An Architecture to Enable Autonomous Control of Spacecraft

Ryan D. May[*]
Vantage Partners, LLC, Cleveland, OH, 44142, USA

Timothy P. Dever[†], James F. Soeder[‡], Patrick J. George[§]
NASA Glenn Research Center, Cleveland, OH, 44135, USA

Paul H. Morris[**], Silvano P. Colombano[††], Jeremy D. Frank[‡‡], Mark A. Schwabacher[§§]
NASA Ames Research Center, Moffett Field, CA, 94035, USA

and

Lui Wang[***], Dennis Lawler[†††]
NASA Johnson Space Center, Houston, TX, 77058, USA

Autonomy is required for manned spacecraft missions distant enough that light-time communication delays make ground-based mission control infeasible. Presently, ground controllers develop a complete schedule of power modes for all spacecraft components based on a large number of factors. The proposed architecture is an early attempt to formalize and automate this process using on-vehicle computation resources. In order to demonstrate this architecture, an autonomous electrical power system controller and vehicle Mission Manager are constructed. These two components are designed to work together in order to plan upcoming load use as well as respond to unanticipated deviations from the plan. The communication protocol was developed using "paper" simulations prior to formally encoding the messages and developing software to implement the required functionality. These software routines exchange data via TCP/IP sockets with the Mission Manager operating at NASA Ames Research Center and the autonomous power controller running at NASA Glenn Research Center. The interconnected systems are tested and shown to be effective at planning the operation of a simulated quasi-steady state spacecraft power system and responding to unexpected disturbances.

Nomenclature

ACAWS = Advanced Caution and Warning System
AD = Activity Dictionary
APC = Autonomous Power Controller
ARC = NASA Ames Research Center, Moffett Field, CA

---

[*] Controls Engineer, 3000 Aerospace Parkway, ryan.d.may@nasa.gov, AIAA Member.
[†] Electrical Engineer, 21000 Brookpark Rd, MS49-8, tdever@nasa.gov.
[‡] Senior Technologist for Power, 21000 Brookpark Rd, MS301-5, james.f.soeder@nasa.gov.
[§] Project Manager, 21000 Brookpark Rd, MS77-4, patrick.j.george@nasa.gov.
[**] Computer Scientist, MS269-1, paul.h.morris@nasa.gov.
[††] Computer Scientist, MS269-2, silvano.p.colombano@nasa.gov.
[‡‡] Group Lead, MS269-3, jeremy.d.frank@nasa.gov.
[§§] Computer Scientist, MS269-1, mark.a.schwabacher@nasa.gov.
[***] 2101 NASA Parkway, lui.wang-1@nasa.gov.
[†††] 2101 NASA Parkway, dennis.g.lawler@nasa.gov.

1

American Institute of Aeronautics and Astronautics

BCDU   =   Battery Charge/Discharge Unit
ECLSS   =   Environmental Control and Life Support Systems
EPS   =   Electric Power System
GRC   =   NASA Glenn Research Center, Cleveland, OH
JSC   =   NASA Johnson Space Center, Houston, TX
MBSU   =   Main Bus Switching Unit
MM   =   Mission Manager
PDU   =   Power Distribution Unit
SNRF   =   Space Network Research Federation
SOC   =   State of Charge (%)
SPIFe   =   Scheduling and Planning Interface for Exploration

I. Introduction

As manned spacecraft venture out from Earth, light-time communication delays will make the current ground-based mission control infeasible.[1] To enable these missions, NASA has begun to investigate an architecture that will enable autonomous control of the spacecraft[1,2] and the electrical power system in particular.[3] In order to accomplish the mission of any crewed spacecraft, the crew schedule must be coordinated with all vehicle subsystems, including Electric Power (EPS), Thermal Management, Communications, Environmental Control and Life Support Systems (ECLSS), and others. Currently, ground controllers develop a complete schedule for all spacecraft components based on system operating mode, crew availability, communications, solar array pointing, and a large number of other factors. This schedule is created by humans on the ground running software tools to analyze the subsystem for which they are responsible.[4,5,6]

For the case of the electric power system onboard the International Space Station, the planning process involves teams at three different locations as described in Ref. 6. The mission operations team develops a mission timeline that schedules vehicle activities required to complete mission objectives. Based on the mission timeline, a team at NASA Glenn Research Center (GRC) uses information about station attitude, orbital path, visiting vehicle traffic, and solar array pointing constraints to estimate the amount of power available at each power convertor, using software called SPACE.[7] This information along with the mission timeline is sent to a team at Boeing that develops a list of loads and their respective power utilization for each point in time along the planning window (typically at one minute intervals). The information is sent back to the team at GRC where the power draw from each load and the solar array power output are used as inputs to the ECAPS software tool.[6] This tool then determines how much power is required to be discharged by the batteries or is available to charge, resulting in the predicted battery state-of-charge (SOC) at each time step. These data are then reviewed to ensure that the battery SOC does not violate constraints and that the power flowing through each device is within acceptable limits. The GRC team can change the network topology (to a limited extent) to avoid constraints. Any remaining discrepancies or issues result in a phone call back to Boeing and a discussion about time-shifting loads to avoid violations. If the two groups are unable to arrive at a feasible solution, the mission operations team is informed and a new mission timeline is planned or constraints on power generation are adjusted. In the worst case, a plan will be flagged as "don't fly." A similar process is conducted by all vehicle subsystems. In addition to solving the planning problem, controllers on the ground monitor telemetry and respond to disturbances and faults that may occur in each of the subsystems.

In order to enable autonomous operations, this functionality must be implemented onboard. Unfortunately, it is neither realistic nor useful to require that the flight crew have expert knowledge of all vehicle subsystems. Thus, software must be designed to respond to these situations in an appropriate manner through an autonomous or semi-autonomous decision-making process.

