A DESIGN FOR AN INTELLIGENT MONITOR AND CONTROLLER FOR SPACE STATION ELECTRICAL POWER USING PARALLEL DISTRIBUTED PROBLEM SOLVING

FT466751

Final Report

Prepared By: Robert A. Morris

Academic Rank: Assistant Professor

University & Department: Florida Institute
                           of Technology
                           Dept. of Computer Science
                           Melbourne, FL 32901

NASA/JSC Directorate: Engineering

Division: Propulsion and Power

Branch: Power

JSC Colleague: Mr. Tom Jeffcoat

Date Submitted: August 21, 1990

Contract Number: NGT-44-005-803

ABSTRACT

This paper describes work leading to the development of an automated electrical power monitor and controller for spacecraft secondary power systems such as the one anticipated for Space Station Freedom. This work integrates current research in parallel processing and distributed models for problem solving, with current research and development being conducted at NASA-Johnson Space Center on computer automation of spacecraft secondary power distribution and control. Although the emphasis here is on secondary power control, the computational model and representational framework is designed with generality, modularity, and extendibility in mind, and hence could be used in the design of a complete spacecraft monitor and control device.

This research is a response to the need for increased autonomy in spacecraft subsystem control. Electrical Power System (EPS) autonomy is defined to be the capability designed into the EPS which permits onboard task execution through monitoring, computation, command, and control functions with minimal or no ground intervention. The reason for EPS autonomy, especially on based missions with technologically sophisticated spacecraft such as Space Station Freedom, is clear. Already, autonomous load scheduling and fault detection and correction have proved to be plausible applications of intelligent (AI) models for decision making. At the same time, there has been a notable lack of success in integrating intelligent techniques into real time software. To make matters worse, the throughput requirements of electrical power control create even more difficult challanges to automation. It has generally been accepted that for spacecraft applications such as control, advanced software must be integrated with hardware and firmware, as exemplified in recent developments in microprocessor technology. In addition, developments in intelligent systems architecture, such as distributed knowledge representation and parallel processing, should be considered.

The starting point for this research is the Advanced Electrical Power Management Techniques for Space Systems (ADEPTS) project at NASA-JSC. The initial success of ADEPTS has demonstrated the viability of integrating knowledge based techniques for monitoring and control into a parallel processing environment. This research proposes to extend the ADEPTS model with three major goals in mind: first to include more intelligent decision-making capabilities; second, to more fully exploit the computational advantages of parallel processing; and third, to make the ADEPTS model for power management more generic, and hence capable of being applied to future space-bound vehicles such as space station. To fulfill these goals, a cognitive model based on cooperative experts is defined.

INTRODUCTION

The emphasis in this report is on defining a set of communicating processes for intelligent spacecraft secondary power distribution and control. The computer hardware and software implementation platform for this work is that of the ADEPTS project at JSC. The Electrical Power System design which was used as the basis for this research is that of Space Station Freedom, although the functionality of the processes defined here generalize to any permanent manned space power control application.

This report is organized as follows. First, the Space Station Electrical Power Subsystem (EPS) hardware to be monitored is described, followed by a set of scenarios describing typical monitor and control activity. Then the parallel distributed problem solving approach to knowledge engineering is introduced. There follows a two-step presentation of the intelligent software design for secondary power control. The first step decomposes the problem of monitoring and control into three primary functions. Each of the primary functions is described in detail, in subsequent sections. This report concludes with suggestions for refinements and embelishments in design specifications.

SPACE STATION SECONDARY ELECTRICAL POWER CONTROL

Monitoring and Control (M&C) involves observing and guiding the behavior of a large physical system towards some objective (ref. 1). Intelligent monitoring and control involves the integration of perceptual, cognitive and reactive knowledge into the monitor and control process. M&C tasks that require such knowledge include: interpretation and prediction of system behavior, reasoned response to important events, diagnosis of exceptional events, and planning long term courses of action.

