

A Human-Centered Approach to Monitoring Complex Dynamic Systems

FINAL REPORT
NASA GRANT NCC2-1220
DECEMBER 31, 2004

DANIEL SHAPIRO
DORRIT BILLMAN
MEI MARKER
PAT LANGLEY

Institute for the Study of Learning and Expertise
2164 Stauton Court, Palo Alto CA 94306 USA
{shapiro, billman, langley}@isle.org

Abstract

This paper reports on a human-centered approach to monitoring complex systems which offers the fidelity of methods that compare predictions against observations with the emphasis on higher-level user feedback found in more inferential techniques. In particular, we cast our work as a decision support tool and we assume that the monitoring system should reason over models that are natural to human users. This has led us to adopt several design tenets: the structure of the model should follow the componential and causal structure of the device; feedback from the system to the user should be tied to the device model; device behaviors should be represented as a set of processes with associated equations; and every process in the model should correspond to some behavior of the device. We claim that these commitments support an intuitive and effective design for a fault detection, isolation, and diagnosis tool. We discuss the knowledge representations, algorithms, and interface capabilities of this system, and we present a scenario of its use that illustrate our claims in the context of monitoring the electric power system on board the International Space Station. We conclude by discussing an evaluation study of the monitoring tool, including its effect on users' diagnostic strategies.

1. Background and Motivation

As engineering technology grows, our ability to employ complex artifacts is limited by our inability to understand them. This difficulty arises for several reasons. First, complex devices – such as high performance aircraft and nuclear power plants – perform complex functions. This requires many interacting parts, which makes it hard to organize knowledge about their structure. Second, complex devices interact with people and the environment to generate a wide variety of possible behaviors, which makes it difficult to discern correct from abnormal operation. Third, complex devices are typically instrumented to produce large quantities of data. This leads to an intricate diagnosis problem, as the information must be analyzed in order to synthesize an integrated picture of what the device is doing, whether it is functioning properly, and, if it is not, how and why it is misbehaving.

The International Space Station is a case in point, and we will use it as an example throughout this paper. Its complexity comes from the fact that it is simultaneously a self-contained habitat, a self-propelled vehicle, and an experimental laboratory. As a result, it contains many subsystems, including power generation, power storage, power distribution, thermal control, atmospheric maintenance, navigation, propulsion, and attitude control, as well as a schedule of demands that often press the Space Station’s performance limits. Its component systems also interact, meaning that a fault (say) in attitude control may show up in power storage. The evidence is contained in an extensive telemetry stream, as the electric power system alone generates approximately 50,000 measurements every ten seconds. Specially trained staff members in Mission Control must examine these data in order to detect and isolate faults. Moreover, the ever-present pressure to reduce costs means that fewer, and less expert, personnel will perform this analysis in the future. NASA clearly needs tools to support the monitoring process.

The monitoring problem, broadly construed to include fault detection, isolation, and repair, has generated a large body of research within AI and related disciplines. Here we distinguish approaches that use background knowledge to infer a relation between observations and faults from approaches that operate by comparing observations with the expectations of a predictive model.

Inference-oriented approaches have grown out of the expert and knowledge-based systems movements. For example, Sermatech’s TIGER system [2] applies production rules to detect anomalous conditions from observations, diagnose/identify the affected system, and abstract the results into a single fault conclusion. This system has been used to monitor motors, generators, turbines, transfer switches, and many other devices. In contrast, Larsson’s [15] GoalArt takes a more model-oriented perspective by describing devices as goal-driven machines and employing a model to map observations into a high-level analysis of system properties. GoalArt has been used to monitor patients in intensive care and is being applied to diagnose faults in nuclear power facilities. In general, inferential mechanisms offer a significant ability to supply analyses at a level of abstraction suitable for human users (i.e., by providing output in terms of lost capabilities vs. anomalous variable values) and their applications are typically cast as decision support tools.

Approaches that compare predictions against observations have emerged from a synthesis of ideas in qualitative physics and operations research. In general, these methods function by finding models

that are consistent with observations and commands; they detect faults when the data are inconsistent with nominal models and they diagnose problems when explicit fault models successfully fit the data. However, the task of matching models to data becomes computationally complex in the presence of multiple device components with multiple behavioral modes. In response, researchers have examined many types of models that capture qualitative behavior [9] [7], semi-quantitative behavior [10] [24], and hybrid continuous/discrete dynamics [21] [22] [12]. These model-based mechanisms also extend to fault repair, viewed as the selection of control inputs that drive a system to a preferred mode. The Livingstone system [28] has had been successful in this regard, as it performed fault detection, isolation, and repair on NASA’s Deep Space One mission. While these techniques provide a great deal of power, the research has focused historically on fully autonomous diagnostic systems. As a result, the underlying representations are instrumental for diagnosis and repair, but not for communicating analyses to human users. We will take this as the main point of departure for our work.

This paper describes a human-centered approach to monitoring complex systems that combines the fidelity of model-based methods with the high-level user feedback associated with inferential techniques. In particular, we cast our work as a decision support tool and we make the commitment, expressed in Langley et al. [14], that such systems should reason over models that are natural to human users. In response, we have adopted four design principles:

- the model structure should follow the componential and causal structure of the device;
- feedback from the system to the user should be tied to the device model;
- device behaviors should be represented in terms of processes with associated equations; and
- every process in the model should correspond to some behavior of the device.

We claim that these commitments have led to an intuitive and effective design for an interactive tool that support fault detection and diagnosis, which we present in the pages that follow.

We begin by introducing our approach in Section 2, then turn to our representational framework for modeling device structure and behavior in Section 3. After this, Section 4 discusses our method for making model-based predictions using this framework, while Section 5 illustrates our system’s fault detection and diagnosis capabilities. Given this background, Section 6 presents the graphical interface that gives users access to this monitoring functionality, including a sample scenario that corresponds to an existing fault in the Space Station’s electric power system. In Section 7, we describe an experiment designed to to evaluate the system’s ability to support diagnostic reasoning in non-expert users. We conclude by placing our work on human-centered monitoring in the context of related efforts, in Section 8, and summarizing the contributions and limitations of our approach, in Section 9.

2. Overview of the Approach

Our approach to monitoring is straightforward. In abstract terms, it detects faults by comparing observations against the predictions of a device model and signaling an alert whenever a significant

discrepancy occurs. Because we have gone to some length to construct accurate models, both manually and via machine learning [3], we can attribute alerts to a fault in the device as opposed to a fault in the model. In addition, because we establish a close correspondence between the structure of the model and the structure of the device, we can localize a problem within the device by associating it with a specific component, subcomponent, property, and behavior within the device model.

In more detail, we characterize devices as hierarchical collections of systems, where each system is composed of processes which act on variables that may belong to further subsystems. Each process contains activation conditions and equations that express causal influences among variables and, since the mathematics supports both algebraic and differential equations, the model makes numeric predictions over time. Our modeling discipline associates device components with systems, device subcomponents with subsystems, device states with variables, and device behaviors with processes. Thus, sets of processes are analogous to the modes of model-based diagnostic techniques and, as the activation conditions gate behaviors, they capture qualitative transitions among modes.

As a whole, a model in this language represents the causal structure of a device and compiles into a directed graph connecting variables by processes. We use this internal form to support fault detection and diagnosis. In particular, a fault is a deviation between a predicted and observed value of a variable that exceeds some threshold. We localize faults by tracing the graph from symptoms towards root causes, and we isolate faults to particular processes (device behaviors) by comparing their expected influences (i.e., their input/output transformations) against the observed input/output relations where possible. We treat fault diagnosis as the task of fitting a model of a faulty component to the data stream.

Our user interface presents fault data via a graphical display that color codes affected elements of the model. The model’s hierarchical structure provides a natural abstraction mechanism in this regard, since the presence, but not the nature, of a fault can be made visible in enclosing systems. Other user-interface operations support navigation through the hierarchical model, navigation through the time history of a monitoring run, navigation through the causal graph connecting symptoms and causes, and projection of both nominal and fault models into the future, together with displays of predicted and observed variable values. The goal is to provide the user with the ability to detect problems, rapidly focus his or her attention, and delve into detail when those problems require special attention.

Before we enter into a detailed discussion of our system, we should point out that, unlike many of the model-based systems referenced earlier, it does not currently perform probabilistic mode identification or propose fault repairs. While this represents a limitation with respect to fully automated techniques, there is reason to believe that attention to probabilistic effects might actually complicate the monitoring process in interactive settings, in comparison with a deterministic (but retractable) approach to prediction and analysis. In addition, our use of a deterministic framework has let us focus on the human-centered aspects of the monitoring problem.

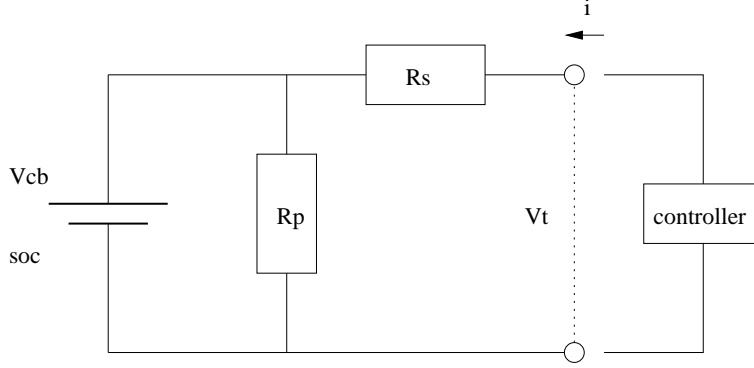


Figure 1: An abstract depiction of a battery model.

3. Modeling Complex Artifacts

In order to model complex devices, we have designed a language for representing their behavior. For clarity, we divide our discussion into two parts that reflect the historical development of the language. The first subsection discusses the primitive elements of the language and their use in modeling simple device behavior. The second adds hierarchy to this picture and shows how to aggregate multiple behaviors into a depiction of abstract devices.

3.1 Quantitative Process Models

The primitive elements of our modeling language represent device behaviors in terms of *variables* and *processes* that act on those variables. A variable encodes device state, like the charge on a battery. A process describes one or more causal relations between input and output variables, such as charging, which increases battery charge in response to an applied current. We state these relations in terms of differential equations (if the process describes change over time) or algebraic equations (if it describes instantaneous effects). A process may also include conditions, expressed as threshold tests on its input variables that identify when the process is active. A process model consists of a set of processes that link observable input variables with observable output variables, possibly through unobserved theoretical terms.

We use a model of the batteries onboard the International Space Station as a simple example. Figure 1 illustrates how an electrical engineer might conceptualize the batteries in a circuit diagram: the battery is a voltage source (V_{cb}) connected to a resistor in parallel (R_p) and in series (R_s). The parallel resistor models spontaneous discharge (as batteries slowly lose their charge over time). The series resistor models internal losses as charge is transferred into and out of the battery.

However, this diagram does not express all of the behavior we might want to capture. In particular, it lacks a quantitative description of currents and voltages in the circuit at any one time, in addition to any representation of the battery's dynamic behavior. For example, its state of charge varies over time and its ability to absorb or dispense charge varies with state of charge. We capture this behavior in the equations attached to processes. Table 3.1 provides a more detailed process

model of the battery in textual notation.¹ It begins by declaring the variables and parameters (constants) involved in the model, then lists the processes representing behavior with their associated conditions and equations. The first four variables – V_{cb} , V_t , soc , and i – describe physical properties of the battery. V_{cb} is a theoretical term corresponding to the voltage of the battery before any losses occur to internal resistances, while V_t represents the actual voltage at the battery terminals. The variable i corresponds to the current, where positive values indicate flow into the battery by convention, and soc represents the total amount of electric charge stored in battery, expressed as a percentage of a nominally full battery. A fifth variable, *charging*, provides a control signal that selects among battery modes. (This battery is either charging or discharging, as the controller never connects it to power sources and loads at the same time.) The battery model also contains three parameters: R_p is the resistance in Ohms of the parallel resistor, R_s represents the battery’s serial resistor, and R_{load} is the effective resistance applied by the controller during discharge.

