CHAPTER 6. ACCELERATOR CONTROL PERFORMANCE

6.1 Performance Criteria

The distributed expert system was applied to the problem of controlling the terminal charging subsystem of the FN Tandem Accelerator (see Chapter 4). This is a desirable test for several reasons:

a) It is presently solved heuristically by a human operator.

b) It is susceptible to an analytical solution (see Figure 4-4).

c) It involves a complex interaction between control variables (charge and corona position).

d) It requires cooperation between processors.

e) It can be compared with similar previous work (Lind and Poehlman 1992).

Terminal charging is also a useful test because it has four clearly measurable performance criteria:

a) Response time to a requested change in terminal voltage.

b) Amount of overshoot of desired terminal voltage.

c) Coarse regulation of desired terminal voltage.¹

d) Maximum rate of change of terminal voltage.

The rate of change statistic is important because increasing the voltage too rapidly may cause sparking within the accelerator.

6.2 Manual Operation

Figure 6-1 is the record of a typical manually operated accelerator run (16 April 1996, Winston Williams, operator). The terminal was charged to 6.4 MV for a period of roughly 30 minutes for "conditioning;" then it was raised to the operating voltage of 7.6 MV. The expanded view in Figure 6-2 shows that it took almost 270 seconds to raise the voltage by 1.2 MV. First the charging voltage was raised, and the terminal reached equilibrium near 7.2 MV (at t=2500 seconds). Then the corona points

^{1.} For fine regulation, a completely separate "terminal stabilizer" subsystem is activated by the operator.



FN960416.001 MANUAL OPERATION (W. WIL

p.58







Figure 6-2. Manual Operation, Enlarged

charging voltage were needed to bring the terminal voltage back to 7.6 MV; the peak voltage was slightly over 7.9 MV. The maximum rate of change was 160 kV/sec at t=665 seconds.

The first automatic control trials made no attempt to adjust corona position. Only the terminal charging was controlled, to achieve a desired terminal voltage.

Terminal Voltage	Low	ОК	High
Action	Increase charge	Do nothing	Reduce charge

Figure 6-3. Control Strategy #1 Heuristic

represented in six rules.² Three are intermediate conclusions about the current terminal voltage:

ChargeUp is true if the terminal voltage is too low.

ChargeDown is true if the terminal voltage is too high.

DeltaMV indicates the change in terminal voltage since the last inferencing cycle (1 second ago).

It is unique in that it returns a numeric result, and that it has a temporal aspect.

Two rules are conclusions which have an output action:

MoreCharge increases the charging setpoint, if the voltage is too low and DeltaMV is too low.

LessCharge decreases the charging setpoint, if the voltage is too high and DeltaMV is too high.³

The output action of these rules is to change the setpoint of a PID controller; this is accomplished by

embedding a "conventional" procedure call in the CONCLUDES phrase of the rule. The sixth rule is a

"safety" rule to prevent excessive rate of change of terminal voltage:

OverCharge decreases the charging setpoint, if DeltaMV is excessively high.

^{2.} In retrospect, these rules could have been given better names.

^{3.} DeltaMV is a signed value, the threshold is set so that any positive value is "too high."

Several additional "rules" perform digital filtering on the sensor inputs, or execute the PID control software for the actuators. This is done by embedding procedure calls (*2ndFilter* or *ApplyPID*) in the evaluator function of each rule. These "I/O rules" are re-evaluated periodically by making them dependent upon the rule *18Hz*, which in turn is fired 18.2 times per second⁴ by a timer routine.

For the first test of this control strategy, the setpoint for both high voltage charging supplies was incremented by 1 unit (0.1 kV) whenever the terminal voltage was more than 2 units (0.02 MV) less from the desired value. As can be seen in Figure 6-4, this caused a large oscillation in charging setpoint, and a smaller oscillation in the terminal voltage, as the system constantly overcorrected. Subsequent tests showed that this could be cured by decreasing the increment to a fractional value,⁵ or by requiring a larger error threshold before making an adjustment. These tests were greatly facilitated by the fact that the increment and threshold were stored in ordinary Forth variables, and that interpretive access to the language was available during testing.

Figure 6-5 shows the greatly improved performance using an increment of 0.5 units (0.05 kV). Generally, a smaller increment gave better regulation, but at the cost of a slower slew rate. Another variant, designated Strategy 1a, incremented the two charging supplies alternately, so as to give finer control of charging voltage.⁶

6.4 Control Strategy #2

The second series of tests attempted simultaneous control of terminal voltage and corona current, using the decision matrix of Figure 4-5. Normal noise fluctuations in the corona current reading caused the system to "hunt" back and forth for a stable position. Also, an error in the control rules caused the system to simultaneously change the charging voltage and corona points position; it was quickly discovered that these two interacting variables should *not* be adjusted together. A modification, Strategy

^{4.} This is the standard clock interrupt of the IBM PC, adopted on the 68HC16s for consistency.

