Opening the Door to Smart Power Management in Small Satellites

Patrick Shriver
Ph.D. Candidate, Aerospace Engineering Sciences, University of Colorado at Boulder
Graduate Research Assistant, NIS-3 Space Data Systems Group, Los Alamos National Laboratory

Faculty Advisors: Professor Scott Palo and Professor Mark Balas
Aerospace Engineering Sciences, University of Colorado at Boulder

Abstract
The objective of this research is to develop a new power management paradigm that will maximize the capabilities of small satellites and therefore help provide cost-effective access to space. Ground-based, mobile processing systems have experienced a similar engineering problem as the small satellite and have developed successful power management methods to provide increased energy and cost savings with improved computing performance. It is envisioned that the new paradigm of smart power for small satellites can utilize these methods to increase autonomy and enable onboard processing. An onboard, processing payload concept and initial software simulation results are discussed.

1. Introduction
The ever-increasing demands and problems unique to the small satellite are presenting a new, complex challenge to the power management system. The engineering problem for small spacecraft is one of limited size, low mass, and tight financial budgets. In order to provide the greatest mission return for the least expense, there is a need to evolve a *smarter* power system that will have a more *autonomous*, *proactive* role in managing the limited power resources.

The engineering problem of the small satellite is mirrored in ground-based, mobile processing systems, particularly driven by the demands placed on laptop computing. System components have been designed with multiple power modes beyond the traditional on/off. On-line *software* control of these hardware states has led to energy and cost savings even with radical improvements in computing capability and performance. A new spacecraft power management paradigm can be developed that utilizes concepts drawn from these existing ground-based, processing technologies. As a major consequence, it is envisioned that such a system will also be able to enable significant onboard processing, which will increase efficient use of downlink bandwidth.

This paper summarizes the motivation behind developing a new power management paradigm; provides an overview of general power management methods and research; portrays a vision of the smartly powered, small satellite; describes work on a processing payload concept with multiple modes of operation and discusses preliminary, software simulation results of this payload concept; and outlines a direction of future work. It is the aim of this paper to open the door to smart power management in small satellites.

2. Motivation
The potential, technological contribution inherent to small spacecraft is the ability to achieve *quick and cost-effective access to space* [1], [2]. In general, small spacecraft have quicker "design-to-launch" times and can utilize smaller, rapidly deployable launch vehicles for faster mission returns. To provide the greatest mission return for the least expense, cost-effectiveness is a tradeoff between mission capabilities and financial resources. The challenge for the small satellite is to maintain cost-effectiveness by maximizing capabilities and performance while meeting difficult design and financial restrictions.

A primary limiting factor on capability and performance is the available power resources. Typically, more ambitious mission goals and enhanced bus capabilities will require more power. The spacecraft must also meet the predicted end of life conditions, which further increases the power requirements. However, there is a direct relationship between available power, size, and mass of the spacecraft. Additional battery cells add more mass; larger photovoltaic surface areas increase both physical size and mass. On the other hand, highly efficient components are usually more expensive and may not be commercially available to meet the desired mission objectives.

Existing, spacecraft power management systems are more *reactive* and operate with relatively little or no onboard intelligence. Careful planning and

significant support from ground-based personnel is required, but this introduces limitations due to communication latency and operations costs. It is predicted that a *smart* power management system can provide enhanced autonomy through awareness of the current power state and *proactive* control of onboard subsystems.

It is known from ground-based power management methods that energy and cost savings can be realized with *proactive* control of devices with multiple operational states. Since microprocessors are now designed with multiple operational modes, a *smart* power management system could enable significant onboard processing even for the power-limited small satellite. This can help increase more efficient use of downlink bandwidth.

The objective of this research is to develop a new approach to the spacecraft power management paradigm that will increase autonomous operations and enable small spacecraft to support advanced, processing payload systems. It is envisioned that this work will help push the envelope of small satellite capabilities and performance to maintain the cutting-edge of *quick and cost-effective access to space*.