The processes in Table 3.1 describe causal interactions between the variables of the battery model. SourceVoltage denotes the behavior of the battery as a voltage source before any losses to internal resistances. This voltage is generally not constant in real power supplies. We model it as a linear function of state of charge, which provides good predictive performance when compared with actual telemetry data [3]. The ChargeTransfer process describes the effect of current flow on battery charge as a differential equation that integrates current flow across time periods (where the scaling constant converts Amperes into state-of-charge units). The process SelfDischarge represents spontaneous battery leakage; like ChargeTransfer, it also has soc as a dependent variable. Whenever two or more processes influence the same variable, we assume the effects are additive.

The last three processes in Table 3.1 characterize the behavior of the battery under an applied load or power source. Here we define charge and discharge as mutually exclusive modes by testing the whether control flag, *charging*, has value 1 or -1 within each condition clause. DischargeCurrent captures the batteries’ ability to supply current under an applied (and internal) load, including an indirect dependence on state of charge (through V_{cb}). DischargeVoltage performs the analogous calculation for voltage. Finally, ChargeVoltage determines battery voltage under an applied current that is selected by the battery controller.

Taken together, these processes specify a causal model of battery function that represents its response to applied current or loads. This particular model has also been tuned to fit telemetry data taken from the batteries onboard the Space Station. The tuning mechanism involves a machine learning method that searches over a space of feasible battery models while balancing accuracy against complexity [3]. Although a discussion of that mechanism is beyond the scope of this paper, it produces a model that can be trusted to represent actual battery behavior for use in monitoring.

3.2 Hierarchical Process Models

Although the process modeling language can capture the behavior of simple devices, the ability to represent and monitor significantly more complex systems requires extensions to the language. In particular, we need methods for organizing our knowledge about devices so that we can appreciate

¹We discuss a graphical representation of process models in Section 6.

Table 1: A simple process model for a battery.

```

Model Battery;
Variables soc, Vcb, Vt, i, charging;
Parameters Rp = 100, Rs = 0.214, Rload = 2.5

Process SourceVoltage;
    equations Vcb = 36.2 + 76.2 * soc;

Process ChargeTransfer;
    equations d[soc,t,1] = 10 * i;

Process SelfDischarge;
    equations d[soc,t,1] = - 10 * Vcb / Rp;

Process DischargeCurrent;
    conditions charging == -1;
    equations i = -Vcb / (Rs + Rload);

Process DischargeVoltage;
    Conditions charging == -1;
    equations Vt = (Vcb * Rload) / (Rs + Rload);

Process ChargeVoltage;
    Conditions charging == 1;
    equations Vt = Vcb + i * Rs;

```

their behavior in more abstract terms. At the same time, our design tenets demand that any language extensions support a mapping from elements of the model onto the structure and function of a device that engineers will find easy to understand.

We address this issue by introducing hierarchy into the representation presented in the previous section to produce a hierarchical process modeling language, which we call HPML. Our approach follows from the observation that all complex devices have a clear functional decomposition (or else it would have been difficult to construct them). In response, we represent a device as a *system* that contains some number of other systems, *variables*, and *processes*, where the processes model interactions among the component systems. These subsystems recursively decompose into similar structures, until the hierarchy ‘bottoms out’ in process models of the form discussed in the previous section.

We illustrate this approach by outlining a model of the Space Station’s electric power system (EPS), which is a self-contained solar power installation. It uses photovoltaic arrays to generate power, batteries to store it, and multiple circuits to deliver power to loads at various points on the Space Station. The EPS maintains these elements in careful balance by routing power; if the supply exceeds demand it stores power in the batteries; if demand exceeds supply (e.g., during eclipse) it

takes power from them. The boundary cases concern extremes of supply and demand that require shutting down generation capacity or shedding load (as in civic power systems).

Because the EPS is a complex device, we represent it via a hierarchy of models. At the top level it is a single system, called the *PowerChannel* (Table 3.2), which acts to govern power flow among three subsystems: power generation (*PowerGen*), power storage (*PowerStore*), and power usage (*Loads*). The *PowerChannel* makes decisions by altering the values of variables that represent targets for power production, power use, and power flow into (or out of) the batteries. These decisions are implemented by processes within the component models.

In more detail, we represent the abstract behavior of the *PowerChannel* via seven processes. The first two determine whether the batteries will act as a source or sink for power. If the solar power system can supply more power than all loads combined, the *condition* field of *ChargeMode* will be true, that process will be active, and the variable *PowerStore.ChargingBat* will be set to 1 (the dot notation associates variables with systems). If demand for power exceeds maximum supply (e.g., during eclipse), *DischargeMode* will be active and the mode indicator set to -1 instead. Processes within the *PowerStore* model check this variable before activating behaviors that trickle charge, taper charge, fully charge, or discharge the batteries.

The next two processes instruct the *PowerGen* system to produce a specific level of solar power. *FullSolarPower* sets the actual output, *PowerGen.SolarPowerOut*, to its maximum possible level provided that the loads and the batteries can collectively absorb that much power. However, if there is an oversupply, *ShuntSolarPower* ‘throttles down’ the generation capacity to produce exactly as much as required. Processes within *PowerGen* calculate the maximum available solar power (*PowerGen.SolarMaxOut*) by tracking the process of pointing the solar panels at the sun as the Space Station moves along its orbit. Other processes implement production targets by reading the *SolarPowerOut* variable and activating or deactivating strings of photovoltaic cells to match the desired supply.

The processes for managing demand are symmetric with those that govern supply. If the power generation and storage systems can collectively support all loads, *ServiceLoad* sets the variable *Loads.ServicedDemand* to the maximum demand. If there is a shortfall, *ShedLoad* forces demand into balance.² Processes within *PowerStore* determine how rapidly the battery can accept or produce current, whereas processes within *Loads* calculate the maximum demand by consulting a load schedule. Other processes implement load shedding via a uniform brown-out policy, but this is an approximation, as the Space Station actually sheds loads in priority order.

The last process within the *PowerChannel* system, *BalancePowerFlow*, determines the amount of power that must flow into, or out of, the *PowerStore* system. It inputs the final values of power supply and demand, and sets *PowerStore.ioBatPower* to their difference, where positive values denote inflow. Processes within *PowerStore* distribute this requirement across three component batteries, modeled as further subsystems that contain dynamic equations for charge transfer, as in Table 3.1.

In summary, the *PowerChannel* system represents EPS behavior in abstract terms. It describes

²EPS load schedulers work hard to avoid this condition, as it is highly disruptive.

power flow instead of propagating currents and voltages, and it suppresses details about how that power is generated, stored, and used. The hierarchical process modeling language also lets us decompose the EPS along natural system boundaries, such that specific devices and more detailed behaviors appear at deeper levels within the model structure. Finally, the language lets us associate processes with recognized device behaviors, preserving the correspondence between the model and the object modeled, which facilitates monitoring.

We note, however, that HPML has several limitations. First, because the language characterizes devices as a casual graphs (where inputs flow through processes towards outputs), it cannot represent feedback within a given time period. This makes it difficult to model control systems (e.g., for pointing the solar panels) and to represent complex equilibria (e.g., the currents and voltages in real circuits), which are typically characterized by in-period feedback and the solution of simultaneous algebraic equations. Next, HPML imposes a reactive structure that predicts the next state from current state and observations, meaning that all past state must be propagated forward via differential equations. (Variables cannot be assigned at time 1 and recalled at time 4.) This reduces convenience, but it enhances the view that devices are simple causal structures. It also lets us define concise mechanisms for performing simulation, fault detection, and diagnosis. More broadly, we note that HPML supplies a deterministic rather than a stochastic framework for modeling device behavior. This limits expressivity, but the simplification supports productive interaction with users, which is our immediate research goal.

4. Simulating Process Models

Given that we can represent a complex device with an HPML model, we can simulate the model to predict values for its variables. Our approach to simulation is straightforward; we compile the model into a causal graph that connects input variables, through processes, to output variables, and we evaluate this graph once for each time period. Since the equations are either algebraic (describing an instantaneous relation among variables) or dynamic (describing a instantaneous rate of change across one time step), the computation is analogous to updating a spreadsheet. The simulator simply sweeps the graph and applies the equations associated with each process whose activation conditions are met. Over multiple time steps, this produces a time series of quantitative predictions for variables in the model, which one can then compare against observations.³ The next few subsections discuss the compilation and simulation of HPML models in more detail.

4.1 Generating Causal Graphs

As introduced in Section 3, every process within an HPML model contains equations that define a causal relationship among variables. Taken together, a set of such processes define a *causal graph*, which is a directed acyclic graph that connects variable and process nodes.

Figure 2 provides the causal graph associated with the battery model in Table 3.1. Here, ovals represent variables, rectangles represent processes, and links indicate data flow. An arc from a

³The details are more complex, in that the simulator invokes an ordinary differential equation solver with adaptive step size to calculate predictions.

Table 2: The top-level of the electric power system model.

```

system PowerChannel;
  components PowerGen, PowerStore, Loads;
  process ChargeMode;
    conditions PowerGen.SolarMaxOut > Loads.MaxDemand;
    equations
      PowerStore.ChargingBat = 1;
  process DischargeMode;
    conditions PowerGen.SolarMaxOut <= Loads.MaxDemand;
    equations
      PowerStore.ChargingBat = -1;
# Generate full power if there is a place to put it
  process FullSolarPower;
    conditions PowerGen.SolarMaxOut <= Loads.MaxDemand + PowerStore.MaxBatPowerIn;
    equations
      PowerGen.SolarPowerOut = PowerGen.SolarMaxOut;
# Reduce power generation if there is nowhere to put it.
  process ShuntSolarPower;
    conditions PowerGen.SolarMaxOut > Loads.MaxDemand + PowerStore.MaxBatPowerIn;
    equations
      PowerGen.SolarPowerOut = Loads.MaxDemand + PowerStore.MaxBatPowerIn;
# Service all loads if the required power is available
  process ServiceLoad;
    conditions Loads.MaxDemand <= PowerGen.SolarMaxOut + PowerStore.MaxBatPowerOut;
    equations
      Loads.ServicedDemand = Loads.MaxDemand;
# Shed load if the required power is not available
  process ShedLoad;
    conditions Loads.MaxDemand > PowerGen.SolarMaxOut + PowerStore.MaxBatPowerOut;
    equations
      Loads.ServicedDemand = PowerGen.SolarMaxOut + PowerStore.MaxBatPowerOut;
  process BalancePowerFlow;
    equations
      PowerStore.ioBatPower = PowerGen.SolarPowerOut - Loads.ServicedDemand;

```

variable to a process implies the variable is an input to the process, whereas an arc from a process to a variable indicates that the variable is an output. Model parameters are not shown. In order to focus on causal relations, this representation also suppresses the control flow information contained in *conditions* fields. As a result, it appears that *ChargeVoltage* and *DischargeVoltage* simultaneously influence Vt , even though those processes are mutually exclusive. In general, multiple influences on a single variable are added.

This particular causal graph describes a dynamic system in which the change in battery state

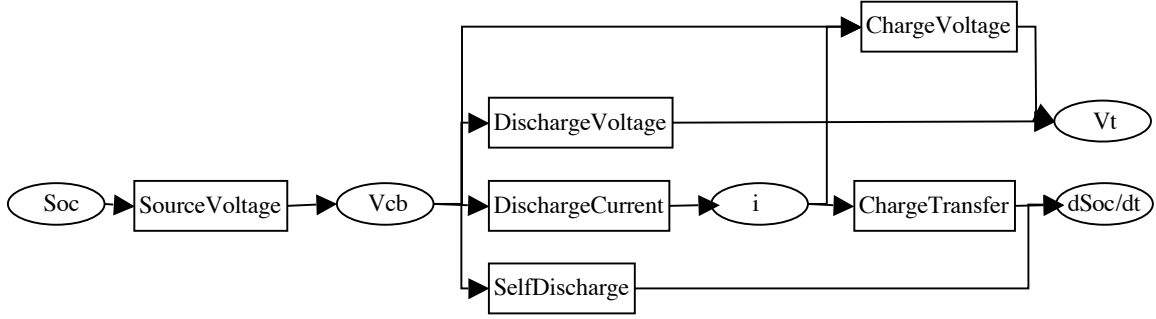


Figure 2: A causal graph for the simple battery model in Table 3.1.

of charge is a function of its current value. The simulator interprets this graph by integrating the equation for $dsoc/dt$ forward between time steps, and adding the result to the value of soc . In light of these semantics, a differential quantity can only appear as the output of a process and an integrated quantity can only be an input. Later figures will reduce the number of variables by mapping both quantities onto a single name (i.e., soc) and depicting differential equations as loops.

