CONF-891094--13

UCRL- 102060
PREPRINT

UCRL--102060

DE90 006775

MAESTRO - A Model and Expert System Tuning Resource for Operators

D. L. Lager, H. R. Brand, W. J. Maurer
F. E. Coffield and F. Chambers

This paper was prepared for submittal to
The International Conference on
Accelerator & Large Experimental
Physics Control Systems
Vancouver, British Columbia
October 30 - November 3, 1989

Lawrence Livermore National Laboratory

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

MAESTRO—A Model And Expert System Tuning Resource for Operators[*]

Darrel L. LAGER, Hal R. BRAND, William J. MAURER,
Fred COFFIELD, and Frank CHAMBERS

Lawrence Livermore National Laboratory, Livermore, California, USA

ABSTRACT

We have developed MAESTRO, a Model And Expert System Tuning Resource for Operators. It provides a unified software environment for optimizing the performance of large, complex machines, in particular the Advanced Test Accelerator and Experimental Test Accelerator at Lawrence Livermore National Laboratory. The system incorporates three approaches to tuning:

- A mouse-based manual interface to select and control magnets and to view displays of machine performance.
- An automation based on "cloning the operator" by implementing the strategies and reasoning used by the operator.
- An automation based on a simulator model which, when accurately matched to the machine, allows downloading of optimal sets of parameters and permits diagnosing errors in the beamline.

The latter two approaches are based on the Artificial Intelligence technique known as Expert Systems.

---

[*] This work was performed jointly under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract W-7405-ENG-48, for the Strategic Defense Initiative Organization and the U. S. Army Strategic Defense Command in support of SDIO/SDC-ATC MIPR No. W43-GBL-D-5007.

I. INTRODUCTION

There is a class of large, complicated machines prevalent in physics research experiments that require expert human control to obtain a satisfactory level of performance. Members of this class include the Advanced Test Accelerator (ATA) and Experimental Test Accelerator (ETA) particle-beam accelerators, the Stanford Linear Accelerator (SLAC), the Superconducting Super Collider (SSC), and laser chains for Laser Isotope Separation (LIS) and Inertial Confinement Fusion (ICF). Many such machines may have several levels of control, with operators responsible for overall global control decisions and closed-loop feedback controllers responsible for relatively small sections of the machine. Over a period of time the operators develop strategies based on rules-of-thumb (heuristics) for *optimizing the performance of the machine.*

As the machines become more complicated and as the performance criteria become more stringent there is a need to provide the operators with an on-line system model to aid in tuning and in failure diagnosis. MAESTRO (Model And Expert System Tuning Resource for Operators) is a software environment that combines physics models of the system and operator heuristics. The MAESTRO acronym was chosen to emphasize the conductor metaphor for unifying and coordinating the activities of control, diagnostics, modelling, and post-run analysis. Traditionally these activities are separate; often different people within different groups in the organization are responsible for each of these aspects. The principal advantage of using one consistent environment is rapid turn-around. Being able to observe the simulation while the machine is running allows the operator to make much more informed control decisions. Also, discrepancies between the model and the machine are readily apparent so that physicists can quickly gain insight into the phenomena that are occurring.

A. Intelligent Assistant

MAESTRO functions as an intelligent assistant for tuning, much as its predecessor, ATES (Accelerator Tuning Expert System) [Lager]. Displays permit the operator to see where MAESTRO is in the tuning processes and allow him to interrupt it so he may do manual tuning. There is a machine interrogation and control interface (MICI) for manually triggering data acquisition and for changing magnet settings by using a mouse. There are menu-selectable tools for characterizing the machine by sweeping focusing and steering magnets through a range of settings and by graphing the resulting beam positions. Finally, MAESTRO provides a very flexible runtime and development environment for rapid prototyping, fast debugging, and online construction of diagnostic codes.