3. Brief Overview of Power Management

The critical relationship between available power, size, and mass is not particular to small spacecraft, but also applies to ground-based mobile systems. As more capabilities are demanded from the system, more power is required. The larger and more massive the system, the less mobile the system becomes. Power management has been a cutting-edge area of research, particularly for computing systems such as laptops[1]. Most of the methods developed to date have focused on *minimizing power consumption*, but a recent area of research is also concerned with *maximizing utility* of available power. It is envisioned that, for the small satellite, a *smart power* management system can draw from these concepts to push the envelope of capability and performance.

[1] Laptops are the significant driving force of today behind the rapid development of sophisticated power management methods. This is in large part due to the high commercial market demands for both lighter and more powerful laptops. However, it is being recognized in research and industry that power management gains can benefit other mobile systems on land, sea, air, or space to improve autonomy and increase lifetime.

3.1 General Methods and Research

Regardless of the application, the power management system attempts to balance the power generation with the load consumption. The two main areas of power management are Static Power Management (SPM) and Dynamic Power Management (DPM) [3]. These two methods are not mutually exclusive and can be applied at all possible levels from the lowest, subsystem components to the higher, system-wide levels.

The SPM approach is applied off-line, either during system design or before a runtime condition to minimize power consumption. Techniques include software such as compilers that compile low power versions of program executables [4] as well as the design of power-efficient interaction between hardware components[2]. Since it is performed off-line, SPM is a conservative method that must satisfy the expected, worst-case power conditions.

In contrast, the DPM approach adjusts the system during runtime. DPM has had significant success in energy and cost savings in laptop and desktop systems while computing performance has continually increased [5], [6]. DPM is capable of managing system components that have multiple modes of operation in addition to the standard "on" and "off." Management methods are based on actual, observed, or assumed knowledge of the system and workload. DPM *policies*[3] are organized according to the following attributes [7], [8], [9]:

- *System Model*- System models[4] can be deterministic with known characteristics and responses, or stochastic with uncertain parameters.
- *Decision Frequency*- A key question is when to implement the policy decisions[5]. Decisions are

[2] An example of hardware SPM for processor technology is the interaction between microprocessor and cache. The microprocessor spends less energy accessing internal cache than other external memory resources. Increasing cache size is therefore an energy-saving form of SPM.

[3] *Policies* are the methods used to implement DPM.

[4] Note that policies are typically designed for a specific type of system model. Heuristic policies are used with deterministic models; and, hence, optimality is not guaranteed but investigated through comparison of different policies. However, stochastic-based policies can guarantee optimal results and provide a more general framework for the system model.

[5] Policy decisions are the choices of when to change the operational mode and which mode to change into.

implemented based on a system clock or on an event basis.

- *Policy Flexibility*- Non-adaptive policies remain fixed during runtime and assume workload conditions are known *a priori*. In contrast, adaptive policies adjust based on the observed workload conditions.

The most commonly used forms of DPM are *timeout*[6] policies, which are now commonplace in cellular phones and Personal Data Assistants (PDAs) [10], as well as desktop and laptop computers. Additional DPM methods include dynamic voltage scaling and processor frequency adjustment [11], but the full application potential for DPM techniques remains largely unexplored [7]. The underlying philosophy of methods developed thus far have been primarily based on minimizing power consumption of system components *without taking into account the status of the power source*[7].

A recently expanding area of power management research is in power-aware systems [12]. Primarily a DPM approach, the premise behind this research is that systems which are aware of their power state, both load demands *and* resource availability, can make better use of the available resources. With power-aware systems, it is possible to not only *minimize power* consumption, but also *maximize use*[8] of the available power. By being aware of the power source and desired objectives, appropriate decisions can be made whether to minimize power consumption, maximize utility, or operate at a tradeoff point between these objectives.

3.2 Existing Satellite Methods

Existing spacecraft power management technologies are based on the similar rationales of SPM and DPM, but have a less *proactive* role in managing the power resources. During design, worst-case conditions are assumed and, during operations, onboard management techniques are relegated to monitoring the power levels and maintaining safe operating thresholds of *over*- and *under*-powered conditions[9] [13], [14]. It is also assumed that subsystems have only two modes of operation: "on" or "off." Although important to the mission objectives, payloads are not crucial for spacecraft survival and are the first to be shut off during power-constrained conditions. Ground operator intervention is then required to perform a system restart.