The algorithm for assembling the causal graph is straightforward. It starts by finding *exogenous* variables that only appear as inputs to some process and considers all processes not in the graph. It adds every process whose inputs are all present, together with its output variables if they had not been added before. This iteration continues until the graph contains all processes and variables in the textual form of the model (success) or until no process can be added (failure). Failure implies that the textual input specified a cyclic graph that has no meaning to the simulator.

Hierarchical process models also specify causal graphs, since the added structure simply partitions the data flow connecting processes and variables. As a result, HPML models compile into flat causal graphs exactly as described above, provided that the algorithm is modified to chase the input and output variables across system boundaries. This causal graph becomes an input to the simulation module.

4.2 Simulation using Causal Graphs

The simulator predicts the values of variables by simulating the causal graph forward in time. More exactly, it generates predictions for *endogenous* variables (intermediate and final outputs of the causal graph) from *exogenous* variables (inputs of the graph). This requires an observed value for every exogenous variable at every time step of the simulation, plus the size of a time step for use in integrating differential equations. Given these data, the simulator sweeps the causal graph once on each time step, from exogenous variables through processes towards final outputs, applying every equation belonging to a process whose activation conditions are met. If more than one process has the same output variable, the simulator adds the corresponding effects, reflecting a modeling commitment to view them as separable influences.

When the simulator encounters differential equations, it invokes a standard integration package [6] to compute the value of the integrated quantity for the next time step. Note that since a differential

equation may depend upon the output of other processes, the rate of change is commonly a function of variables observed on that time step.

It is important to notice that the simulator often has access to both a predicted and an observed value for endogenous variables, since observed values may be present in the telemetry stream. This creates a choice in source data. If the simulator uses predicted values to make further predictions, it generates estimates that only depend upon the exogenous data. If the simulator relies on observed values for intermediate variables instead, it generates more informed predictions but loses the causal dependency among predicted values. Each of these approaches has advantages. We rely on predicted rather than observed values during fault detection, as this lends an intuitive consistency to the system's expectations on any single time step. However, we rely on observation over prediction during fault diagnosis, where the goal is to isolate problems to the affected processes. We also rely on observed data over predicted values to determine which processes are active, regardless of context.

5. Monitoring Complex Devices

We have used the ability to simulate device models to construct a decision aid for fault detection and diagnosis, which we refer to here as the *monitor*. This tool helps the user to observe the status of the monitored device, detect anomalies and track them to their source, and anticipate future hazards that are a consequence of current events. In overview, the monitor operates by comparing a time stream of observations from the monitored device against the simulator's predictions. Whenever there is a significant mismatch, it raises an alert and attempts to associate the fault with particular variables, processes, and systems throughout the model hierarchy. It presents the results to the user via a graphical interface that supplies operations for navigating the hierarchy, navigating the underlying causal graph (to move from symptoms towards causes), navigating the history of monitoring system output, and projecting the model into the future to supply an early warning facility.

This section discusses the monitor's primitive operations for detecting anomalous conditions in variables and faults in processes, for isolating faults, and for projecting model behavior into the future. Later sections will discuss the graphical interface and the use of these capabilities to diagnose fault scenarios.

5.1 Detecting Anomalies in Variables

The monitoring system supports two methods of associating anomalies with variables. The first compares observed values against *absolute* limits and generates an alert if the value falls outside of an allowable range. This implements a standard threshold test. In the second method, the simulator generates a prediction based on exogenous inputs and the monitor compares it against the observed value. The monitor signals a fault if the error lies outside a predefined tolerance, expressed as a numeric range, or as a percentage of the observed value. We call this a *relative* threshold.

The monitoring system provides the user with a convenient checkbox-style interface for selecting

variables to monitor and for setting both absolute and relative thresholds. For example, the user might set an absolute threshold that triggers an alarm if battery goes below a 65% state of charge, in recognition of the fact that there is a real safety risk if the Space Station enters the night side of its orbit with less charge. In contrast, the user might set a relative threshold on total power demand, reflecting an expectation that the scheduled demand will cover a fairly wide dynamic range, making an absolute threshold impractical.

5.2 Detecting Faults in Processes

Surprisingly, the monitor can detect faults in *processes* as well as anomalies in variable values. It does this by comparing the predicted input/output mapping of a process against the observed input/output relation and flagging any significant difference. If such a difference exists, it serves as a smoking gun for fault isolation because it names a specific device behavior that is in error. However, the relevant data must be available. This only occurs if the process is *measurable* according to the definition:

A process P is measurable if and only if (i) all of its inputs are observable, (ii) at least one output is observable, and (iii) no other active processes contribute to the prediction of the observable output.

This definition is straightforward with the exception of condition (iii), which is designed to exclude situations where some other process affects the variable(s) in question, making it difficult to discover that P is in error. When the definition is satisfied, we can unambiguously declare that P is malfunctioning if $|o_{predicted} - o_{observed}| > k$, where k is a fixed threshold, o is a measurable output variable, and $o_{predicted}$ is computed by passing observed values of the input variables through P . At this point, we can take advantage of the fact that HPML fosters a close association between the model and the device to ascribe the process fault to a device behavior.

If multiple processes do affect a given output variable, o , we can sometimes treat the set of processes that influence o as a *compound process* and apply a relaxed version of the same definition. Here a compound process is measurable if its combined inputs are all observable, as well as at least one of its combined outputs. Again, no other processes can influence the chosen output variable. In this case, the compound process is malfunctioning when its predicted output (calculated from its input data) is sufficiently different from its observed output.

As an example, consider Figure 2 and assume that the variables *soc*, *i*, and *dsoc* are observable. When the processes *ChargeTransfer* and *SelfDischarge* are active, both influence *dsoc*, so neither process is measurable on its own. However, if we combine them with *SourceVoltage* and *DischargeCurrent*, the resulting compound process *is* measurable. In particular, if the value predicted for *dsoc* as a function of *soc* in this time period is significantly different from its observed value, the compound process as a whole is malfunctioning. We cannot say which component process is at fault, but we have isolated the problem to a set of possible candidates, which we call an *ambiguity region*.

5.3 Isolating Faults

The monitor uses its ability to detect process faults to support a more general fault isolation capability. Here the task is to search the causal graph for the smallest *measurable* compound process that explains a variable fault. The algorithm for doing this starts with an observed variable, o , that violates a monitoring threshold (relative or absolute), and traces the causal graph backwards towards exogenous variables while growing the ambiguity region. The algorithm stops when it can observe all of the inputs to a connected subset of the causal graph. In the limit, the ambiguity region will contain all processes between the exogenous variables and o . At this point, the system can compute a prediction for o as a function of these observed inputs and compare it to the observed value of o . If the difference is outside of the tolerance, k , we say that the ambiguity region *explains* the fault.

We have found it necessary to restrict this algorithm to prevent it from generating overly large ambiguity regions. We do this by constraining the length of the longest causal chain in an ambiguity region to eight processes (this is a tunable parameter). This approach gives the monitor considerable freedom to trace problems from symptoms to causes, while restricting attention to problems the user can plausibly correct.

We note that a number of other systems cast fault isolation as a search through a causal graph [22] [1], but ours seems unusual in its ability to employ quantitative rather than qualitative predictions. This difference should increase the reliability of our isolation results, given an accurate predictive model.

5.4 Projecting Models into the Future

The monitor also has the capability to generate an early warning of possible faults by projecting current state in to the future. We should contrast this with the ability to detect faults as they happen by performing tests on observations. The projection mechanism employs the simulation engine discussed in Section 4 to generate a deterministic prediction of future events, but since this process moves beyond the available data, we make one modification: the models posit observations for exogenous variables over the projection interval. This ‘closes’ the simulation, giving it the input data necessary to run.

The model developer has access to all of HPML to create predictors for exogenous variables, under the proviso that the new models are themselves closed. That is, its processes can define the value of an exogenous variable as a function of any variables in the rest of the model, specifically in the form of a differential equation (which is the only means in HPML of introducing dependencies across time periods). For example, the process:

```
process SetGimbalPosition;  
  equations d[position, t, 1] = 4;
```

models an assumption that the gimbal responsible for pointing the solar panels towards the sun will advance at four degrees per minute (where its position formerly had to be observed as an

exogenous variable). The projection system can then predict the consequences of this assumption for the behavior of the electric power system as a whole.

Since projection looks into the future, the monitor cannot compare predicted data against observed values. However, it can compare projected data against fixed thresholds to detect likely failures. Since the projections are quantitative, users can detect features such as the first moment at which battery charge will redline by dipping below 65 percent of its full capacity.

Any mechanism for projecting the current situation into the future will necessarily introduce assumptions, and thus expose issues of uncertainty and inaccuracy in the underlying model. As a result, the scope of the projections should be bounded in time. To this end, we introduce user-defined parameters that declare how many time steps to run a projection and how often to update that projection. We use one orbit (93 minutes) as the default value for this look ahead, and we update the projection twice in this interval, following our discussions with NASA personnel.⁴

5.5 Diagnosing Faults

Our capability for fault diagnosis builds on the projection facility described above. In particular, we let the user employ any model of a faulty system as the basis for projection, and compare those data against the actual observations as they arrive. For example, if the user suspects a failed NiH battery cell, he/she can select and project the consequences of that specific fault model forward in time, creating predictions for many system variables. If later telemetry broadly agrees with those predictions, the diagnosis was correct.

This approach to fault diagnosis requires a means for finding and selecting appropriate fault models. One response (dating back to the 1980s [26]) is to develop a library of explicit fault models with tools that let users browse, combine, and apply such models. We have investigated a related approach that employs machine learning to automatically construct fault models that best fit the data. This mechanism operates by searching over plausible replacements for each process in the ambiguity region output by the fault isolation facility and tuning its free parameters to the observed data. We have shown that this form of model revision is feasible in other domains [25] [4], and it appears extensible to fault diagnosis.

In summary, our approach to monitoring complex devices employs causal process models to offer a variety of core capabilities, including mechanisms for anomaly detection in variables and in processes, fault isolation, projection, and fault diagnosis. The following section shows how we embed these capabilities within a graphical interface to provide users with a powerful tool for interactive monitoring.

6. The Graphical User Interface

The monitor's graphical interface provides users with a rapid way to understand the health of complex devices. Here, the key design challenge has been to maintain a sense of simplicity while

⁴Alan Crocker, Electric Power System flight monitor, personal communication, 2002.

supporting a daunting task characterized by a complex model, large volumes of data, multiple analysis methods, and potentially cascading alerts.

Our design relies on the hierarchical structure of HPML as the organizing principle. In particular, we tie all alerts and data displays to the graphical model to provide the user with a constant sense of orientation within the device, and we employ the model hierarchy as a means of hiding and summarizing detail. This implies that we are focusing on informed, but non-expert, users who understand the basic function and organization of the device but who still need some reminder of its structure. All totaled, the interface must display the model’s structure, associate alerts with that structure, display time-varying data, and supply tools for navigating the data and the model, in addition to operators for invoking the various monitoring functions. We discuss each of these component capabilities below.

6.1 Graphical Model Display

The monitor displays HPML models via a two-dimensional layout that preserves both causal and hierarchical structure. We illustrate this in Figure 3, using the top-level model for the Space Station’s electric power system (the *PowerChannel*). This display corresponds to the textual representation in Table 3.2. Here rectangular nodes denote processes, oval nodes denote variables, and long rectangles with rounded corners (and internal structure) denote subsystems. As before, an edge from a variable to a process indicates that the variable is an input to the process, while an edge from a process to a variable indicates that the variable is an output. Data flow goes from left to right in the diagram, while the subsystems are arranged vertically down the page.