Secondary electrical power distribution hardware on space station comprises the following elements: DC-to-DC converter units (DDCU), which output 120 volts of DC power, Secondary and Remote Power Distribution Assemblies (SPDA, RPDA) which contain a collection of Remote Power Controller Modules (RPCM) for switching, 1553 buses for transmitting data, power lines, and a set of loads. These elements are arranged in a hierarchical framework to allow for more effective hardware control.

Space station M&C offers challenges which do not arise on either ground-based system control or control on short-term missions such as Space Shuttle. In particular, electrical power will be a limited resource, requiring scheduling on the basis of load priority. In addition, EPS hardware maintenance will need to be performed on potentially dangerous electrical equipment by relatively inexperienced crew members. Finally, the 30-year duration of space station makes continuous ground monitoring

economically and logistically inefficient. Each of these
considerations strongly suggest a need for an increased role for
automation.

Current design for space station involves dividing the
overall EPS M&C process into a hierarchical configuration of
tiers (Figure 1). Tier I consists of Operations Management
System (OMS) control, including software and hardware for
overall ground and onboard management of Space Station
operations. Tier II consists of a set of Standard Data
Processors (SDP) containing 80386- or 486-based processing
capabilities, running concurrently to perform functions such as
Guidance, Navigation, and Control (GN&C), data management (DMS),
storage unit control, and Node and ITA management. Currently,
the Node/ITA SDP and DMS SDP are jointly responsible for Tier II
secondary power control. Communication among the Tier II
managers and the OMA are via a Fiberoptic Distributed Data
Interconnect (FDDI) bus. A Tier III secondary power controller
(SPC) provides local control and protection for the various
Space Station subsystems. Each SPC, as currently defined, will
consist of a Multiplexer-Demultiplexer (MDM), and functions for
controlling input/output and self-monitoring. Finally, Tier IV
control consists of a hierarchy of Secondary Power Distribution
Assemblies (SPDAs) and Remote Power Distribution Assemblies
(RPDAs) which are collections of Remote Power Controller Modules
(RPCMs). The RPCM switches and distributes 120-Vdc power via a
set of unidirectional, solid-state switches. Each RPCM provides
deadfacing features, remote manual off capabilities, hardware
setpoint programmability and built-in testing features (ref. 2).
Each RPCM is assigned a 1553b data bus which responds to
external control signals and transmits data and status
information within 200 ms of application of control power. These
local features of the RPCMs allow for a real time response to
fault conditions.

Any proposed advanced automation for space station must
integrate with standard control firmware and software. In
particular, this report describes software that would naturally
migrate to Tier III MDM operations, although certain functions
could more properly be assigned to the Tier II level. Certain
constraints have been placed on application software at this
level, due to two primary factors: the desire to preserve the
main function of the MDM, which is high throughput of sensor and
effector i/o to the higher control levels, and to maintain
compatibility with MDM software. The primary reason for adding
intelligence to the MDM is to cut 1553 bus traffic to the higher
Tiers (viz., to the SDPs and OMA) (ref. 3). These constraints,
among others, are major factors influencing the design decisions
made in this report.

