REACTIVE CONTROL AND REASONING ASSISTANCE FOR SCIENTIFIC LABORATORY INSTRUMENTS

David E. Thompson
Richard Levinson
Peter Robinson

NASA Ames Research Center
Mail Stop 269-2
Moffett Field, CA 94035

ABSTRACT

Scientific laboratory instruments that are involved in chemical or physical sample identification frequently require substantial human preparation, attention, and interactive control during their operation. Successful real-time analysis of incoming data that supports such interactive control requires (1) a clear recognition of variance of the data from expected results and (2) rapid diagnosis of possible alternative hypotheses which might explain the variance. Such analysis then aids in decisions about modifying the experiment protocol, as well as being a goal itself. This paper reports on a collaborative project at the NASA Ames Research Center between artificial intelligence researchers and planetary microbial ecologists. Our team is currently engaged in developing software that autonomously controls science laboratory instruments and that provides data analysis of the real-time data in support of dynamic refinement of the experiment control. The first two instruments to which this technology has been applied are a differential thermal analyzer (DTA) and a gas chromatograph (GC). Coupled together, they form a new geochemistry and microbial analysis tool that is capable of rapid identification of the organic and mineralogical constituents in soils. The thermal decomposition of the minerals and organics, and the attendant release of evolved gases, provides data about the structural and molecular chemistry of the soil samples.

INTRODUCTION

Over the past two years, researchers at NASA Ames have developed a new scientific laboratory instrument and have implemented intelligent control and analysis software to support the operations and data analysis of this new instrument. In particular, the authors, as researchers in artificial intelligence, have worked in close collaboration with two Ames microbial ecologists, Rocco Mancinelli and Lisa White, to affect this development. This paper focusses on the intelligent software technology part of that project and its potential generalization to other scientific laboratory instruments. Scientific laboratory instruments that are involved in chemical or physical sample identification frequently require substantial human preparation, attention, and interactive control during their operation. Our software is intended to alleviate the user of much of this attention and interactive control. Successful real-time analysis of incoming data that supports such interactive control requires (1) a clear recognition of variance of the data from expected results and (2) rapid diagnosis of possible alternative hypotheses which might explain the variance. Data analysis is a goal in its own right; real-time analysis, however, can support decisions about modifying the experiment protocol. Thus, the software both reactively controls science laboratory instruments and provides data analysis in support of dynamic refinement of the experiment control. The first two instruments to which this technology has been applied are a differential thermal analyzer (DTA) and a gas chromatograph (GC). Coupled together, they form a new geochemistry and microbial analysis tool that is capable of rapid, robust identification of the organic and mineralogical constituents in soils. The thermal decomposition of the minerals and organics, and the attendant release of evolved gases, provides data about the structural and molecular chemistry of the soil samples. The details of this new tool are provided in the following section.

The coupling of these analysis systems results in a more detailed characterization of the minerals and organics in the soil samples than has previously been available; their combined use has also required the development of new reasoning expertise, detailing how the data or results of the two types of analyses interrelate. This expertise has been gained through the construction of an integrated DTA-GC instrument itself, through development of the control and reasoning software in synchrony with this construction, and finally through the use of the system on soils and mixtures whose chemical decompositions provide clear examples of the interplay between thermophysical and chemical processes.

The DTA-GC software has been implemented in terms of three development levels. Level 1 represents functionality of the system as a reactive (that is, non-planning) controller. It requires the operation of the sensory perception, analysis, and control components, and the system reacts to evolved gas events by recognizing the event as an increase in oven pressure, and then exercising the GC sampling protocol. Both DTA and GC data is analyzed. At level 2, a predictive control loop is added by introducing an experiment planner, but all the components still operate sequentially. This means that in this first phase of operation (levels 1 and 2), the system is capable of controlling a single experiment run and then reasoning about the data after the run has completed. The received data is matched against encoded representations of data in mineral library records. The matches form explanations of the observed features in the data, and represent a best-guess identification of mineral and organic content. This explanation and identification can then be used to suggest follow-up experiment protocol needed to resolve any ambiguities in the identification. At level 3, all of the components can operate in parallel. This phase of implementation is a transition to operations in an interrupt mode, exploiting parallel reasoning, planning, and execution, whereby the system carries out partial matches of data with the library records while the data is "coming in" from a run. It thus allows re-programming of an experiment profile during that run, based on expectations of identification, if there are deadline limits. The status of our system with respect to each of these development levels is discussed in a later section of this paper.

The team is engaged in establishing performance criteria and evaluation standards for all the software, yielding performance metrics which can guide our extensions and help empirically determine the meaning of 'improvements' to the system. We are particularly interested in exploring the necessary trade-offs between speed of analysis and fidelity of analysis: accuracy in reporting identification versus speed and economy of the representation, noting discrimination errors. These metrics are currently being established only for the DTA-GC instrument, without consideration to generalization of the system to other instruments. The more improvements we make to affect robust control and reasoning, the more we will understand the possibilities for generalization of this software to other science instruments.

As a second prototype, it is our intention to apply the software system to a multistage bioreactor being developed at Ames during this next year. Our particular bioreactor will be used in part to evaluate the microbial paleo-environmental conditions and constraints that are implicitly represented in inhabited soil and mineral samples studied through DTA-GC. It will also be used separately to study the nutrient and environmental characteristics of natural carbon and nitrogen cycling in the Earth system. This work therefore supports NASA's interest in the role of nutrient cycles in Global Change studies, in the effects of planetary physico-chemical environments on early evolution of life, and in controlled ecological life support systems. The multistage bioreactor requires far more extensive reactive control and reasoning assistance during its operation than has been found necessary under DTA-GC, so this extension will help guide further software enhancements.

