

An Adaptive Noise Cancelling System Used* for Beam Control at the Stanford Linear Accelerator Center

T. Himel, S. Allison, P. Grossberg, L. Hendrickson, R. Sass, H. Shoaee
Stanford Linear Accelerator Center, Stanford University, Stanford, CA 94309 USA

Abstract

The SLAC Linear Collider now has a total of twenty-four beam-steering feedback loops used to keep the electron and positron beams on their desired trajectories. Seven of these loops measure and control the same beam as it proceeds down the linac through the arcs to the final focus. Ideally each loop should correct only for disturbances that occur between it and the immediate upstream loop. In fact, in the original system each loop corrected for all upstream disturbances. This resulted in undesirable over-correction and ringing. We added MIMO (Multiple Input Multiple Output) adaptive noise cancellers to separate the signal we wish to correct from disturbances further upstream. This adaptive control improved performance in the 1992 run.

I. INTRODUCTION

The SLC presently has twenty-four steering feedback loops running [1]. Seven of these loops are placed one after the other along the linac.

A typical loop measures and controls eight states: the position and angle of the electron beam in both the horizontal and vertical directions and the same for positrons. The loop measures these states using ten beam position monitors (BPMs). Each monitor gives the horizontal and vertical position for electrons and positrons. Hence, there are a total of forty measurements.

Each feedback loop is designed using our knowledge of accelerator optics and the state-space formalism of control theory. The linear quadratic Gaussian (LQG) method is used to design optimum filters to minimize the rms disturbance seen in the beam. Since there is a fair amount of white noise in the incoming beam disturbance, this filter averages measurements of about six beam pulses. Hence the typical loop corrects most of a step change in six pulses.

A problem exists with the system as described so far. Seven loops in a row examine the same beam. Figure 1 depicts the beam trajectory in the region of two of these loops. Figure 1a shows the trajectory on the first pulse after a sudden disturbance (such as an operator adjusting a dipole magnet strength) upstream of the two loops. The plot of transverse beam position as a function of distance along the linac shows the sine-like trajectory caused by the focusing quadrupole lenses. At this time, the loops have not made a correction. Figure 1b shows the trajectory on the next pulse. To keep this example simple, the loops were set to completely fix an error detected in one pulse instead of in six. The first loop completely corrected the original disturbance. The second loop also made a correction, *which was unnecessary* because the first loop corrected for the disturbance. Of course, on the next pulse the second loop would correct its error but the damage has been done, the loops have overshot the mark for a pulse. The problem gets much worse with seven loops in a row. The overshoot can be reduced by having each loop respond more slowly but the system still overshoots and then rings for many pulses. The system is stable and the ringing gradually dies out, but the overall response of the loops is not optimal, hence the beam positions and angles have a larger rms than need be.

The proper solution is to have each loop correct only for disturbances which happen between it and the next upstream loop. This would completely eliminate the overshooting caused by multiple loops correcting for the same disturbance.

ADDING A MIMO ADAPTIVE NOISE CANCELER

An individual loop (say loop $n+1$) has only a few local BPMs to detect disturbances in the beam. It has no way to tell how far upstream the disturbance occurred. Since we want loop $n+1$ to correct for disturbances downstream of loop n , but not upstream, the upstream disturbances can be thought of as noise.

* Work supported by Department of Energy contract DE-AC03-76SF00515.

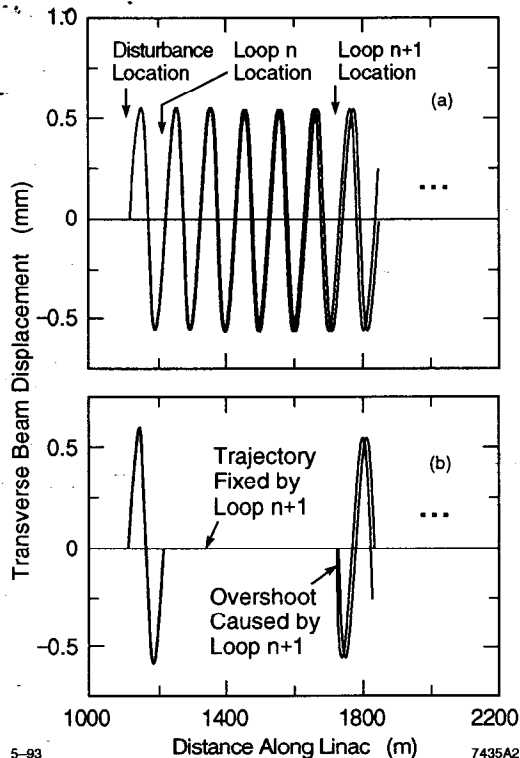


Figure 1. Feedback's response to a disturbance. The beam trajectory shown is on the first pulse (a) and second pulse (b) after a sudden disturbance is introduced. The response of the two feedback loops shows the need for the adaptive noise cancelling system.

Hence an adaptive noise canceler can be used to solve our problem.