This representation shows all of the variables and processes that belong to the enclosing system and all of the causal relations between the systems it contains. At the same time, it suppresses details associated with the structure of subsystems. However, the user can bring up those details by double-clicking on the subsystem’s box (e.g., the *PowerStore* system) to open an analogous display in a separate window.

In addition to communicating system structure, the display also animates that structure as the user compares the model against telemetry. At each time step, active processes are fully displayed with incoming and outgoing edges, while the inactive processes are shaded and their corresponding edges are removed. For example, the display in Figure 3 corresponds to a state where the batteries are charging, and the system is servicing all loads while shunting some solar panels. This animation may change in the next time step, as motivated by the data.

6.2 Alert Propagation and Display

The monitoring system displays status by color-coding variables, processes, and systems. It uses green to signal OK, yellow for warnings, and red to indicate severe alerts in accordance with people’s preconceptions. In particular, it displays the oval representing a variable in green if it satisfies all monitoring tolerances, yellow if the comparison between its observed and predicted values lies outside of a *relative* tolerance, and red if the value exceeds an *absolute* threshold (see Section 5.1). It displays the variable in a dull olive color if it has no attached monitoring tolerances

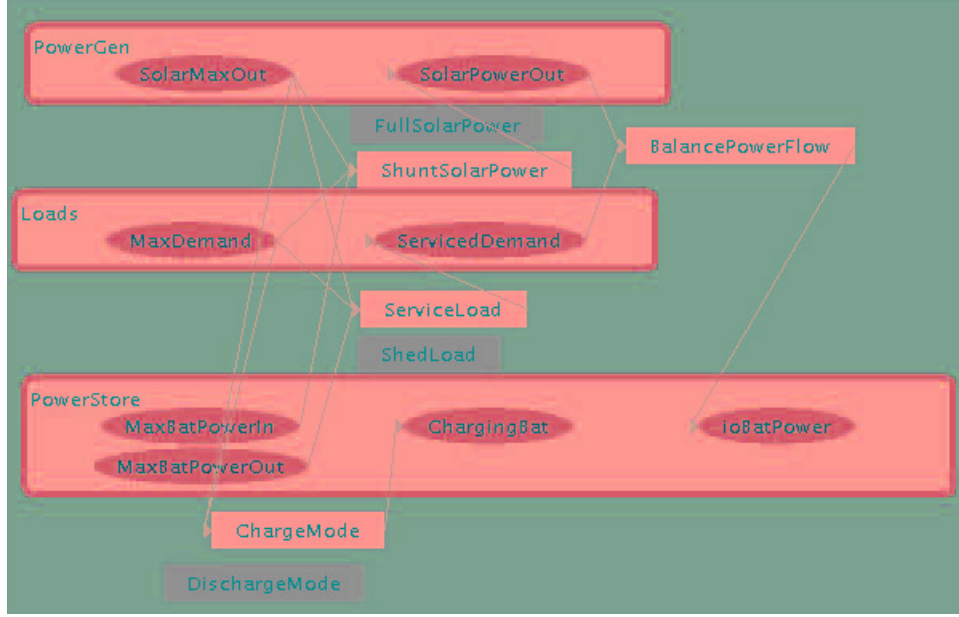


Figure 3: Graphical display for the top level of the electric power system.

(although note that the variable may still be observable in the telemetry stream). The system draws a process boundary in yellow if it is a part of an ambiguity region associated with a variable fault, as discussed in Section 5.2.

The monitor color codes component systems if they contain faults in variables and/or processes. This provides an abstraction principle: a system inherits the most severe color found in any of its components, recursively down the subsystem hierarchy. As a result, a system remains uncolored if it is fault free, it will be bounded in yellow if there are relative threshold violations or process faults (but not absolute threshold violations), and it will be bounded in red if it (or its subsystems) contain any absolute threshold violations.

Figure 4 illustrates this principle by showing the status of the *PowerChannel* system during a monitoring run. This display indicates that the *PowerGen* and *Loads* systems are fault-free, whereas *PowerStore* contains a severe fault at some location. If the user wants to see precisely where, he/she can drill down through the model hierarchy by double-clicking on systems. This leads to a display like that in Figure 5, showing that the state of charge in *Battery3* is out of bounds, where *Battery3* is a subsystem of *PowerStore*, which is a subsystem of *PowerChannel*.

6.3 Data Displays

In addition, the monitoring system contains a graphing package that lets the user plot the predicted and observed value of any variable over time by double-clicking on its icon. Figure 5 provides a simple example concerning the state of charge on the battery. It illustrates that the observed value of *soc* deviates from the value predicted by the model by integrating current flow forward from

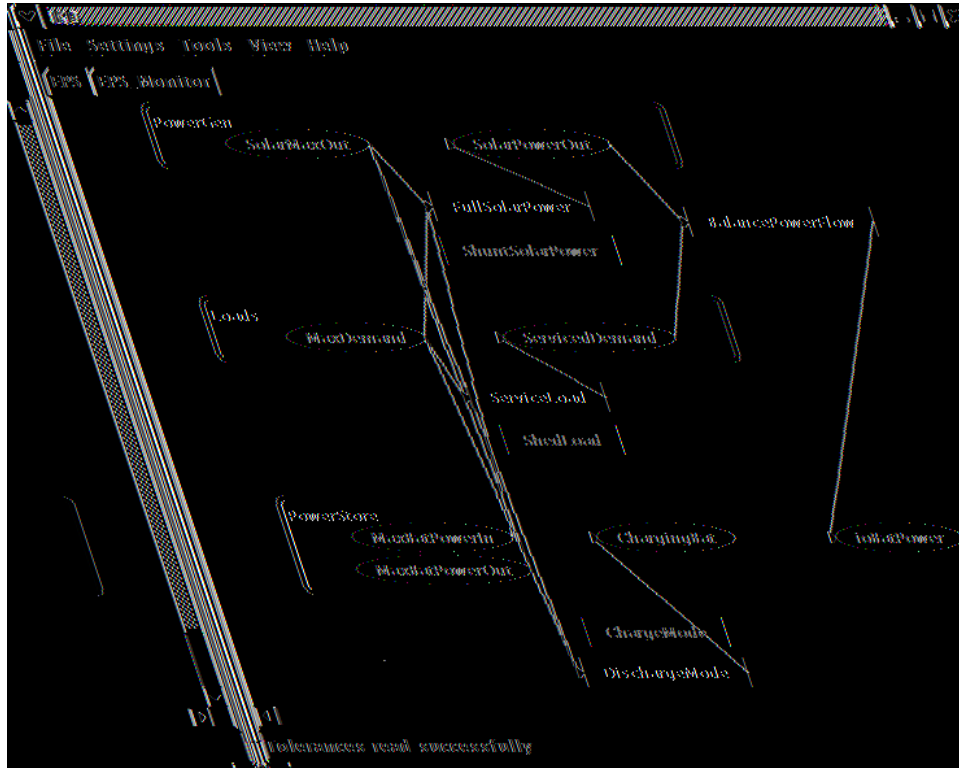


Figure 4: The PowerChannel showing a fault during a monitoring run.

time zero. The user can also view a third curve, not shown in the figure, that projects the value of the indicated variable into the future under various fault models.

6.4 Navigation

The monitor provides several navigation tools to help the user detect and diagnose problems. These fall into three categories, for navigating the hierarchical device model, the time history of alerts, and the causal graph containing alert data. The operators for navigating the model are quite simple; they let the user open subsystems by double-clicking on their icons (creating a display of that system's contents in a separate window), choose open models (by clicking on windows), and move to any enclosing model from a point deeper in the tree.

The operators for navigating the time history of alerts treat that history as a sequence of snapshots depicting device health over time. In this framework, time zero is the beginning of the telemetry stream, *now* is the time index of the most recent telemetry, and any index later than *now* is in advance of the data, and thus solely associated with model projections. (In actuality, the data in our scenarios comes from a file rather than a real-time feed, so data is available at all time indices. However, we maintain the fiction that *now* corresponds to the most recent telemetry for the purpose of exposition.) With this background, our navigation metaphor follows the control functions of a DVD player. Reading left to right in the top corner of Figure 6, the user can play

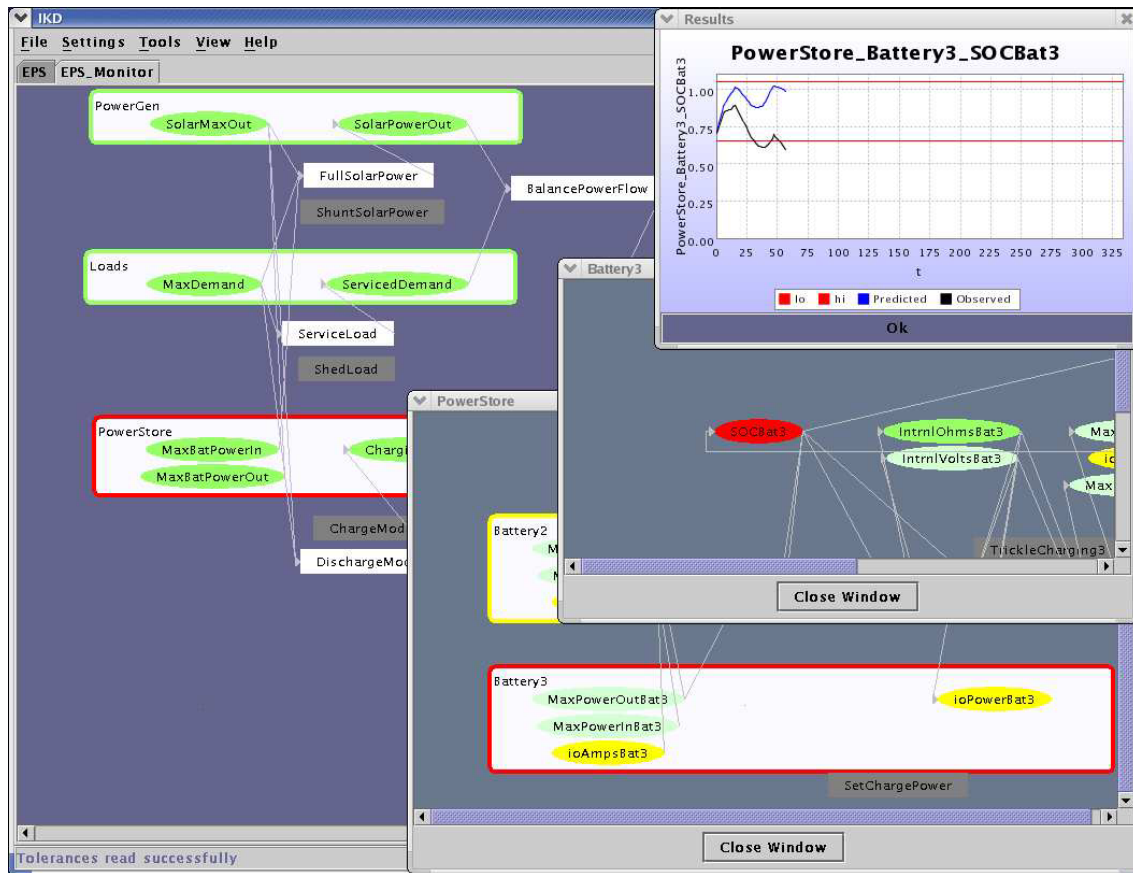


Figure 5: Drilling down through a model hierarchy to find a fault.

back the time-line (at three frames per second), move forward one timestep, move backward one timestep, pause, and stop the automated playback. The user can also drag the slider to any time of interest, in which case the variable *now* is set to that new time. In response, the monitor updates the graphs, and color coding within all open models to their appropriate values.

The operators for navigating the causal graph let the user shift attention to systems that are causally upstream or downstream of a fault within a given time index. In particular, the user can select the forward (or backward) causal navigation option by right-clicking on a variable. The monitor will respond by finding the the next (or previous) process, opening the model that contains it, and highlighting that process on the screen. An analogous operation navigates from processes to adjacent variables. Simple extensions to this facility will find variables and processes flagged with faults by searching over multiple links in the causal graph. This will provide a rapid means of managing fault cascades, for example, to locate the causally earliest evidence of a problem.

