/N-63-CR
134753
16 P.

JPL Publication 87-40

Predictive Monitoring Research: Summary of the PREMON System

Richard J. Doyle
Suzanne M. Sellers
David J. Atkinson

(NASA-CR-182687) PREDICTIVE MONITORING RESEARCH: SUMMARY OF THE PREMON SYSTEM (Jet Propulsion Lab.) 16 p CSCL 09B

N88-20905

Unclas
G3/63 0134753

December 15, 1987

NASA
National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

JPL Publication 87-40

Predictive Monitoring Research: Summary of the PREMON System

Richard J. Doyle
Suzanne M. Sellers
David J. Atkinson

December 15, 1987

NASA

National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

ABSTRACT

Traditional approaches to monitoring are proving inadequate in the face of two important issues: the dynamic adjustment of expectations about sensor values when the behavior of the device is too complex to enumerate beforehand, and the selective but effective interpretation of sensor readings when the number of sensors becomes overwhelming. Our system addresses these issues by building an explicit model of a device and applying common-sense theories of physics to model causality in the device. The resulting causal simulation of the device supports planning decisions about how to efficiently yet reliably utilize a limited number of sensors to verify correct operation of the device.

Table of Contents

PRECEDING PAGE BLANK NOT FILMED

PAGE IV INTENTIONALLY BLANK

TECHNICAL REPORT STANDARD TITLE PAGE

1. Report No. JPL PUB 87-40	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle PREDICTIVE MONITORING RESEARCH: SUMMARY OF THE PREMON SYSTEM		5. Report Date December 15, 1987 6. Performing Organization Code
7. Author(s) Richard J. Doyle, Suzanne M. Sellers, David J. Atkinson		8. Performing Organization Report No.
9. Performing Organization Name and Address JET PROPULSION LABORATORY California Institute of Technology 4800 Oak Grove Drive Pasadena, California 91109		10. Work Unit No. RE65 CY-549-01-00-05-59 11. Contract or Grant No. NAS7-918 13. Type of Report and Period Covered
12. Sponsoring Agency Name and Address NATIONAL AERONAUTICS AND SPACE ADMINISTRATION Washington, D.C. 20546		14. Sponsoring Agency Code
15. Supplementary Notes		

16. Abstract

Traditional approaches to monitoring are proving inadequate in the face of two important issues: the dynamic adjustment of expectations about sensor values when the behavior of the device is too complex to ennumerate beforehand, and the selective but effective interpretation of sensor readings when the number of sensors becomes overwhelming. Our system addresses these issues by building an explicit model of a device and applying common-sense theories of physics to model causality in the device. The resulting causal simulation of the device supports planning decisions about how to efficiently yet reliably utilize a limited number of sensors to verify correct operation of the device.

17. Key Words (Selected by Author(s)) Spacecraft testing Qualitative simulation Sensor planning Causal modeling Common-sense reasoning Artificial intelligence	18. Distribution Statement UNCLASSIFIED--UNLIMITED

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
UNCLASSIFIED	UNCLASSIFIED	14	

JPL 0184 R 9/83

Overview

PREMON is a predictive monitoring system being developed by JPL's Artificial Intelligence Group in the Computer Science and Applications Section. The system is funded through the NASA Systems Autonomy Core Research Program of the Office of Aeronautics and Space Technology. Our overall, long-range goal is to design and implement a software prototype that will demonstrate cooperation (via a blackboard architecture) between knowledge-based systems for the problem of monitoring a complex device. The prototype will consist of three modules: a *causal simulator*, a *sensor verification planner*, and a *sensor interpreter*.

System Architecture

Currently, PREMON consists of the *causal simulator* only. The causal simulator maintains a causal model of the device being monitored. Our initial device domain is the 25-foot space simulator at JPL, a vacuum chamber used for testing spacecraft in an environment similar to that found in space. The causal model is based on a representation designed by Richard Doyle at the Massachusetts Institute of Technology. Doyle's representation is an extension to one developed by Kenneth Forbus [Forbus 84]. The causal model simulates the future behavior of the device based on the device's current state or a hypothetical state. A valuable attribute of the causal simulator is that it can be used to generate and reason about alternative scenarios of future device behavior. The need to explore alternative scenarios arises from the desire to resolve uncertainty about the device's observed state, the desire to investigate the effects of potential operator inputs to the device, or the desire to detect possible inaccuracies in the causal model itself.