B. Tuning

Tuning is the art and science of optimizing the performance of a machine for a desired experiment. The art is in knowing that a certain bump on a waveform indicates that a particular power supply is off line, or that having a certain waveshape at the end of the injector gives poorer performance through the accelerator but much better performance in the wiggler.

As a science, tuning may be viewed as a multivariate nonlinear optimization of multiple cost functions from multiple measurements. Typical systems involve tens to hundreds of interacting control variables, five to ten cost functions, and three or four types of measurements. Some of these measurements are readily available for computer manipulation (for example, digitized oscilloscope traces) and others are not (for example, audio information from microphones placed near the machine). Systems are often under-diagnosed so that a single measurement is affected by several control parameters. Multiple data sets must then be taken to decouple the effects of the individual controls.

We have chosen the domain of tuning particle beam accelerators to implement the concepts embodied in MAESTRO as an example of a complex system that can benefit from this approach.

II. TUNING APPROACHES

MAESTRO blends three distinct tuning approaches to achieve better performance than possible from any one approach alone. A trade-off among these approaches can be made as the machine is modified or grows more complex, as the operators learn more of its idiosyncrasies, or as more is understood about its physics.

The first approach is "cloning the operator." We encode as faithfully as possible the strategies and reasoning followed by the operator. This was the dominant approach taken for tuning ATA since it was not possible within the allotted time to develop an accurate model for the beamline. A disadvantage of using only this approach is that operators may not be able to develop successful strategies for the far more complex machines being designed, for example the SSC.

The second approach is model-based tuning, where an on-line simulator of the system is coupled with the system's control and diagnostics subsystems. This approach has the advantage that a complete set of tuning parameters can be determined using the model and then simply downloaded onto the real machine at a substantial time savings. The disadvantage is that it may not be economically feasible to develop a sufficiently accurate simulator.

The third approach is to tune manually, but to provide the operator with more powerful tools for tuning the machine. This takes the form of displays derived from the raw data the operator is normally presented, and different interfaces making it easier to control the machine. For most applications the manual approach is needed because the operator will almost always have more knowledge of the machine than is economically feasible to encode into a computer program. He can take into account information that may not be readily accessible to the computer, such as the "sound" the machine makes when it is not quite running right. He may also be able to instantly diagnose a failure because for example he remembers what happened when that same situation arose on a machine he was tuning 20 years previous.

The goal is to achieve a blend of these three approaches that minimizes the time required to tune and maximizes the time available for performing physics experiments. Each of these approaches is discussed below.

A. Cloning the Operator

This was the dominant approach taken for tuning ATA since it was not possible within the allotted time to develop an accurate model for the beamline. The fundamental goal followed by the expert was to tune for a zero; that is, tune to put the beam on the center of the pipe. Given that goal, it was unnecessary to be concerned with rotation, tilt, and calibration-factor errors in the beamline components. He only needed to be qualitatively concerned with coupling errors between the vertical and horizontal steering magnets.

We concentrated our efforts on tuning the transport section of ATA. The tuning process was to follow a global strategy by selectively applying local strategies involving two to four components in the section. The local strategies were implemented using the Monitored Decision Script (MDS) representation we developed for ATES [Lager]. The script part of the MDS represented the step-by-step procedure followed by the operator, the decision part represented the choices made when there was anomalous behavior by the machine, and the monitor part emulated the behavior on the part of the operator when he would periodically "pop his head up" and check the overall status of the machine, especially those portions affecting the section being tuned.

Our initial approach for implementing the local strategies was to decompose the problem into two simpler ones, namely to center the beam in X, and then center in Y. We had the capability to account for the relative rotations of the components and their various calibration factors. Unfortunately, when we attempted tuning ATA in February 1989, there was coupling between the steering magnets that we could not accommodate, so we modified the local strategies to more closely follow the operator's procedure. There is now an initial phase where a pair of steerers are chosen, one nominally steering vertically, the other horizontally, and an experiment is performed to diagnose which most affects the X position and which affects Y. These steerers are used to center the beam by iteratively halving the error in X, then in Y, using a uniplex adaptive optimization algorithm [King]. With this approach, we were able to successfully center the beam.