To summarize, the intelligent secondary power system
controller proposed here, if used on space station, would
monitor, evaluate and control performance from the DDCU units
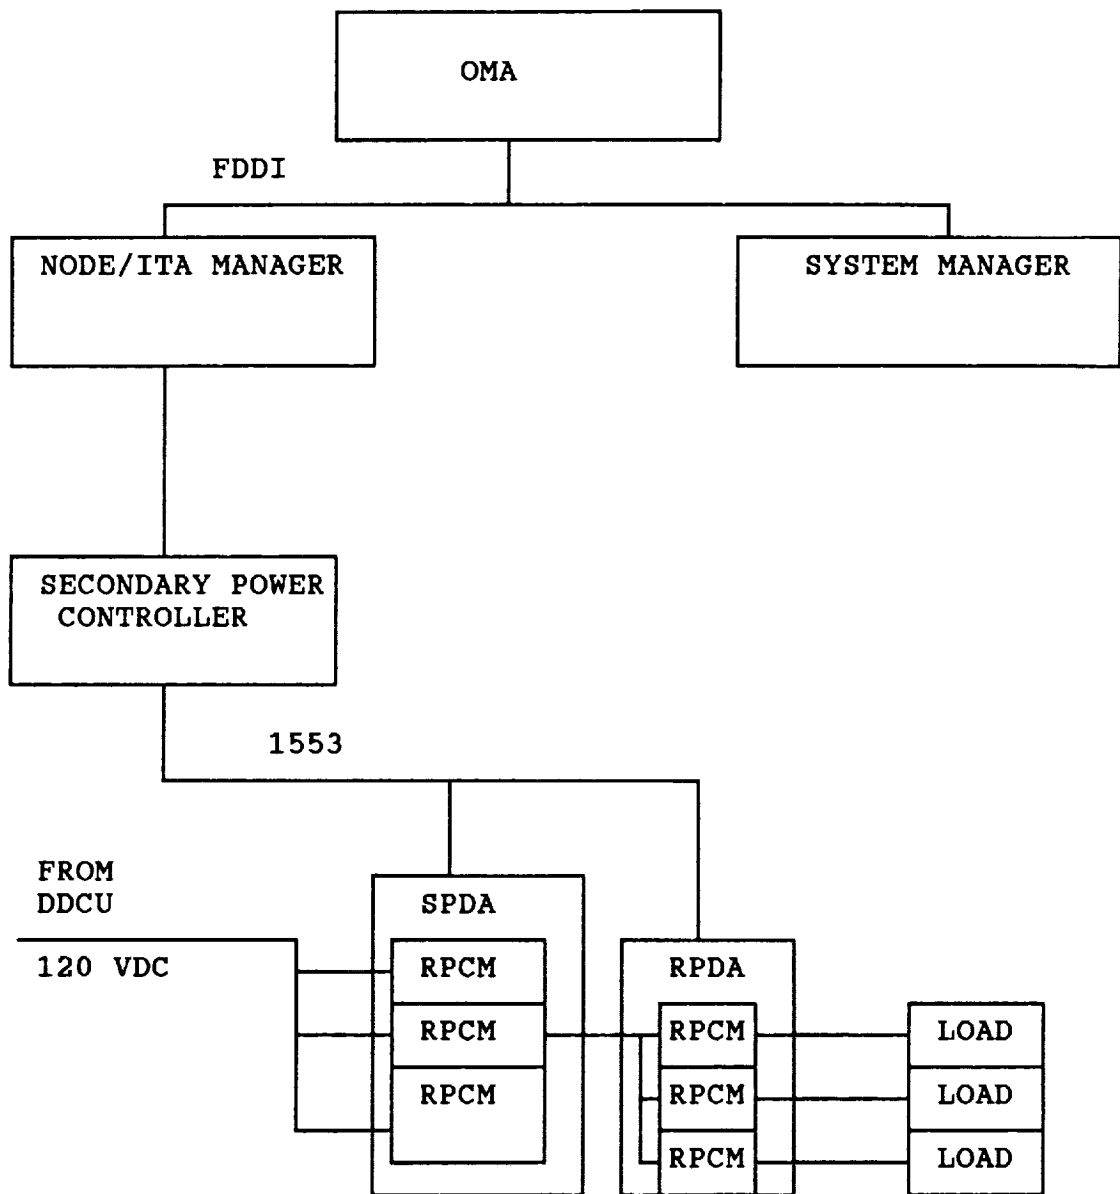to the interface to the loads. It should be viewed as an

FIGURE 1. SPACE STATION POWER DISTRIBUTION
AND CONTROL HIERARCHY

extension of Tier II and Tier III control operations. This controller would therefore have one instantiation for each space station node or element.

SCENARIOS FOR CONTROL

In this section we present four scenarios involving the secondary power distribution hardware for which M&C activities are required. These scenarios are based on a four-fold classification of power system states into normal secure, normal insecure, emergency, and restorative. These states are defined in terms of two sets of constraints: load constraints and operating constraints (ref. 4). Load constraints are requirements for power to the loads. Operating constraints are upper and lower limits on voltage, current, and other parameters of the system.