The long range commercial potential for the DTA-GC instrument itself is primarily for use as an analysis tool in laboratories (or in the field) that require rapid identification of solid samples without the need for refined wet-chemistry or scanning calorimetry. Additionally, the intelligent software developed for DTA-GC provides further commercial potential as a generic predictive/reactive control and reasoning architecture that can assist scientists in critical control and analysis decisions, and can allow for instrument operations and electronic-linked analysis under remote or hostile conditions.

BACKGROUND ON DTA AND GC, AND THE COUPLED SYSTEM

A differential thermal analyzer is an unpressurized programmable oven -- it heats up mineral or other solid samples at a controlled rate, from ambient temperature and pressure to 1200 degrees C. The heating causes the minerals to undergo chemical and structural changes. These changes include phase transformations in the mineral structure, melting, oxidation, nucleation and crystal reorganization, or simple breaking of chemical bonds and release of gases that are either physically adsorbed interstitially or are chemically bonded in the lattice structure of the particular minerals. Any organics that are contained in the sample of course undergo similar decompositions with attendant release of gas from the residue. Any substance put into the DTA oven will produce particular changes upon heating depending on its chemistry and crystal structure, thus allowing for partial identification of the substance. The temperature changes in the sample are measured against the temperature of an inert reference. The resulting difference in temperature, at ambient pressure, is proportional to the energy utilized or released in the sample during these thermophysical events. Hence, the changes in the sample are recorded by the DTA as "difference features" in the data stream. Any event which utilizes energy is *endothermic*. The sample then appears as "cooler" than the reference, and

thus produces "valleys" in the data stream. Events that release energy are *exothermic* and show up as "peaks" in the data stream. The character of such a thermal event, such as its duration, intensity, onset temperature, and whether it is endo- or exothermic, is indicative of the mineral structure, proportion, and content in the sample, but it is not unambiguously diagnostic for identification. It is important to realize that DTA by itself does not provide any chemical information except from inference. Furthermore, because DTA measures only *temperature* differences, there is no direct measure of the actual heat involved in a given reaction, so no information is revealed concerning heat capacities of the minerals. Nor can one definitively measure relative proportions of different minerals contained in the sample cup, rather only their presence or absence. There is also no guarantee that presence of trace minerals will be detected unless significant amounts of these are contained in the sample to yield a signal, but of course then one does not know whether their presence is significant or only at trace levels in the parent sample. Certain variations in silicate structure, notably in clays, will not show up as differences in DTA signatures, yet this structure can be quite important since it controls the availability of lattice sites for certain ionic replacements or even preferential locations for organic compounds. DTA also provides no information on the grain size distribution of the sample or of the parent rock. All these points notwithstanding, DTA is an extremely robust thermal analysis instrument which faithfully identifies the presence or absence of diagnostic thermal events from which detailed mineral structure can be inferred, and from which much headway can be made concerning inference about processes and reaction pathways. This makes it an ideal system to be coupled with other analysis techniques, and it also may be made capable of operating in field conditions outside of a well-supplied laboratory, or even on planetary surfaces.

A gas chromatograph essentially consists of a column of material through which gas mixtures flow for purposes of constituent identification, plus a detector that quantifies the gases as they flow out of the column. When a gas mixture flows through the column, the individual gas compounds diffusively separate due to their differing affinities for the material packed in the column, and thus the compounds can be identified chemically according to their relative flow rates. This identification is at the molecular level, not elemental or ionic level; GC provides chemical information, not molecular weight information as might a mass spectrometer. The GC gives total proportional volume of gas compounds eluted through the column during one diffusion event or gas injection. By sending the sample through both polar and nonpolar columns and detectors in parallel, that are separately calibrated for particular gas compounds, the normally varying retention times can be compressed so that all the data becomes available at roughly the same time (on the order of minutes) without the various gases interfering with or masking each other. This is especially important when trace gases are sought, because the high resolution on the column necessary to detect the trace gas signatures would be swamped by even a small amount of water being eluted through the column.

The coupled DTA-GC instrument is itself a new research tool. By coupling a DTA to a GC, the scientist can determine both structural and evolved gas chemistry of a single sample. When both sources of information are combined, a more complete and less ambiguous characterization results. Typically, GCs have a pyrolitic "front-end", and rapidly heat up an entire sample. By using a programmable oven, the samples are heated slowly so that when gases are released during a thermophysical event, that release temperature is recorded and the gases are temperature-(time-) stamped according to when in the experiment run they evolved off the sample. Hence, if one observes carbon dioxide gas coming off at around 350 degrees C, then one knows it is from decomposition of organics and not from decomposition of a calcium carbonate like limestone since the limestone decomposition and its release of CO_2 occurs at near 600 degrees C. Thus, decisions can be made as to the amount and type of minerals that are present in the sample, and one can discriminate between the gases from minerals and from organics. If the sample is an unknown, then its DTA and GC "signatures" are compared by our software to characterizations in the database, along with geochemical domain knowledge and with expectations generated by the system. The system generates a set of hypotheses about what the sample contains, and it suggests and controls variations in the experiment run that will help to eliminate alternative hypotheses. Such a system can perform analysis either for target minerals and organics or for toxic compounds, and it can both verify expectations and suggest presence or absence of unanticipated species. The additional information available from the coupled DTA-GC system enhances independent DTA structural information or molecular chemistry from the GC. It also contributes to elucidation of reaction pathways and provides gas volume proportions, but not unambiguous mineral proportions. Were mineral proportions available, then one could map the DTA-GC information back to parent rock or even to geologic environment information. However, the problem remains that from knowing only presence or absence of minerals, only disjunctive possibilities of parent rocks or environments can be known. Rocks are identified not only by their chemistry but by the exact proportions and occurrence of those chemical constituents so that several vastly different kinds of rocks may still have identical chemistry. Furthermore, there is a critical sampling issue concerning whether the distribution of minerals or substance in the DTA sample cup is representative of the proportionate distribution in the parent rock. This indicates that in order to make the DTA-GC system into a functioning geologic analysis assistant, a different class of information is needed, specifically that concerning grain size of species and proportionate occurrence.