The *sensor verification planner* will have the capability to use causal simulator output, in the form of device behavior envisionments, to generate a plan to utilize sensor resources during device operation. The development of this module is motivated by limitations in current monitoring systems with respect to sensors, e.g.: limitations on the bandwidth of data that may be acquired, limitations on the available sensors or their configuration, and limitations on human attention to relevant sensor data. Current systems lack a way to automatically maximize acquisition of the most important sensor data. The most important sensor data is that which provides the most feedback on the processes critical to evaluating the functionality of the device <u>in the current operating context.</u> By careful analysis of the active dependencies at each stage in the operation of the device, we can establish which physical components are most directly supporting the stability of the device state, or are contributing to changes in the device state. Sensor data reflecting the state of those components would be the most important. A theory of relative importance of sensor data with respect to operating context will be developed and tested in the implementation of the sensor verification planner.

The *sensor interpreter* will use output from the sensor verification planner, in the form of a sensor plan, qualitative parameter value expectations, and observed parameter values, to determine whether predictions about the behavior of the device being monitored match its actual behavior. The process will be model-driven.

Approach to Causal Modeling and Qualitative Simulation

In building the PREMON's causal simulator module, we are concerned directly with the distinguishable states of physical systems and how they change from one state to another. We take a process-oriented, rather than a device-oriented, approach to determining the causality in physical systems. Device-orientation considers first the behaviors of the components of a device, then how those behaviors are propagated through the device. Process-orientation, on the other hand, considers first what processes support causality in a device, then which components participate in those processes.

The qualitative simulation technique we are using is based on research by Kenneth Forbus. It is known as limit analysis, it determines which processes are active at any given time in a physical system, and it propagates values accordingly. As they change, the values of quantities may reach thresholds ("limits"), resulting in state changes in the system where active processes become inactive or vice versa.

The causal modelling capability that we are developing is driven by knowledge of the kinds of mechanisms existing in physical systems. Examples of these causal mechanisms are fluid flow, mechanical coupling, expansion, gravity, and physical contact.

As part of this effort, Richard Doyle has developed a representation for causality in physical systems and a vocabulary of causal mechanisms describable within it. The representation distinguishes four categories of causal mechanisms:

(1) Propagations: causally related events at different sites, structurally linked through *media* and potentially disrupted by *barriers*. Example: Fluid Flow
(2) Transformations: causally related events at a single site, typically involving some kind of energy transformation. Example: Expansion
(3) Field Interactions: events due to interaction with a field. Example: Gravity
(4) Thresholds: continuous changes which lead to abrupt changes in physical or geometrical relations. Example: Contact

The current causal simulator works as follows: Given an initial state of a physical device and a specification of state changes to simulate input from device sensors (we are not currently connected to the space simulator sensors), the causal simulator uses a common-sense theory of causal mechanisms in physical devices to generate a simulation of the future behavior of the device.

To support simulation, a qualitative calculus has been developed which encodes knowledge of how qualitative values combine under negation, addition, and multiplication.

Research Issues

Qualitative Simulation

Doyle's representation for causality directly describes physical relations which enable and disable processes. A research goal for PREMON is to give this causal modeling system and simulator the capability to identify and reason about thresholds that establish and remove media and barriers and thus alter the active process structure of a device.

Most of the existing qualitative simulators use a simplified notion of temporal integration. They merely note the sign of the derivative of a quantity and find the next qualitative value in the indicated direction. Forbus goes a bit further by using the magnitude, as well as the sign of the derivative. Our simulator uses a more powerful notion of temporal integration that includes the duration of the interval during which a quantity's derivative is non-zero. Therefore, our system can reason about the result of turning on a heating or chilling component based on the length of time the component is active. This is a seemingly straightforward inference, yet the more complete notion of temporal integration embedded in a simulator is needed to make it. We will continue this line of research.

Currently, PREMON simulates future behavior of the device deducible from singular causation, e.g., when the chiller is turned on, a negative temperature rate of change is propagated through the gaseous nitrogen cooling circuit for the space simulator mirror. We are researching methods of integrating multiple causes of events into the simulation of future device behavior. For example, it is desirable to integrate changes in the mirror temperature resulting from heat loss to the liquid-nitrogen-cooled walls, heat gain from the ceiling of the chamber, as well as temperature changes internal to the mirror cooling circuit.

Sensor Planning

The key to determining which components are most appropriate for monitoring is the establishment of the relative importance of the operation of that component at each instant of time in the dynamic behavior of the device as a whole. We believe that a deep, qualitative model of the behavior of the device yields information about the causal dependencies between different active processes. Development of methods for elaborating these dependencies and including them in a planning process are key objectives of our research.

Sensor Interpretation

Substantial domain knowledge is required to detect departures from nominal behavior when there are no active expectations; knowledge about how other events or data may *differ* from what is nominally true is needed. We suspect that this knowledge is fundamentally the same as the underlying qualitative causal model. Identifying this knowledge and devising indexing and accessing strategies are objectives of our research.

In addition, we have the problem that the data being sensed can have a variety of features that make the process of symbolically matching it against expectations difficult. The computational complexity of interpreting sensor data increases with the variety and number of events that must be recognized. A match against expectations is fundamentally an inexact process called "partial matching." We propose to investigate heuristic techniques for resolving uncertainty in the matching process and to investigate improved approaches to knowledge organization.