^{5.} Fractional values are accumulated until a change of one whole unit occurs. This "unit," 0.1 kV, is the smallest change which can be output to the charging power supplies.

^{6.} Incrementing one supply by one unit (10 kV) yields half the effective increase of incrementing both supplies by one unit.





p.62

Figure 6-4. Strategy #1 Oscillation



FN960418.002 STRATEGY 1 INC=0.5 THR=2

p.63

2a, solved the noise problem by adding hysterisis (in the form of a dual threshold) to the *CoronaLoad-low* and *CoronaLoad-high* rules. The rules were rewritten so that *only* corona position *or* charging voltage is adjusted -- never both together.

In addition to simultaneous control of two interacting variables, this was the first strategy to employ two cooperating processors. CPU "0D" measured terminal voltage and controlled the charging supplies with eight rules, exporting the facts *TermMV-high* and *TermMV-low* to the other processor. CPU "0C" measured corona current and controlled corona position with five rules, exporting the facts *CoronaLoad-high* and *CoronaLoad-low*.

Two new safety rules were added. One stopped charging if a terminal voltage limit was reached. The other stopped charging if the terminal voltage suddenly dropped more an 0.5 MV -- evidence that a spark has occurred in the generator tank.⁷

6.5 Control Strategy #3

The combination of corona and charge adjustment can cause a large overshoot of terminal voltage; this is undesirable near the desired voltage. Also, a large increment and large threshold will reach the desired voltage sooner, but small increment and threshold will yield better regulation. This

Terminal Voltage	Low	Slightly low	ОК	Slightly high	High
Corona Current					
Low	Large charge increase	Small charge increase	Do nothing	Small charge decrease	Extend points
ОК	Large charge Increase	Small charge increase	Do nothing	Small charge decrease	Large charge decease
High	Retract points	Small charge increase	Do nothing	Small charge decrease	Large charge decrease

Figure 6-6. Control Strategy #3 Heuristic

^{7.} The accelerator formerly employed an electronic spark detector, but now only this heuristic method is used.

suggests using one strategy when far from the desired voltage, and another when near. This hybrid approach was implemented as Strategy #3.

Appendix E gives the program listing of the rules for strategy #3.

Good results were obtained using a "large" increment of 0.5 unit, and a "small" increment of 0.2 unit. Previous tests revealed that a corona adjustment could cause a jump of 0.2 MV; therefore, when less than 0.2 MV from the desired voltage, the "fine adjustment" rules take effect. The result is shown in Figure 6-7.

6.6 Discussion of Results

The terminal voltage control experiments are summarized in Figure 6-8. In addition to the automatic control strategies just described, five test runs were performed with human operators to provide comparison data.

Two principal figures of merit are voltage regulation (RMS error), and maximum rate of change of terminal voltage (slew rate). Both of these statistics are somewhat compromised by noise in the terminal voltmeter reading. To estimate this noise, two tests were performed using the independent "beam stabilizer" circuit, which regulates the terminal voltage to within 2 kV (Cairns, Greene, and Kuehner 1974). These suggest about a 10 kV RMS error in the terminal voltage reading, corresponding to one count of the external digital voltmeter.

The "refined" expert control strategies exhibited measured variations of about 15 kV RMS. Assuming that voltmeter noise and terminal voltage fluctuations are statistically independent⁸, this corresponds to a regulation of approximately 11 kV RMS. "Unregulated" manual operation exhibits variations of 20 to 43 kV RMS.

Since slew rate is computed as the difference between successive measurements, it is strongly affected by random fluctuations. It is reasonable to assume, over a long series of samples, that the largest difference will be exaggerated by one count (10 kV); this accords with casual observation of the voltmeter

^{8.} In this case, the total variance is the sum of the two independent variances; RMS is the square root of the variance.