REQUIREMENTS FOR ANALYSIS AND CONTROL

The kinds of science instruments we are concerned with may be broadly classified as ones in which control decisions are made reactively, in real-time, based on incoming data. For example, we run the DTA-GC under mild vacuum so that release of gases from the sample during decomposition may be recognized by pressure sensors, thereby immediately triggering or changing the GC sampling strategy for that run. This reactive control notwithstanding, the DTA-GC currently operates best in a mode where decisions on sample identification are delayed until all relevant data has been acquired so that as much uncertainty as possible is eliminated before analysis. In this section, we discuss the capabilities required to support intelligent analysis and control.

The DTA-GC application requires *sensory perception* capabilities. We define these simply as the ability to acquire information about the external world via sensors. The system must interpret real-time DTA, GC, and pressure signals from the hardware sensors. These sensors provide results in the form of voltage streams that are typically plotted graphically and then visually interpreted by humans. Because our system operates semi-autonomously, it needs some signal processing capabilities for recognizing peak and valley features in the voltage streams. Even though it does not require graphical representation for its decisions, we provide graphical display of the data for use by the attending scientist. Furthermore, the system must address a form of limited perception since it is never certain which events will be encountered during the heating process. This uncertainty is compounded by signal/noise or figure/background discrimination tasks. For example, it may be difficult to discriminate between, or assign semantic meaning to, a single "valley" signal versus two "peaks". Thus, some heuristics are necessary to bias such decisions.

The application requires *data analysis* capabilities, which we define as any processing or reasoning over data that was acquired through sensory perception. The results of DTA-GC data analysis is a set of hypotheses that postulate mineral combinations that could be contained in the sample. When a single observed event can be explained by two different minerals because they both have events in the same temperature range, multiple hypotheses are produced. The result is a set of competing hypotheses that represent an ambiguous model of the unknown soil. Because this is the first combined "DTA-GC" system, the only experts on the analysis of this combined data are our microbial ecologists, who themselves are learning about the new system. However, experts in DTA and in GC separately often employ a variety of heuristic knowledge when they choose between alternative hypotheses or explanations. We need to model the expert's reasoning process using a high level language so that our results will make sense to these scientists. Ideally, the scientists should also be able to develop and maintain the knowledge base themselves. This need for a high-level knowledge-based representation combined with heuristic search are the typical motivations for expert system techniques. Since a given observation may not perfectly match the generalized characterization in our mineral library, the use of probabilistic techniques for assignment of matches is also needed. Further, belief revision techniques are motivated due to the system's limited perception and incremental data acquisition in an uncertain world.

This application requires *planning* capabilities, which we define as the ability to select actions by performing "look ahead" or "predictive" search. Because the constituents of the soil sample and its in situ environment are unknown, an appropriate set of experiments to validate or clarify identification cannot be designed in advance. Therefore, the system must perform on-line planning in order to design experiments based on knowledge gained. Also, since competing hypotheses will often exist, the system must take actions aimed at clarifying ambiguities. For example, consider a case in which the first sample run indicates only that gas evolved somewhere between the temperatures of 200 and 700 degrees. The data analysis results from that run might then induce two competing hypotheses: one assuming the gas was produced at 300 degrees and another assuming it happened from a different event at 600 degrees. A simple follow-up experiment on a second sample might collect gas only between 200 and 400 degrees. If the gas were again detected in that smaller interval, then the second hypothesis could be eliminated.

The use of planning techniques is further motivated by the need to contend with limited resources. In a remote planetary setting, the system might not always have enough time or soil sample for a complete second run. Therefore, the planner must reason about resources in order to choose its best experiment design strategy. For example, in the lab, a complete experiment involves heating the reference and sample up to 1200 degrees C at a rate of 10 degrees/minute, thus taking about two hours. If the system were to have only one hour in which to clarify ambiguities that occur at 1000 degrees, there would not be enough time for a complete second run. The planner could choose a strategy that uses a much faster heating rate to "skip" data collection in the first 900 degrees, stop and come to thermal equilibrium, then proceed at the desired 10 degrees/minute for data collection in the critical section. When there is not even enough time or soil for a partial second run, the planner might instead choose between strategies that seek to clarify the results by simply analyzing the data differently without requiring the hardware. In particular, it

356

could rerun the analysis in order to (1) look for masking effects between two decomposition events that occur in similar temperature ranges, (2) clarify matches under different prior probability assignments of mineral groupings in the Bayes net, or (3) look for possible alternative assignments of endotherm/exotherm features in the data due to "single valley versus two peak" ambiguities or other figure/background assignments. These actions also represent an experiment, even though no science hardware is involved. The knowledge representation used to model these strategies needs to be a high-level language so that scientists can develop such critical strategies themselves. Additionally, the language must support heuristic search techniques, and it must be procedurally expressive enough to represent the conditional and iterative control required for encoding arbitrarily complex strategies. One additional point: recall that the planner designs experiments based on the results of data analysis, which often contain competing hypotheses. However, those hypotheses may change at any time as unexpected exothermic, endothermic, or gas-release events are observed. Thus, the planner must operate in an uncertain and changing environment. In order to plan appropriate experiments in a changing world, the planner must be able to incorporate asynchronous sensor reports into its search process.

Finally, this application requires *real-time control* capabilities, which we define as the ability to take actions in bounded time. Our system must perform real-time control in order to react to unexpected thermal events, and to capture gas produced while heating the sample. Although the system cannot be certain in advance whether these events will occur or when, it must respond within seconds of their detection. If the planner cannot produce a plan within the available time, the controller must still operate with some intelligence. Thus, it must be able to generate experiments reactively by instantiating a design strategy according to heuristics that do not involve time consuming look-ahead search.