The architecture presented in this paper is an early attempt to formalize and automate the planning process using on-vehicle computation resources. Additionally, the architecture is developed to be able to respond to disturbances in the plan or changes to plan constraints. For this work, the architecture will be applied to the electrical power system to serve as an example of how the autonomous architecture could be implemented.

Section II provides a high-level overview of the proposed autonomous control architecture for the spacecraft. The third section describes how the Autonomous Power Controller and the Mission Manager interact to develop a plan and respond to associated disturbances. This is followed by a brief discussion of the paper simulation test cases as well as results from the software testing. Finally, Section V provides a conclusion and a discussion of the future direction of the work.

American Institute of Aeronautics and Astronautics

II.   Proposed Autonomous Control Architecture

The proposed vehicle communication architecture is shown in Figure 1 based on Reference 8. The Mission Manager (MM) is responsible for coordinating the vehicle subsystems and the crew resources. Each of the vehicle subsystems is autonomous in the sense that it controls all related spacecraft components in order to meet the objectives set forth by the mission manager. Similar architectures have been proposed previously, e.g. Reference. 9.

The vehicle-level optimal planning problem is highly complex; however, the overall problem can be solved via a 'divide and conquer' strategy. This strategy both allows loosely coupled subsystems to be managed independently, and allows lower level details to be addressed at the subsystem level. Therefore, one of the goals of this system is to limit the amount of subsystem detail that must be known by the Mission Manager. Thus, the system architecture is distributed in nature with information about state, dynamics, and constraints being held at the lowest level possible. Nevertheless, the Mission Manager must be designed with sufficient understanding of the constraints on subsystems in order to generate plans that respect the subsystem constraints.
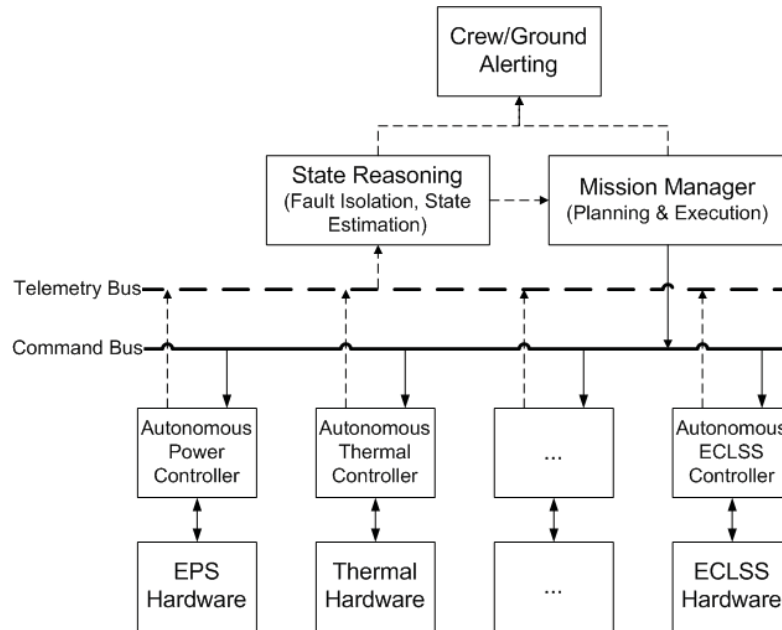

Figure 1. Communication architecture of notional autonomous spacecraft.

A. Vehicle-Level Task Overview

The Mission Manager is responsible for strategic planning in order to accomplish mission objectives. These objectives are accomplished by performing activities that collectively, over time, achieve the goals of the mission. Each activity, in turn, requires one or more subsystems of the spacecraft to perform a variety of functions. Each step of the activity plan is then scheduled to particular time slots in order to minimize the amount of time needed to achieve the goal while ensuring that any time-based constraints are observed. The MM uses automated planning and scheduling to plan and schedule the high-level operation of equipment on board the spacecraft, such as the Electrical Power System and the ECLSS, which can then be operated in an automated fashion. It also plans and schedules the activities of the crew members. The MM will support either fully automated planning and scheduling, or "mixed initiative" planning and scheduling, in which a crew member collaborates with the automated system to produce a plan or a schedule.

The MM also ensures that the plan is executed correctly by monitoring activity progress in real-time. If an unexpected deviation from the prescribed schedule occurs, the MM will work to adjust the schedule to maximize the mission objectives achieved within the new operating environment. For example, if an activity takes longer to complete than had been expected other activities may have to be delayed to accommodate the change in crew and subsystem availability. Further, these delays may result in constraint violations in later activities that need to be rescheduled.

American Institute of Aeronautics and Astronautics

State Reasoning determines the current state of the system based on the information available from the spacecraft and subsystems. Examples of the system state to be determined include status of switches, component modes, temperatures, pressures, and so on. A subset of the state reasoning task is the detection of off-nominal conditions: is there evidence of something wrong? If so, is it an anomaly (a component is not behaving the way it usually does) or a failure (unacceptable system performance) caused by a fault? If there is a fault, can it be determined uniquely from the available information, or is it necessary to perform diagnostic activities to isolate the faults? This determination may require the use of information from multiple subsystems (provided by the subsystem controllers) in order to determine the root cause of a failure that may have symptoms across multiple subsystems. In some cases, the diagnostic process will be fully automated. In other cases, it will require crew involvement in a "guided troubleshooting" process. For some failures, there will be multiple candidate faults that are consistent with available sensor data. In this case, the automated system will produce an "ambiguous diagnosis", which is a list of candidate failures, all of which are consistent with available sensor data. The system will then ask the crew to collect additional data to help the system disambiguate the diagnosis.