The ground operations crew bears the more *proactive* role in managing the power resources, but is hampered by communication latency. There are limited opportunities for ground contacts, and the contacts tend to be of short duration. Thus, operators are not privy to continuous health status information and can only quickly scan a limited number of parameters to determine the immediate status of the spacecraft. As a result, ground operators typically issue commands based on delayed information. Additionally, essential supervision and management activities significantly increase with the number of spacecraft to support. Constellations of spacecraft require increased operator involvement and larger support staff.

A need for change in the power management paradigm is being recognized and supported within the small satellite community, but remains relatively unexplored. In [15], the need for an intelligent power system is outlined for a constellation of nanosatellites. A more integrated CubeSat power system is described in [16], which doesn't utilize Direct Energy Transfer or Peak Power Tracker systems. A smarter, more integrated approach to power management *is* possible with currently emerging technologies and new power management strategies.

---

[6] After exceeding a specified idle duration, the device is powered down into a lower power state.

[7] During implementation, battery voltage *is* monitored so that the end voltage condition is not exceeded. At which point, the system is shut down. If the battery has not reached the end voltage condition, policy knowledge is based solely on the device characteristics and load demands. A prediction, or awareness, of available power resources is not taken into account.

[8] Maximizing use is appropriate for certain instances in mobile systems. It can be thought of as maximizing utility or reward to achieve the operational objectives.

[9] This refers to the use of Direct Energy Transfer or Peak Power Tracking systems to shunt excess power or adjust the photovoltaic operating voltage to match load demands, respectively. Additionally, onboard logic and switching systems turn subsystems on/off to prevent draining the batteries beyond safe operational limits.

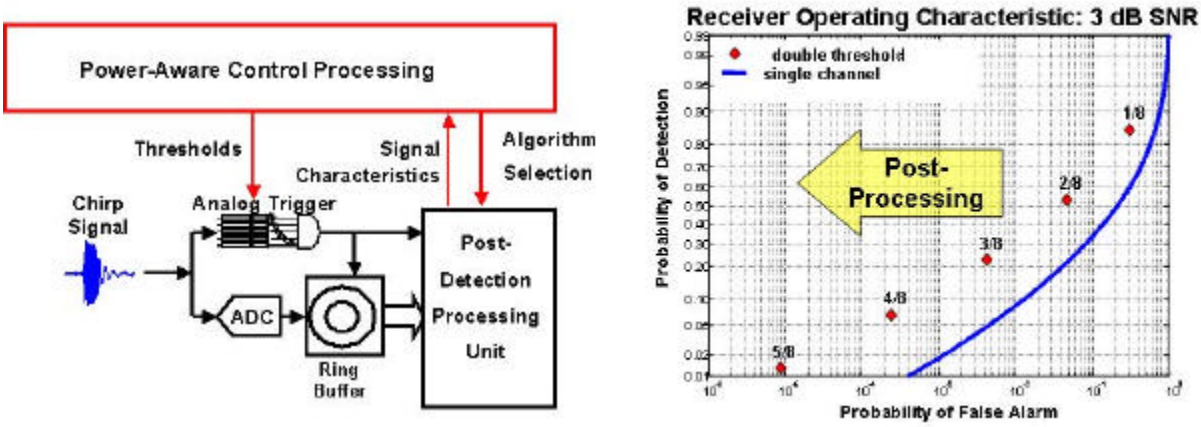

Figure 1. Payload Concept: The diagram on the left represents the conceptual data flow diagram of processing the chirp signals. Post-trigger processing algorithm decisions are to be made based on the available resources. The chart on the right displays the trigger box Receiver Operating Characteristic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

3.3 A New Paradigm

It is envisioned that the next-generation of small spacecraft will not only incorporate more power-efficient devices but also *smarter* power management systems to meet the ever-increasing demands on capability and performance. For the smartly-powered, small satellite of the future, the ground operator will outline a high-level plan of operation, and the satellite will determine, through use of DPM and power-aware concepts, the best management of available power resources to accomplish that plan.

