feature extraction through accurate parameter estimation and have determined a method to improve detection performance.

3. PAYLOAD CONCEPT

With support from the DARPA PAC/C program, we have developed an intelligent onboard processing payload concept using an open application familiar to LANL staff with a potential impact on next-generation payload designs. It is our vision that APM will help enable onboard processing technologies for future spacecraft. The application selected is the Fast On-Orbit Recording of Transient Events (FORTÉ) satellite mission. FORTÉ, funded by the Department of Energy, was launched in August of 1997 and is a joint LANL project with Sandia National Laboratory. Our concept is described in literature,[10] but for convenience, we briefly summarize the concept in this section.

3.1. FORTÉ Mission Application

A primary objective of FORTÉ is to detect the Radio-Frequency (RF) signal of lightning events in the Earth's atmosphere.[11] As received on-orbit, the RF lightning signal is a "chirp" waveform amidst a noise of anthropogenic signals and background cosmic ray particles. The chirp signal is a result of the frequency dispersion

[‡]It should be noted that *power-aware* does not necessarily imply *power-efficiency*. Instead, *power-aware* systems are concerned primarily with the *maximum utilization* of available power resources, which, depending upon the application, may or may not include *power minimization*.

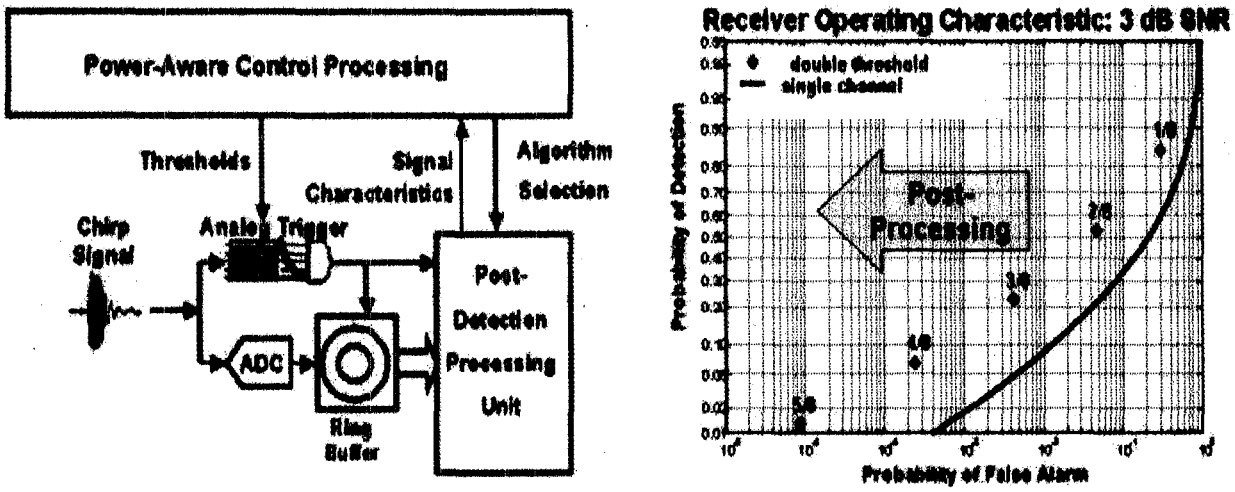

Figure 1. Payload Concept: *The diagram on the left represents the conceptual data flow diagram of processing the chirp signals.* Post-trigger processing algorithm decisions are to be made based on the available resources. The chart on the right displays the trigger box Receiver Operating Characteristic Curve. The post-trigger processing can reduce the probability of false alarms.

experienced during propagation through the ionosphere. An analog trigger box provides multiple channels of sub-band filters that attempt to detect the presence of a lightning event. A detection trigger occurs when N of M channels break threshold to satisfy the predetermined criterion. FORTÉ does not have the capability to process this data onboard and, hence, stores only raw data for downlink. Since the threshold criterion is preset, the receiver's operating point remains fixed and a certain probability of false alarms must be accepted for a desired probability of detection. As discussed in the next section, post-trigger processing techniques can further reduce the amount of false alarms through increased expenditure of energy. This application concept is depicted in Figure 1.