It is the responsibility of the Sensor Interpreter to use the sensor plan to configure the use of sensors during device operation. Mechanisms for sequencing the execution of a sensor plan with the corresponding expectations about sensor data need to be developed. This is an additional focus of our research.

Choice of Initial Domain

Although we found the 25-foot space simulator to be the best initial domain for our research, the following also have merit as future candidate domains: Deep Space Network Tracking Station, Mars Rover, and Space Station. We sought a domain with the following characteristics:

(1) *Numerous and diverse mechanisms and sensors.* We want to address research issues in the context of a problem domain of significant complexity to demonstrate the efficacy of new, as well as extant, theories.

(2) *Multiple operating modes.* Multiple modes present a problem for current monitoring systems in the form of false alarms. False alarms result from the inability of pre-defined nominal operating ranges for sensors to correctly represent alarm states <u>within each operating context</u>. We wish to demonstrate the ability to reduce false alarms by dynamically shifting alarm ranges of sensors to reflect what is and is not an alarm state specific to each operating mode.

(3) *Requirement for real-time response.* Our performance goal is to simulate the device state, plan sensor usage, and interpret sensor data in complex systems such that anomalies are detected in real time.

(4) *Burden of interpretation on human operators.* PREMON is intended to automate the global interpretation of sensor readings and, in some cases, ameliorate the need for round-the-clock human supervision.

(5) *Difficulty of comprehensive sensor interpretation.* When sensors become too numerous, the task of synthesizing a global picture of the state of the system becomes overwhelming. We are exploring sensor planning--sampling some subset of the available sensors based on the evaluation of the relevance of particular sensor readings.

The space simulator matches our current criteria better than the other domains we considered. It contains a diverse set of mechanisms and sensors. There are several different environments achievable within the chamber. These target environments represent the various operating modes of the simulator. Nominal sensor readings for each mode are different. Detection of anomalies during the process of reaching a target environment, and while maintaining it, must occur in real time. Round-the-clock human supervision is required. Sensor planning is appropriate for the space simulator and is implicit in the established procedures. These procedures require operators to monitor only the sensors relevant to the current operating mode.

The three domains other than the space simulator represent excellent longer-term targets. The DSN tracking stations offer the best support for testing and gradually phasing in new technology. The Mars Rover scenario represents the most challenging monitoring task, and almost demands a new approach to monitoring. Finally, the Space Station Thermal Management System, with its greater complexity, would prove the most thorough test of our proposed approach to monitoring.

Reference

Forbus, Kenneth D.
Qualitative Process Theory.
Technical Report #789, Artificial Intelligence Laboratory, Massachusetts Institute of Technology. Cambridge, MA. July 1984.

Bibliography

Atkinson, D., James, M., Porta, H., Doyle, R.
Autonomous Task Level Control of a Robot.
In Proceedings of Robotics and Expert Systems, 2nd Workshop, Instrument Society of America. NASA/Johnson Space Center. Houston, TX. June 1986.

Bell, Colin
Resource Management in Automated Planning.
AIAI TR #8, Artificial Intelligence Applications Institute, University of Edinburgh. United Kingdom. July 1985.

Brooks, Rodney.
Symbolic Error Analysis and Robot Planning.
A.I. Memo #685, Artificial Intelligence Laboratory, Massachusetts Institute of Technology. Cambridge, MA. September 1982.

Corkill, Daniel D., Gallagher, Kevin Q., and Murray, Kelly E.
GBB: A Generic Blackboard Development System.
In Proceedings of the Fifth National Conference on Artificial Intelligence. Philadelphia, PA. 1986.

Cullingford, Richard E.
Script Application: Computer Understanding of Newspaper Stories.
Research Report #116, Department of Computer Science, Yale University. New Haven, CT. January 1978.

de Kleer, Johan., Seely Brown, John.
Mental Models of Physical Mechanisms.
Unpublished Manuscript, personal communication.
XEROX PARC, Cognitive and Instructional Sciences. Palo Alto, CA. December 1980.

DeJong, Gerald F.
Skimming Stories in Real Time: An Experiment in Integrated Understanding.
Research Report #158, Department of Computer Science, Yale University. New Haven, CT. May 1979.

Donald, Bruce.
Robot Motion Planning with Uncertainty in the Geometric Models of the Robot and Environment: A Formal Framework for Error Detection and Recovery.
In Proceedings of the IEEE International Conference on Robotics and Automation. San Francisco, CA. 1986.

Doyle, Richard J.
Constructing and Refining Causal Explanations from an Inconsistent Domain Theory.
In Proceedings of the Fifth National Conference on Artificial Intelligence. Philadelphia, PA. 1986.