6.5 An Illustrated Monitoring Scenario

The monitor's palette of primitive capabilities should help people perform a variety of fault detection and isolation tasks. In particular, it should let users:



Figure 6: The control bar for navigating the time series of monitoring data.

- detect problems in the operation of a complex device,
- localize symptoms to components of the device model,
- track the development of those symptoms over time,
- trace symptoms to their causes within a given time step, and
- identify the specific processes (and corresponding device behaviors) that are malfunctioning.

We illustrate how these capabilities come together by walking through a hand-crafted scenario concerning a fault that has actually occurred on board the Space Station.

This scenario begins with the electric power system in some distress, at a time when two of its major components contain serious faults (*PowerGen* and *PowerStore*), and its third displays a warning. The left side of Figure 7 illustrates this case. Faced with this ambiguity, the user arbitrarily decides to drill down into the *PowerStore* system, producing the right side of Figure 7, and then to expand one of its three battery subsystems, all of which indicate faults. This generates the bottom illustration in Figure 7, which shows that *soc* is red-lined, but that all of its active processes appear to be functioning properly. This situation repeats in the other two batteries (not shown). The fact that all three batteries simultaneously contain variable faults suggests that these anomalies are symptoms, not causes. In response, the user traces the causal connections from *SOCBat3* back towards the source of the problem.

The user can employ the monitor’s causal navigation facility to trace the data flow from *SOCBat3* through *NormalCharging3*, towards inputs of the *Battery3* model (variables with no incoming arcs). Alternatively, the user can exploit the hierarchical structure to recognize that *ioPowerBat3* represents power flow into, or out of *Battery3*, and that *PowerStore.ioBatPower* computes that quantity for the *PowerStore* subsystem as a whole. This chain of reasoning leads back to the *PowerChannel* model (the left side of Figure 7). Here, *ioBatPower* leads to *BalancePowerFlow*, *SolarPowerOut*, and ultimately, to the input variable, *SolarMaxOut*. We show the user employing causal navigation to find the predecessor of this variable.

This path leads into the *PowerGen* system, as it is causally upstream of power storage as a whole. Figure 8 depicts the results. Here, the user has chased *NormalSupply* (the predecessor of *SolarMaxOut*) into the *BetaGimbal* system, which models the controller for the rotating joint that points the solar panels at the sun. The *BetaGimbal* system is the source of the error, as evidenced by the yellow highlight on the *GimbalResponse* process, which declares that its expected

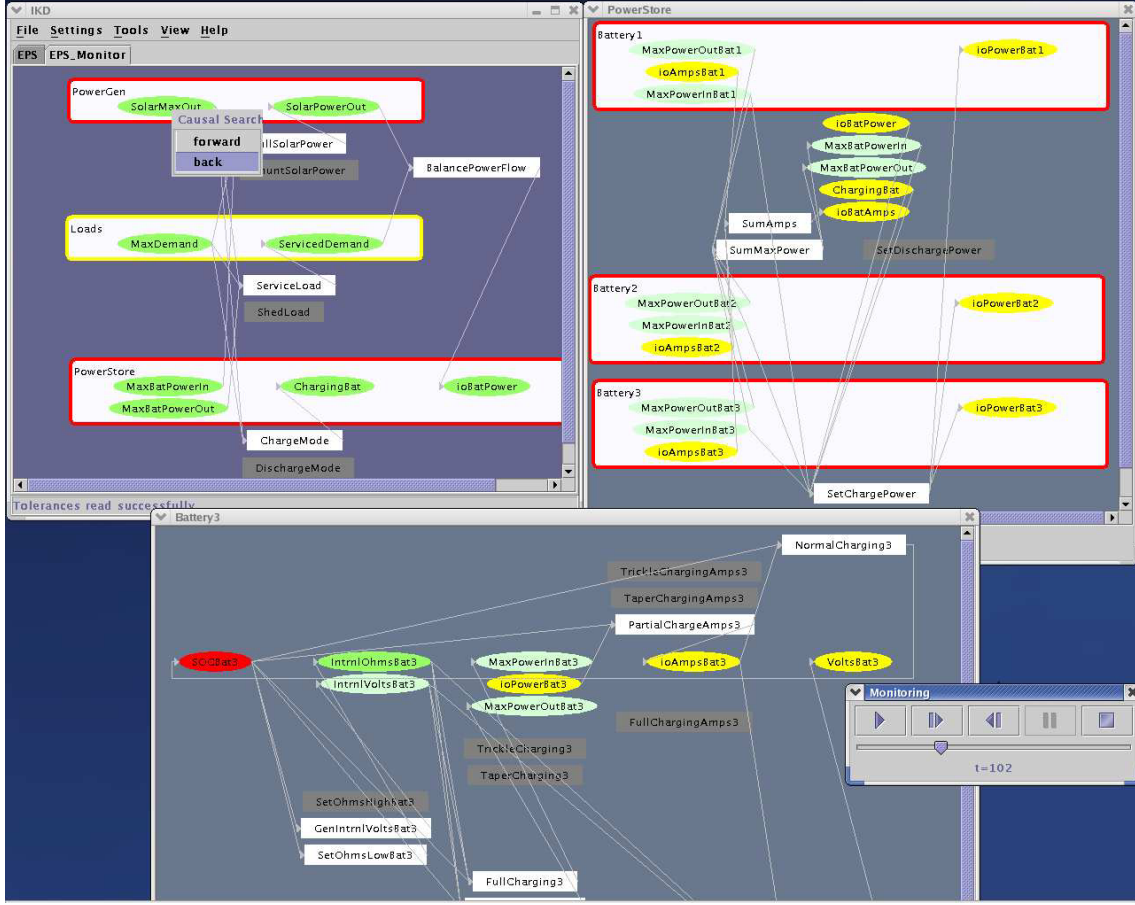


Figure 7: The PowerChannel system with simultaneous faults.

i/o mapping does not equal the observed behavior. Something is wrong with the gimbal joint; its position (integrated over time) is not responding appropriately to the applied current.

The user calls up three graphs to clarify the nature of the problem. The data in Figure 9 tells a convincing story that the gimbal joint is frozen in place. In particular, the predicted value for the supplied current is constant (corresponding to a constant rotation speed), while the observed current follows a sawtooth pattern, in lock step with the day/night cycle shown in the variable *SolarPowerOut*. This suggests that the gimbal controller is trying to rotate the panels in response to an increased pointing error, first one way, then another. Given that the battery discharged throughout the day cycle, the gimbal also froze with the panels at a dysfunctional angle to the sun.

7. Experimental Evaluation of the Monitor

Although the preceding scenario showed how one might employ the monitor's capabilities to diagnose a fault, we have also studied its behavior in practice with our target audience of non-expert users. The key open questions were whether the monitor's novel features affect diagnostic reasoning and which features supply advantage (or disadvantage) relative to existing interactive diagnostic

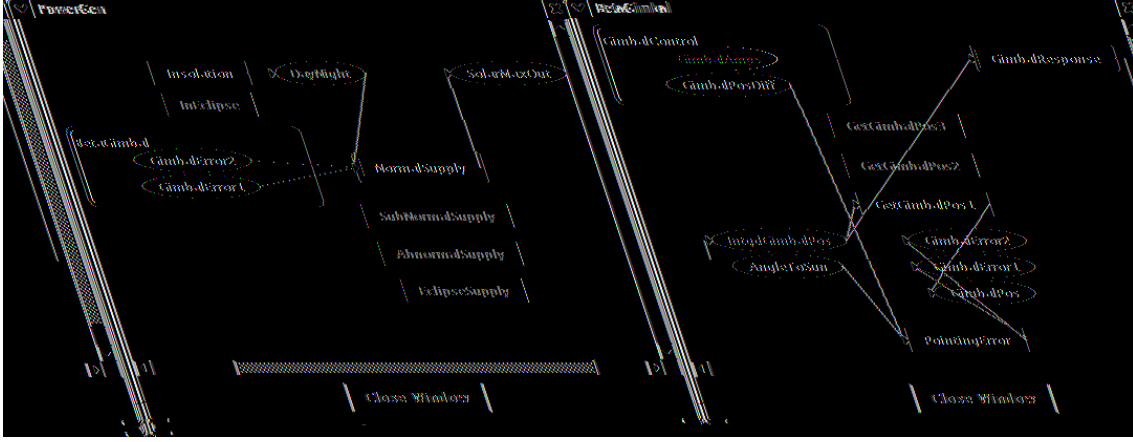


Figure 8: The gimbal controller and gimbal systems.

techniques. We focused on the first question, as we judged the second too difficult to address in the context of this project.⁵ In particular, we conducted an exploratory study that recorded monitoring and diagnosis strategies in the presence or absence of hierarchical model structure. Then we mined the records for regularities about how people used the tool: which features they relied upon in both conditions, which they found frustrating or ignored, and whether the protocols indicated a difference in reasoning style or diagnostic accuracy across the hierarchical and non-hierarchical conditions. We describe this study in some detail below.

7.1 Experimental Method

7.1.1 STIMULI AND DESIGN

In order to study the impact of the monitor on diagnostic strategies, we needed a realistic model of a non-trivial device interacting with its environment, plus data describing a series of malfunctions. In response, we developed a rather complex model of the Space Station’s electric power system; it contained 16 systems, 24 distinct variables, and 79 processes that reflected the components, observable features, and behaviors of the device. (The displays in this report are drawn from this model.) On average, the model hierarchy was four levels deep, and to enhance realism, some components had been tuned to fit Space Station telemetry data, but most had not. As a result, the monitor could not predict actual telemetry values, and the scenarios employed synthetic data.

We also developed six fault scenarios by altering processes within our model of the electric power system, which we presented to subjects as monitoring tasks. Two scenarios involved faults in each of the three subsystems. They were presented in the following order:

- a recurrent shadow reduces output from the solar panel.

⁵The second question calls for a lesion experiment on a complex system, while measuring the performance of multiple users as they pursue a non-trivial cognitive task. Individual variations in diagnostic strategies would very likely govern the results, unless we ran a prohibitive number of subjects.

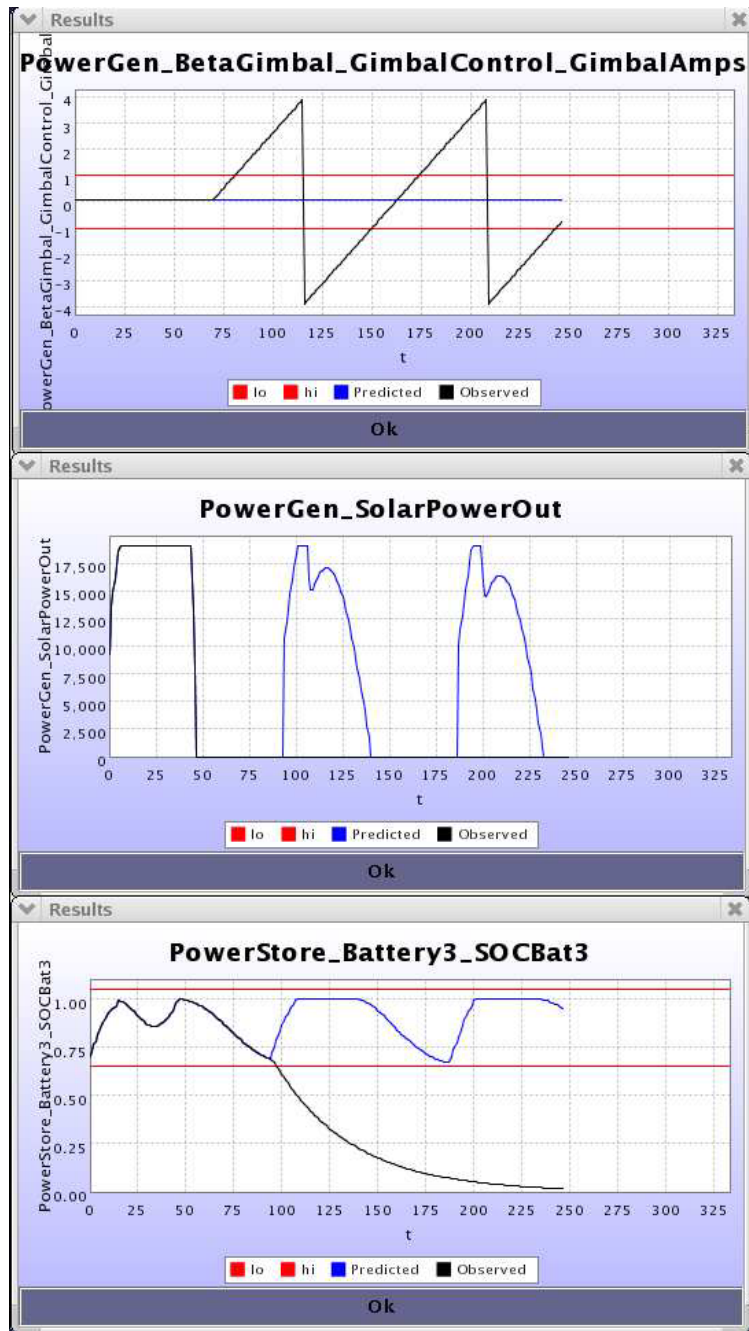


Figure 9: Data indicating a stuck gimbal fault.