A block diagram of the cascading of information from one loop to the next is shown in Figure 2. The bold lines represent information carried by the beam and the bold boxes represent transfer functions which are part of the plant (accelerator). The non-bold items represent items implemented as part of our feedback system.

The line in the upper left labeled "Positions, angles at loop n " represents the eight states. Since loop n is responsible for maintaining these states at their desired set points (which are typically zero since we want the beam to move in a straight line down the center of the linac), as far as loop $n+1$ is concerned, these states are noise. Loop n reads some BPMs and calculates the positions and angles from their readings. It uses the numbers for its own feedback loop, and sends them via a communications link (labeled "Measured positions, angles at loop n ") to loop $n+1$, that uses them as its noise reference signal for its adaptive noise canceler.

Similar information is carried to loop $n+1$ by the beam itself. Between the two loops, the beam executes a betatron oscillation so that positions and angles transform into each other. This is represented by the box labeled "Transport from n to $n+1$," and represents the accelerator, dynamics between the two loops. It is very important to note that our problem is static;

the transport of this beam pulse does not depend on the positions and angles of the previous beam pulse. Hence, the box can be represented as a simple 8×8 matrix.

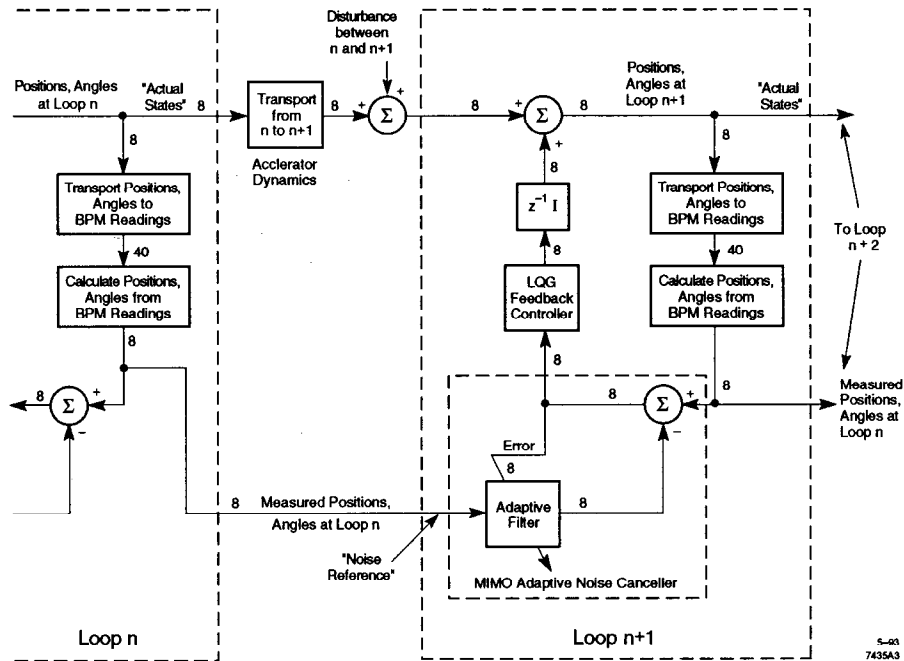
In addition to the simple transport of the beam, an additional "Disturbance between n and $n+1$ " may be added. This disturbance could be due to a klystron tripping off or an operator adjusting a magnet. Loop $n+1$ is intended to correct this kind of disturbance so that it corresponds to the signal that we want the noise canceler to extract.

The last box that needs an explanation is the "LQG Feedback Controller." This box represents the controller feedback loop $n+1$. The controller now takes as its input the output of the MIMO adaptive noise canceler, which represents our best estimate of the "Disturbance between n and $n+1$." That is precisely what we want loop $n+1$ to correct. The output of the controller controls the dipole magnets that steer the beam between n and $n+1$. Hence its output is shown summed into the positions and angles of the beam transported from loop n .

In summary, before the implementation of the adaptive noise canceler, the series of seven feedback loops over-corrected for deviations in the position and angle of the beam because each feedback loop acted independently, and all feedback loops applied a correction for the same disturbance. MIMO adaptive noise cancelers allow each loop to separate disturbances that happen immediately upstream from those that occur upstream of the previous loop. This action cures the over-correction problem.

ADAPTIVE CALCULATION

Before delving into the details of the adaptive calculation, it is worthwhile to ask why adaptation is necessary at all. What is varying? The box labeled "Transport from n to $n+1$ " in Figure 2 is what varies. It accounts for the sine-like trajectory, caused by the focusing magnets, that the beam follows as it travels down the accelerator. For example, if loop $n+1$ is 90° of the betatron (sine-like) oscillation downstream of loop n , then a position offset at loop n becomes an angle at loop $n+1$, and an angle transforms into a position. The transformation is critically dependent on the number of betatron oscillations between the loops. This is parameterized as the *phase advance* where 360° of phase advance corresponds to one full oscillation. Figure 1 shows two loops separated by $5 \times 360^\circ$ of phase advance, the average for the loops in the SLC. The dotted line in Figure 1a shows a betatron oscillation where the focusing strength is incorrect by 1 percent, an error typical of the real linac. Note that the position and angle at the second loop are quite different due to the 1 percent error. This significant variation of the "Transport from n to $n+1$ " forces the use of an adaptive method for the noise canceler.