NORMAL SECURE STATE. In a normal secure state, both operating and load constraints are met. While in this state, load schedules are being dispatched and executed. Status information from the RPCMs and other hardware is monitored and analyzed to check for potential fault conditions.

NORMAL INSECURE STATE. In this state, both load and operating constraints are met by the power system, but there is a likelihood of a transition to an emergency state. Some violations caused by a contingency cannot be corrected without load constraint violations. Preventive rescheduling of loads may be needed to return the system to a secure state.

EMERGENCY STATE. Transition to this state is caused by major violations of operating constraints. Redundant power is currently being used if available, but there is serious loss or degradation of power. While in this state, remedial action or corrective rescheduling is needed to return the system to a normal state.

RESTORATIVE STATE. In this state, only the load constraints are violated, typically because of a need for load shedding and reconfiguration. Actions are required to restore power to the effected loads.

The functionality of the intelligent power controller is based on the need to perform the activities described in these scenarios. It is clear that a number of perceptual, cognitive, and responsive capabilities are involved. When performed by mission control personnel, spacecraft M&C is a cooperative distributed activity performed by a hierarchical committee of flight controllers led by the ``front room'' controllers, with support and backup by power system experts in the ``back rooms.'' In addition, the power system hardware will be

geographically distributed on space station, making distributed, localized on-board automated control a logical approach to intelligent automation. For these reasons, we propose a software model of this activity in terms of a collection of processing agents. This model is presented in the next section.

GENERIC CONTROL MODEL USING DISTRIBUTED PROBLEM SOLVING AND PARALLELISM

Distributed problem solving involves cooperation by a decentralized and loosely connected group of processing units called agents (ref. 5). The problem solving is cooperative in that communication of data and task sharing are required in order to come up with a solution. The problem solving is decentralized in that knowledge, computational resources, control, and data are distributed.

Many distributed problem-solving architectures have appeared in recent years. Among these are blackboard systems, contract net systems, and multi-agent planning systems. These systems are based upon metaphors of social, economic, or biological organization. There are many dimensions along which to categorize these systems. These include:
1. system metaphor (individual, society, committee);
2. grain size (how much processing power is in each of the smallest units);
3. system size (small, medium, large);
4. agent characteristics (controlled or independent, restricted or unrestricted in their resources, interactions simple or complex); and
5. techniques for attaining results (problem solving methods).

Space limitations forbid a detailed discussion of each approach and methodology. Rather, we proceed to present a specific distributed model for intelligent monitoring and control.

For reasons relating to problem domain, implementation, and space station design constraints, the model proposed here is based on the idea of a small, loosely-coupled group of semi-autonomous agents cooperating to maintain the proper distribution of electrical power to each of a set of loads. The actual number of agents will depend on factors such as desired throughput, the number of independent subtasks of the primary tasks that can be discerned, and hardware restrictions. These factors will emerge as the proposed model is refined and instantiated in further research. For now, the proposed model recognizes at least three separate communicating agents.

