Abductive Model Refinement for Accelerator Control

Carl Stern, William Klein, George Luger, Mike Kroupa

Vista Control Systems, Inc.,
134 B Eastgate Drive, Los Alamos, New Mexico 87544
email stern@cs.unm.edu

Abstract. Many aspects of accelerator control require a complex procedure that includes planning, control, and re-evaluation of the process model. As control actions are performed new information is obtained from the system which allows the model to be adjusted. In many cases, observed errors in the model suggest certain control actions for gathering new information used for further refining the model. The process of comparing predicted with observed behavior to produce testable hypotheses for adjusting the predictive model is called *abductive model refinement.* This paper describes our ideas for applying abductive model refinement to beamline tuning tasks, including minimum steering through a set of quadrupole lenses and developing a waist at a specified location in a beamline.

TRADITIONAL AND KNOWLEDGE BASED CONTROL

The development of a theory and methodology for model refinement comprises one element of a larger research project: the development of an hierarchically structured architecture for distributed adaptive control. Our approach, described elsewhere (1), integrates knowledge based methods, including the explicit representation of control knowledge and control models, to support an adaptive capability. The view of control upon which this architecture is based is significantly different from the traditional view.

Traditional control theory views process and controller as strongly separated components of a control system. The process takes inputs from the environment and the controller and produces outputs which are then operated on by the controller to produce new inputs. Missing from this model of control is a view of the process and controller together as comprising a mutually dependent system with potentially time-variant behavior. Even in adaptive systems, where the controller adjusts its response function to minimize system error, the control system lacks the ability to recognize that a particular control method can no longer function successfully and to modify its internal representations and control algorithms accordingly. This is an acceptable model for control in simple stable systems where predictable low order functions can be minimized by traditional control methods (e.g., PID, FAM, etc.). *Figure 1a* illustrates the traditional control paradigm.

Stern, C.R., Klein, W.B., Luger, G.F., and Kroupa M. "Abductive model refinement for accelerator control" *Computational Accelerator Physics Conference 1996, American Institute of Physics*. Williamsburg, VA (1996).

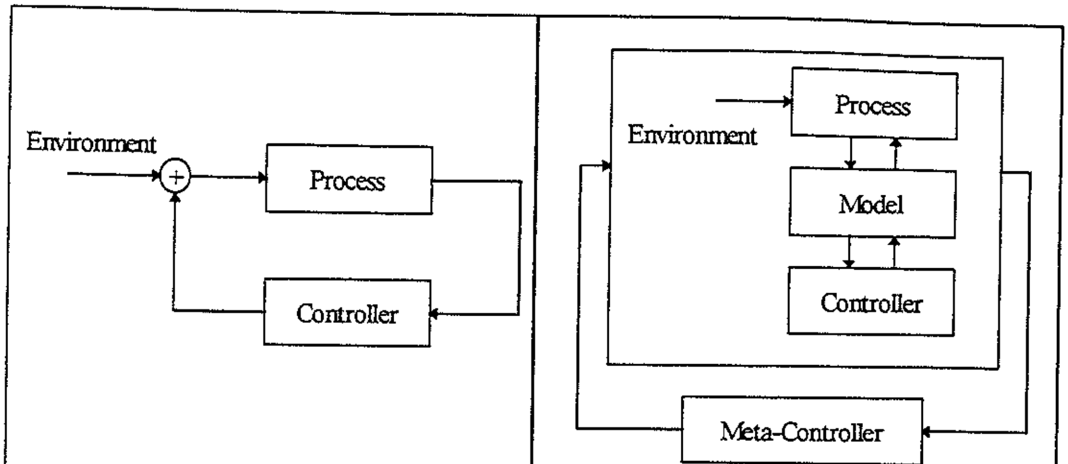

Figure 1a. **Figure 1b.**

The whole system view of control includes structures for modifying internal control processes based on information about both the process and the controller. An important aspect of this paradigm is the use meta-level control in conjunction with an explicit process model. The model is built to reflect both the state of the physical system and the relationships between the physical process, the control hardware, and the control software. The model provides a shared representational framework between process and controller and facilitates intelligent decision making and reasoning. *Figure 1b* illustrates the whole system view of control.

In the whole system view, it is also useful to employ confidence metrics encoding the reliability or effectiveness of control methods relative to observable states of the system. Knowledge based control systems can then select control algorithms by matching knowledge about the system's current state against a set of available control algorithms indexed by their estimated effectiveness under the current set of conditions.

From our viewpoint, model refinement is not a persistent process continually tracking the system's current state. Rather, it is triggered by the inability of the system to satisfy a goal using the current model. This may occur because of the failure of attempted control actions or due to lack of confidence in the effectiveness of potential actions. In many cases, it may be sufficient to update model data by making new system measurements. In other cases, however, failure to achieve a goal is directly correlated to a mismatch between the model and the process. Abductive model refinement is then used to diagnose and correct the deviation.