After the system produces an unambiguous diagnosis, it will perform failure consequence assessment. It will determine which components of the vehicle have been impacted by the failure, and what capabilities have been lost. It will communicate this information to the automated planning and scheduling system, which will then replan and reschedule the rest of the mission to either repair the failed component (moving other items in the schedule to make room for the repair), or to accomplish as many mission objectives as possible given the loss of capability.

The state of the vehicle, the validity of the current plan, and the presence of either anomalies or faults must all be reported both to ground controllers and to the crew. This allows the crew and/or ground controllers to participate in mixed-initiative planning and guided troubleshooting.

B. Subsystem Task Overview

Each of the vehicle subsystems is responsible for a large number of tasks as best understood by the respective subject matter experts. The purpose of this work is not to dictate how each of these subsystems is to be controlled, rather to develop an architecture to enable long time-scale planning and operations level control to be implemented in an autonomous and coordinated manner. The vehicle Mission Manager is responsible for strategic planning and coordination of all subsystems whereas the autonomous subsystem controllers are responsible for operating the subsystem in the manner most consistent with the strategic plan. The short time-scale control tasks are handled by low level controllers in each of the subsystems with the responsibility for ensuring system safety and operability. An example of such a control architecture for the electric power system is presented in Ref. 10.

One of the central tenets of this architecture is that operational details should be visible to the lowest level system possible. For example, the vehicle Mission Manager needs to be aware of crew resources and mission objectives; however, information such as the impedance of distribution lines should be hidden and only available to the Autonomous Power Controller. The communication between the subsystem controllers and the MM are phrased in terms of system capability. For the power system, this capability is the ability to provide power to loads. Capabilities are expressed as constraints on the Mission Manager's activity planning problem. This enables the Mission Manager to solve the vehicle planning problem as an optimization problem with each of the subsystems providing relevant constraints.

In addition to nominal planning operations, the autonomous subsystem controllers must be able to correctly and promptly respond to off-nominal events. It is envisioned that a diagnostic and recovery system would be built into each of the subsystems, based on the technology developed at Ames for the Advanced Caution and Warning System (ACAWS).[2] ACAWS is a fault management tool that provides the following capabilities:
- Dynamic and interactive graphical representations of spacecraft systems
- Systems modeling
- Automated diagnostic analysis and root cause identification
- System and mission impact assessment
- Procedure and flight rule identification
- Interaction with other tools to help spacecraft operators (both flight controllers and crew) understand and respond to anomalies more effectively.

These capabilities are implemented in four major modules: a) anomaly detection, b) fault detection and isolation, c) system effects, and d) graphical user interface. The current implementation of this system was designed for interaction with crew and flight controllers, with a view towards enhancement of automation capabilities. It was successfully tested on the power system for the Deep Space Habitat during a Mission Operations Test in September 2012. While the existing implementation is a centralized system in that ACAWS operates at the vehicle level and

American Institute of Aeronautics and Astronautics

has intimate knowledge of all vehicle subsystems, work to develop a distributed implementation is underway. This new implementation would be integrated into each subsystem to enable failure detection, isolation, and effects analysis at the subsystem level while sharing information that concerns the vehicle state and capabilities with the vehicle level ACAWS. This would enable subsystems to respond to internal faults without requiring coordination with the vehicle, as well as enabling coordination for failures that extend through multiple subsystems.

III. Interface with Autonomous Electric Power System

In order to develop and test this system, an example Autonomous Power Controller (APC) and Mission Manager are constructed. It is anticipated that the other subsystems would have a similar communication scheme, modified to account for the particular aspects of interest of each subsystem. The primary purpose of the interaction between the MM and the APC is to determine which loads in the power system should be operational at each point in time. This information allows the power system to configure itself to most efficiently and safely deliver the necessary power. It also enables the Mission Manager to make decisions about crew time, vehicle orientation, and long-term planning without requiring a detailed understanding of the power system.

The interaction between the two systems is composed of a planning process as well as event-driven replanning. The messages and data communicated between these two systems were developed through a series of "paper simulations" conducted by system experts using email and manual optimization/planning tools. A number of scenarios (described in Ref. 11) were considered to ensure that the communication scheme is robust enough to handle situations that would be encountered in operation. Once the paper simulations were complete, an interface document was developed[11] and the protocol was implemented in software for further testing.

A. Planning

At the initialization of the planning process, the Mission Manager informs the Autonomous Power Controller of the planning period start time, end time, and time step. In addition, the Mission Manager will send all information necessary for the power system to determine the power availability of the various power generation components. For solar arrays, information such as pointing angle, local traffic, shadowing information, and orbital information are sent. The APC then uses this data as well as its knowledge of the system status (faults, degradation, etc.) to compute the power availability at each time step of the planning window. The system also computes the peak power and energy available (by utilizing storage systems off-nominally) and passes this information back to the Mission Manager.

The MM then uses the power availability, list of desired crew activities to perform, and information from other subsystems to develop a proposed vehicle plan. One component of this plan is a schedule of load power utilization and priority. This schedule is derived from the orderings of the crew activities and the equipment used by each activity. This information (for each load on the system) is then communicated to the APC, where power flow calculations and contingency analyses are conducted to determine the feasibility of the proposed load plan at each time slice of the planning window. For any point at which the plan is found to be infeasible, the APC will determine which of the load profiles are in conflict and inform the Mission Manager. The MM can then determine which activities are in conflict and propose resolutions. The planning process will repeat until the Mission Manager states that the proposed plan should be implemented. These messages must be constructed in a manner such that the MM can determine how to develop a feasible plan.

Finally, the Mission Manager informs the APC that the previous plan will be implemented at the associated start time or of its intention to develop a new plan. As discussed earlier, the MM can ask the APC to implement any plan, regardless of feasibility. However, if the APC is tasked with implementing an infeasible load plan, load will be shed by the low-level control laws to force the system to remain within operational constraints.