Durfee, Edmund H., and Lesser, Victor R.
Incremental Planning to Control a Blackboard-based Problem Solver.
In Proceedings of the Fifth National Conference on Artificial Intelligence. Philadelphia, PA. 1986.

Erdmann, Michael.
Using Backprojections for Fine Motion Planning with Uncertainty.
In Proceedings of the IEEE International Conference on Robotics and Automation. St. Louis, MO. 1985.

Fikes, R.E., Nilsson, N.J.
STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving.
Artificial Intelligence Journal 2(3-4), 1971.

Fikes, R.E., Hart, P.E., Nilsson, N.J.
New Directions in Robot Problem Solving.
Machine Intelligence 7, 1972.

Forbus, Kenneth D.
Measurement Interpretation in Qualitative Process Theory.
In Proceedings of the Eighth International Joint Conference on Artificial Intelligence. Cambridge, MA. 1983.

Friedman, Leonard.
Diagnosis Combining Empirical and Design Knowledge.
(JPL Internal Document D-1328), Jet Propulsion Laboratory. Pasadena, CA. December 1983.

Gini, Maria., Doshi, Rajkumar S., Garber, Sharon., Gluch, Marc., Smith, Richard., Zualkernain, Imran.
Symbolic Reasoning as a Basis For Automatic Error Recovery in Robots.
Technical Report 85-24, University of Minnesota. St. Paul, MN. July 1985.

Hayes-Roth, B.
A Blackboard Architecture for Control.
Artificial Intelligence Journal 26(3), 1985.

Hayes-Roth, F.
The Role of Partial and Best Matches in Knowledge Systems.
In Waterman, D.A., and Hayes-Roth, F. (Eds.),
Pattern-directed Inference Systems.
New York: Academic Press. 1978.

Kuipers, B.J.
Qualitative Simulation.
Artificial Intelligence Journal 29(3), 1986.

Miller, David P.
Planning by Search through Simulations.
Ph.D. thesis, Yale University. New Haven, CT. October 1985.

Moravec, Hans P., and the Mobile Robot Laboratory Staff.
Towards Autonomous Vehicles.
In *Autonomous Mobile Robots: Annual Report 1985.*
Technical Report CMU-RI-TR-86-4, Mobile Robot Laboratory,
Carnegie Mellon University. Pittsburgh, PA. February 1985.

Nii, H. Penny, et al.
Signal-to-Symbol Transformation: HASP/SIAP Case Study.
The AI Magazine. Spring 1982. pp. 23-35.

Nii, H. Penny, Feigenbaum, E.A.
Rule-based Understanding of Signals.
In Waterman, D.A., and Hayes-Roth, F. (Eds.),
Pattern-directed Inference Systems.
New York: Academic Press. 1978.

Porta, Harry.
Dynamic Replanning.
In Proceedings of Robotics and Expert Systems, 2nd Workshop.
Instrument Society of America. NASA/Johnson Space Center.
Houston, TX. June 1986.

Rieger, Chuck., and Grinberg, Milt.
The Declarative Representation and Procedural Simuation of Causality in Physical Mechanisms.
In Proceedings of the Fifth International Joint Conference on
Artificial Intelligence. Cambridge, MA. 1977.

Riesbeck, Christopher K.
An Expectation-driven Production System for Natural Language Understanding.
In Waterman, D.A., and Hayes-Roth, F. (Eds.),
Pattern-directed Inference Systems.
New York: Academic Press. 1978.

Sacerdoti, Earl.
A Structure for Plans and Behaviour.
Elsevier North-Holland Inc., 1977.

Tate, Austin.
Planning and Condition Monitoring in a FMS.
Technical Report AIAI TR #2, Artificial Intelligence Applications Institute, University of Edinburgh. United Kingdom. July 1984.

Van Baalen, Jeffrey.
Exception Handling in a Robot Planning System.
Presented at the IEEE Workshop on Principles of Knowledge-based Systems, Denver, CO. December 1984.

Wallace, Richard., et al.
First Results in Robot Road-following.
In *Autonomous Mobile Robots: Annual Report 1985.*
Technical Report CMU-RI-TR-86-4, Mobile Robot Laboratory, Carnegie Mellon University. Pittsburgh, PA. February 1985.

Weld, Daniel S.
Switching Between Discrete and Continuous Process Models To Predict Genetic Activity.
Technical Report #793, Artificial Intelligence Laboratory, Massachusetts Institute of Technology. Cambridge, MA. May 1984.

Wilkins, David.
Domain Independent Planning: Representation and Plan Generation.
Technical Note #266, SRI. Menlo Park, CA. August 1982.

Wilkins, David.
Recovering From Execution Errors in SIPE.
Technical Note #346, SRI. Menlo Park, CA. January 1985.