The global tuning strategy is driven by the overriding constraint, "don't put the beam into the wall." As soon as the machine is powered up the operator checks the beginning of each major section to see if the beam is there. If not, he does coarse steering in the previous section to get the beam through it. Once he has the beam to the tuning dump, he goes back to the beginning of the transport section and meticulously centers the beam while monitoring the current reaching the tuning dump. For the meticulous centering, he uses more sophisticated local strategies based on pairs of steering magnets, where the first member of the pair is used to position the beam at the center of the second member. The second is then used to remove any angular offset in the beam.

We developed an inference engine [Charniak] to implement the global strategy. The philosophy is to choose and execute an appropriate local strategy, based on the present state of the machine. The inference engine operates by repeatedly "matching," "selecting," and "executing" MDSs until the desired machine state has been achieved. Since each MDS has a pre-condition field defining the state of the machine when it is valid to be executed, and a post-condition field defining the state of the machine if it successfully executes, the "match" part of the cycle consists of scanning all the MDSs to find which ones match the present state of the machine. The "select" part of the cycle chooses one of the MDSs to execute if there is more than one match. Then the selected MDS is "executed." The actions of the MDS change the state of the machine, and the cycle is repeated.

B. Model-Based Tuning

Model-based tuning promises to put more science into the art of tuning leading to a more rigorous understanding of the machines. The model can enable operators to "see" the effects of their tuning in the regions between sensors. Given sufficiently accurate models, it should be possible to determine a set of parameters that will change the present state of the machine to a desired (i.e., tuned) state in a single step. With less accurate models, iteration will be necessary, but will typically converge rapidly. Experiments can be performed off-line using the model to determine, for example, the effects of adding new components to the beamline at far less expense than involving the real machine.

MAESTRO directly couples the beamline model to the machine representation in order to make all the present magnet fields, locations, calibration factors, rotation angles, and tilt angles directly available to the model. A nonlinear least-squares algorithm performs least-squares fits to data sets of beam positions. Fit parameters include magnet field strengths and component displacements, rotations, and tilts. The art (and artificial intelligence) is in choosing the proper subset of parameters to be fit and in choosing between alternative fits when more than one explanation of the data is possible.

Model-based tuning requires two distinct phases. The first is "commissioning" [Lee], where the simulator is matched with the real machine. This phase forces proper bookkeeping because it requires all the components to be properly characterized. Any rotations, tilts, and miscalibrations must be eliminated from the machine or incorporated into the model during this phase so that the simulator and real machine produce the same beam trajectory. We have developed a variety of tools for commissioning. Steering sweeps measure the beam position as first the horizontal steerer is swept through a range of values, then the vertical steerer is swept. This produces the cross-shaped pattern in Fig. 1. Focus sweeps measure the beam position as the field in a focusing magnet is varied over a range. This produces the spiral-shaped pattern in Fig. 2.

We have begun the commissioning on the smallest subset of ETA possible, one magnet followed immediately by a beam-bug in the injector section of the machine. Once these components have been characterized, we will commission the remaining injector magnets and then proceed down the accelerator, beginning with the first 10-cell set. Proceeding in this fashion insures that the smallest number of noncommissioned devices influence a measured value. A typical commissioning fit for a steering sweep is shown in Fig. 3

where the small squares are measured beam positions and the large squares are model predictions.