COGNITIVE FOUNDATIONS

Human experts in many domains exhibit the ability to effectively improve their model of a given environment through exploratory action. This involves the use of an initial rough model as a basis for planned testing and interaction followed by a process of evaluation. The attempt to explain discrepancies between expectations and observations generates a new understanding of the environment that results in a "refined model." In many cases, this process must be repeated a number of times before a satisfactory model of the environment is found. This cycle of exploration and model adaptation is called *abductive model refinement* because it depends on abductive reasoning, i.e. reasoning that attempts to *explain* the source of the differential between expectations and experience.

In our study of expert human performance in the area of particle accelerator control we have encountered precisely this pattern of problem solving. For example, we have found that accelerator physicists commissioning a new or reconfigured beam line will try to understand the behavior of the system by conducting a carefully designed sequence of experiments. Typically this involves an attempt to produce some standard set of beam conditions at specific locations. They employ an initial model of the accelerator beam line in conjunction with a software modeling code such as TRANSPORT to compute some configuration of magnet field strengths that is expected to produce the desired beam condition. This prediction then serves as the basis of an experiment to verify the expected beam condition based on the computed magnet settings. If the experimental finding fails to fall within an acceptable range of accuracy, this is in fact a very useful result. The physicist then analyzes the results in order to generate hypotheses regarding the source of the discrepancy between prediction and observation. This in general leads to a cycle of experimentation and explanatory hypotheses that eventually results in an improved beam line model.

REPRESENTATIONAL AND ALGORITHMIC ISSUES

We are currently developing a set of representational schemes and algorithms necessary to support a computational implementation of abductive model refinement in the context of accelerator control. Active study and analysis of the problem solving of expert accelerator physicists [1] has played a key role in shaping our current approach. We are also collaborating with these same physicists in the development of mathematical methods for diagnostic estimation. What follows is a brief account of our current insights and achievements.

Implementing a Model Refinement Algorithm

Abductive model refinement includes the following sequence of interrelated steps (although not necessarily in the given order):

1. Recognition of a discrepancy between model-based prediction and observation;
2. Generation of hypotheses for explaining this discrepancy;

[1] Most notably Andrew Jason at LANL and Xijie Wang at BNL.

3. Planning and execution of sequences of actions to test explanatory hypotheses;
4. Modification of a model based on a verified hypothesis.

Each of these steps presupposes a complex set of capabilities. Developing a computational implementation of the model refinement process has required a detailed understanding of patterns of reasoning involved in the second, third, and fourth steps in particular. This has entailed a number of related research goals:

- the application or modification of artificial intelligence planning methods to model the design of useful experiments;
- the application of abductive reasoning to generate plausible hypotheses that explain observed discrepancies;
- a form of model-based reasoning in which conjectured model revisions are tested for consistency with known data and causal relationships.

The high degree of complexity of the model refinement task necessitates the use of certain simplifying assumptions. We follow the traditional model based reasoning approach in assuming that certain dimensions of the causal connection structure of the physical system are fixed and known (2). These define a fixed causal framework for reasoning without which too many diagnostic possibilities and combinations would have to be considered.

Following the traditional representation used in model based reasoning, we consider a model of a physical system to be a set of elements together with a causal connection structure between elements. Each element has a set of inputs and outputs where each output is defined as a transfer function over the inputs together with the element's internal state. The elements together with their causal connections constitute a network as seen in *Figure 2*.

An issue of great concern for model refinement in complex systems such as particle accelerators arises from the dependencies between beliefs. An inference regarding some property of an element, such as a remnant field or misalignment in a magnet, often depends upon a good deal of assumed knowledge about other elements. If these assumptions are incorrect or of limited accuracy, the inference itself is tainted. Since incremental model refinement involves the construction of chains of inferences, it is necessary to keep track of the evidence for inferred beliefs as well as the degree of certainty and accuracy that such evidence warrants. For this reason we cache inferences, recording the dependencies between assumptions and conclusions. We use a network representation similar to that used by a Justification Based Truth Maintenance Systems (JTMS) architecture (3) to organize cached inferences. This allows us to immediately determine the set of beliefs to which a revised belief is linked and reevaluate the status of those dependent beliefs accordingly.