4. Advanced Processing Payload Concept

A payload concept based on the Fast On-Orbit Recording of Transient Events (FORTÉ) satellite mission was developed [17] and provides an example of a processing payload with multiple modes of operation. A primary objective of FORTÉ is to detect the Radio-Frequency (RF) signal of lightning events in the Earth's atmosphere [18]. 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 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. This application concept is depicted in Figure 1.

4.1 Algorithm Power Modulation (APM)

The 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)[10], 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. 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 [19].

Through further analysis, the relationship between energy usage and reduction in the probability of false alarms for each algorithm was determined for a

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

given operating point of the trigger box [20]. The results of this analysis are reproduced in Table 1.

Table 1. Post-Processing False Alarm Performance

Algorithm	Probability of False Alarm
Least Mean Squares	3.58%
Maximum Likelihood	3.74%
Software Trigger Box	2.01%
Match Filter Bank	0.00%

5. Payload Concept Simulation

A software simulation of the payload concept has been developed to help quantify the performance of APM. The simulation is accomplished through the following procedure:

- The first step simulates the FORTÉ orbit using the *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 module to determine the event rate at specified latitude and longitude positions.
- 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[11]. 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. A ring buffer module keeps track of the number of events processed or "lost." The maximum capacity of the simulated ring buffer is 150Mb[12].

Lightning event rates are taken from data observed by the FORTÉ optical lightning sensor. As a result some lightning events may not be reflected in the data due to clouds in the field of view. The data provided comprises 15 months of data between 1998 and 2002. The monthly variation in lightning has been summarized into 3 months of seasonal activity. However, for the initial simulation results, only one eclipse period of lightning data is used during Northern Hemisphere Summer.

5.1 Battery Model

The battery model is based on the Maxell ICR18650G manufacturer cell data [21]. 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 this work, a model based on linear interpolation and extrapolation of the manufacturer data can yield reasonable accuracy.
- *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. The simulation assumes a maximum capacity at the beginning of the eclipse period.

[11] For evaluation and comparison purposes, the user has the option of implementing the APM decision scheduler or to specify one of the LMS, ML, ST, or MF algorithms to execute.

[12] 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.

- *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 [22]. However, for simplification, temperature effects are neglected.

Using the cell data, the battery was sized according to the power requirements of the PPC750 processor and the estimated orbital eclipse duration. Only power consumed by the microprocessor is considered. It is understood there would also be memory, cache, voltage supply, etc., in order to make practical use the processor, but, for the preliminary results, the power consumption and inefficiencies of these components is neglected. Table 2 lists the parameters involved in sizing the battery.

Table 2. Battery Sizing Parameters

PPC750 Nominal Voltage	2.5 – 2.7 V [23]
PPC750 Typical Full-On Power Consumption	5.7 W [23]
Maximum Eclipse Duration	35 min. [14]
Cell Depth-Of-Discharge	80%
Required Battery Capacity	1600mAh

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 [23], only one cell is necessary per processor.

The model takes, as input, the increment of capacity drained and the discharge rate over the one second simulation time step. 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 2 illustrates the battery model under constant rates of discharge. Qualitatively, these curves are typical of battery properties and match those of the manufacturer data. Since access is not available to this battery cell and a discharge test bench, a quantitative validation of this model is not possible.