3.2. Algorithm Power Modulation Solution

Our approach to this problem has been to develop a suite of signal processing algorithms that can be run on a multi-processor system. The algorithms can be executed independently to estimate the parameters of Total Electron Content (TEC) and Time-Of-Arrival (TOA) from simulated chirp signals. Each algorithm has an associated level of estimation accuracy and energy consumption. The chosen algorithms include a Least-Mean-Squares (LMS), Maximum Likelihood (ML), Software Trigger (ST)[§], and a bank of Matched Filters (MF). An algorithm power experiment was performed on a PPC750 266MHz test-bench provided by the Jet Propulsion Laboratory as part of the PAC/C effort. The result was a 10^6 order of magnitude difference in energy usage between the four algorithms. These four algorithms have been exercised via Monte Carlo testing with the simulated signals. Using the Root-Mean-Squared (RMS) error as a metric of performance, these performance values were correlated with the energy measurements and outline a decaying exponential profile with an increase in energy expended.[12]

Through further analysis, the relationship between energy usage and reduction in the probability of false alarms for each algorithm was determined for a given operating point of the trigger box.[13] The results of this analysis are depicted in Figure 2. It is shown that the MF algorithm can reduce the probability of false alarms to near zero with a significantly increased cost in energy usage.

[§]The ST algorithm performs multiple, short FFTs on the signal to estimate TEC and TOA.

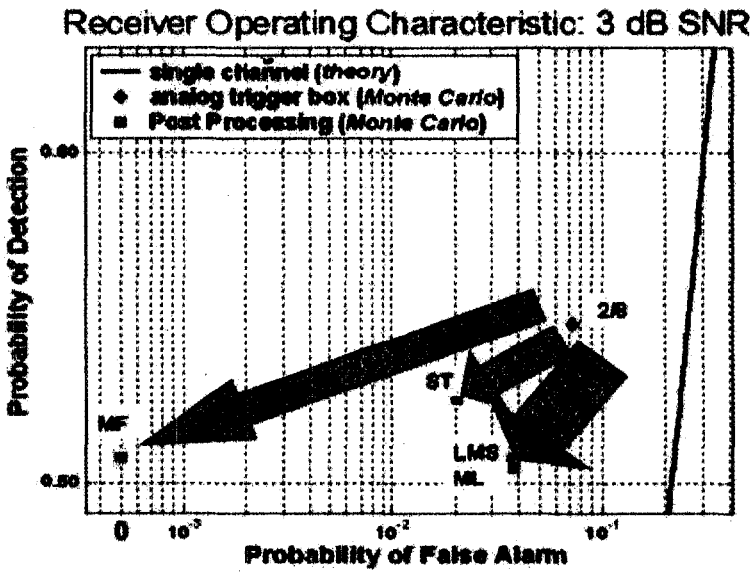

Figure 2. Probability of False Alarm Reduction: This chart illustrates the performance of the signal processing algorithms to reduce the probability of false alarms through expenditure of energy.

4. SIMULATION DESCRIPTION

The simulation is accomplished through a two step process. The first step simulates the FORTÉ orbit using the comercially available *Satellite Tool Kit*®program. In this step, the FORTÉ latitude and longitude locations, for a given time period, are fed into a lightning event rate model to determine the event rate at a specified latitude and longitude. The second step simulates the payload operation given the event rates, which is written in *C++*. As a first-generation of the simulation, only one processor is considered in the payload operation. The operation algorithm steps through the event rates and power is consumed from the battery model based on the chosen algorithm properties. Since, in spacecraft, the highest typical resolution of monitoring the battery state is one second, power is only drawn from the battery every second in the simulation. Additionally, only one algorithm is executed during this interval. For evaluation and comparison purposes, the user has the option of implementing the APM decision scheduler (see section 4.3) or to specify one of the LMS, ML, ST, or MF algorithms to execute. An average rate of discharge over this time step is determined based on the number of times a given algorithm executes and the corresponding charge consumption. When an algorithm is not executing, charge is drained at an operational background rate of 1800mA, as observed in the JPL power experiment (see section 3.2) when the algorithms were not executing on the PPC750 processor. A ring buffer module keeps track of the number of events residing in memory. The maximum capacity of the simulated ring buffer 150Mb[¶]. The following subsections briefly discuss the main modules of the simulation, which include the event rate model, battery model, and the APM decision scheduler.