- a known, periodic load kicks in at the wrong time.
- a gimbal joint is stuck and cannot move.
- a battery spontaneously discharges at an accelerated rate.
- a planned schedule of loads is unsatisfiable.

- an aging battery cannot hold as much charge as it once did.

We used two versions of the interface: the hierarchical design illustrated above and one that presented a flat, non-hierarchical model. The flat model contained the same processes and variables, but without a supervening system/subsystem structure. The layout for the flat condition spread across six screens and included a single, system-wide status indicator that, like those in the hierarchical condition, turned red or yellow if any component within had that status. Figures 5, 7, and 8 presented examples of the hierarchical layout, whereas Figure *** shows a portion of the flat layout.

7.1.2 USERS AND PROCEDURE

We recruited subjects by sending electronic mail to students in Stanford engineering classes and posting notices on engineering bulletin boards. We used engineering graduates and undergraduates because they should have some familiarity with models of dynamical systems and complex electrical devices. We felt these users offered a reasonable surrogate for our target audience of informed, but non-expert, diagnosticians. We included 12 users in the study, with six in each condition, although several more served in an earlier pilot experiment.

The procedure involved giving a tutorial to subjects, presenting them with the six fault scenarios to be solved with one of the two interfaces, and giving them a final debriefing with exposure to the alternative interface. We videotaped over the user's shoulder during the training and fault scenarios, capturing the screen, some gestures, and their utterances.

The session lasted three hours. Presenting the tutorial required about an hour, solving the six diagnostic problems took from 1.5 to 2 hours, and exploring the cross-condition layout lasted from a few minutes to half an hour, as time was available. Users took a standard break after the first or second task, according to the user's choice, and we let them take a break after any problem. The monitoring system is a complex tool that was simulating a very complex system. We developed the tutorial after pilot studies, which identified the information and activities needed in training to let users act effectively and reason with the monitor's tools.

7.1.3 TRAINING

After an introduction to the entire session, the tutorial began with an overview of the Space Station's electrical power system. This was organized to cover broad information about the Station and its three components: power generation, loads, and power storage. The experimenter read information to the user in four sections and asked him to summarize and report back the main information from each one. After this, the tutorial introduced aspects of the PowerMonitor: how structure of the electrical system is represented, variable name conventions, and navigating the structure by scrolling and causal links; dynamics of the system through running the scenario and displaying variables' values over time; alerts and other information useful in diagnosis. We used a scenario in which all aspects of the station electrical system operated normally, and the subject practiced finding and viewing each aspect of the PowerMonitor as we described it. After the experimenter had

presented all the information and the user had explored it, the latter explained and demonstrated answers to a series of questions that reviewed the presented information.

In both phases of training, the subject was questioned further if his first response did not include all the target information and, if needed, the experimenter resupplied information that the subject had not retained. In general, users did very well in answering the questions; some realized after the first question that considerable detail was expected and shifted their strategies accordingly.

In the final part of the tutorial, the subject solved a series of reasoning problems about normal operations, which gave him practice interpreting graphs and using them to reason about the state of the system and about relations among variables. This section of the tutorial varied most in time, in user activity, and in amount of scaffolding provided by the experimenter. Our intent was to ensure all participants had a roughly comparable understanding of the system at the beginning of the next stage, while still keeping the user-experimenter interaction as standardized as possible.

7.1.4 DIAGNOSTIC TRIALS

We instructed each subject that he would view a series of scenarios that involved problems with the power system and that, for each, the subject should detect, diagnose, and report on the fault. We told each user what counts as a problem (any thing not as predicted), what he should include in the diagnosis or explanation (the causes and effects that generated the unpredicted and possibly damaging states), and how he should give the report. During the first two phases, users spoke aloud; during the report phase they were asked to give their explanation of the fault, state the root cause, state its effects, provide advice about actions to repair, compensate for, or prevent the problem, and rate its severity on a numeric scale.

We asked subjects to give their explanation first, as a relatively broad task. We found that many users continued to problem solve while giving their report and this first, open question let them do that. We also asked them to state the root cause and effects separately to get a cleaner, user-provided specification than if we had tried to extract this from the initial explanation. The final rating of severity ranged from 1, meaning that something was not as predicted but had no negative impact, to 5, meaning it could be fatal. The user had access to the computer while making the report. Participants frequently used the system while making their initial explanation (from just using the cursor to point at variables to opening systems and variables while testing a hypothesis). However, they rarely used the computer for the remaining four items in the report. After the subject provided the verbal report, we moved the computer away and gave him a check list of candidate descriptions.

Users talked aloud as they detected and diagnosed the faults. The experimenter sat beside the user and watched the screen with the him, but made only three types of intervention. First, if the subject had not said anything or seemed inaudible, the experimenter asked “What are you thinking?” or “Could you talk louder?” Second, if the user seemed confused about the interface (e.g., he was not clicking fast enough to get a double click), the experimenter tried to clear up the problem. Third, since time spent before concluding varied considerably, the experimenter urged the participant to conclude if he had been working more than 25 minutes on a task. (This happened

N times; in m of them, the user was looking for more information that the system provided, or trying to go outside of the system’s instrumentation.)

All participants completed the six monitoring scenarios. Depending on the remaining time, the experimenter included as many of three activities as possible: payment and debriefing; view and comment on the alternative system; instruction on different aspects and problem solving of one or two problems on the alternative interface, which involved two novel scenarios. The procedure was the same as before, except that the experimenter provided no check list.

7.2 Experimental Results

Our experimental results provide evidence on the overall viability of the monitoring system for novice users and document the role of several key design features. Evidence for the system viability comes from the results of the check list responses and from more qualitative analysis of the first scenario. We consider several design features: hierarchical organization into subsystems, inclusion of variable information, and process status information. Evidence for the benefits of hierarchical layout comes from comparison of performance between users with hierarchical versus flat presentation and from observations about how subjects used the high-level window. The roles of other design aspects are illustrated with examples from the process-tracing for the Shadowed Panels scenario. We first present the check list data, then findings on the accuracy of solutions to the Shadowed Panel scenario, and finally findings on user processes.

Users found the task to be highly engaging and they took their role seriously. Many users placed themselves in the role of astronauts, as indicated by comments like “it is beautiful up here” or talking about “sending down to Houston” for new batteries. Users frequently talked in the first person, about “I” or “we” having a problem or taking actions. The scenarios were complex and difficult for these generally knowledgeable but inexperienced users. Many problem-solving trajectories were possible and the observed behaviors were complex and variable across users and tasks. The data collected from our 12 users is rich and provides a seedbed for generating hypotheses more than a testbed for evaluating them.

7.2.1 CHECK LIST DATA ON PROBLEM DESCRIPTIONS

The check list provided a structured set of alternatives to assess users’ reasoning in a more supported task. It offered a basis for qualitative understanding of what problem descriptions users attributed to the situation and for comparing between conditions. Two experimenters with advice of a third rated the check list items as 1, 2, 3, or not applicable, as the users had done. This was used to categorize descriptions as False, True, and Near Miss, the latter referring to ones that were closely related to the actual problem, consistent with much of the data, but not in fact part of the scenario. These were items which participants might think plausible, be unsure about because they did not have a complete enough understanding, or consider for which the data did not allow certain determination.

To compare conditions, we dropped the descriptions that no one selected and compared conditions on the remaining items. We provide inferential as well as descriptive statistics. Because we had

so few users per condition, and these were highly variable, we did analyses by item ⁶ rather than by subject, i.e., using the quasi-F test (Clark, ****). We believe inferential statistics are helpful because they provide a consistent criterion for which differences should be taken most seriously. Figure 10 shows the average number of users (out of six) who marked items ‘1’, that is, who indicated that they were certain the description held for the scenario. Effect of item type was highly significant $F(2, 64) = 34.2$, $p \ll .001$, $Eta2 = .517$, reflecting the users’ strong discrimination among item types. The effect of condition was not significant, reflecting that neither group was significantly more inclined to include items, regardless of type.

Of greatest interest is the significant interaction of condition and item type, MANOVA $F(2, 64) = 5.82$, $p = .005$, $Eta2 = .154$. Users in the hierarchical condition showed much greater discrimination between the True and Close Items than did users in the flat condition. This suggests that the hierarchical users had a more detailed or accurate model of the situation, letting them make finer discriminations. This is supported by t tests, which showed different response to True vs. Close items by hierarchical condition users (mean True = 3.50 differs from mean Close = 1.47 $t(41) = 3.94$, $p < .001$) but not by flat condition users (mean True = 2.67, mean Close = 2.05 $t(41) = 1.254$, $p = .22$). Further, hierarchical users rated True items as ‘1’ significantly more often than did flat condition users, $t(23) = 2.286$, $p = .03$.

We analyzed the data in two additional, complementary ways using a second and third dependent variable, and we found the same pattern as with the measure based on descriptions rated ‘1’. One of these was the number of users who gave the item a rating of 1, 2, or 3. Here we found the same effects and a MANOVA $F(2, 64) = 3.77$, $p = .028$ for the interaction of condition with item type, and we observed the same pattern in t tests. The third dependent variable was a weighted ‘certainty’ measure: a score of 3 meant most certain (a response of 1), 2 meant moderate certainty (a response of 2), 1 (a response of 3) meant low belief, and a score of 0 meant the item was not checked as describing the scenario. In this case, we found the same effects: a MANOVA $F(2, 64) = 6.00$, $p = .004$ for the interaction and the same pattern in t tests.

7.2.2 SOLUTION QUALITY AND CONDITION COMPARISONS

We designed the “Shadowed Panels” scenario to be particularly challenging, in that its solution requires “thinking outside the box” in a way people normally find very difficult. The situation violates expectations that a problem results from a failure within the power system, since the power system monitor is the topic of the study and what subjects are using. Correct solution of this scenario requires a thorough understanding of the situation that is deep enough to trust unexpected conclusions. If the hierarchical interface fosters more complete or coherent models, such benefits might be visible on this task. Note that, although this scenario is labeled “Shadowed

⁶Formally, this means we are generalizing over how these two groups of users would respond across other items of the same type, rather than over how other users would respond to these items. We believe it is valuable to utilize a consistent evaluation criterion when identifying which differences in the data to take most seriously. Averaging across participants for a given item allows a relatively large number of items for comparative statistics, while averaging across items leaves only six users to compare between groups. Further, we can compare performance on each item in both conditions, while no user was assessed in both conditions. Thus, in addition to allowing tests on a larger sample, analysis by item allows repeated measures and paired t tests.