An overview of the communication necessary to produce a load plan is shown in Table 1. The complete details of the implementation are given in Ref. 11 and an example is presented in Section IV.C.

American Institute of Aeronautics and Astronautics

Table 1. Summary of communication between the MM and APC to develop a load schedule.

Initiator	Receiver	Purpose	Content
MM	APC	Initiate planning	Planning period start time, end time, time step
APC	MM	Send power constraints	Nominal and peak power for each time step, peak energy available for planning period, other power constraints
MM	APC	Send load schedule	State and priority of each load for each time step
APC	MM	Send feasibility	Communicate detected conflicts and the associated time step.
MM	APC	Re-plan or implement	Intention to re-plan the load schedule or inform the APC that the plan should be implemented

B. Event-Driven Replanning

In addition to the planning sequence, the system must be able to respond to off-nominal events and disturbances. For the purposes of this work, only events that disturb the power system plan or planning parameters are considered. However, it is anticipated that the basic architecture will enable the system to respond to faults and other unanticipated events in any subsystem.

As summarized in Table 2, the communication necessary to appropriately handle these events is inherently unplanned and will be initiated by the subsystem (in this case the power system) that detects the disturbance. Upon detection, the Autonomous Power Control will determine the impact of the disturbance on its capability to operate as specified during the planning process. This information, which is key to determining if a replan is necessary in order to achieve the vehicle's mission objectives, is passed to the MM. The MM then decides whether a replan is necessary; if so, the planning process is started over and the disturbance is accounted for as a new constraint.

Table 2. Summary of communication between the MM and APC to determine the response to a power system disturbance.

Initiator	Receiver	Purpose	Content
APC	MM	Inform of disturbance	Disturbance type, expected impact of disturbance
MM	APC	Decision	Take no action or notify of replan

For example, consider the case in which an activity runs longer than had been planned. In this situation, system loads that the power system had expected to be off are still operating. When this condition is detected, the APC runs calculations to determine if and when the deviation will impact system operation should the loads continue to draw power. If the loads are large or the system loading is high, problems can arise such as excessive discharge of storage devices or there may be a violation of a maximum power flow constraint. Information about the impact of these disturbances is sent to the Mission Manger to facilitate decision making.

As with the development of the planning process, a number of paper simulations were conducted of various classes of disturbances in order to validate the choice of message. Complete details of the test cases and final implementation are included in Ref. 11.

It is worth reiterating that the autonomous control in the subsystems is constructed in a hierarchical manner based on time-to-respond. Thus, if there is some component failure or disturbance, the lowest level of the control architecture will respond in order to maintain system safety and stability while maximizing operational capability. The communication between the MM and the autonomous subsystems is not intended to handle these situations; instead it is tasked with planning and long time-scale operation.

American Institute of Aeronautics and Astronautics

IV. System Simulation

In order to determine the feasibility of this architecture and to validate its capability to respond to both disturbances and changing constraints in an appropriate manner, a computer-based simulation is conducted. The system is configured as shown in Figure 2. The Mission Manager runs on a computer at NASA Ames Research Center in Moffett Field, California, and the Autonomous Power Controller is executed at NASA Glenn Research Center in Cleveland, Ohio. For ease of testing during this research effort, the two systems communicate via TCP Sockets over the Space Network Research Federation (SNRF) dedicated network. In order to test the communication a notional vehicle and power system are constructed. The power system of this vehicle is simulated using a quasi-steady state, power flow-based, simulation running on a computer at NASA GRC. Visualization software for all systems is provided to offer insight into the operation state and decision making process. Future studies will likely be conducted using hardware in the NASA Johnson Space Center's (JSC) Deep Space Vehicle electrical testbed. As a proof of concept that this hardware can be controlled over a high latency network connection, this initial study will utilize the Autonomous Power Controller to control power relays on the testbed via UDP messages over the SNRF network (typical round-trip latency exceeding 100ms).


Figure 2. Architecture of the software simulation. Blue components are located at NASA ARC, green at NASA GRC, and red at NASA JSC.

The following subsections describe the vehicle and power system under study, the test cases considered, and the results from the software simulation study.

A. System under Study

The vehicle and associated power architecture considered in this work form a notional architecture for a prospective deep-space vehicle. A schematic of the electrical power system is shown in Figure 3. This architecture is based on the results of studies[12] conducted for vehicles designed for deep space operations, therefore requiring autonomous control.

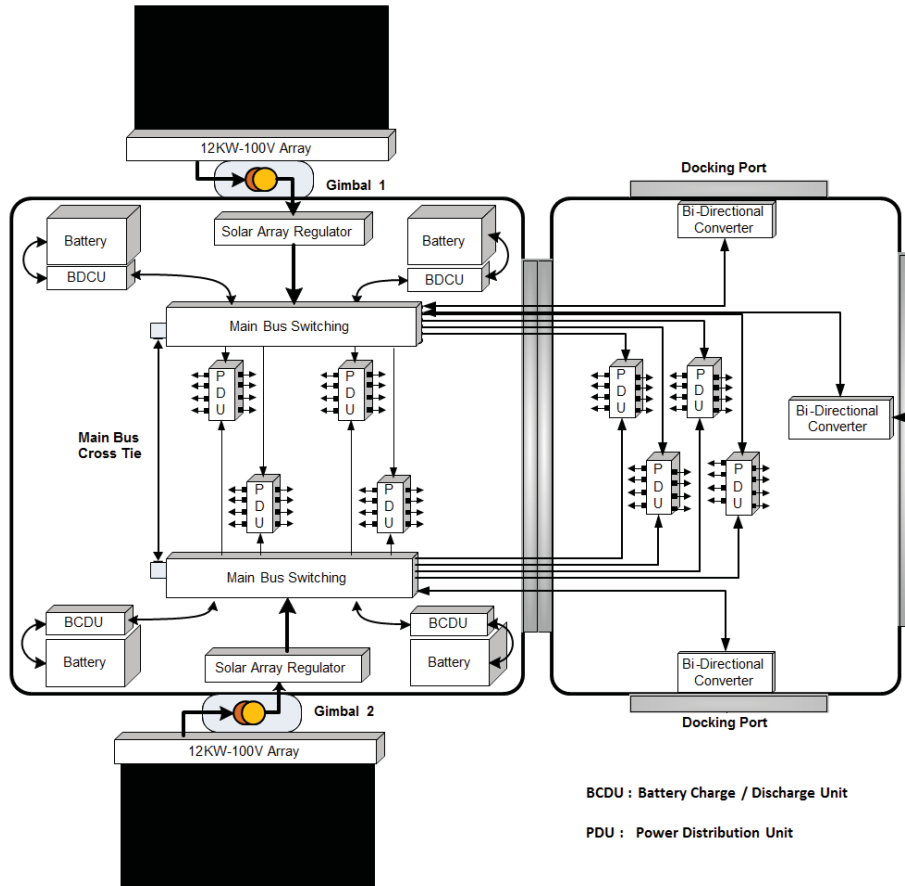American Institute of Aeronautics and Astronautics