In summary, DTA-GC needs to combine a mineralogical expert system with integrated sensory perception, probabilistic data analysis, planning, and control. The next section describes our architecture and its components in terms of the software engineering and artificial intelligence techniques we have applied to these requirements.

THE DTA-GC SOFTWARE ARCHITECTURE

A simplified view of our software architecture is illustrated in Figure 1. It consists of three elements: a hardware relay, an analysis component, and a control component. The 'hardware relay' is responsible for sending effector commands to the hardware, and for receiving sensor reports from the hardware. The 'analysis' components provide the sensory perception capabilities that acquire information via hardware sensors, and the data analysis capabilities which reason about the sensory data. The 'control' components provide both the experiment planning and the real-time control capabilities. The software is written in LISP and C, and operates on a Sparc 2 Sun Workstation. The system accepts scientific goals and a time limit as input, includes both reactive and predictive control loops, and produces analytical results. The reactive control loop, indicated by the solid arrows in the figure, selects actions in bounded time by matching sensor readings against condition-action "reflex" rules. The predictive control loop, indicated by the dashed arrows, involves sending the analysis results to the experiment planner. The planner searches through a space of experiment design procedures either for a useful follow-up experiment, for modifications to the current experiment, or for analyzing the data differently. A successful search produces a new experiment in the form of condition-action rules to be passed to the experiment controller. We now briefly discuss each of the five software components in Figure 1, and the techniques we have used to address the requirements described in the previous section.

The job of the *hardware relay* is to receive sensor readings and transmit effector commands to the hardware. The DTA-GC hardware includes a programmable DTA oven, two GC columns and detectors, two pressure sensors, and four valves which control the gas flow between the DTA and the GC. The hardware relay currently receives nine real-time data streams from the hardware sensors, and it can transmit over 100 distinct effector commands to the hardware. All of the these instruments communicate with our Sparc 2 through a General Purpose Instrument Bus (GPIB), the IEEE-488 standard for byte serial, bit parallel interface. To facilitate this communication, we have developed a general LISP/GPIB interface written in C.

The job of the *sensory perception* component is to identify the qualitative features in the DT, pressure, and GC signals. We use a "Scale Space Filtering" technique originally developed by Witkin [5,6] for use in image processing domains. This technique detects peaks and valleys in a curve by convolving Gaussian filters of varying standard deviation with the input signal. As the size of the filter increases, the convolved signal becomes increasingly smoothed. Hence, the points of inflection that remain after applying the largest filters correspond to the most prominent variations in the input signal. Points of inflection at varying filter scales are then grouped into scale-space
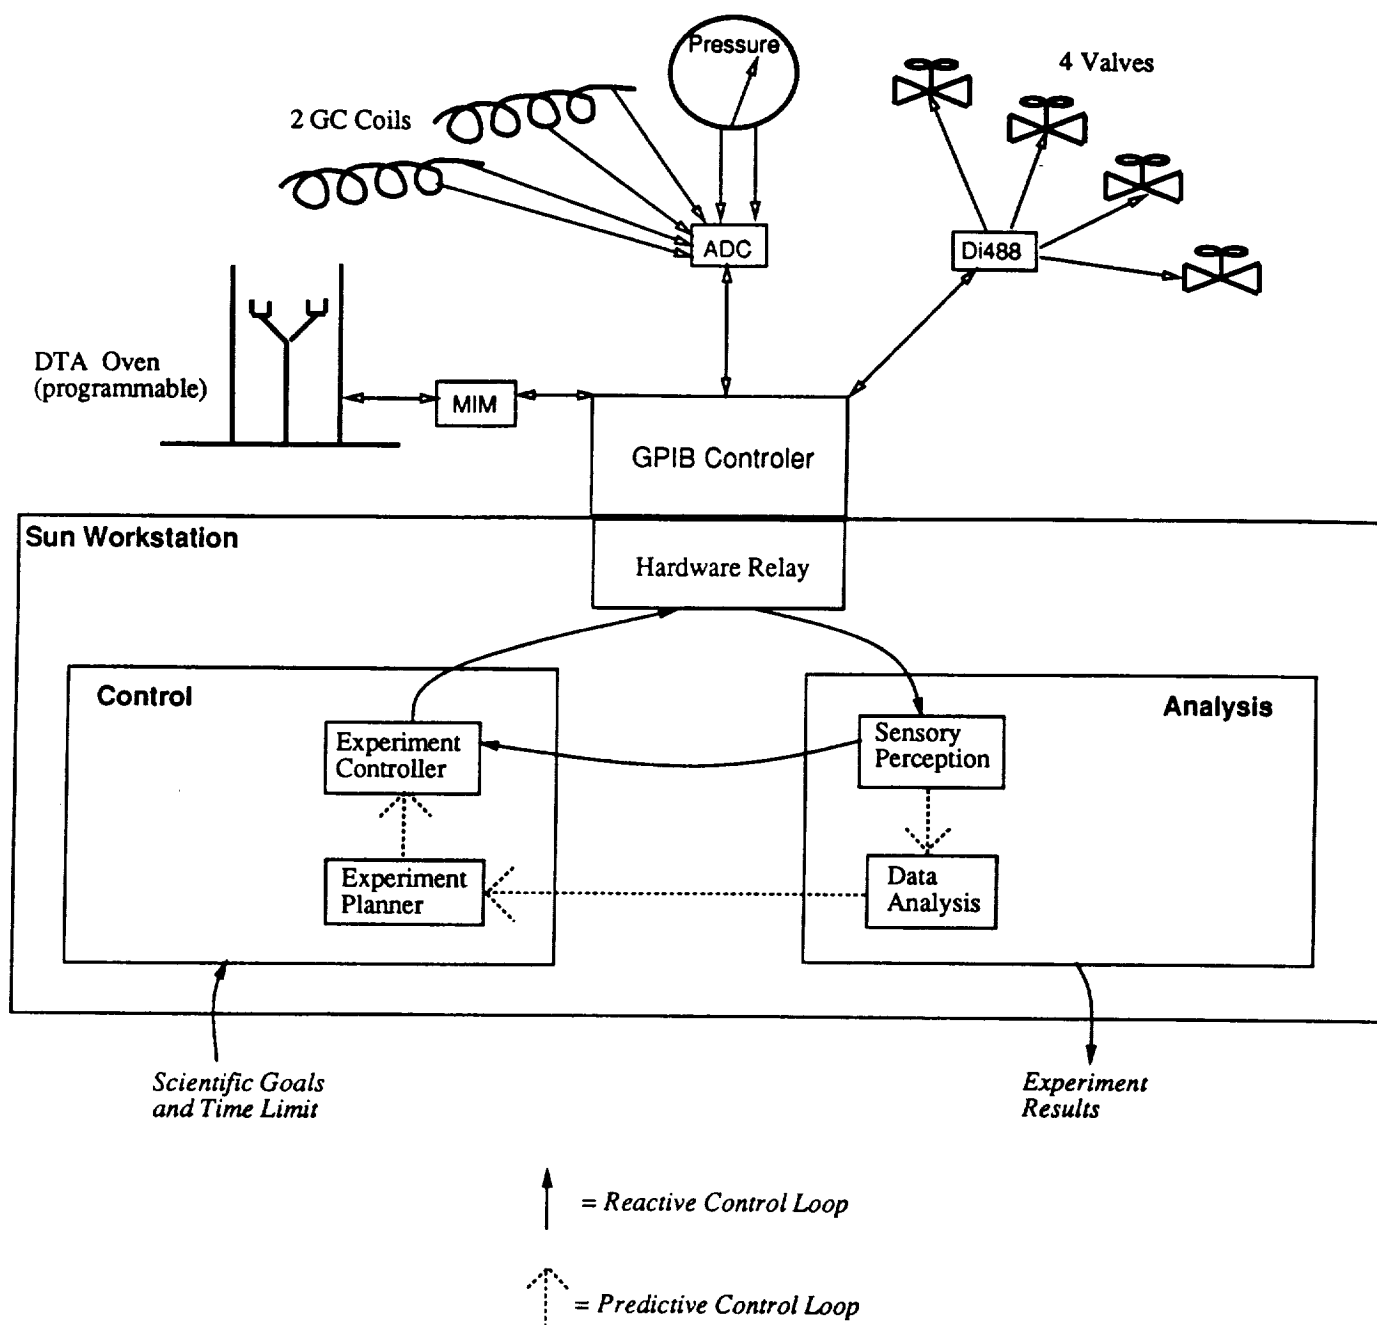
FIGURE 1. The DTA-GC Architecture