Table 3: Reported and selected (check list) causes of the problem.

| | Incorrect Subsystem | Correct Subsystem | | | | |
|-------------------------|---------------------|-------------------|----------------------------|------------------------------|----------------------|--------------------------|
| | | Nonspecific only | Loss of input: block light | Loss of input: panel alignm. | Control: shunting | Other (sensor; multiple) |
| Hierarchical report | 0 / 6 | 1 | 3 (1 or array | 1 (or circuit from panels) | 1 | |
| Hierarchical check list | 0 / 6 | 1 + 2 | 2 | 1 | 1 write in | |
| Final report | 1 / 6 (load) | 0 | 1 | | 3 (1 or maybe align. | 1 (panel: mech. failure |
| Flat check list | 1 / 6 | 1 + 2 | 2 | | | 1 (sensor) |

Panel”, the simulation data are consistent with any external factor that briefly reduces sunlight over the panels, and we judged any diagnosis consistent with this interpretation as correct.

Participants were generally successful in localizing the cause of the problem for the Shadowed Panel scenario, with 11 or 12 users correctly identifying the root cause as a generation fault and only one incorrectly identifying the cause as excess demand. We have no clear difficulty metric for the scenarios, but localization to the generation system is a significant accomplishment. Problems in the storage and generation systems appeared simultaneously; and many individual fault symptoms could result from faults originating in any of the subsystems.

Table 3 shows the categories of causes from the explanations generated in the verbal reports on the first line for each condition; it shows the cause selected as best explanation from the check list on the second line. If a user generated a second possible cause, x, we have listed this in parentheses “(or x)”. Of considerable interest, five users endorsed reduction of input to the solar panels as the cause. Four (3/6 of the hierarchical condition users) attributed this to an ‘out of system’ cause and generated the explanation on their own: shadowed panels, something blocking the sun, or the space shuttle itself getting misaligned (flat condition).

There are several hints of greater accuracy in the hierarchical condition: every hierarchical user localized the cause within the generation system, four users attributed the cause to reduced input, and each of these four generated this explanation on his own rather than just selecting the item from the check list. There is more variation in the flat condition, with solutions distributed across shunting, mechanical failure in the panels, gimble sensor error, and excessive loads. Further, users in the hierarchical condition seem to have deeper and more stable explanations, with less change between generated and check list accounts. Informal reading of the protocol reports suggests there was more direct, systematic convergence on a solution in the hierarchical condition, with fewer

changes back and forth in the user’s hypotheses. These possible differences between conditions should be followed up both by further analyses on these data and by more targeted experiments.

Our general impression was that hierarchical condition users had richer, more complete models of the problem situation, as well as a stronger sense of closure and more stability in what they believed. However, it is very difficult to identify clear measures of model quality given the very different solution activities and reporting styles that we observed across subjects.

Careful study of the protocols revealed one subproblem that gave us a very sensitive index of the sophistication of the user’s model: providing the correct analysis for complex, Day 3 behavior that solar power generation is greater (not less) than predicted and that this ‘overgeneration’ compensates for the undergeneration of solar power on Day 2, which resulted in more-than-predicted draw down of the batteries. The alternative models included either not noticing or analyzing this less critical departure from normal at all, or considering it a separate problem, say caused by another issue in regulating shunting. Because all users’ attention was primarily focused on the more serious Day 2 faults, this is a very difficult aspect of the scenario.

Three of the six hierarchical users reached the correct and complete analysis of this subproblem, with two of them also noting the overgeneration but providing different explanations. In contrast, no user in the flat condition had the correct model: three never noted the discrepancy at all (either by cursor pointing or by comment), two noted it but provided no explanation, and one provided an incomplete account. Developing the correct and complete model depended on comparisons between predicted and observed values across distinct variables, an activity described below.

We believe that this approach to analyzing the quality of situation models developed by users will be important in more extensive studies of how computer tools support humans in complex problem-solving tasks.

7.2.3 IMPACT OF DESIGN FEATURES ON USER PROCESS

We also examined how participants used key features of the monitoring system to support their problem solving. Although these features combine synergistically, we presented the various aspects separately using examples taken from the study.

Hierarchical Organization. Users in the hierarchical condition could navigate by opening subsystem windows and all users drew on this strategy. Predominantly, users in both conditions opened a variable graph either because it was color flagged as a problem or because it was linked to a flagged variable. However, variables in the highest, system-level window are never flagged red or yellow. Interestingly, four of the six hierarchical users opened unflagged, high-level variables from the system window, apparently with a proactive goal of monitoring or understanding the system rather than reacting to a particular problem variable. When hierarchical users were informally shown the flat layout after the experiment, some commented that they missed having the more informative high-level alert from the hierarchical condition, which localized the fault to the relevant subsystem.

Temporal Aspects of Variables. Both interfaces let users get more information about a variable than just its indication of green, yellow, or red status. The ability to plot a variable’s trajectory over time also helped problem solving. Users checked the day/night variable to establish the overall

pattern of activity for the power system and they used the battery-charging graph similarly to track the system's high-level power flow. They also used the shape of these function to make very specific inferences. For example, one user used the step-function contour of the SolarPowerOut graph (at the point when both sunlight and solar generation drop) to reason that the probable cause of the change was something outside the system: "Here, at the beginning it goes as expected, and then suddenly, it drops. (pause) Things usually don't happen like this, like, it doesn't suddenly go into a right angle. It must be some kind of external thing."

Comparing Predicted and Observed Values. The availability of predicted values, in addition to actual ones, supported a number of reasoning activities. Users examined such pairs of trajectories to identify when when a variable diverged from its predicted value and the point when 'things return to normal', thus bounding the scope of a problem. This facility also helped users determine the nature of departure from a normal value, thus constraining the nature of the analysis task. For example, immediately after solar power drops, one user opened IOBatAmps and noted "Here's a spike here [plays scenario] ... it's lower than expected." One particularly interesting case involves the examination of SolarPowerOut when the generation on the third day is higher than expected, in order to compensate for the battery discharge on the second day. One user selected SolarPowerOut, looked at Day 3, started to say it was again too low, did nearly a double take, and corrected himself to say the power generation was now too high.

Comparison of actual to predicted values also served the very important function of letting users cleanly test and reject hypotheses. For example, after determining that generation was lower than it should be, several users hypothesized that the gimble system might be responsible and checked the gimble variables. Finding that the actual values matched those predicted was a sufficient and compelling basis for rejecting the hypothesis that panel alignment by the gimbles was responsible for the anomaly. A few users also checked load variables to reject the hypothesis that excess demand was contributing to the problem.

Comparing Values of Different Variables. We also frequently observed subjects display multiple graphs of variable at the same time in order to compare their trajectories. Evidence for this comparison comes from both the verbal protocols and instance of pointing that indices efforts to align corresponding parts of the graphs.

Many subjects related a reference variable to a second variable to develop a more integrated or coherent model of what was happening. Several users viewed the day/night graph to interpret what was happening in other graphs, such as SolarPowerout. The variable BatteryCharging was also used this way, as were the SOC (battery state of charge) graphs. In addition, users compared the variables of analogous components to see if a fault were local to one part or general to the system. For example, many users selected one variable, such as state of charge, for each of the three batteries. If the function looked the same for all three, they concluded that the problem was not specific to any one battery, but originated outside and upstream of them. Users then monitored a variable from just one battery to track all three. By the same logic, several users also did this for the two power string variables.

Comparing Relations Across Variables. The ability to display information about variables supports

a still more powerful type of reasoning that is critical to understanding the system’s causal structure. Users often compared how and when one variable departed from its predicted value with the same details about a second variable, which let them make complicated inferences about causal dynamics. Six users clearly did this for multiple variable sets (five in the hierarchical condition but only one in the flat). Two additional (flat) users made clear comparisons between day/night and another variable. Two users (flat condition) never made multi-variable function comparisons and two users’ activity was ambiguous but did not clearly show comparison. To score behavior as variable comparison, the user needed to relate one graph to the other, either by explicitly pointing between corresponding points on the two plots or by explicitly relating the two variables verbally.

Subjects’ comparison of predicted-actual profiles across variables appeared to support several types of reasoning. One involved determining which variable deviated from the predicted trajectory first, and hence was a candidate cause of the later change. Users also made comparisons to better understand and reason about compensatory relations along variables. For instance, after one subject had plotted SOCBattery1, SOCBattery2, and SolarPowerOut through the third day cycle, he said “There was a deficit in solar power out [points day 2]. But here we have surplus [Points day 3; pause] that could cure [points SOC day 3] the problem of battery to go back to its original predicted level.”

Use of Process Information. Subjects seemed to have more difficulty using information about processes than they did information about variables, and they did not always make the appropriate inferences about processes. Specifically, users might attribute a fault to a process even when that process was not flagged yellow. For example, several subjects concluded that the fault lay in shunting because process “shuntSolarPower” was upstream of the problematic variable “SolarPoweOut”. This conclusion is suspect because the process was green, indicating that the expected input and output relations were being maintained. In contrast, the correct operation of this process was the critical piece of evidence for one user to hypothesize that the cause must lie outside the system itself. This subject had plotted SolarPowerOut, IOBatAmps, and SOCBat1, and aligned them at the right of screen. After checking the gimble system and concluding that it was fine, he said: “They [the processes] are not lighting up either, uh, providing output for a given input. So, [sighs, pause] bad input equals bad output. Right input. [very long pause] All I can say is they’re not getting enough sun. At this point.” This reasoning, one of the most sophisticated pieces we observed, critically exploits the information available about processes.

As we have seen, the combination of information about variables and processes provides a very rich, organized body of content the user can query. The study produced many examples of very sophisticated reasoning that draw upon and depended on these tools. However, subjects could also miss the relevant variables, draw conclusions where they should have consulted additional data, and, perhaps most pervasively, lose track of where they were in the problem-solving process. The data we collected should let us identify a number of possible improvements to the interface, as well as providing evidence for the utility of basic design.

8. Related Research

Although this paper has shown how to employ model-based techniques to develop an interactive monitoring system, many other ideas in the fault detection, isolation and diagnosis arena relate to our work. In this section, we discuss methods for handling uncertainty, supplying goal-oriented feedback to users, and extending diagnosis into repair, along with other languages for representing and simulating device models.

Although our work assumes that device modes are fully determined by available observations, many efforts in model-based monitoring take mode identification under uncertainty as the central issue. In this perspective, a device’s behavioral mode changes in response to hidden state that must be inferred from uncertain measurements and control inputs over time. For example, Williams et al. [27] employ a Bayesian updating scheme to compute the most likely model consistent with observations and commands. They represent device behaviors and transitions with hidden Markov models, augment them with constraints that capture system dynamics within each mode, and then compute each modes’ posterior likelihood from its prior distribution, the command input, and the available observations. This approach has the benefit that it naturally encompasses diagnosis, in that mode identification can select a fault mode, but it is computationally complex.

Researchers have developed a number of techniques for controlling the implied search, which, in principle, encompasses the space of all feasible modes for all device components, to find the most likely vector of modes at each instant in time. Williams et al. consider modes in descending order of prior likelihood in the context of an anytime algorithm. Dearden and Clancy [8] employ particle filters to sample possible modes probabilistically, while being sure to include the most interesting ones. In contrast, Narasimhan et al. [22] [23] decompose the problem by employing quantitative data to track device behaviors across mode transitions (assuming a linear dynamic and compensating for noise with Kalman filters), qualitative reasoning to isolate faults to causally preceding modes, and quantitative parameter estimation techniques to find the candidate that best fits the numeric observations. Other approaches encounter this search problem in different guises: Lawesson et al. [17] seek fault modes that entail other alerts by reasoning about models stated in a process algebra, while Jones et al. [13] search for logical constraints that, when removed, restore consistency between predictions and observed behavior. This relates to the seminal work by de Kleer et al. [7], who describe observations and component behavior in logic and map diagnosis onto the problem of assigning truth values (OK/Not OK) to component descriptions.

In summary, the task of mode identification in the presence of uncertainty entails unavoidable search problems. As our research expands to consider this task, we will encounter the same issues in an interactive setting. This suggests alternative strategies that rely on the user to prioritize candidate modes or that present only the most probable fault hypothesis. If the user rejects a candidate, the can shift attention to the next most likely interpretation. In general, the notion of presenting a single, but retractable, hypothesis may provide an effective tool for encapsulating the uncertainty of mode identification when interacting with people.