Figure 3. Power system architecture for a notional deep space vehicle.

The notional vehicle consists of a Power Module (shown on the left) and a Habitation Module (on the right). The Power Module can provide approximately 16kW average power to the user loads. It is characterized by two independent power channels, with multi-level cross strapping. Each channel includes a 12kW solar array and lithium ion batteries with a capacity of 110 Amp-hours per channel. The large arrays and batteries enable missions in environments with an eclipse/insolation period of 30/60 minutes. Each independent power channel operates at 120V and is compatible with SAE specification AS 5698.[13] Power from each solar array is fed to a regulator that provides conditioned power to the Main Bus Switching Unit (MBSU). From the MBSU the power is used to charge the batteries during insolation and is distributed to the Power Distribution Units (PDUs) which feed the user loads. For redundancy and reliability considerations the MBSUs can be cross-strapped, and each PDU can be fed from either power channel. Finally, the MBSU can also feed bi-directional converters that can provide or draw power from the three Habitat Module docking ports. For the purpose of this initial work, the docking ports are assumed to be unused.

In order to conduct load planning exercises, it is necessary to have a description of the loads on the vehicle. Since the vehicle is in the early development stages, a set of thirty notional loads and their corresponding power consumption was arbitrarily constructed. The list of the loads, their power consumption, and the subsystem to which they belong is shown in Table 3. As shown, most of the loads are actually in an "idle" state (still drawing power) when turned "off;" thus all of the loads can be in one of two states. This is a simplification from the more general case in which every load can draw a variable amount of power between a maximum and minimum level. This assumption simplifies the planning problem as it enables the use of discrete rather than continuous decision variables.

American Institute of Aeronautics and Astronautics

Table 3. Notional loads and their associated "load off" and "load on" power consumption levels.

Subsystem	Load Name	Load ID	Power Consumption (W)	
			Load Off	Load On
Communications	Array 1	1	1.5	350.0
	Array 2	2	1.0	196.8
Crew Health Care	Exercise 1	3	2.0	652.4
	Exercise 2	4	2.0	300.0
	Exercise 3	5	3.0	652.4
	Exercise 4	6	3.0	100.0
Crew Support	Water Processing	7	5.0	150.0
	CO2 Processing	8	20.0	100.0
	Lighting	9	-	200.0
Command and Data Handling	Data Acq. 1	10	6.0	107.8
	Data Acq. 2	11	4.0	109.2
Extra-Vehicular Activities	Suit Prep	12	-	2,800.0
	Airlock Cycle	13	2.0	980.0
	EVA Lighting	14	1.7	420.0
Crew Off-Duty	Entertainment	15	5.0	126.0
	Food Prep	16	15.0	197.4
Solar Array	Gimbal 1	17	-	125.0
	Gimbal 2	18	-	125.0
Guidance, Navigation & Control	Moment Dump 1	19	5.0	346.5
	Moment Dump 2	20	5.0	346.5
	Moment Dump 3	21	5.0	346.5
	Moment Dump 4	22	5.0	346.5
	Thruster Valves	23	-	70.0
Mechanical	Propellant Maint.	24	56.4	280.0
Science	Express Rack	25	30.0	100.0
	Science 1	26	1.0	700.0
	Science 2	27	1.0	350.0
	Science 3	28	1.0	1,400.0
	Science 4	29	1.0	2,100.0
	Science 5	30	-	420.0

In the spacecraft, power is needed to support a crew activity plan – a timeline describing how the resources of the craft are to be utilized over time to achieve the vehicle's mission objectives. This plan has to simultaneously satisfy temporal, state, and resource constraints, including constraints arising from the power system itself. Most activities considered require one crew member; some require two. The activities may also utilize equipment of various types which must be in working order for the activities to be successful. In general, temporal constraints include activity durations and activity precedence orderings. The precedences may be directly imposed for science or engineering reasons, or may be inserted by the planner in order to enforce state and resource constraints. For example, an activity that supplies a precondition for another activity (e.g. turning on an experiment rack which provides power for experiment hardware inside the rack) will need to precede it. Other precedences may resolve activity exclusions that arise from the limited availability of discrete resources, such as crew members or instruments. For example, if two different activities require the exclusive attention of a particular crew member, then they cannot happen simultaneously and a planner/scheduler must choose an ordering for them. In general, the ordering decisions must be judiciously selected to satisfy multiple constraints.