contours. The first derivative of the signal and its trend is used to determine whether a given contour group is a peak or a valley; it also aids in determining a degree of belief associated with the contour according to the probability that a feature observed really is a thermophysical event. This belief attribute helps to address the inherent perceptual uncertainty in our domain generated by signal/noise or figure/background discrimination issues, as well as the use of a sparse set of Gaussian filters. See [3] for a more complete description of our sensory perception component.

In the DTA-GC system, *data analysis* corresponds to generating hypotheses that postulate mineral combinations contained in the soil sample. We generate hypotheses through a two step method: Bayesian classification and heuristic search.

The classifier uses a Bayes tree to probabilistically match observations against events associated with known minerals in its library. The library contains knowledge of thermal and gas evolution events for over 30 classes of minerals including clays, carbonates, and salts. The classifier defines a Bayes tree for each mineral. Each child of a root mineral node defines a process node such as 'phase transition' or 'chemical reaction' which is produced by heating the minerals. Each of these process nodes has a terminal child node which corresponds to specific mineral decomposition events. These mineral event nodes test observations for membership in a class of endotherm, exotherm, or gas events that occur within a given temperature range. The classifier uses the probabilities generated during sensory perception to assign probabilities to the terminal nodes in the Bayes trees. Using the conditional probability links from mineral nodes to process nodes and from process nodes to mineral-event nodes, a standard Bayes tree propagation algorithm [4] is used to deduce the probabilities at all non-terminal nodes. The minerals are then ranked according to their associated degrees of belief. Here the belief attribute helps to address domain uncertainty by indicating the probability that the observation really is an instance of mineral decomposition event. Two issues arise with the output of the classifier. First, since the mineral events in our library may overlap in temperature range, the classifier may match a single observation to multiple mineral events, thus increasing the belief in both minerals based on the same piece of evidence. For example, both types of clays, montmorillonite and kaolinite, will match a single observed exotherm at 1000 °C. Second, each mineral model may only account for a subset of the total observations. Thus, another procedure is required to provide global explanations for the entire set of observations. In order to address these two issues, the classifier output is passed to an explainer that has the job of constructing systematic explanations for the set of observations as a whole.