The second or "operational" phase is actual accelerator tuning and diagnosis. With the simulation model and accelerator in relative agreement, tuning can be done by using the nonlinear parameter estimator to select the optimum tuning parameters for the injected beam after its launch conditions have been estimated. Obviously, a very large number of tuning parameters are available. Initially, an experienced operator will select a subset of these for tuning, and will iterate the process with different subsets. Eventually, the operator's knowledge and experience will be encoded into an expert system that will select the subset of tuning parameters and will decide how much iteration is necessary

When the simulation model and the accelerator go out of relative agreement, the model can be used to find those regions within the accelerator where there is still agreement. For a single component failure, there will be two regions of agreement roughly surrounding the failed component. From this information, a set of suspect components can be constructed. With the aid of an expert operator (or eventually, an expert system), a number of possible failures can be proposed. The nonlinear parameter estimator is then used to estimate both the magnitude of the proposed error and the improvement in the model's predictive ability, given the proposed error and the data currently available from the accelerator sensors. With this information, many proposed errors can be rejected because either the magnitude of the error is beyond reasonable limits, and/or the improvement in the model by the addition of the error is too small. In the ideal case, only one reasonable error remains. In other cases, one is normally left with only a few possibilities. Beam redirection experiments can then be performed to further isolate the accelerator-model discrepancy. Selection of these experiments will initially be done by experienced operators but will eventually be included within the expert system's capabilities. When the problem is finally found, then the accelerator can be fixed and/or the model can be updated with the (fit) error information, and tuning can proceed as before.

C. Manual Interface

The manual interface for MAESTRO is shown in the screen dump in Fig. 4. The interface consists of several windows that become visible on the screen as necessary. The machine interrogation and control interface (MICI) window is the main one for interacting with the system. It consists of two panes, with the upper one for displaying the output from the simulator, showing the horizontal beam position vs location along the beamline as a thick line and the vertical position as a thin line. The lower pane is a set of icons depicting the components in the beamline. The locations of the icons and their shape are derived from the information describing the beamline in the MAESTRO knowledge base. Components added to the beamline are automatically included in the MICI display, once the information has been added to the knowledge base.

The operator controls the fields in the magnets by clicking a mouse button as a cursor is positioned over an icon. Clicking over a vertical steering magnet icon, for example, causes power supply control windows to become visible below the MICI window, as in Fig. 4. Clicking the mouse as the cursor is positioned in those windows causes the current in the appropriate power supply to be incremented or decremented.

The scope display window is used to control and display data acquisition. Clicking the mouse over the appropriate label in the window causes data to be acquired and displayed as a set of three traces in the window, giving the time histories of the beam current, *X* position, and *Y* position. The actual location of the beam within the pipe is derived from the trace data, after accounting for sensor misalignment and is displayed in the circular "bulls-eye" displays in the upper part of the scope display window.

Other windows display additional derived data. A position display window shows a value vs position down the beamline. A "bug-walk," for example, makes position displays that show peak current, *X* position, and *Y* position at the locations of the beam-bug position monitors along the beamline. Similarly, a "magnet-walk" shows the fields at the centers of the magnets vs their location in the beamline. Three magnet walks appear in Fig. 4, showing the solenoidal fields (labeled PD-B), the horizontal steering fields (labeled PD-H), and the vertical steering fields (labeled PD-V) vs distance down beamline.

There are also windows for displaying historical data. Clicking the mouse over an icon for a beam bug causes a "shot history" window to appear. By clicking the mouse over buttons on the window, the operator can view all the oscilloscope traces acquired for that bug for that day's run.

III. FUTURE

We are planning to extend the capabilities of the shot history mechanism to improve its flexibility for manipulating not only past shot data but also past machine configuration data. We want to have the capability to redo the signal processing with different control parameters, for example. We also want to make queries, such as: "during the last three months what was the highest current magnitude measured when the machine had the long collimator installed?" The approach we are taking is to develop an unstructured database based on the Artificial Intelligence representation scheme knows as a *semantic network*.

We are also developing the capability to acquire image data from cameras that observe the beam as it strikes foils inserted into the beamline. We will apply image-understanding techniques to determine the position and focus of the beam for making tuning decisions.