Partitioning work activity should insure an equal distribution of processing labor; otherwise, the benefits of the distributed approach is lost. The organizational topology of our control model consists of a two-level decomposition of tasks (Figure 2). On the first level, the problem of secondary power

```
                    ┌─────────────────────────┐
                    │  INTELLIGENT            │
                    │  POWER CONTROL          │
                    └───────────┬─────────────┘
                                │
         ┌──────────────────────┼──────────────────────┐
┌─────────────────────┐ ┌─────────────────────┐ ┌─────────────────────┐
│  FAULT MANAGER      │ │  POWER MANAGER      │ │  LOAD MANAGER       │
│                     │ │                     │ │                     │
│ *FAULT DETECTION    │ │ *MONITORING         │ │ *LOCAL PLANNING     │
│ *FAULT ISOLATION    │ │ *SWITCHING          │ │ *HEALTH ANALYSIS    │
│ *FAULT RECOVERY     │ │ *RECONFIGURATION    │ │                     │
└─────────────────────┘ └─────────────────────┘ └─────────────────────┘
```
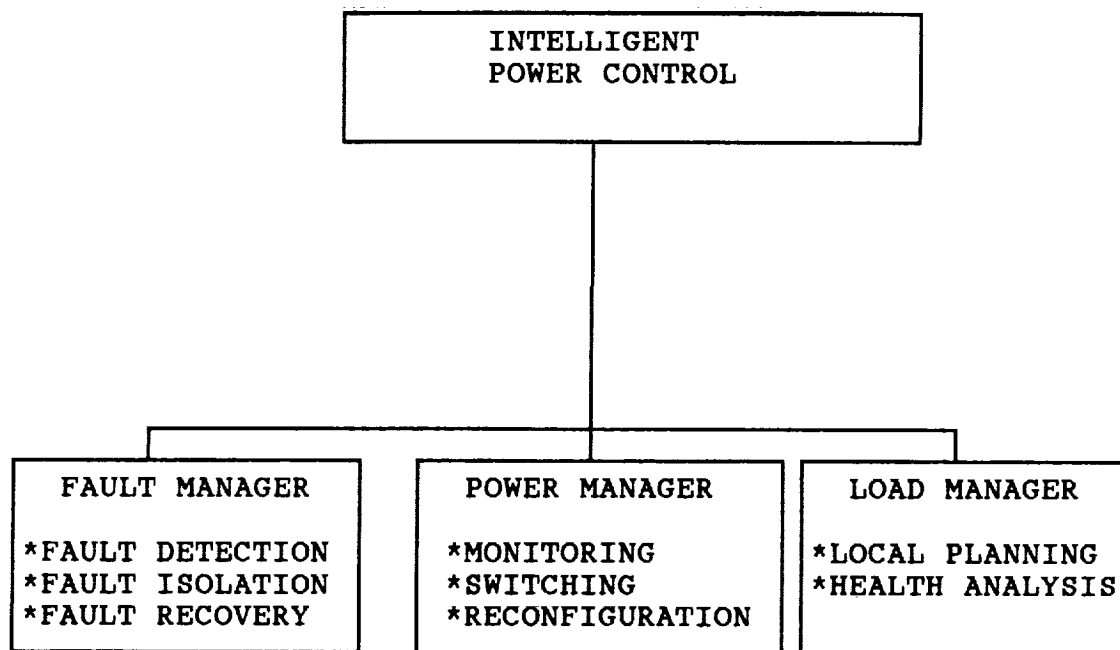
FIGURE 2. POWER CONTROL TASK PARTITIONING

control is divided into three primary tasks: power management, fault management, and load management. This division is based on a decomposition of monitoring and control functions in terms of a three-fold temporal partitioning of the events that make up the behavior of the power system. Events that occurred in the past include exceptional or urgent events that require diagnosis, and are of interest to the fault manager. Events that are presently occurring include the flow of electrical power and relay switching. These events need to be monitored and controlled, and will be the assigned task of the power manager. Finally, future events are those which need to be planned or predicted; in particular, loads need to be scheduled and possible unhealthy conditions predicted. This will be the job of the load manager.

On the second level, each primary task is subdivided into a collection of perceptual and cognitive operations. For example, the fault manager consists of procedures for gathering and evaluating voltage and current information in order to perform diagnostics. Each of these subfunctions are described further in the next section.

As mentioned, in this report, we will view each primary task as an intelligent agent. Further refinements of this model may result in assigning to each primary task a network of agents, depending on whether certain subtasks within a primary task may be viewed as independent perceptual or cognitive activities. Furthermore, each primary agent can be viewed as either operating concurrently on a separate processor, or as sharing processing resources. In the discussion that follows, we will assume for simplicity that the agents run continuously on their own processors. This assumption allows for a consideration of the potential for parallelism in the monitor and control process.