The explainer is a general purpose inference engine that uses the local matches provided by the classifier to construct explanations or hypotheses for the set of observations as a whole. Each explanation contains a set of distinct mappings from each observation to a unique mineral decomposition event. This is done by reasoning about the matches provided by the classifier. The classifier can match a single observation to two different mineral events, or it can match a single mineral event to two different observations. Each of these cases produces disjunctive explanations. Thus, in our above example, one explanation will match the exotherm to the kaolinite decomposition event while another explanation matches it to the montmorillonite decomposition event. More disjunction is introduced to model cases where an observation is left unexplained. The explainer searches through this space of alternative explanations with the aid of a heuristic control function that combines multiple scoring dimensions. This heuristic is a form of Occam's Razor which prefers explanations that minimize the number of minerals used, the number of unmatched observations, and the number of unobserved events, while maximizing the combined probabilistic beliefs of the observations and the mineral events. The explainer currently uses two very simple hypothesis generation rules. The first rule defines a search space that matches each set of observations to a distinct set of classifications. The second rule completes the search space by allowing observations to remain unexplained. Even for simple examples, these rules can produce many distinct explanations. This ability to automatically and systematically construct and evaluate so many alternative, yet viable, explanations can provide a benefit to the human expert who may not be so rigorous in exploring alternatives. Our system includes closed loop control, which enables the system to design and perform its own experiments. The primary output of *data analysis* is a set of explanations, termed the 'result'.

The integration of *planning* and *control* components in this architecture is based on Drummond's Entropy Reduction Engine (ERE) [1,2]. We chose the ERE approach because it has the benefit that the controller operates independently from the planner so that real-time control is not dependent on the more expensive search behavior of the planner. Our system differs from ERE primarily in the style of search used by the planner component. Our planner generates a task decomposition space, whereas their planner generates a state-space search.

The experiment *controller* is a rule-based system that matches sensory enablement conditions to GPIB effector commands. Its job is to control the laboratory equipment in real-time according to a set of Experiment Control Rules (ECRs) that are either provided by the scientist or synthesized by the experiment planner. Our controller is based on

the "Reactor" and "Situated Control Rule (SCR)" elements of the ERE architecture. Under this approach, the controller operates in a perpetual sense-act cycle, executing rules that function as quick reflexes to provide the reactive control capabilities of the system. In the DTA-GC system, the controller must be able to react to unexpected thermal and gas events within seconds of their detection in order to properly analyze them.

Although many types of low-level commands can be sent to the DTA-GC instrument, we have defined three abstract operations that characterize our required experiment control behavior. These commands are 'record', 'skip', and 'sniff'. 'Record' causes the oven to heat up at the regular rate of 10 degrees/minute, during which time data is collected. 'Skip' causes the oven to heat up quickly, during which time data is not collected. 'Sniff' causes gas to be passed to the GCs for analysis, and then reconfigure the valve system for acquiring the next event.

The job of the experiment *planner* is to produce an experiment that clarifies the ambiguous results of a current or a previous run. A 'clear result' contains only one explanation that explains all observations; this rarely occurs. More often, the result contains multiple explanations that use different minerals to explain the same observation. Additionally, the result often contains observations that cannot be explained, and events that were expected but not observed. These cases represent three distinct forms of ambiguity. The planner searches through a task decomposition space to generate a set of Experiment Control Rules (ECRs) that might clarify the given ambiguities. First, the experiment planner selects which ambiguities to clarify using heuristics that consider ambiguity type and resource availability. The planner then chooses among hypotheses that postulate experimental, chemical, sensory, or modelling causes for each ambiguity. Next the planner selects a strategy for proving the hypotheses. General strategies include designing a second run that skips uninteresting temperature intervals, modifying the current run, or modifying the data analysis procedure alone. Lower-level strategies produce specific ECRs by selecting specific temperature intervals for 'skipping', 'recording', or 'sniffing'. Experiment plans that do not violate resource constraints are passed to the controller.

The planner is implemented in *Propel*, a general-purpose language that we have designed to be procedurally expressive enough to represent real-world procedures, while maintaining the benefits of heuristic search. *Propel* procedures allow subgoals and other choice points to be embedded within the conditional and iterative control constructs of a LISP-like language. These procedures are used to represent our experiment design strategies. The *Propel* interpreter generates disjunctive experiment plans by heuristically searching through the task-decomposition space that is defined by these strategies. Even though *Propel* was primarily designed for search, our system performs closed loop control by actually executing the experiments it designs, and analyzing the results.

To address our deadline management requirements, the planner must ensure that results are returned within the given time limit. The planner first estimates the available computation time by subtracting an initial estimate of required execution time from the given time limit. During simultaneous planning and execution, this estimate of computation time is adjusted according to the projected durations of developing plans. If a plan is found within the available computation time, then it is passed to the controller for execution. Otherwise, the controller could begin execution of a default experiment, or it could reactively instantiate an experiment design strategy. This is facilitated by the *Propel* strategy representation which can be instantiated in bounded time using predetermined heuristics. This type of action representation, which can be used by both the planner and the controller, allows for a tighter integration between planning and execution.

Since the planner must operate in a changing environment, we developed a mechanism called 'dynamic dependencies' that integrates asynchronous perception and analysis into the planner's search process. With our mechanism, the planner performs dependency analysis on the projection paths to identify external conditions on which its plans rely. The analysis component is informed about these plan assumptions so that it can notify the planner as soon as their status changes. The planner can then adjust its search control to favor plans that are based on new beliefs instead of continuing to develop plans that are based on obsolete assumptions. This technique allows DTA-GC to break the typical planning system assumption that the world does not change during the planning process. This "static world assumption" does not hold when the system is planning changes to the current experiment. Performing dependency analysis on our procedurally expressive experiment strategies is a difficult task.

STATUS AND TRANSFER PROSPECTS

Much time has been spent building the coupled DTA-GC instrument hardware itself and our LISP/GPIB interface to it. We have also focussed extensively on building up the mineral library used by the Bayes classifier by

running the DTA-GC on known samples; further work has concentrated on the sensory perception and explainer components, and on the development of *Propel*, the experiment method language, and its reactive dependency mechanism. The status of our system can be presented in terms of the three development levels described at the outset; work has progressed at all levels.