IV. SUMMARY

The MAESTRO software environment was developed to function as an intelligent assistant to an operator in tuning complex systems such as particle-beam accelerators. It incorporates three approaches to tuning. The "cloning the operator" approach uses an inference engine and the MDS representation to encode the strategies and reasoning followed by the operator. The model-based approach makes use of a beamline simulator and a nonlinear least-squares parameter estimator first to match the model with the machine and then to determine optimum tuning parameters after computing the beam launch conditions. The third approach lets the operator manually perform tuning and provides him with displays that easily let him determine the machine status. Finally, a history mechanism lets the operator view past data to compare the present tune with ones previously obtained.

ACKNOWLEDGEMENTS

The authors gratefully acknowledge the contributions of Doyle Rogers and John Clark, the tuning experts. Their ability to make order out of chaos is truly amazing.

REFERENCES

Charniak, Eugene; Riesbeck, Christopher K.; and McDermott, Drew V., *Artificial Intelligence Programming*, Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1980.

King, Paul G. and Demming, Stanley N., "UNIPLEX: Single Factor Optimization of Response in the Presence of Error," *Analytical Chemistry*, 46, 1476-1481, (1974).

Lager, Darrel L.; Brand, Hal R.; Maurer, William J.; Searfus, Robert M.; and Hernandez, Jose E.; "An Expert System for Tuning Particle-Beam Accelerators," Proceedings of SPIE- The International Society for Optical Engineering, Vol. 1095, Applications of Artificial Intelligence VII, 28-30 March, 1989.

Lee, Martin and Clearwater, Scott, "GOLD: Integration of Model-Based Control Systems with Artificial Intelligence and Workstations," Workshop on Model-Based Accelerator Controls, Brookhaven National Laboratory, Upton, New York, 17 August 1987.

Figure Captions

Figure 1. Variation of X and Y beam position as the vertical and horizontal steering magnets are swept through a range of values.

Figure 2. Variation of X and Y beam position as the solenoidal field is varied through a range of values.

Figure 3. Comparison of fit between model and measured data for a steering sweep. The small squares indicate measured values, the large squares indicate model predictions.

Figure 4. Screen dump of the windows associated with the manual interface to MAESTRO. These windows include: (a) simulated beam position vs distance down beamline with the X position shown by the thin line and Y position by the thick line; (b) icons depicting the components according to their positions in the beamline; (c) power supply control window; (d) scope display showing the I, X, and Y waveforms from the selected beam bugs; and (e) field strength vs distance down the beamline for the solenoids, horizontal steerers, and vertical steerers.

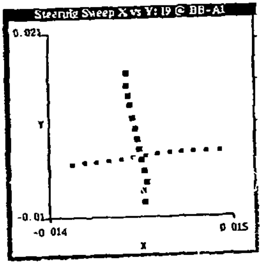

Figure 1.

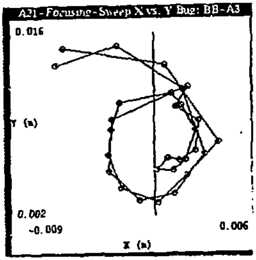

Figure 2.