From the power perspective, each activity requires a set of user loads that need to be "on" for the duration of the activity. If two or more activities overlap, the load requirements for the overlap period are merged (without double-counting). Some loads use a small amount of power even when they are off, so each load contributes some amount of power usage during each time step. The power system requirements, as communicated to the MM, add several types of numeric constraints that must be respected by the activity plan:

- Absolute upper limits on the total power available ("peak power"). These vary from time-step to time-step.
- Absolute upper limits on the power summed over designated combinations of user loads. These apply to specific intervals of time-steps or all time steps.

American Institute of Aeronautics and Astronautics

- Nominal upper limits on the total power usage that is tighter than the absolute limits. The actual usage is permitted to exceed these limits for individual time steps provided that the total energy consumed above the nominal limits is within a "cumulative energy" allowance that applies to the whole plan. However, no credit is given for time-steps where the usage is under the nominal amount. Thus, the resource impact cannot be described as a simple increment or decrement of total energy.

As an example of the effect on these constraints on the activity plan, suppose we have a constraint for all time steps that the power used by loads 1, 19, and 23 together cannot sum to more than 360 Watts. Such a constraint may arise because they use a power distribution path that is currently degraded. Based on the power utilization shown in Table 3, this constraint means that loads 19 and 23 cannot be on at the same time because their combined "on" usage would exceed 360 watts. Consider a maneuver activity that utilizes load numbers 17-24. Since the maneuver activity requires both loads 19 and 23, there is no time step where the maneuver activity can satisfy the power constraint under current circumstances, so it must be removed from the plan. However, if the constraint was due to a temporary limitation, such as an eclipse, it might apply only to a subset of time steps. In that case, the activity could be scheduled outside of that period. More generally, the load constraints might prevent activities of two or more particular types from overlapping during specified periods. This type of constraint stands in contrast with the peak energy restriction which tends to require a leveling of the power loads over time.

All of the constraints listed above are dynamic (varying from plan session to plan session). This introduces a challenge for automatic planning software. Another important aspect of this application is that disturbances may occur during execution, in which case short-fuse time-limited re-planning is required. This precludes extensive search and requires an "any-time" approach where a "good enough" solution is available quickly, for example where constraint violations are resolved by discarding low-priority tasks.

In an activity planning framework, whether manual, automatic, or mixed-initiative, the properties (requirements and effects) of activities are specified by means of an Activity Dictionary (AD) or activity model. The AD provides the information needed for the planner to detect and eliminate constraint violations. Information in the AD is static in the sense that it does not vary from one planning session to another. If constraints are dynamic, i.e. varying from plan to plan, a more complex representation is needed that may involve variables that are defined in the AD but may be set to specific values at planning time. Our experiments involved 26 activity types and plans that used most of the types. The activities required between one and nine loads to be "on" (most used 2-3 loads). Typical plans involved about 20 time-steps, each of 30 minutes duration. These plans were feasible to construct manually for the paper simulation stage of development. However, working with plans of longer duration or finer granularity would require substantial automated assistance.

Two existing tools, SPIFe[14] and EUROPA,[15,16] have been used for previous NASA activity planning applications, including International Space Station scheduling[14] and Mars rover planning[15,16]. SPIFe is designed for mixed-initiative (allowing human operators to collaborate with the automated planner) plan editing using a graphical user interface that automatically detects and reports plan violations. However, it relies on the user to fix plan violations by moving or deleting activities. EUROPA can automatically construct plans and fix violations of temporal, state, and certain resource violations. However, its native numeric resource capabilities are limited to the simple increment/decrement type. SPIFe and EUROPA have been integrated for some applications to provide a combination of manual and automatic fixing.[17]

Since a fully automatic interaction with the MM is desired for this project, EUROPA would be the more suitable tool. However, the new types of numeric constraints are challenging due to their dynamic nature and the fact that they do not necessarily observe the limitation to delta-type resource impacts. A significant amount of customization is needed for either EUROPA or SPIFe to fix violations involving these resource types automatically. In particular, the resource flaws need to be mapped to temporal decisions that can resolve the flaws. Given the limited development time for the demonstration, we have initially concentrated on developing a standalone module that focuses on satisfying the power constraints in the activity plan. We expect this experience to be instructive for devising specialized power constraint methods that can later be integrated into the EUROPA automatic planner.

The standalone module uses a greedy algorithm to schedule activities. First, all the activities are placed in a list according to priority. Then, while activities remain on the list, the algorithm does the following:

- Select and remove the leftmost (highest priority) activity A from the list.
- For each time step *t*, determine whether starting A at *t* is excluded because it would violate an absolute constraint at some time step in the duration of A, taking into account the activities already scheduled. Starting A at *t* may also be excluded if there is insufficient time in the remaining time steps. If all start times for A are excluded, remove A from the plan.

10

- Otherwise, for each non-excluded time step *t,* compute the excess (above-nominal) energy usage of the total plan that results from starting A at *t*. If, for every such *t*, this exceeds the maximum excess energy allowed, remove A from the plan.
- Otherwise choose a non-excluded time step *t* that minimizes the excess energy usage, and schedule A so that it starts at *t*. (This will leave the most energy "room" for subsequent activities.)

This algorithm performed better than the manual construction as done during the paper simulation because the final plan uses less excess energy and includes more activities. A future EUROPA model has the potential to fit even more activities into a plan because it would allow the repositioning of previously placed activities to facilitate later activities. In addition, it could take into account a wider variety of constraints that are essential for planning at the activity level. However, the increased capabilities would involve a more complete and more complex search that would take more time and might require the development of additional search control methods.

B. Test Cases Considered

During the paper simulation stage of development, the following power system constraints were considered:
- Battery state-of-charge must remain within prescribed limits (30%-90%)
- The solar array output is based on orbital location (12kW during insolation and 0kW during eclipse)
- The amount of power flowing through a PDU is limited (multiple PDUs with different power levels)
- The amount of current flowing through a distribution line is limited
- The amount of current the batteries are able to source is limited

These were tested both as constraints on the entire planning window and constraints limited to specific time periods within the planning horizon. It was found that any of these constraints can be formulated as a constraint on load power consumption (i.e. the sum of power used by loads A, B, and D must be less than some limit), and that the communication sequence described in Table 1 is sufficient to handle any of these situations.