In April 1992, the first level of functionality for the reactive control loop of DTA-GC was successfully demonstrated and turned over to the scientists. Since then, the mineral library and classifier has been enhanced to include characterizations of over 30 classes of minerals, and the experiment control language has been greatly expanded. Currently, the system can execute default Experiment Control Rules which heat a sample slowly while monitoring the incoming DTA, GC, and pressure data. If the pressure in the oven reaches an assigned threshold, our system automatically reacts by evacuating the gas into the GC for analysis, and then it prepares for the next gas event. The operation of each component at this level, sensory perception, data analysis, and experiment controller, functions well. The sensory perception component is implemented, but we are exploring alternative methods, especially to identify the *onset* of DTA events rather than the *peak* of DTA events. Even though "peaks" are easier to identify, because DTA peak amplitudes may shift due to the amount of material present, we must focus on onset temperatures of events. This of course allows us to map our events directly to the traditional melting point or phase change literature on minerals. Our LISP/GPIB interface and hardware relay currently forwards all sensor data to signal processing, but it will soon perform data filtering so that only "significant" sensor data is relayed for evaluation. The data analysis component has been implemented and produces explanations, but the rules and heuristics it uses need to be tuned through additional knowledge engineering efforts. Capturing this knowledge is necessarily slow since no one has previously performed computer analysis of DTA data, let alone fusion of that data with asynchronous GC data. We intend to continue addressing issues of identifying, representing, and modelling thermophysical interactions between decomposing mineral combinations that tend to obscure data and hence confuse the classifier. Finally, the experiment controller has been implemented. We have demonstrated the ability to react to detected gas events to within one second. Since the controller is a rule-system, it has been straightforward to implement. However, the current default Experiment Control Rules have turned out to be rather brittle and as yet provide little coverage for unexpected events. Thus, we will be developing a more robust set of default ECRs through knowledge engineering efforts as we learn more about the system through our two collaborators' usage.

The second level primarily has involved the introduction of the experiment planner component and the development of better modelling and heuristic control techniques for data analysis. This level consists of *serial* predictive control. At this level, the planner can suggest follow-up runs that could produce better explanations. The experiment planner has been prototyped but needs further development. In particular, the *Propel* language for representing and searching through experiment strategies is implemented, but the knowledge engineering of these strategies has just begun. Since the DTA-GC is a new instrument, there are no existing strategies, and our experts will first have to develop them. At this level, we also introduce deadline limits into the problem. The deadline management mechanism has been partially designed but has not been completely implemented.

At the third level, termed *parallel* predictive control, all components operate in parallel, and this is the phase in which the dynamic dependency mechanism is required. The dynamic dependency mechanism is not fully implemented, but development has begun. We are currently converting the original ERE state-space search approach to work for *Propel*'s task-decomposition space.

We feel that our work on the DTA-GC system will yield several self-contained and general technological components that could transfer easily to other applications. Our general architecture, characterized by dual reactive and predictive control loops, can be applied to any scientific instrument that requires real-time control in conjunction with autonomous design of experiments that clarify previous results. Our LISP/GPIB interface is also a general tool that can be used by any LISP-based system to communicate with any of the more than 4000 instruments that use the GPIB protocol. On the data analysis side, the scale-space filtering, Bayesian classification, and development of disjunctive explanations are general techniques that could easily be instantiated for other applications. Our contribution has been primarily in the linking of these capabilities and producing code to affect unified data analysis. We have implemented these techniques as linked, general tools that could be instantiated for other applications. Such an architecture could also allow for more model-based analysis. The explainer itself is a standard production rule system that uses domain specific rules. The *Propel* language, which we use to represent and interpret the experiment design strategies, is specifically designed to be a general tool that can be transferred to many applications. *Propel* can be used by any application that requires control procedures to be represented in a heuristic search framework. *Propel* procedures can be reasoned with by a planner, and also executed directly by the controller. This allows the controller to execute

procedures rather than simple if-then rules. This is a feature that can be used by a variety of real-time applications where execution of a control procedure may have to begin before it has been completely instantiated by the planner.

CONCLUSION

We have described an architecture designed to autonomously control a new geochemistry instrument. The system functions as an instance of a general class of autonomous scientific instruments, that integrate sensory perception, data analysis, experiment planning, and experiment control. We have described how these components function and how they interact to provide autonomous control of the DTA-GC instrument. The architecture itself is now being used as we extend the system to other instruments. The system we have described represents a synergy between AI applications and AI techniques. The DTA-GC application has stimulated the development of techniques for the integration of perception, planning, and control, which in turn allow us to tackle new real-world applications that are even more ambitious.

ACKNOWLEDGEMENTS

REFERENCES

[1]   Bresina, J., and M. Drummond (1990) Integrating Planning and Reaction: A Preliminary Report. Proc. AAAI Spring Symposium Series. Stanford University, CA.

[2]   Drummond, M., and J. Bresina (1990) Anytime Synthetic Projection: Maximizing the Probability of Goal Satisfaction. Proc. 8th AAAI, Boston, MA.

[3]   Kulkarni, D., K. Kutulakos, and P. Robinson (1992) Data Analysis using Scale-Space Filtering and Bayesian Probabilistic Reasoning. Computers and Chemistry, vol.16, no.1.

[4]   Pearl, J. (1988) Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann Publ.

[5]   Witkin, A. (1983) Scale-Space Filtering. Proc. 8th Int. Joint Conf. Artificial Intelligence, Karlsruhe, Germany.