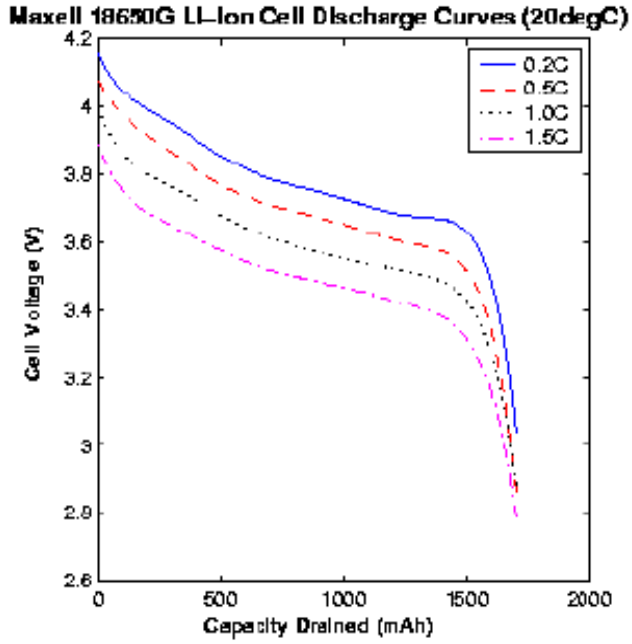

Figure 2. Battery Model Discharge Curves: This figure illustrates the battery model discharge properties under constant rates of discharge. Note that the temperature is constant at 20°C.

5.2 APM Decision Scheduler

Initially, a *heuristic, power-aware* 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:

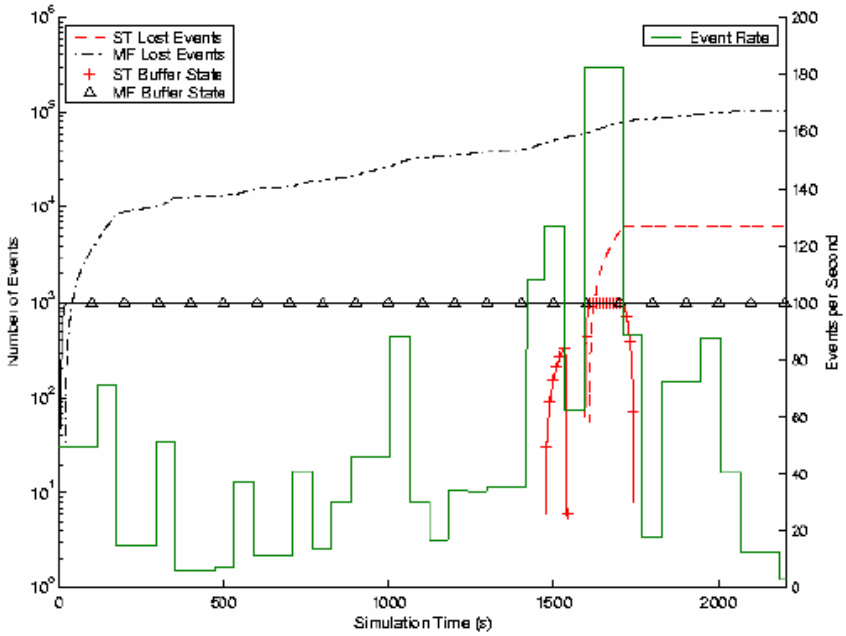

Figure 3. Event Processing Performance of ST and MF Algorithms: This figure shows the performance of the ST and MF algorithms in processing events in the ring buffer. Both the number of events in the ring buffer and the events lost over the simulation time are shown. For comparison, a plot of the event rate is also illustrated.

- 0%-25% of the maximum capacity, run MF
- 25%-50% of the maximum capacity, run ST
- 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. In this case, the ring buffer will continue to fill up with events for a non-zero event rate.

5.3 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 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 previously, 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 3 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 [17].

Table 3. 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

In comparison to the data in Table 3, 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 it can only execute twice in the