Two sorts of parallelism can be distinguished: in one, a set of necessary and independent tasks are identified and processed concurrently (AND-parallelism). The other (OR-parallelism), is the result of the non-determinism involved in intelligent problem solving, i.e., in the potential for solving the same problem in a number of different ways. As will be shown, both kinds of parallelism are inherent to the problem of monitoring and control.

As mentioned, in choosing this preliminary model, a number of constraints have been applied. Some of these involve integrating distributed conventional control hardware and software with intelligence. Other constraints involve the nature of the process of spacecraft monitoring and control itself, which, when solved by humans, typically involves a small group working cooperatively on the problem. Finally, the desire to implement the model using the transputer environment, as envisioned by the ADEPTS designers, was a consideration in constructing the model.

PARTITIONING KNOWLEDGE AND EXPERTISE AMONG LOCAL AGENTS

This section contains a generic description of each of the three primary agents in terms of four characteristics: the agent's functional components; responsibilities; resource requirements (software, hardware, and communication); and interactions (i.e., which agents talk to which others, and what do they exchange). In addition, there is a brief informal discussion of the potential for AND- and OR-parallelism within each agent's activity. Finally, we describe briefly how each of these generic functions could be instantiated on space station.

Local Fault Management

Responsibilities. A fault manager is responsible for isolating, identifying, and aiding in the recovery from faults involving power system components. The problem of fault diagnosis can be stated thusly: given a set of symptoms, find a set of faults to account for them. The problem of fault identification is complicated by the fact that different faults may share symptoms, and also by virtue of the possibility of multiple faults. The amount of reasoning skills required by fault diagnostic software is inversely proportional to the amount of sensor data to be processed. Therefore, by adding an intelligent fault manager to system software, data processing can be potentially reduced.

Functional Components. Every intelligent fault manager consists of the following:
1. A procedure for gathering data
2. A simululator of the behavior of the artifact
3. A diagnostic problem solver, consisting of:
   i. A procedure for recognizing discrepencies
   ii. A procedure for generating hypotheses
   iii. A procedure for testing hypotheses

To perform fault management, data must be constantly gathered and tested against the expectations generated by the simulator. When a discrepency is identified, a symptom set is generated. The model is consulted to generate hypotheses about the cause of the discrepency. Each hypothesis is tested by observing further data, or by injecting the hypothesis into the simulated model of the system, generating simulated symptoms which are further matched to the actual data. This procedure is repeated until a unique hypothesis is generated.

Resource Requirements. Input to this agent will consist of voltage and current data from remote controller modules. The knowledge sources of this agent include the model of the components of the subsystem under its control, including the buses, switchgear, and loads. Its primary local data are current symptom sets.

Interactions. This agent can perform its major functions locally. Diagnostic information, i.e., data identifying faulty components; the nature of the fault; and recommendations for recovery procedures; should be communicated to the local power manager, as well as the load manager. The load manager must use this information in plan generation. In addition, the human user of the system should be alerted and recommendations for recovery enumerated.

Potential for Parallelism. AND-parallelism is inherent in the relative independence of the subtasks that are performed by the fault manager. For example, the data analysis and problem solving phases can be run concurrently. The fault manager activity also allows for some degree of OR-parallelism, e.g., when multiple hypotheses are generated from symptom sets. Each hypothesis can be tested separately against current data. This parallelism could be useful if multiple fault detection is required by the fault manager.

Space Station Instantiation. Distributed diagnosis is particularly suited to space station design. The point of local diagnosis is to reduce the flow of status data from the component controllers to the upper level managers. Furthermore, the purpose of intelligent automated diagnosis is to reduce the number of points that need to be sensed in order to make a diagnosis, thus reducing the data processing requirements.

Power Manager

Responsibilities. The power manager is responsible for maintaining the proper flow of electrical power to the loads. This involves controlling the allocation of power resources to meet the current system load. The cognitive requirements for the power manager consist in the ability to compute changes in configuration of the system on the basis of load priority, which includes the ability to determine which loads must be shed. It is also envisioned that the power manager will share the local monitoring tasks with the fault manager.