4.1. Lightning Event Rate Model

Lightning occurrence and lightning type is a function of terrain, concentrations of the cloud nucleus condensation, and temperature. Lightning can traverse between clouds, from cloud to ground and from ground to cloud. The data used in developing the lightning model discussed in this paper does not distinguish between these different types of lightning. The data used in this simulation was recorded by the optical sensor on FORTÉ. As a result, some lightning may be missed because of clouds in the field of view. The primary emphasis on FORTÉ is to

[¶]FORTÉ has a variable sampling rate of the lightning chirp signals. In our simulation, we have chosen a typical chirp size of 150Kb. The simulated ring buffer can therefore hold a maximum of 1,000 events.

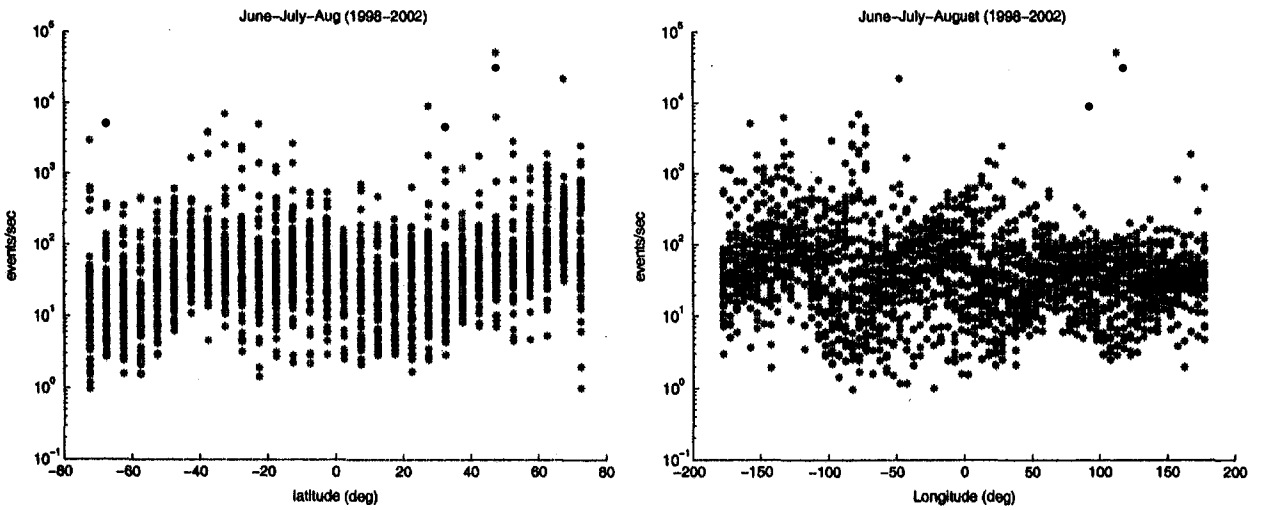

Figure 3. Northern Hemisphere Summer Event Rates: This figure shows the number of FORTÉ optical sensor triggers as a function of latitude and longitude for the Northern Hemisphere Summer season.

measure impulsive electromagnetic pulses due to lightning within a noise environment dominated by continuous-wave carriers such as television and FM stations.[11] The optical sensors on FORTÉ augment the RF system on FORTÉ in characterizing lightning events.

The satellite optical sensor is a fast-time response photodiode detector with a 80° field of view. Since FORTÉ is in a near-circular orbit at ~825km altitude, the optical sensor viewing area is ~1200km of Earth's diameter. The photodiode detector is amplitude-threshold triggered with a noise riding threshold. A lightning event is assumed to have occured when a trigger is recorded by the sensor. When trigger rates exceeded ~500/sec the presence of high energy particle hits was assumed and lightning signals are assumed non-existent in our model.

The lightning event rate model for our current application was developed from the FORTÉ lightning data set that comprises fifteen months of data between 1998 and 2002. The monthly variation in lightning has been summarized 3 months at a time to produce 4 sets of seasonal global lightning activity. December, January, February was designated as winter; March, April, May constitute spring; June, July, August represent summer; September, October, November compose lightning activity in Fall[‖]. Principal component analysis was next applied to detect if there was a predominant lightning pattern data set. The current data did not reveal any dominant pattern. This supports the basic claim that lightning is a function of season and terrain. Figure 3 show the lightning events (triggers recorded by the photodiode detector) as a function of latitude and longitude for the Northern Hemisphere Summer. A subset of these event rates, reflecting one orbital eclipse period during Summer, is used in the preliminary simulation results.