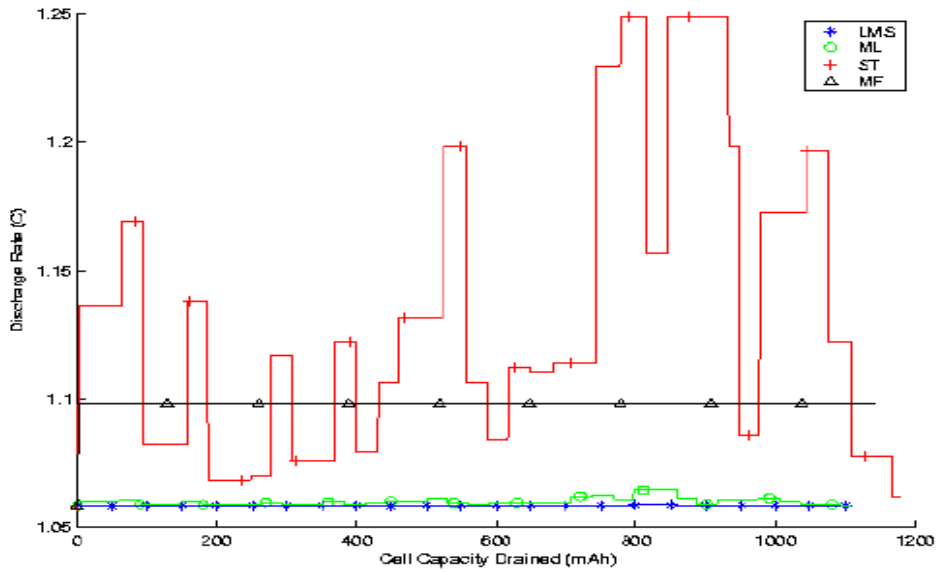

Figure 4. Discharge Rates: This figure depicts the discharge rates of LMS, ML, ST, and MF algorithms. The ST algorithm shows the largest fluctuation in discharge rate.

one second interval. 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 ML algorithm when the number of events in the ring buffer exceeds 500.

Performance of the ST and MF algorithms in processing events is illustrated in Figure 3. 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 4 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 4, 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 5.

Table 4 summarizes the final results of the simulation runs illustrating the DOD, number of events lost, and the average probability of false alarms in the processed data for each algorithm. This shows the tradeoffs in using each of the different methods.

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 Table 4, 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. There is a tradeoff point for ST and MF in the amount of capacity consumed based on the number of executions. Due to the decision procedure used, APM selects the worst-case, i.e., it

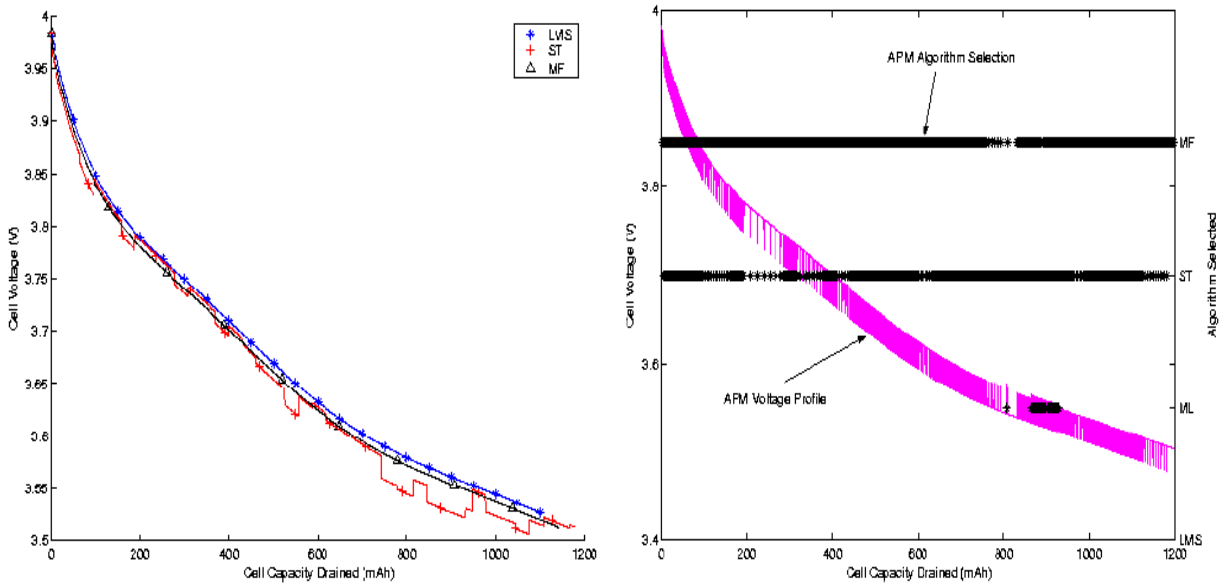

Figure 5. Voltage Profile: This figure illustrates the voltage profiles versus the capacity drained during simulation. Note that the chart on the right also illustrates APM switching between ML, ST, and MF algorithms. Due to the fast execution of ML, APM did not need to execute LMS.