Functional Components. The activities of the power manager can be broken down into the following functions:
1. Configuration Processor
2. Resource Controller
3. Health Monitor

The configuration processor computes switch settings for the remote power controllers based on current load demands and load priorities. These load demands may either result from normal execution of a load schedule, or by an exceptional event such as a component fault or addition of a load. For example, if power to a critical load has been switched automatically to a redundant path, and if a part of that path has already been fully allocated to other loads, then the power manager will compute which loads to shed along the redundant path.

16-11

The other functions of the power manager may consist of more conventional algorithmic software. The resource controller distributes commands to remote controllers (RPCMs) for switch settings. The health monitor will periodically perform analysis of the data from the remote power controllers.

Resource Requirements. Like the fault manager, the power manager will contain component and behavior knowledge of the power system. In addition, it will require at least two small databases: one, consisting of current switchgear configuration information, which will be updated as load requirements change; the other, a temporal database, which will allow for the maintenance of health information. Finally, it must contain load priority data, in order to compute power allocation.

Interactions. The power manager will interact both with the fault manager, as described above, and the load manager. The load manager will request information about current power availability in order to verify or update a load schedule. In addition, configuration updates will be reported to users.

Potential For Parallelism. The three power manager subtasks are relatively independent, which suggests a source of AND-parallelism. In particular, the monitoring task can be performed concurrently with the other tasks. OR-parallelism may be found in the computation of configurations.

Space Station Instantiation. The power manager will allow for more load shedding decisions and health monitoring decisions to be made locally, thus reducing the volume of OMA-directed activity.

Load Manager.

Responsibilities. Scheduling loads can be a highly complex problem. Distributing this responsibility by localizing the processing would lessen the bottleneck associated with generating load schedules. A promising approach to partitioning the problem solving knowledge and control is by distinguishing between the static and dynamic aspects of scheduling. Static scheduling involves taking a set of loads and constraints and generating a consistent, usable schedule. Parameters and constraints which are used to develop a static schedule involve global knowledge of each EPS function, from generation and storage to distribution. As such, the design and implementation of the static scheduling procedure would occur on Tier I or Tier II level control. Consequently, this problem is beyond the current scope of the subsystem controller being designed here.

The dynamic scheduling procedure has the function of modification of proposed schedules on the basis of the addition or deletion of resource or load constraints. These new constraints may arise as the result of damage to EPS equipment, additional load requirements, or other resource constraints such as crew or instrument availability. This, we propose, is a function that can be localized and distributed to a Tier III-

type controller. This dynamic load scheduler will be called a (local) load manager.

Functional Components. A load manager will verify and modify schedules and plans based on current resources. The need for such a function arises due to expected frequent changes in resource availability. Schedule verification consists of ensuring that the power and other resource expectations demanded by the static schedule can be currently met. For example, the static schedule may assume that a certain activity A, can be assigned an interval of time T with power requirements P, instrument I and crew members C D and E. The contingency handler will verify that no changes to any of these resource requirements has occured since the static schedule was generated. If current resources do not match load requests, the load manager will attempt to modify the schedule in order to satisfy the requests.

Resource Requirements. The data and knowledge sources for this module include a load profile database for each load, a resource database of currently scheduled resources, and a database consisting of current schedule information.

Interactions. Rescheduling may involve reconfiguration and load shedding. Therefore, it will be necessary for the load manager to request data and share tasks with the power manager. It will also receive and send schedule information from/to human planners.

Potential for Parallelism. Scheduling can be viewed as a search problem. Parallel versions of depth-first search algorithms are being developed, which are essentially examples of exploiting OR-parallelism.

Space Station Instantiation. Because of the complexity of space station planning, it is less clear to what extent a local planner can be employed on Tier III-level software. It is likely that automated planning on station would be a Tier I or II activity. It was suggested that a Tier III load manager could be an ``intelligent interface'' to human planners, notifying them that planned loads can or cannot be powered, and perhaps suggesting contingency short-term schedules.