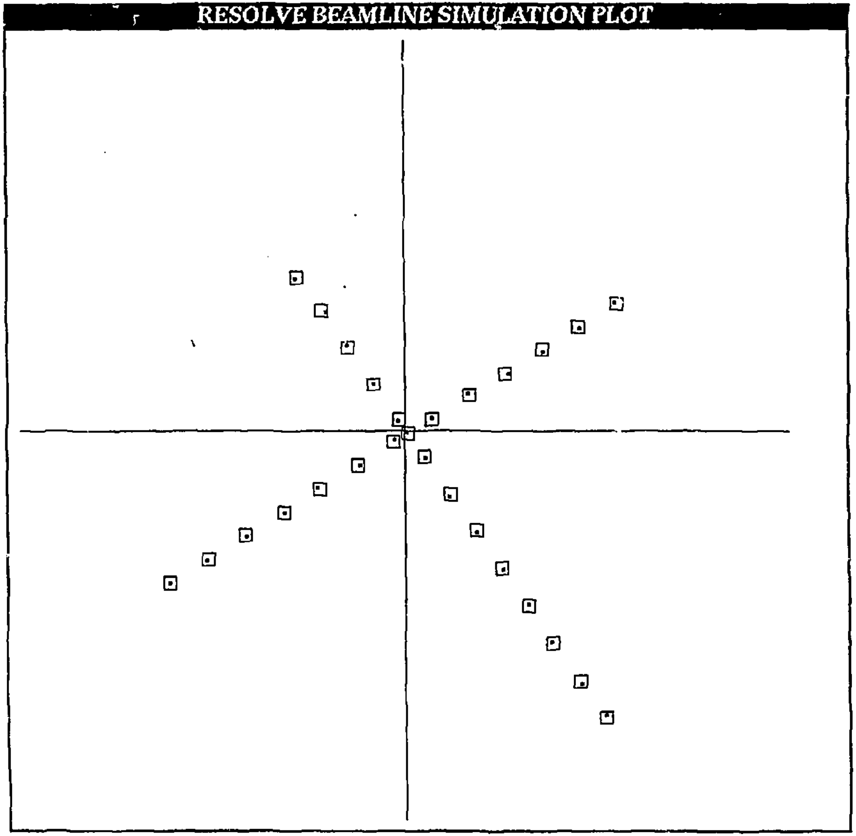

Figure 3.

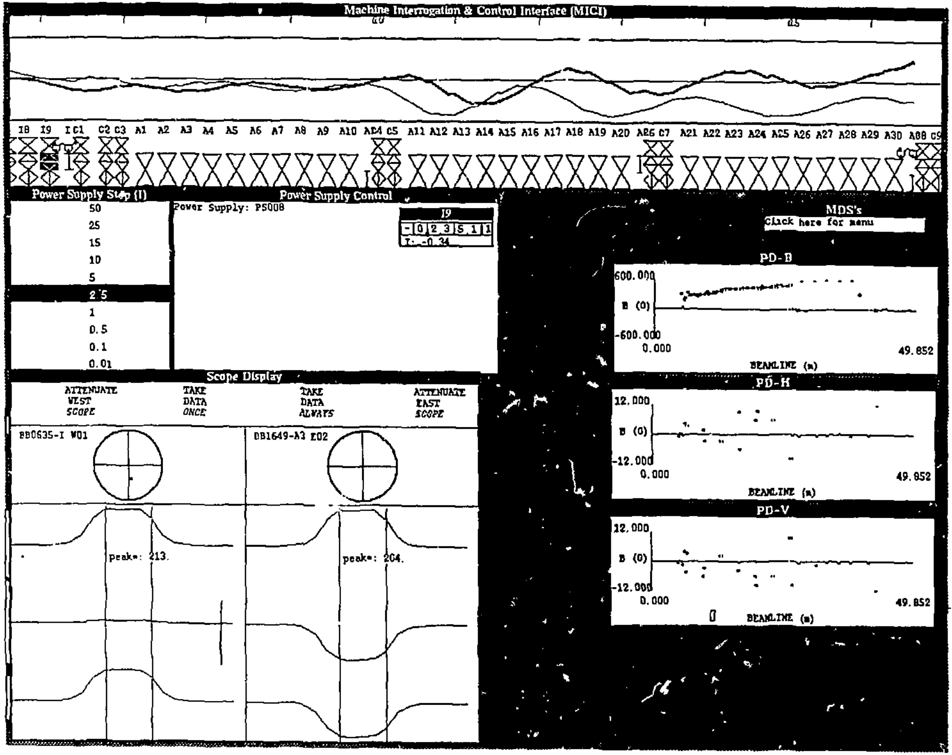

Power Supply: PS008

Figure 4.