4.2. Battery Model

The battery model is based on the Maxell ICR18650G manufacturer cell data.[14] This is a Lithium-Ion cylindrical cell with a nominal voltage of 3.6V and nominal capacity of 1700mAh. The following considerations were made in developing the battery model:

discharging The discharge curves are based on the discharge characteristics given in the manufacturer data. Linear interpolation and extrapolation is used to determine points not on the given curves. Cells are nonlinear in nature, but, for our work, a model based on linear interpolation and extrapolation of the manufacturer data can yield reasonable accuracy.[15]

‖Note that these seasons correspond to those in the Northern Hemisphere.

charging The charging characteristics are neglected in this first version of the simulation. Normally, the spacecraft batteries are recharged during sunlight periods immediately following the eclipse periods. The amount of battery capacity at the beginning of an eclipse period is therefore a function of the capacity drained during the previous eclipse period and the amount charged during sunlight. Our simulation assumes a maximum capacity at the beginning of the eclipse period.

cycle life The long-term effects of discharge-charge cycles are neglected. At this time, the simulation is not used to evaluate battery operation life.

temperature Temperature has a significant impact on battery performance. As temperature decreases, the battery capacity during discharge also decreases.[16] However, for simplification, temperature effects are neglected.

Table 1. Battery Sizing Parameters

PPC750 Nominal Voltage	2.5V - 2.7V[8]
PPC750 Typical Full-On Power Consumption	5.7W[8]
Maximum Eclipse Duration	35min.[17]
Cell Depth-Of-Discharge (DOD)	80%
Required Battery Capacity	1600mAh

Using the cell data, the battery was sized according to the power requirements of the PPC750 processor and the estimated orbital eclipse duration. We only consider power consumed by the microprocessor. It is understood there would also be memory, cache, voltage supply, etc., in order to make practical use the processor, but, for our preliminary results, we are neglecting the power consumption and inefficiencies of these components. Table 1 lists the parameters involved in sizing the battery. The required battery capacity is obtained by dividing the average capacity drained during eclipse, 1280mAh, by the DOD. This is the capacity of the battery required to meet the processor full-on power consumption during eclipse. Since one cell provides a nominal voltage of 3.6V and 1700mAh,[14] only one cell is necesary per processor.

The model takes, as input, the increment of capacity drained and the discharge rate over the one second simulation time step (as described in section 4). The capacity increment is subtracted off from the total battery capacity. The battery voltage is then interpolated or extrapolated from the total battery capacity and discharge rate. The model then returns both the total capacity left and the voltage level. Figure 4 illustrates the battery model under constant rates of discharge. Qualitatively, these curves are typical of battery properties[16] and match those of the manufacturer data. Since we do not have access to this battery cell and a discharge testbench, a quantitative validation of this model is not possible.

4.3. APM Decision Scheduler

Initially, we have taken a heuristic approach to developing the decision scheduler that selects a given algorithm to execute. The decision procedure first determines a set of algorithms that can execute in the available power over the next one second interval. For each of the algorithm properties, a comparison is made of the predicted battery capacity drained over the next interval and the resultant voltage. If the predicted capacity does not exceed the maximum battery capacity (1700mAh) or the end voltage (3V), the algorithm is placed in the possible set of choices. A second decision determines which of the algorithms in the set can execute under the time constraint predicted by the amount of events in the ring buffer over the next interval. If the algorithm is in the set and the predicted number of events is:

0%-25% of the maximum capacity, run MF

25%-50% of the maximum capacity, run ST

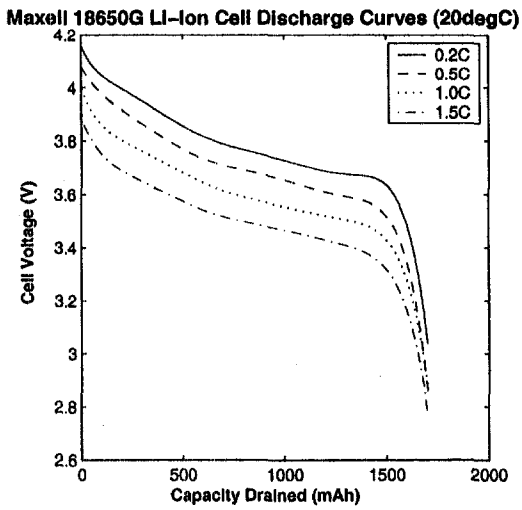

Figure 4. Battery Model Discharge Curves: This figure illustrates the battery model discharge properties under constant rates of discharge. Note that the temperature is constant at 20°C.

50%-75% of the maximum capacity, run ML

75%-100% of the maximum capacity, run LMS

If there is not enough predicted battery capacity or voltage, no algorithms are executed and the processor is placed in the "idle" state (1800mA). In this case, the ring buffer will continue to fill up with events for a non-zero event rate.

5. SIMULATION RESULTS

The simulation is run through five test case scenarios: four of which execute only one specified algorithm and the fifth scenario uses APM (see section 4.3) to switch between the algorithms. Each test case uses the same event rate data and simulates a FORTÉ eclipse period of approximately 36 minutes. During this time, the event rates range from approximately 2 events/sec to 182 events/sec. As mentioned in section 4, the battery is drained at every one second interval, so an average discharge rate is calculated over this one second time step. The discharge rate depends upon the number of executions each algorithm performs in this time step. The number of events processed also depends on the maximum number of possible executions. Table 2 depicts the time duration of each algorithm and corresponding number of maximum possible executions within a one second interval. The time duration of each algorithm was determined during the power measurement test on the Jet Propulsion Laboratory test-bench.[10]

Table 2. Algorithm Duration Properties

Algorithm	Execution Duration	Maximum Number of Executions Possible
LMS	3.4μs	294,117
ML	183μs	5,464
ST	8.34ms	119
MF	470ms	2

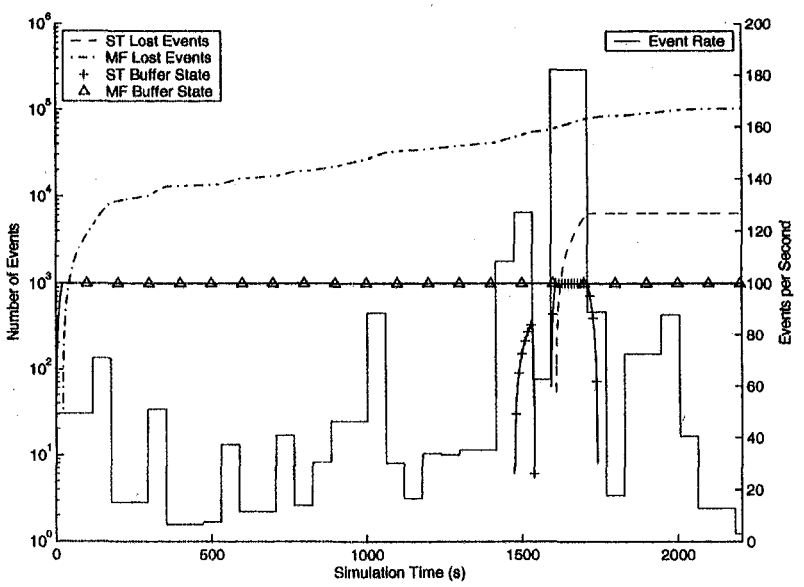

Figure 5. Event Processing Performance of ST and MF Algorithms: This figure shows the performance of the ST and MF algorithms in processing events in the ring buffer. Both the number of events in the ring buffer and the events lost over the simulation time are shown. For comparison, a plot of the event rate is also illustrated.

In comparison to the data in Table 2, it is apparent that the LMS and ML algorithms should be able to process *all* events as they occur in the one second interval. However, the ST and MF algorithms will have difficulty processing events in this interval. ST will begin to have difficulty when the event rate exceeds 119 events/sec. At this point, events will begin to fill memory space in the ring buffer. MF should have difficulty during the majority of the simulation run since. Events are counted as "lost" when the number of events in the ring buffer exceeds 1,000**. It should be noted that APM executes the faster LMS and ML algorithms when the number of events in the ring buffer exceeds 500 (see section 4.3).