Also, during the paper simulation a number of power system disturbances were studied, including:
- Load "on" for too long
- Load "on" for too short
- Actual load consumption is less than estimated value
- Actual load consumption is greater than estimated value
- Solar array output power is less than estimated

These disturbances were examined both with and without an impact to future operation within the plan horizon. The load "on" for too short of a time period is an interesting case in that the decrease in load consumption may provide an opportunity for the MM to add additional activities (and thus loads) to a new plan.

For this paper, two of the above scenarios are evaluated: 1) nominal operation and 2) a load persists in consuming power past the point at which it was scheduled to terminate. Nominal operation must take into account the typical insolation/eclipse cycles experienced while in orbit. The scenario in which a load is "on" for too long is a plan disturbance in which an activity (or part of an activity) exceeds the allotted period of time. This may be due to the crew member encountering unexpected difficulties in performing the desired operation or due to some other unanticipated event.

C. Example Simulation Results

It is infeasible to present the results of each of the tests conducted in this paper, thus a representative test case is shown. For this case, the planning period extends from 00:00 until 10:30 with 30 minute intervals. The reply from the power system informs the MM of the nominal and peak power available (as shown in Figure 4) along with the fact that there is 8kW of peak energy available. In addition, the APC informs the MM of additional load constraints:
- At each interval over the entire planning period the total power drawn by loads 1, 19, and 23 must not exceed 360W. (PDU 1)
- At each interval over the entire planning period the total power drawn by loads 3, 7, 26, and 28 must not exceed 1.5kW. (PDU 5)
- At each interval during 04:00 – 0:600 the total power drawn by loads 4, 8, 25, 27, and 29 must not exceed 60W. (PDU 6)

These constraints represent limitations on the power throughput of the power distribution units (PDUs) in the system. The last constraint is only valid for a short period of the overall planning window, representing a situation where the component has a period of anticipated degraded performance (e.g. due to planned maintenance).

American Institute of Aeronautics and Astronautics

The Mission Manager is then responsible for developing a load schedule that meets each of the power constraints, load power constraints, and peak energy constraint. The resulting schedule is shown in Figure 4 as an aggregate power usage over time. We can clearly see that the proposed schedule is always less than the peak power available. Integrating the scheduled load power for the periods in which the power consumed is higher than nominal (00:00 - 02:30 and 03:30 – 04:00) yields a total of 7.369kWh used of the 8kW available.

Figure 4. Nominal and peak power availability over the planning period of interest as provided by the APC. The aggregate power of the scheduled loads is shown in red.

The aggregate load power usage on the three constrained PDUs is plotted in Figure 5 along with the time-varying upper limit (shown as dotted lines). Again, the scheduled load power usage falls within the allowed range.


Figure 5. The aggregate scheduled load power usage on PDUs 1, 5, and 6 (solid lines) along with the power constraint (dotted lines).

The other constraint of interest is the battery state-of-charge. For this case, the SOC is constrained to lie within the range of 30% - 90%. If the total amount of load is less than the nominal power available, the battery SOC is designed to remain within 60% - 90%. For this example, the initial SOC on both batteries is 60% and the resulting

12

American Institute of Aeronautics and Astronautics

profile is shown in Figure 6. As desired, the SOC always remains above the 30% limit, although it gets close at 03:00 after a period of sustained peak power usage. Towards the end of the simulation, the SOC is limited to 90% by the simulated hardware controller. Therefore all power system constraints are observed by the agreed upon load schedule. Overall, a total of 16 different activities were scheduled during the 10.5 hour planning period.
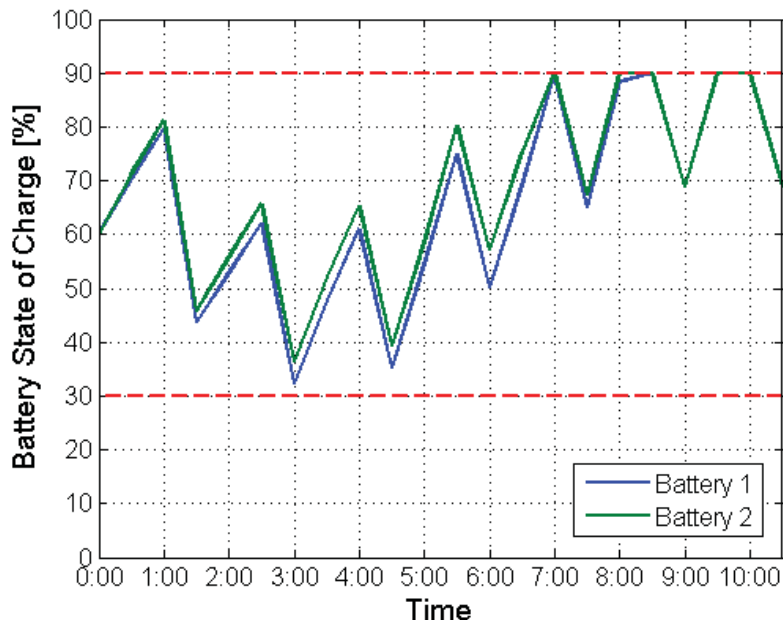

Figure 6. State-of-charge for the two batteries in the electrical power system. The SOC limits are shown as red dashed lines.