Summary of Preliminary Model

The model described here is in a preliminary stage of development. More is required in order to define the knowledge sources of each agent, the communication and synchronization mechanism among the primary agents, and the potential for concurrency among the agents' subfunctions.

A useful metaphor for distributed power control is human organization. To avoid control bottlenecks among upper level managers, it is useful to distribute intelligent decision making among lower-level managers. On the other hand, too much local control could result in the lack of a common direction among the different local units. The issue of the amount of local

automated control will need to be addressed in applying the generic model proposed here to space station.

SUMMARY AND FUTURE DIRECTIONS

This paper has sketched a design for an intelligent secondary power controller, based on current Space Station Freedom secondary power distribution hardware specifications, as well as on current developments in cooperative problem solving methodology. The model has arisen from conversations with station designers of EPS hardware, the OMS, data management system (DMS), and application software. These developers were helpful in detailing the functionality of the different software modules, and suggesting avenues for intelligent automation.

A number of improvements and extensions of this work are needed. First, as mentioned, the model is in need of significant further development. A prototype of the model should be built and tested, and ported to the ADEPTS transputer implementation environment. Currently, an implementation of a prototype of this model is being developed.

Second, more agents should be added as needed. For example, a separate data pre-processing agent could monitor sensors and translate and filter sensed data to the appropriate manager. In addition, an intelligent user interface could be added for displaying data to users, such as crew members, with varying expertise in the problem domain.

Third, more functional autonomy could be added to each agent. Each agent could possess more sophisticated cognitive and decision-making capabilities. Our design favored a more modest approach to agent autonomy, in line with current opinions regarding SPC MDM processing capabilities. As early prototypes are tested and validated, more functionality can be added.

Fourth, a more detailed investigation of the potential for AND- and OR-parallelism should be undertaken. Finally, the committee of agents defined in this report could be extendned to include more global (i.e., Tier I or Tier II) control operations, such as static scheduling and user interfacing. With few modifications, the generic model could also be instantiated for primary as well as secondary control operations.

It is believed that the model defined here could be useful to space station designers in the automation of space station power control. Instantiating the generic model for space station will imply tailoring the roles of each manager to the requirements of a Tier III or Tier II controller. As the model proposed here and the space station automation process progress, the proposed software could ``migrate'' on board in later growth stages. Finally, the overall model advanced here can be applied to future automation projects such as for lunar and Mars bases.

ACRONYMS

ADEPTS	Advanced Electrical Power Management Techniques for Space Systems
DDCU	DC-to-DC Converter Units
DMS	Data Management System
FDDI	Fiberoptic Distributed Data Interconnect
FDIR	Fault Detection, Isolation and Recovery
ITA	Integrated Truss Assembly
M&C	Monitoring and Control
MDM	Multiplexer/Demultiplexer
OMA	Operations Management Applications
RPCM	Remote Power Controller Module
RPDA	Remote Power Distribution Assembly
SDP	Standard Data Processor
SPC	Secondary Power Controller
SPDA	Secondary Power Distribution Assembly

REFERENCES

1. Hayes-Roth, B.; Washington, R.; Hewett, R.; Hewett, M.; and Seiver, A.: Intelligent Monitoring and Control, (Proceedings of the International Conference on Artificial Intelligence, IJCAI-89), Morgan Kaufmann Publishers, Inc., 1989, pp. 243-249.

2. Interface Development Document: Remote Power Controller Module Standard (Preliminary), Space Station Freedom Electrical Power System, Work Package 04 Document, SSP 44004, March 20, 1990.

3. Systems Engineering and Integration Data: Multiplexer/ Demultiplexer (MDM) DMS PDR#3 Support Data, Work Package 2 (WP-2) MDC H4778, March, 1990.

4. Talukdar, S.; Cardozo, E.; Leao, L.: Toast: The Power System Operator's Assistant, Proceedings of the IEEE, Vol. , No. July, 1986, pp. 53-59.

5. Gasser, L.; Huhns, M.H.: Distributed Artificial Intelligence Volume II, Morgan Kaufmann Pub., Inc., 1989.