Performance of the ST and MF algorithms in processing events is illustrated in Figure 5. As expected, the MF algorithm cannot keep up with the high initial event rate and begins to lose events almost immediately. At most, MF can only process two events in the one second interval. The ST algorithm provides better performance until the event rate exceeds 119 events/sec. At this time, events begin to fill up in the ring buffer. A brief drop in the event rate near 1500 sec allows the ST algorithm to process all events. However, a sharp rise in the event rate to 180 events/sec causes the ST algorithm to once again fall behind in processing and start losing events.

The number of executions of a given algorithm has an effect on the capacity drained and discharge rate over the simulation time step. For a given algorithm, the greater the number of executions, the more capacity drained, and the higher the discharge rate. Figure 6 depicts the discharge rate versus the capacity drained during simulation. Both LMS and ML have relatively small changes in the discharge rates due to their fast execution times. The MF algorithm has no change in its discharge rate since it always executes twice in the one second time step. By contrast, the ST algorithm causes large changes in the discharge rate due to its varying rates of execution and power consumption.

From the results of Figure 6, it is expected that the LMS, ML, and MF during discharge will exhibit relatively smooth voltage profiles while ST should exhibit visible changes in the voltage levels. Additionally, APM should exhibit more abrupt changes in voltage to reflect the switch between different algorithms. The simulation does exhibit these properties as illustrated in Figure 7.

**Note that the simulation assumes that once an event is processed, it is removed from the ring buffer. Thus, the ring buffer is simulated as a *temporary* memory buffer for the incoming data sets.

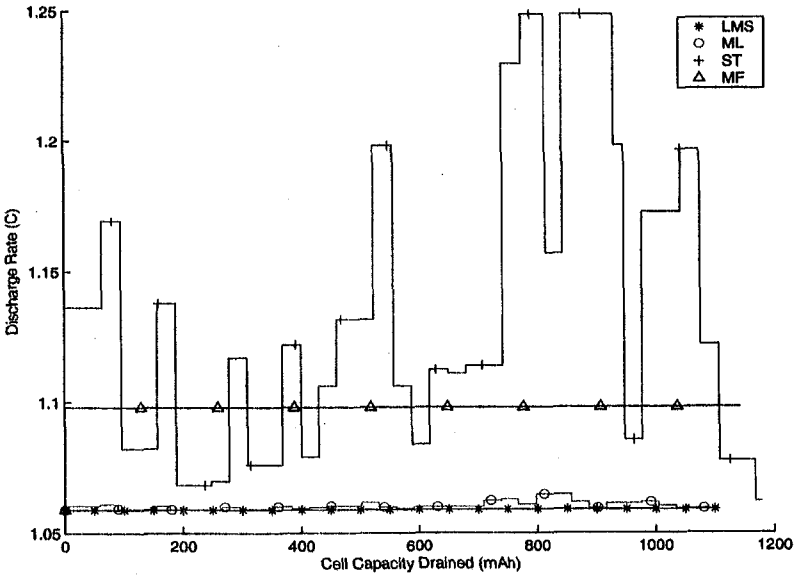

Figure 6. Discharge Rates: This figure depicts the discharge rates of LMS, ML, ST, and MF algorithms. The ST algorithm shows the largest fluctuation in discharge rate.