It is interesting to observe that while most of the available peak energy was used by the load schedule (7.37 of 8.00kWh) the batteries are at the maximum SOC at the end of the planning period. The simple reason for this is that the algorithm presented here does not allow the Mission Manager to earn an "energy credit" for using less than the nominal amount of power, as is done after 04:00 in this example. The amount of peak energy available is simply computed based on the initial SOC:

$$\text{Peak Energy Available} = (SOC_{t=0} - SOC_{min}) / 100 * \text{Voltage} * \text{Capacity}$$

Additionally, this value will not change regardless of the length of the planning window. This approach is inherently over-conservative as it guarantees that if peak power is used immediately at the start of the planning window the lower SOC limit will not be violated.

An investigation was performed to determine what is necessary to allow the MM to earn credit by using less than nominal power. It was found that the MM would need to know the initial SOC, maximum limit, and minimum limit for all storage devices, as well as the bus voltage and capacity of each battery. Additionally, the MM would need to make assumptions about how each of the batteries is used during the window (i.e. how much load each battery is supporting). It was felt that placing all of this information as well as EPS operational assumptions within the MM violated the separation of information precept that has driven much of the development. However, there is a significant potential for improvement here and mechanisms to enable this will be investigated in future work.

V. Conclusion & Future Work

An architecture to enable the automated control of spacecraft has been developed composed of a set of vehicle control functions that work to coordinate and monitor autonomous subsystems. This approach has been applied to the electrical power system for the purpose of developing a load schedule that maximizes the vehicle's resources to meet mission objectives and serves as an input the autonomous power controller for operation of the power subsystem. The communication protocol between the two systems has been developed and tested against a variety of nominal operational conditions and found to result in operation that observes all appropriate constraints. One such test case was presented here. Additionally, the communication algorithm was modified to enable the vehicle to respond to a selection of unanticipated disturbances to a currently executing plan and react in an appropriate manner.

13

This approach is intended to serve as a model upon which other subsystems can be constructed and interfaced with the autonomous vehicle-level control functions.

Significant work remains to continue the development of this technology. Of particular interest is the development of a mechanism to respond to both known system faults as well as unanticipated or multiple faults. It also remains to be seen if any convergence problems will occur when the vehicle planning problem is implemented across multiple subsystems. Further, many of the implementation challenges remain to be tackled, including minimization of communication bandwidth, robustness to communication noise, robustness to communication delay and asynchrony, as well as hardening against malicious users.

Acknowledgments

References

[1]J. Frank, L. Spirkovska, R. McCann, L. Wang, K. Pohlkamp, L. Morin, "Autonomous Mission Operations," 2013 IEEE Aerospace Conference, Big Sky, MT, March 2-9, 2013.

[2]S. Colombano, L. Sprikovska, V. Baskaran, G. Aaseng, R. McCann, J. Ossenfort, I. Smith, D. Iverson, and M. Schwabacher, "A system for fault management and fault consequences analysis for NASA's Deep Space Habitat," AIAA SPACE 2013 Conference and Exposition, September 2013.

[3]Soeder, J.F., Beach, R., McNelis, N., McNelis, A., Dever, T., May, R.D., "Overview of Intelligent Power Systems Development for Human Deep Space Exploration", to be published at IECEC, Cleveland, OH, July 28-30, 2014.

[4]J. Frank. When Plans are Executed By Mice and Men. Proceedings of the IEEE Aerospace Conference, 2010.

[5]E. E. Smith and D. J. Korsmeyer, "Intelligent Systems Technologies for Human Space Exploration Mission Operations" 2011 IEEE Fourth International Conference on Space Mission Challenges for Information Technology, pp.169-176, 2011.

[6]J. Fincannon, A. Delleur, R. Green, J. Hojnicki, "Load-Following Power Timeline Analyses for the International Space Station," Energy Conversion Engineering Conference, 1996. IECEC 96., Proceedings of the 31st, August 1996. doi: 10.1109/IECEC.1996.552868

[7]J.S. Hojnicki, "Computer Code Analyzes Electric Power System Performance," *Research and Technology 1991*, NASA TM-105320, p130-131, 1991.

[8]D. Lawler, L. Wang, et al, "Habitat Demonstration Unit Core Avionics Software," Johnson Space Center Biennial Research Report, pp. 138-139, 2011.

[9]H. Stetson, D. Deitsch, C. Cruzen, A. Haddock, "Autonomous Operations Onboard the International Space Station," IEEE Aerospace Conference, IEEE, Big Sky, MT, 2007.

[10]May, R.D., Loparo, K.A., "The Use of Software Agents for Autonomous Control of a DC Space Power System," to be published at IEEE EnergyTech 2014, Cleveland, OH, July 28-30, 2014.

[11]T.P. Dever, R.D. May, P.H. Morris, "Autonomous Spacecraft Communication Interface for Load Planning," NASA TM-2014-218362.

[12]Deep Space Vehicle Power Architecture Description Document and Deep Space Habitat Test Bed Evolution, AES-AMPS-RPT-002, January 15, 2013.

[13]SAE Space Power Standard AS5698, SAE Aerospace, 2/18/2012.

[14]J.J. Marquez, M. Ludowise, M. McCurdy, M., and J. Li, "Evolving from Planning and Scheduling to Real-Time Operations Support: Design Challenges," In Proceedings of 40th International Conference on Environmental Systems. Barcelona, Spain, 2010.

[15]Bresina, J., Jonsson, A., Morris, P., and Rajan, K. "Activity Planning for the Mars Exploration Rovers," In Proceedings of 14th International Conference on Automated Planning and Scheduling (ICAPS-05), 40-49, Menlo Park, CA: AAAI Press, 2005.

[16]Bresina, J., and Morris, P., "Mixed-Initiative Planning in Space Mission Operations". AI Magazine 20:2, 2007.

[17]P.H. Morris and J.L. Bresina, "Active and passive constraint enforcement for activity planning," International Symposium on Artificial Intelligence, Robotics and Automation in Space (ISAIRAS-08), Pasadena, CA, 2008.

American Institute of Aeronautics and Astronautics