[6]   Witkin, A. (1987) Scale-Space Methods; in Encyclopedia of Artificial Intelligence. Edited by S. Shapiro. Wiley.
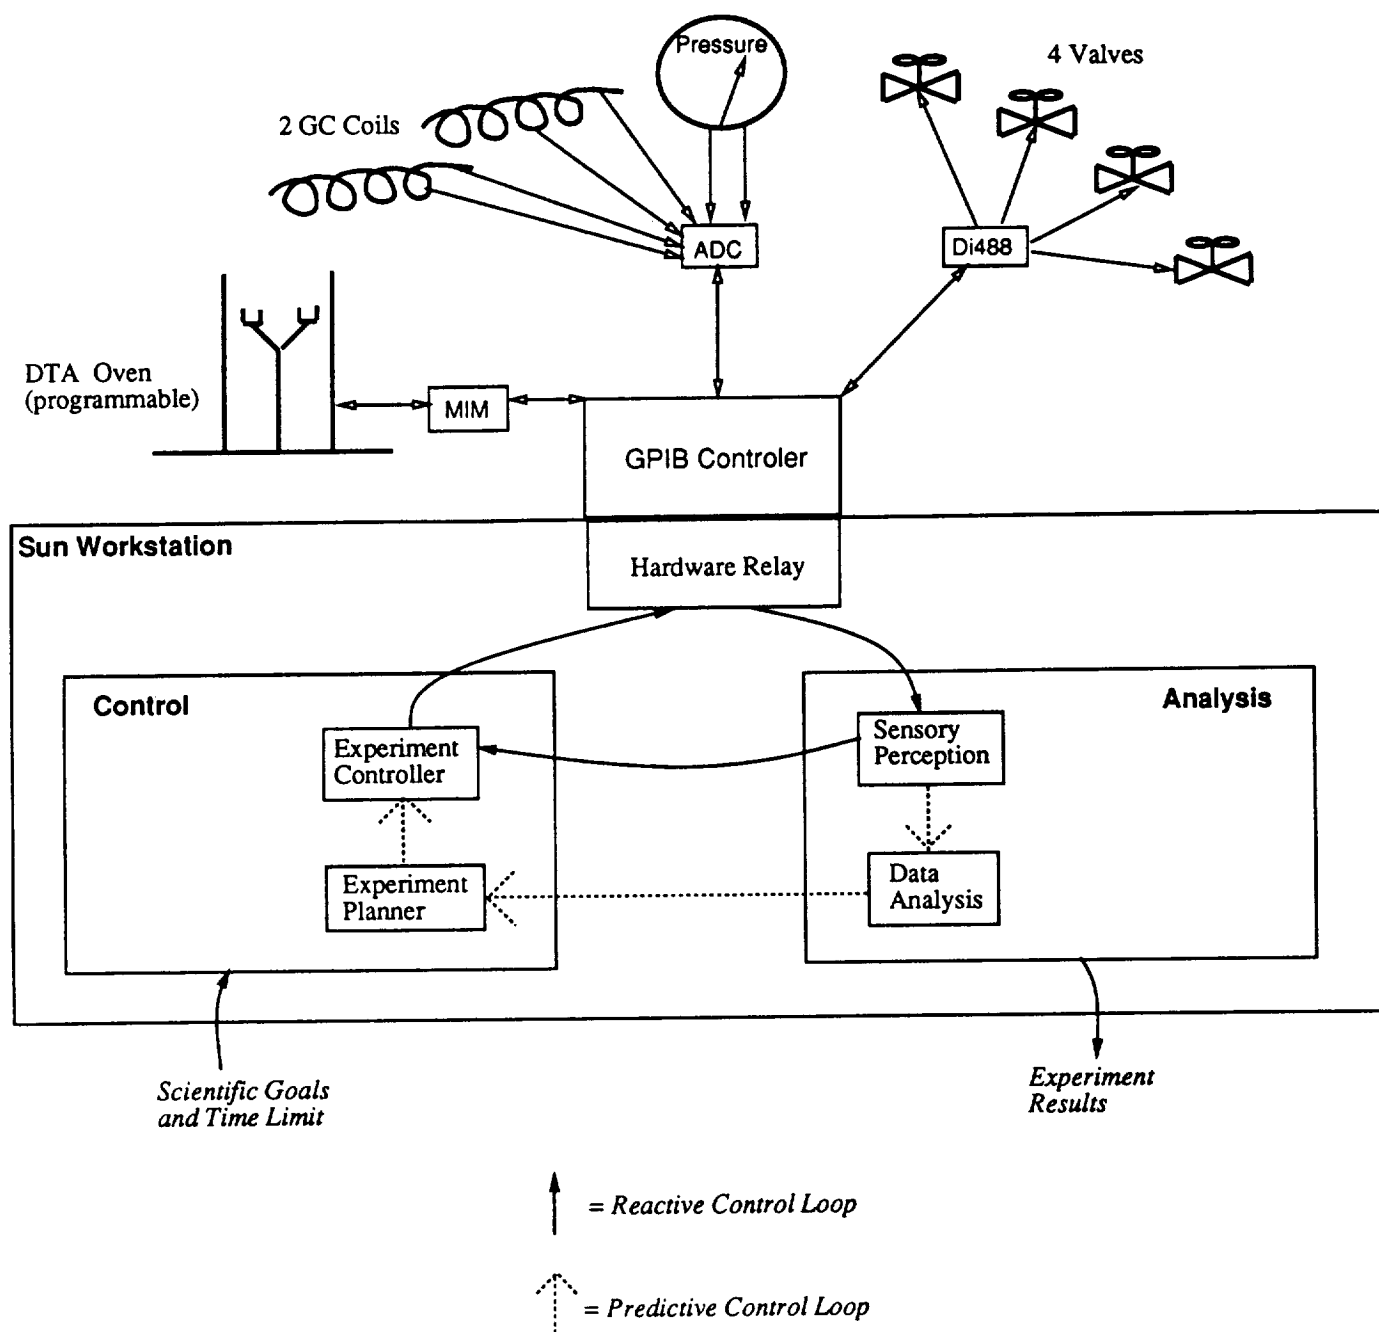
FIGURE 1. The DTA-GC Architecture

363