2. Adaptive MIMO noise canceler added to the typical feedback loop.

The updates of the weights in the adaptive filter are made using the Sequential Regression (SER) algorithm [2]. The equations used in the SER algorithm are explained in Reference [2]. Instead, we will examine our experiences which lead to the use of this algorithm and mention minor modifications that were made to ensure it would be robust.

We started out using the least-mean-square (LMS) algorithm for updating the weights (matrix elements). This is the simplest algorithm, is very fast computationally, and has been successfully used in many applications. In the design phase of the project, simulations were done to check our algorithms. Two problems turned up.


- As explained in reference [2], the LMS method is only stable if the learning rate is less than the inverse of the largest eigenvalue of the input correlation matrix. For our problem it is the natural jitter of the beam due to magnet supply fluctuations and klystron problems that cause the variations of the positions and angles and hence the information from which the adaptation is done. The amplitude of this jitter can easily change by an order of magnitude in a short time which in turn means the eigenvalues change by that amount. If the jitter increased too much the LMS method would become unstable and our feedback system would malfunction, making things still worse. To ensure this didn't happen we would have to pick a learning rate much smaller than the typical optimum value which would result in a very slow convergence.

- The LMS method has a different convergence rate for each eigenmode which depends on the corresponding eigenvalue. Unless we carefully scaled our inputs, the convergence of the slowest eigenmode would be much less than optimum.

At the cost of added complexity and CPU time, the SER algorithm avoids the above problems. Basically it adaptively estimates the inverse of the input correlation matrix. This is used to scale the inputs so that all the eigenvalues of the correlation matrix of the scaled inputs are equal to 1.

Even with the SER method, the calculation of the weights becomes unstable for a short time if the beam jitter suddenly increases. During the time it takes for the estimate of the inverse of the input correlation matrix to converge to the new value, the weights run away. This problem was found in simulation along with the solution: not to update the weights if the inverse correlation matrix is receiving large updates.

After testing the algorithms with the computer simulation we implemented them in the SLC control system. The software was tested and debugged using a hardware accelerator simulator capable of mimicking the accelerator response to 3 simple loops, each having one beam position monitor and one dipole corrector. In this environment the proper operation of both the method of cascading to reduce overshoot and the adaptive learning to the beam transport matrix were verified. Finally (about six months into the project) the time had come to use it on the real accelerator.



5-93
7435A4
Figure 3. Response of a chain of six feedback loops to a sudden disturbance in the incoming beam. In part (a) adaptive noise cancelling is off so there is a ringing caused by the over correction of many loops. In part (b) adaptive noise cancelling is on, so the whole chain of loops responds like a single loop. In fact, the first loop did all the work to correct the beam and the downstream loops did virtually nothing.

EXPERIENCE ON THE REAL ACCELERATOR

First we turned on just the adaptive algorithm. The results were not used to control the beam. After confirming that the matrices had converged to reasonable values, we turned on the noise cancelling system. As shown in Figure 3 the response to a step disturbance in the beam trajectory was greatly improved with the startup of the adaptive noise-cancelling system.

Over the next few weeks we varied the learning rate to find the optimum value that would allow the adaptation to converge rapidly without having too much noise introduced by the adaptive process. We settled on a learning rate of 0.001 and an adaptive update rate of 10 Hz. A convergence time of about 100 seconds resulted. The system ran for several days with learning rates of 0.1 and 0.01 and was completely stable, but with these higher learning rates more random noise showed in the adaptive matrix elements.

The adaptive noise-cancelling addition to the fast feedback system has been running stably in seven locations on the SLAC linear collider for over six months. Probably the best measure of its robustness and stability is that operators have made *no* middle of the night phone calls asking for help to recover from a problem. In fact there have been no significant problems with the system. Adaptive noise cancelling has significantly improved the performance of our feedback systems and helped us achieve our goals of accelerating two beams over a distance of three kilometers, pointing the beams at each other, and then colliding them head on so they pass through each other even though they have a radius of only $2 \mu\text{m}$ at the collision point.

In fact we have received an unexpected bonus from the adaptive calculation. The adaptive weights can be interpreted as measurements of the beam transport matrix from one loop to the next. These measurements are recorded on disk and can be displayed. Such data shows a typical variation of over 30 degrees which is about 1 percent of the total phase advance between the two loops. We have made many checks and convinced ourselves that this variation is caused by a real variation in the focusing strengths in the linac (typically due to rf phase and energy changes). Accelerator physicists are using this data to identify and try to fix the cause of the changes in focusing strength. This would make a still more stable accelerator.

REFERENCES

- [1] F. Rouse et al., "A Database-Driven Fast Feedback System for the Stanford Linear Collider," *Nuc. Instrum. Meth.* A316:343-350, 1992.
- [2] B. Widrow and S. Sterns, *Adaptive Signal Processing*. New Jersey: Prentice-Hall, 1985.