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[13].

For the decision procedures used in the initial simulation, the benefit APM provides is in the low average probability of false alarms in the processed data and the zero loss of events. Since APM primarily uses MF and ST algorithms, the false alarm performance is greater than running LMS or ML alone. This is an indication of APM's ability to reduce the number of false alarms that will wastefully use onboard memory space and improve the quality of data downlinked to the ground. It is a more efficient use of bandwidth. 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.

Table 4. Summary of Simulation Results

Algorithm	DOD	Events Lost	Average Probability of False Alarms
LMS	64.9%	0	3.58 %
ML	64.9%	0	3.74 %
ST	69.3%	6,316	2.01 %
MF	67.3%	102,722	0.00 %
APM	**70.4%**	**0**	**0.74 %**

[13] 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.

6. Future Direction

Now that the first generation of the simulation concept is developed, future work on the payload concept includes the following:

- *decision process*- The APM decision process used in the preliminary simulation runs was a simple process, used primarily to evaluate the simulation behavior. Investigation into more sophisticated decision processes is warranted.
- *lightning event rates*- Provide a more stochastic mechanism for determining the event rates, based on the FORTÉ data. This would help ensure that future decision procedures do not adapt to a specific event rate set.
- *battery model*- Adding in charging characteristics and temperature effects would make the model more realistic and allow for simulation beyond eclipse conditions.

During the course of this work, it was realized that a two-tiered approach to power management would help support a more modular bus architecture. In this concept, smart power management would be applied at the both the system level and the payload level. This could conceivably help reduce mission costs further and increase capabilities by supporting a wider range of payloads. The system-level power management work is being targeted towards an approximately 50W average spacecraft application based on the AeroAstro Bitsy bus.

7. Conclusions

There is a need for a new power management paradigm in small satellites to provide a more integrated and cost-effective solution for small satellites. It is envisioned that a more *proactive* power management system will enhance autonomy and enable significant onboard processing.

In this paper, a payload concept is described that has multiple modes of processing operation. The preliminary results have shown that a power-aware, heuristic process, APM, utilized more battery capacity, but did not loose events. Additionally, APM provided a low probability of false alarms in the processed data.

8. Acknowledgements

The processing payload concept and simulation has been sponsored by the Defense Advanced Research Projects Agency, DARPA, through the Air Force Research Laboratory, USAF, under agreement number F30602-00-2-0548. The system-level power management work is being sponsored by a National Aeronautics and Space Administration, NASA, Small Business Technology Transfer, STTR, grant through AeroAstro, Inc.

Thanks are due to Professor Scott Palo, Professor Mark Balas, and Professor Dirk Grunwald of the University of Colorado at Boulder; Mr. Raymond Zenick of AeroAstro, Inc.; and Mr. Philip Lyman of Ball Aerospace Corp. for support and guidance in this work. This work has also been made possible through the efforts of the LANL NIS-3 Power Aware Computing team: Dr. Maya Gokhale, Dr. Scott Briles, and Dr. Jayashree Harikumar. Thanks are also due to the LANL FORTÉ team for supplying details of FORTÉ operation and to Dr. Tracy Light, LANL NIS-1, for providing the lightning event rate data.

9. REFERENCES

[1] Fleeter, Rick. The Logic of Microspace. El Segundo: Microcosm Press, 2000.

[2] Wertz, James R. and Wiley J. Larson. Reducing Space Mission Cost. El Segundo: Microcosm Press, 1996.

[3] Lu, Yung-Hsiang and Giovanni De Micheli. "Comparing System-Level Power Management Policies." IEEE Design and Test of Computers. 2001.

[4] Kandemir, Mahmut, N. Vijaykrishnan, and Mary Jane Irwin. "Compiler Optimizations for Low Power Systems." Power Aware Computing. Eds. Robert Graybill and Rami Melhem. New York: Kluwer Academic/Plenum Publishers, 2002. Ch. 10: 191-210.