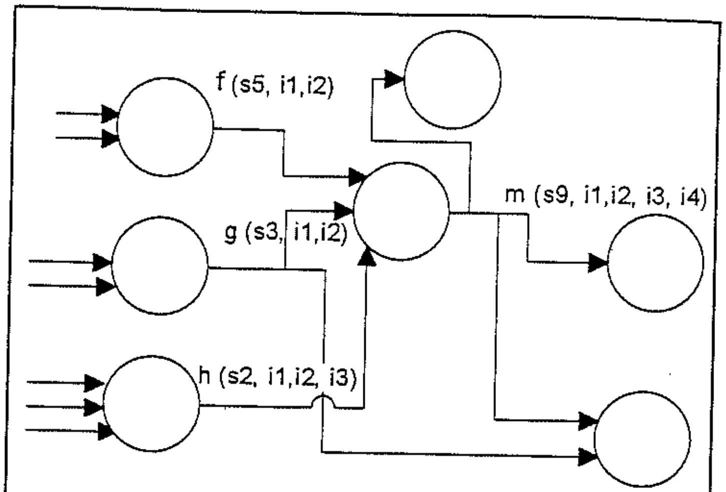

Figure 2. A causal network.

Two key issues in model revision present significant challenges. The first, the problem of cyclic dependencies in beliefs, poses a problem from the algorithmic point of view. Algorithms that update the numerical probabilities or confidence levels associated with beliefs in a belief network behave erratically in the presence of cycles, e.g., when belief A supports belief B, B supports C, and C supports A. When A supports C and C also indirectly supports A this can lead to an incorrect calculation of the evidential weights supporting both A and C. Unfortunately, belief cycles tend to permeate reasoning about models.

A second more fundamental issue is related to the problem of multiple source errors. A recognized discrepancy between prediction and observation can derive from a combination of errors in a model, or even worse, a sequence of causal interactions that magnifies minor inaccuracies into significant errors. Finding explanations for multiple source errors is not only difficult for humans; it also poses a serious computational challenge for most abduction algorithms.

We do not believe that there is a general solution to the second problem, but rather heuristics that are effective in finding explanations in many cases. Our current approach is to focus on knowledge engineering with expert accelerator physicists in order to discover such heuristics.

One heuristic that we are currently exploring is to refine the beamline model by identifying "islands" in the beamline that can be measured and studied in relative isolation, i.e., in a way that is minimally dependent on assumptions about the rest of the beamline. We have a identified a few experimental methods that support such independent calibration techniques. Once islands of high accuracy and confidence are constructed, the generally strategy is to cautiously extend them, verifying them by generating predictions through the use of the usual beam propagation and fitting algorithms and then testing those predictions against actual beam measurements.

Consider a simple example of a method that implements the strategy described above, i.e., to begin from and gradually extend islands in the model that have been calibrated with a high degree of accuracy. In this example we begin by measuring the beam itself using a method that assumes nothing about the rest of the model. We start with a location in the beamline that contains two screens with a drift space of length L between them. This is illustrated in the right half of *Figure 3*. Working in one dimension, we then produce a minimum spot size at the second screen, measuring both *Xs* and *Xs'* under this condition.

Using the matrix for a drift space and setting $\delta\sigma 11'/\delta\sigma 12 = 0$ corresponding to a minimum spot size at the second screen, we can determine a number of useful parameters: $\sigma 11$ at the first screen as a linear function of $\sigma 12$, beam emittance (ε), as well as the position of the beam waist. Varying the current to quads α and β (left side of *Figure 3*) and applying curve fitting to measured beam sizes on the first screen allows us to calculate actual field strengths as a function of current as well as remnant fields in α and β.

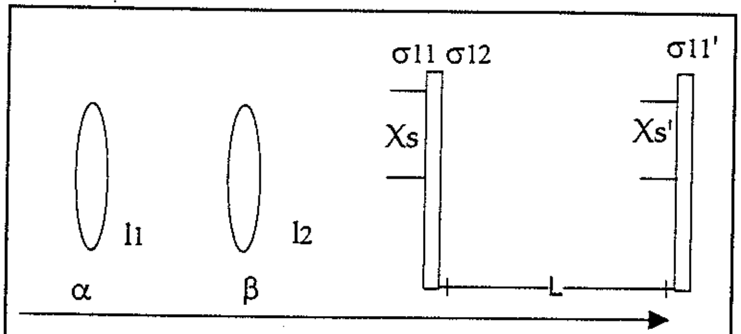

Figure 3

CONCLUSION

We are currently planning to test our ideas in abductive model refinement through a series of tests at Brookhaven National Laboratory's Accelerator Test Facility. There is reason to believe that the most difficult of these tests, production of a specified waist condition in an undulator cavity of a Free Electron Laser, can only be accomplished by refining the current beamline model. Future work involves the extension and generalization of our current system into a general architecture for accelerator control.

REFERENCES

1. Klein, W.B., Westervelt, R. T., & Luger, G.F., *Journal of Intelligent & Fuzzy Systems*, in press, (1996).
2. Davis, R., *Artificial Intelligence*, 24, 99-94 (1983).
3. Doyle, J., *Artificial Intelligence*, 12, 231-272 (1979).