Figure 6-7. Strategy #3

Experiment Number	Control Strategy	Terminal Voltage Threshold	Belt Charge Increment	Belt "Fine" Increment	Corona Current Threshold	Regulation (RMS error)	Maximum Slew Rate
960416.001	manual					13.8 kV [1]	160 kV/s
960417.001	manual					n/a	70 kV/s
960501.001	manual					44.2 kV	150 kV/s
960501.002	manual					10.9 kV [1]	n/a
960501.003	manual					21.8 kV	180 kV/s
960416.002	1	0.02 MV	0.1 kV			20.8 kV	60 kV/s
960416.003	1	0.04 MV	0.1 kV			24.2 kV	50 kV/s
960418.001	1	0.02 MV	0.02 kV			16.0 kV	50 kV/s
960418.002	1	0.02 MV	0.05 kV			15.3 kV	40 kV/s
960418.003	1 [2]	0.02 MV	0.05 kV			11.6 kV	40 kV/s
960419.001	1a	0.02 MV	0.02 kV			12.7 kV	60 kV/s
960419.002	1a	0.02 MV	0.05 kV			14.3 kV	50 kV/s
960426.001	2	0.02 MV	0.05 kV		5 uA	n/a	70 kV/s
960426.002	2a	0.02 MV	0.05 kV		10 / 7 uA	11.8 kV	50 kV/s
960426.003	2a	0.02 MV	0.05 kV		5 / 2 uA	15.6 kV	90 kV/s
960430.001	3	0.02 MV	0.05 kV	0.01 kV	5 / 2 uA	15.7 kV	50 kV/s
960502.001	3	0.02 MV	0.05 kV	0.02 kV	5 / 2 uA	15.6 kV	50 kV/s
960502.002	3	0.02 MV	0.1 kV	0.02 kV	5 / 2 uA	17.1 kV	60 kV/s

Note 1. For these tests the automatic terminal voltage stabilizer was activated.

Note 2. This was a duplicate of the previous test, to observe the effect on terminal voltage of manual changes in the corona points position

Figure 6-8. Summary of Terminal Control Trials

reading. Then the "better" control strategies limit slew rate to about 40 kV/second⁹, compared with 60 to 170 kV/second for different human operators.

Comparing the experimental runs of Figure 6-1 and Figure 6-7 shows that the expert system is faster to reach a new voltage (80 vs. 250 seconds, for a 1.5 MV change), and has less overshoot (30 kV vs. 300 kV). These are also important figures of merit for terminal control. Some test runs were recorded in which a human operator was faster than the expert system to set a new voltage. This usually involved the human operator using his judgment to exceed the recommended maximum rate of increase of terminal voltage; presumably, this knowledge could be incorporated into additional rules which allow the expert system to use a higher slew rate under specified conditions (such as terminal voltage below 5 MV). Invariably, though, the human operators overshoot the desired terminal voltage by substantially more than the expert system.

6.7 Observations

It is noteworthy that control strategy #3 was programmed in approximately 30 minutes. This "rapid prototyping" is a tribute both to the ease of writing rule-based control algorithms, and the advantage of implementing the rule base as an extension of the programming language (thereby allowing all of the interactive programming tools to be applied to the knowledge base, as well). Control strategy #2 required roughly half a day to develop, but this was largely the "design" phase, in consultation with accelerator lab staff. Once a strategy is designed, actually writing the rules to embody that strategy is a quick and straightforward task.

All of the expert control strategies operated with a "major inferencing cycle" of one second (that is, all of the control rules were evaluated once per second). This was felt to be a reasonable choice for initial experiments, since it accords roughly with a human reaction time. It would be instructive to explore control strategies with a longer inferencing cycle, to discover a lower bound on expert system response time. Reducing the cycle time would reduce the inferencing load on the processors, and thus

^{9.} The 90 kV/s recorded for experiment 960426.003 is probably a statistical fluke, given the repeatability of the slew rate limits for all other test runs.

permit larger and more complex problems to be handled. It may also be useful to run different parts of the knowledge base on different schedules: perhaps charging current can be adjusted every five seconds, but safety rules such as terminal overvoltage should be checked every second.

Corona current and terminal voltage are strongly coupled, and should properly be controlled by the same CPU. That they were separated here was an accident of the input/output assignment -- a fortuitous accident, because it offered a very effective test of two processors cooperating on one problem.

This may well be carried to a greater extreme in future control systems. One can readily visualize a network of more and smaller CPUs, with -- for instance -- one CPU located at the corona points, dedicated solely to measuring corona current and controlling corona position.¹⁰ The extent to which the problem can be distributed is limited mainly by the amount of interprocessor communication required, which in turn depends on how many facts must be exported and imported. A problem for future research will be to discover where the tradeoff between inferencing (each rule in a different processor) and communication (all rules in one processor) is optimum.

^{10.} From an installation standpoint this is advantageous, because it can greatly reduce the amount of cabling required in the physical plant.