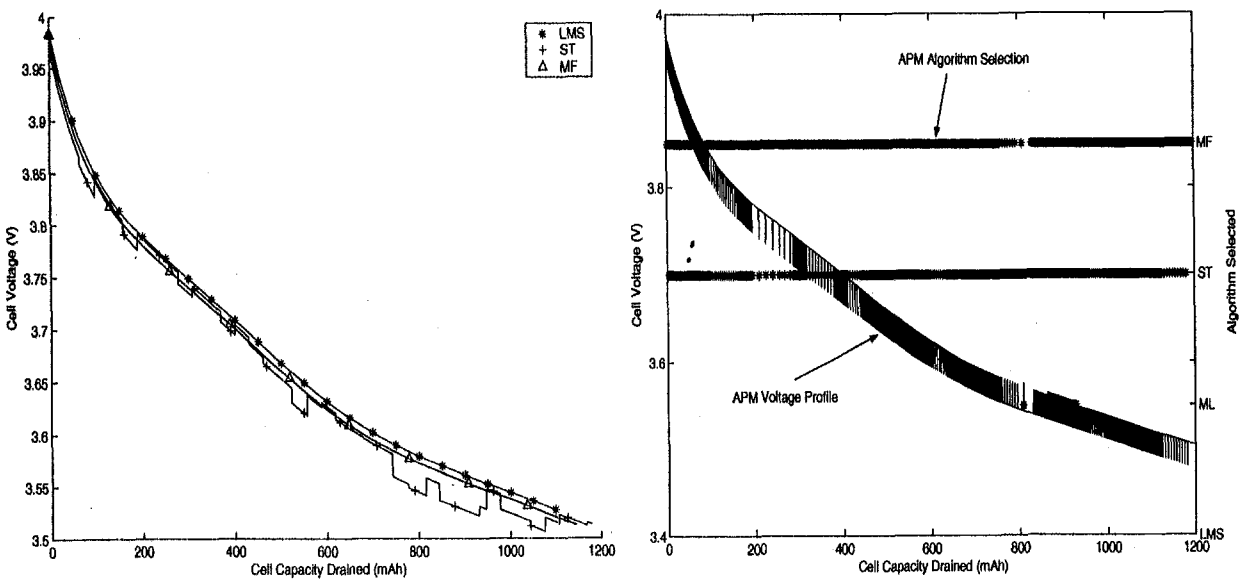

Figure 7. Voltage Profile: This figure illustrates the voltage profiles versus the capacity drained during simulation. Note that the chart on the right also illustrates APM switching between ML, ST, and MF algorithms. Due to the fast execution of ML, APM did not need to execute LMS.

One might expect that, since APM switches between algorithms, the capacity drained by APM would be in a mid-range when compared to the capacity drained by all the other algorithms. Through examination of Figure 7, one can see that APM actually consumes more capacity than any of the other algorithms. This is due to APM predominantly running the ST and MF algorithms. When the ST algorithm executes approximately 25 times in the one second time step, the capacity drained is equivalent to that drained by executing MF twice in the time step. When ST executes more than 25 times, the capacity drained is greater than the drain experienced by executing the MF only twice. Likewise, when ST executes less than 25 times, the capacity drained is less than that drained by executing MF twice. Strictly in terms of capacity drained, APM appears to select the worst case, i.e., it selects ST when ST executes enough times to drain more capacity than MF and selects MF when MF drains more capacity than ST. Thus, APM drains more capacity than either ST or MF alone by selecting them during these worst-case conditions[††].

This provides an example of how a *power-aware* system is concerned with the *maximum utilization* of the available power resources and not necessarily concerned with *power minimization.* The benefit APM provides is in the accuracy of the estimated parameters, the reduction in false alarms, and the number of processed events. Since APM primarily uses MF and ST algorithms, the parameter estimation accuracy is greater than running LMS or ML alone. Through running MF and ST, APM also reduces the probability of false alarms. This helps to reduce the number of false alarms that will wastefully use onboard memory space and improve the quality of data downlinked to the ground. Additionally, APM does not loose any events as with the MF or ST algorithms alone. Thus, APM does make better *use* of the available power in processing events.

6. CONCLUSIONS

As instrument technology advances, the need for onboard processing is becoming more apparent. Onboard processing can help to reduce the required high data rates, and thereby mission costs, reduce data latency, and support onboard science analysis to provide better quality of data to the end-user. However, most spacecraft do not have the capability to provide significant onboard processing in large part due to the conventional power management paradigm. We seek to enable onboard processing through development of intelligent payload systems that can dynamically adjust their operational mode by being aware of the currently available power resources. We use the concept of APM to select a signal processing algorithm, from a suite of possible algorithms, based on the algorithm performance, incoming event rate, and available power.

We have built a simplified software simulation based on the mission application of the FORTÉ satellite to study lightning events in the Earth's atmosphere. The objective of our processing concept is to process lightning events as received on-orbit. The simulation is composed of three main modules including the battery power source, event rate, and decision scheduler. We ran the simulation through two types of test case scenarios. The benchmark case is in executing a single algorithm to process the lightning signals. It is shown that APM consumes more battery capacity than in the benchmark cases, but, unlike the benchmark cases, APM balances power *utilization* with improved parameter estimation and enhanced detection performance while minimizing the number of lost events.