Systems that infer failures from behavioral measures can supply users with a highly desirable form of abstract feedback about device faults. Although expert systems can fulfill this role [2],

Larsson’s [16] work on diagnosis (recently incorporated into the commercial system, GoalArt [15]) performs this function in a particularly interesting way. He takes advantage of two properties of multi-level flow models, which represent device behavior in terms of a efforts and flows, much like hybrid bond graphs [20]. First, the primitive elements of such models support causal analysis of stylized faults, such as, “a *low capacity* alert on a *storage* object will cause a *low flow* alert in a downstream *transport*”. Similar rules suffice to restrict alarm cascades and isolate primal faults. Second, multi-level flow models link these functional descriptions into a hierarchy of goals and subgoals, as in saying that a cooling system’s *purpose* is to maintain desired water flow and that its ability to transport water through a reactor vessel depends upon the subgoals of having electrical supply and a cooled pump. This produces a direct connection from observations to device goals, enabling high-level feedback from primitive faults.

There are several interesting points of comparison between this work and our own. First, where GoalArt exploits the semantics of specific functional primitives to reduce alarm cascades, we trace faults through the causal graph. It is not yet clear if either mechanism is more powerful. However, GoalArt can use expectations from its primitives to question sensor values, while we must take sensor data as ground truth. A second point is that our work supplies more detail on the relation between device structure and function, while GoalArt adds a connection between function and goals. The implication is that we can improve the quality of our system’s output by incorporating teleological structure into our representation language, so that it can simultaneously localize faults to processes or functional subsystems and identify failed goals.

Although our work does not extend into fault repair, it is important to note that other monitoring systems do tackle this problem. For example, many deployed systems employ fault trees to map a branching series of yes/no questions into repair actions stored in the leaves. A more sophisticated approach views repair in control theoretic terms. For example, Williams and Nayak [28] describe a combinatorial optimization technique that drives a state transition model to a desired configuration, as applied in the context of NASA’s Deep Space One mission. Williams et al. [27] elevate this merger of monitoring and control into a proposal for developing fault-aware systems. Here reactive control programs set mode configuration targets, while fault detection, diagnosis, and repair methods implement that behavior plan.

Looking beyond other monitoring techniques, we should acknowledge the relation between our modeling formalism and other representational languages. In particular, the core notion of describing device behavior in terms of processes that are active under qualitative conditions comes from Forbus’ [11] work on qualitative process theory, although we incorporate more quantitative predictions. Work by Bobrow et al. [5] is also quite relevant, as they propose a compositional modeling language that shares many features with our own, including a part-of hierarchy, typed variables, and processes with activation conditions whose behavior can be described by both algebraic and differential equations. However, their work emphasizes issues of declarative knowledge representation, as opposed to causal accounts of behavior, and is less concerned about establishing a tight coupling between elements of the model and the modeled system or device. In other words, although the languages overlap, our research follows a very different theme.

Finally, we should note graphical environment, CONFIG [18] [19], that lets users compose and

simulate models of complex devices. This system supports discrete-event simulation with capabilities for continuous-time models in which models are expressed as deterministic state-transition systems. CONFIG has been employed as the central engine in a number of NASA simulation tests. Relative to our language, it offers more diverse model types and a richer capability for simulating models forward, but pays less attention to hierarchical model structure.

9. Concluding Remarks

In closing, we should consider some general issues that have emerged from our research on interactive monitoring. One surprising result is that a single representation of knowledge – hierarchically organized sets of quantitative process models – can support not only anomaly detection, but also fault isolation and diagnosis. This indicates that process models have substantial practical power and suggests they should be considered for other tasks as well. We believe the source of this power lies in the direct mapping between model structure and device structure, which leads directly to the formalism’s transparency and to the monitoring framework’s simplicity.

A second result is that we have been able to generate high-level user feedback on the basis of a functional description of device behavior, even though such feedback is typically associated with goal-oriented methods. Rather than view these two approaches as dichotomous, our work suggests that they can be productively merged. In particular, if the abstract vocabulary for describing faults is what carries value to users, we should consider augmenting our hierarchical but mechanical perspective with a layer of modeling that describes artifacts in purposive terms and use the new layer to interpret monitoring alerts. More broadly, we may be able to represent both physical and man-made systems as process models with functional and teleological interpretations.

Despite the advantages of hierarchical process models for monitoring tasks, there remain many ways in which we can extend the framework to improve it further. One limitation of the current scheme is that it cannot encode directly when a device or subsystem is operating in different modes. We can mimic mode behavior by placing conditions on processes, but some modes can activate or deactivate entire subsystems, and operating modes play such a central role in engineering that they may deserve special treatment.

Another drawback is the current framework’s inability to handle uncertainty. Even when accurate deterministic models are available, noise in sensors can produce uncertain predictions, and even well-established models have only limited accuracy. This suggests that we extend the modeling framework to support confidence intervals on both process parameters and sensor values, which will require more sophisticated methods for simulation to take this ambiguity into account. Such extensions will bring it closer to the capabilities of other model-based monitoring systems but still retain the interpretability that is one of its greatest strengths.

As mentioned earlier, we have done initial work using machine learning to revise our model of the electric power grid, but we hope to expand on this idea in our future efforts. Here we plan to extend the environment to include libraries of generic processes and subsystems that can serve as building blocks for new models. We will also develop methods for revising hand-crafted process models from the knowledge in these libraries and from observations. We envision not automated techniques for

model revision, but rather ones that let the user direct high-level aspects of the search, with the system handling low-level details like eliminating inconsistent candidates and estimating parameters from the data.

In parallel with these extensions, we should expand upon our efforts to evaluate the monitoring system with human test subjects. Earlier, we hypothesized that the system’s hierarchical display and support for causal analyses would improve user performance on diagnostic tasks. However, we need to test these predictions in controlled experiments, ideally via a lesion study that measures diagnostic accuracy and speed as system capabilities are removed. It would also be useful to extend that comparison to include a system like the one currently employed by NASA to monitor the electric power system on board the Space Station, which provides very few cues about EPS structure, as it is intended for expert users. If these studies produce encouraging results, we should consider paths for transitioning our human-centered approach to monitoring into operational use.

Acknowledgements

The research reported in this paper was funded by Grant NCC 2-1220 from NASA Ames Research Center. We thank Rick Alena at NASA Ames for many helpful conversations and for supplying a great deal of background knowledge concerning the electric power system onboard the International Space Station. We also thank Daryl Fletcher, Charles Lee, and Stephen Bay for their assistance in examining Station telemetry, and Javier Sánchez for his contributions to initial design and implementation of the process modeling environment.

References

- [1] G. Aaseng, K. Cavanaugh, and S. Deb. An intelligent remote monitoring solution for the international space station. In *Proceedings of the IEEE Aero Conference*, 2002.
- [2] Sermatech Intelligent Applications. Tiger: Gas turbine condition monitoring, <http://www.intapp.co.uk>.
- [3] S. Bay, D. Shapiro, and P. Langley. Revising engineering models: Combining computational discovery with knowledge. In *Proceedings of the Thirteenth European Conference on Machine Learning*, 2002.
- [4] S. Bay, J. Shrager, A. Pohorille, and P. Langley. Revising regulatory networks: From expression data to linear causal models. *Journal of Biomedical Informatics*, 35:289–297, 2003.
- [5] D. Bobrow, B. Falkenhainer, A. Farquhar, R. Fikes, K. Forbus, T. Gruber, Y. Iwasaki, and B. Kuipers. A compositional modeling language. In *Proceedings of The Tenth International Workshop on Qualitative Reasoning*, 1996.
- [6] S.̃. Cohen and A.̃. Hindmarsh. Ccode, a stiff/nonstiff ODE solver in C. *Computers in Physics*, 10:138–143, 1996.

- [7] J. de Kleer, A. Mackworth, and R. Reiter. Characterizing diagnoses and systems. *Artificial Intelligence*, 56:197–222, 1992.
- [8] R. Dearden and D. Clancy. Particle filters for real-time fault detection in planetary rovers. In *Proceedings of the 13th International Workshop on Principles of Diagnosis (DX-2002)*. <http://www.dbai.tuwien.ac.at/user/dx2002/>, 2002.
- [9] R.J. Doyle, S.M. Sellers, and D.J. Atkinson. A focused, context-sensitive approach to monitoring. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1989.
- [10] D. Dvorak. Monitoring and diagnosis of continuous dynamic systems using semiquantitative simulation. *PhD thesis, University of Texas, Austin, TX*, 1992.
- [11] K. Forbus. Qualitative process theory. *Artificial Intelligence*, 24:85–168, 1984.
- [12] M. Hofbaur and B. Williams. Mode estimation of probabilistic hybrid systems. In *International Conference on Hybrid Systems: Computation and Control*, 2002.
- [13] C. Jones, G. Bond, and P. Lawrence. Consistency-based fault isolation for uncertain systems with applications to quantitative dynamic models. In *Proceedings of the 13th International Workshop on Principles of Diagnosis (DX-2002)*. <http://www.dbai.tuwien.ac.at/user/dx2002/>, 2002.
- [14] P. Langley, J. Sanchez, L. Todorovski, and S. Dzeroski. Inducing process models from continuous data. In *Proceedings of the Nineteenth International Conference on Machine Learning*, pages 347–354. Morgan Kaufmann, 2002.
- [15] J. Larsson. Goalart, <http://www.goalart.com>.
- [16] J. Larsson. Diagnosis based on explicit means-end models. *Artificial Intelligence*, 80(1):29–83, 1996.
- [17] D. Lawesson, U. Nilsson, and I. Klein. Fault isolation using process algebra models. In *Proceedings of the 13th International Workshop on Principles of Diagnosis (DX-2002)*. <http://www.dbai.tuwien.ac.at/user/dx2002/>, 2002.
- [18] J. Malin. Using hybrid modeling for testing intelligent software for lunar-mars closed life support. *JOM-e*, 51(9), <http://www.tms.org/pubs/journals/JOM/9909/Malin/Malin-9909.html>, 1999.
- [19] J. Malin, D. Ryan, and L. Fleming. Configintegrated engineering of systems and their operation. In *Proceedings of the Fourth National Technology Transfer Conference (NASA Conference Publication CP-3249)*, pages 97 – 104, 1993.
- [20] P. Mosterman. Hybrsim - a modelling and simulation environment for hybrid bond graphs. *Proceedings of the I MECH E Part I Journal of Systems & Control Engineering*, 216:35–46, 2002.

- [21] P. Mosterman and G. Biswas. Monitoring, prediction, and fault isolation in dynamic physical systems. In *Proceedings of the American Association for Artificial Intelligence AAAI*, 1997.
- [22] S. Narasimhan, B. Biswas, G. Karsai, and T. Szemethy. Hybrid modeling and diagnosis in the real world: A case study. In *Proceedings of the 13th International Workshop on Principles of Diagnosis (DX-2002)*. <http://www.dbai.tuwien.ac.at/user/dx2002/>, 2002.
- [23] S. Narasimhan, G. Biswas, and G. Karsai. An integrated approach to diagnosis of complex hybrid systems. In *15th Annual International Symposium on AeroSense (Component and Systems Diagnostics, Prognosis, and Health Management)*, pages 275–286, 2001.
- [24] B. Rinner and B. Kuipers. Monitoring piecewise continuous behaviors by refining semi-quantitative trackers. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, 1999.
- [25] G. Saito, P. Langley, S. Bay, and K. Arrigo. Discovering ecosystem models from time-series data. In *Proceedings of the Sixth International Conference on Discovery Science*, in press.
- [26] D. Shapiro. Sniffer: A system that understands bugs. *Master’s Thesis: Massachusetts Institute of Technology, Artificial Intelligence Laboratory. AI Memo 638*, 1981.
- [27] B. Williams, M. Ingham, S. Chung, P. Elliott, M. Hofbaur, and G. Sullivan. Model-based programming of fault-aware systems. *AI-Magazine*, Fall, 2003.
- [28] B. Williams and P. Nayak. A model-based approach to reactive self-configuring systems. In *Proceedings of the American Association for Artificial Intelligence*, pages 971–978, 1996.