[5] Slawsby, Alex, "Taking Charge: Trends in Mobile Device Power Consumption," 2002. On-line bulletin available at http://www.intel.com/ebusiness/pdf/prod/related_mobile/ar023804.pdf

[6] "PC Energy-Efficiency Trends and Technologies." Intel, 2002. On-line article at http://www.intel.com/ebusiness/pdf/prod/related_mobile/ar024103.pdf

[7] Benini, Luca, Alessandro Bogliolo, and Giovanni De Micheli. "A Survey of Design Techniques for System-Level Dynamic Power Management." IEEE Transactions on Very Large Scale Integration (VLSI) Systems, 8 (3), 2000.

[8] Ramanathan, Dinesh and Rajesh Gupta. "System Level Online Power Management Algorithms." Proceedings of Design Automation and Test in Europe, 2000.

[9] Šimunic, Tajana, Luca Benini, Peter Glynn, and Giovanni De Micheli, "Event-Driven Power Management." IEEE Transactions on Computer Aided Design of Integrated Circuits and Systems, 20 (7), 2001.

[10] Poole, Bill, "RF 'Sleep Modes' can Improve Battery Life in Wireless Phones." Planet Analog, 2003. On-line article available at http://www.planetanalog.com/features/power/OEG20030502S0037

[11] Šimunic, Tajana, Luca Benini, Andrea Acquaviva, Peter Glynn, Giovanni De Micheli. "Dynamic Voltage Scaling and Power Management for Portable Systems." Proceedings of the 38th Conference on Design Automation. 2001.

[12] http://www.darpa.mil/ipto/research/pacc/

[13] McDermott, Joseph, K. "Power." Space Mission Analysis and Design. Eds. James R. Wertz and Wiley J. Larson. El Segundo: Microcosm Press, 3rd ed., 1999. Ch. 11.4: 407-427.

[14] Sullivan, Ralph M. "Space Power Systems." Fundamentals of Space Systems. Eds. Vincent L. Pisacane and Robert C. Moore. New York: Oxford University Press, 1994. Ch. 6: 337-431.

[15] Johnson, Michael A., Robert G. Beaman, Joseph A. Mica, Walter F. Truszkowski, Michael L. Rilee, and David E. Simm. "Nanosat Intelligent Power System Development." Proceedings of the Second International Conference on Integrated Micro/Nanotechnology for Space Applications. 1999.

[16] Schaffner, Jake A. "The Electronic System Design, Analysis, Integration, and Construction of the Cal Poly State University CP1 CubeSat." Proceedings of the 16th Annual AIAA/USU Conference on Small Satellites. 2002.

[17] Shriver, Patrick M., Maya B. Gokhale, Scott D. Briles, Dong-In Kang, Michael Cai, Kevin McCabe, Stephen P. Crago, and Jinwoo Suh. "A Power-Aware, Satellite-Based Parallel Signal Processing Scheme," Power Aware Computing. Eds. Robert Graybill and Rami Melhem. New York: Kluwer Academic/Plenum Publishers, 2002. Ch. 13: 243-259.

[18] Roussel-Dupré, Diane, Phillip Klingner, Leslie Carlson, Robert Dingler, and Diana Esch-Mosher. "Four Years of Operations and Results with FORTÉ." Proceedings of the AIAA Space 2001 Conference. AIAA-2001-4627.

[19] Shriver, Patrick, Scott Briles, Jayashree Harikumar, and Maya Gokhale. "A Power-Aware Approach to Processing Payload Design." Proceedings of the Government Microcircuit and Critical Technology Conference. 2003.

[20] Briles, Scott, Patrick Shriver, Maya Gokhale, and Jayashree Harikumar. "Power-Aware Improvement in Signal Detection." Proceedings of the International Signal Processing Conference, 2003.

[21] Maxell Lithium Ion Rechargeable Batteries. Product information brochure.

[22] Linden, David and Thomas B. Reddy, eds. Handbook of Batteries. New York: McGraw-Hill, 3rd ed., 2002.

[23] PowerPC740 and PowerPC750 Embedded Microprocessor: Hardware Specifications. Available on-line at http://www-3.ibm.com/chips/techlib/techlib.nsf/techdocs/8CAAF8D73C47F5B587256A0900630F73/$file/7xxp8t_em_hw2-99.pdf