With APM, a satellite onboard payload processor is sensitive to the mission tasks and power availability. The payload processor operates in a power aware manner by determining the status of the battery capacity and the size of the data ring buffer before processing an event. Both, the ring buffer and the battery have a finite capacity and consequently decisions to process events with a particular algorithm is a function of both parameters. If events are not processed quickly, the data in the ring buffer can be overwritten and lost. At the same time, the battery charge level is a function of time and tasks performed. A battery can reverse polarity and destruct if it goes beyond a specified end voltage. Under such conditions on a spacecraft, the fault protection system usually switches off the subsystems to protect the battery.

[††]Although APM drains more capacity, it does not go below the specified end voltage or rated capacity. APM still operates within the safe limits of the battery.

ACKNOWLEDGMENTS

This effort is sponsored by Defense Advanced Research Projects Agency (DARPA) through the Air Force Research Laboratory, USAF, under agreement number F30602-00-2-0548. We are grateful to Mr. Raymond Zenick of AeroAstro, Inc. for providing the manufacturer battery cell information and to Mr. Philip Lyman of Ball Aerospace Corporation for providing guidance on development of the battery model. Thanks are also due to the LANL FORTÉ team for supplying details of FORTÉ operation and to Dr. Tracy Light for providing the lightning event rate data.

REFERENCES

1. S. M. Schweizer and J. M. F. Moura, "Efficient detection in hyperspectral imagery," *IEEE Transactions on Image Processing* **10**, pp. 584–597, 2001.
2. W. L. Morgan and G. D. Gordon, *Communications Satellite Handbook*, John Wiley & Sons, New York, New York, 1989.
3. J. R. Wertz and W. J. Larson, eds., *Space Mission Analysis and Design*, Microcosm Press, El Segundo, California, 3rd ed., 1999.
4. J. R. Wertz and W. J. Larson, *Reducing Space Mission Cost*, Microcosm Press, El Segundo, California, 1996.
5. S. Cook and J. Harsanyi, "Real-time data processing onboard remote sensor platforms," in *Proceedings of the Earth Science Technology Conference*, 2001.
6. M. A. Figueriedo, P. H. Stakern, *et al.*, "Extending NASA's data processing to spacecraft." http://klabs.org/richcontent/Papers/NASA_Data_Processing.pdf, 1999.
7. Established through personal communication with satellite operators at Los Alamos National Laboratory, United States Air Force, and Laboratory for Atmospheric and Space Physics.
8. "PowerPC740 and PowerPC750 embedded microprocessor: Hardware specifications." http://www-3.ibm.com/chips/techlib/techlib.nsf/techdocs/8CAAF8D73C47F5B587256A0900630F73/$file/7xxp8t_em_hw2-99.pdf.
9. http://www.darpa.mil/ipto/research/pacc/index.html.
10. P. M. Shriver, M. B. Gokhale, *et al.*, "A power-aware, satellite-based parallel signal processing scheme," in *Power Aware Computing*, R. Graybill and R. Melhem, eds., pp. 243 – 259, Kluwer Academic/Plenum Publishers, 2002.
11. D. Roussel-Dupré, P. Klinger, *et al.*, "Four years of operations and results with FORTÉ," in *Proceedings of the AIAA Space 2001 Conference*, (AIAA-2001-4627), 2001.
12. P. Shriver, S. Briles, *et al.*, "A power-aware approach to processing payload design." To be published in Proceedings of the Government Microcircuit and Critical Technology Conference, 2003.
13. S. Briles, P. Shriver, *et al.*, "Power-aware improvement in signal detection." To be published in Proceedings of the International Signal Processing Conference, 2003.
14. "Maxell lithium ion rechargeable batteries." Product information brochure.
15. Established through personal communication with Mr. Philip Lyman, a power systems engineer with Ball Aerospace Corporation.
16. D. Linden and T. B. Reddy, eds., *Handbook of Batteries*, McGraw-Hill, New York, New York, 3rd ed., 2002.
17. V. L. Pisacane and R. C. Moore, eds., *Fundamentals of Space Systems*, p. 346. Oxford University Press, 1994.