Adjustment of Trimming-Coil Currents

H. G. Blosser

We have heard a good deal of discussion from several speakers on various ways for adjusting trimming-coil currents. I want to discuss still another way for handling this problem, based on field measurements and least-square fitting procedures, with a computer being used of course to do the actual calculating. The method is essentially a closed loop; measurements are plugged in and the answers grind out. There is little need for any sort of judgment or physical intuition on the part of the operator, such as is required in various of the other methods of trimming-coil adjustment.

We were led to this system by experience obtained in the design of the Oak Ridge electron model, and let me begin then by describing some of the pertinent experience from the design of that model. The most elementary approach to the trimming coil problem is to present your operator with a panel of knobs and a meter and tell him to adjust the knobs to peak the meter. If there are only a few knobs, this works fine.

In the Oak Ridge electron model we had 21 knobs, each of which controlled the current in a pair of circular trimming-coils; the currents had already been set fairly well by a computer program. The computed settings assumed that all the coils in the model were perfect, but we knew from the performance that the actual coils weren't perfect. The beam wiggles around; there are clearly median-plane errors, and that sort of thing. One ought then to be able to find a considerably better arrangement of the trimming-coil currents, an arrangement which would compensate for the fabrication errors the computed settings did not consider. On the basis of these arguments several of us have tried knob twiddling for a good many hours. We watched the beam vs radius and tried to find a setting of the knobs which would give a lower threshold voltage. No change in the measured threshold was ever achieved.

This experience clearly demonstrated to us a fact which everyone else probably knew to begin with; if you have a large number of interacting variables and no detailed information on the relationship of the variables to the quantity which you are trying to improve, then you have a real jungle and its practically impossible to improve things simply by empirical adjustment of the variables. And, of course, if the particles make a large number of turns (of the order of a thousand) in a cyclotron, you must have many variables in order to control the field with sufficient accuracy.

At another stage in the design of the analogue we tried an improved system of empirical adjustment, again without a great deal of success. In this system we used the curve plotter attachment on the Oak Ridge computer to give us a continuous display of the error in average field vs radius. To change the current in a coil we had simply to type a revised number into the computer and the revised error profile would pop up immediately on the curve plotter. We had then much the same sort of problem as you would have if you observed the phase; this is certainly much better than simply looking at the threshold. You can tell at what radius the field is in error and work to improve things at that radius. We spent a number of hours of computer time with this procedure and found we could get the errors down to 1/2%

or better fairly readily, but to go further (we were trying to get to 0.05%), it got extremely tedious.

At about this stage we decided to try a standard least-squares technique and it worked so well that we promptly dropped all the other methods. In 5 to 10 min. on the computer we would get solutions which fitted by about a factor of ten better than the tolerance we were trying to achieve. With a few precautions, which I'll describe in a minute, the required currents would be quite reasonable.

The same system can be adopted fairly straight-forwardly to the trimming of magnets for medium energy cyclotrons. Let the total field be the sum of the main magnet field and the trimming-coil corrections, then we have:

$$B_{total}(r,\theta) = B_{main}(r,\theta) + \sum_{j=1}^{n} I_j B_j(r,\theta). \tag{1}$$

I've assumed that the fields of the trimming-coils are linear, that they're given simply by a shape factor times the current. For corrections of the order of 1 or 2% this linear assumption holds to good accuracy. Also, non-linear current terms could be included without a great deal of additional complexity. For now, though, lets assume that Eq. (1) holds and, of course, always we're interested in trimming only the average field since this is quantity which controls the isochronism.

Then Eq. (1) becomes

$$\overline{B}_{total}(r) = \overline{B}_{main}(r) + \sum_{n=1}^{j} I_j \overline{B_j}(r), \tag{2}$$

which is now in a form such that a standard least-squares procedure can be applied to find the currents, I, which most nearly produce the desired B_{total}. It is also, of course, a trivial extension to include a weighting factor for the turn density if desired. Almost all computers have in fact a standard subroutine for handling this problem and they usually include the use of a weighting gactor as an option.

The most difficult part of the trimming is determining the functions $\overline{B}_{main}(r)$ and $\overline{B}_j(r)$, because all of these functions are certainly going to vary due to saturation effects as the main magnet field is changed to provide for variable energy. The way we propose to handle this is by an extensive series of field measurements before the cyclotron starts operation. We would measure the main field at a sequence of field strengths with perhaps 1-kilogauss increments over the range where we planned to operate, and we would measure the incremental field of each trimming-coil at a similar sequence of main magnet excitations; probably it would be adequate to take coarser steps in this series, say 2-kilogauss. If we use eight trimming-coils and had a range of, say, 10 kilogauss this would mean around 50 field scans. We would do this once and store the information in a permanent storage. To calculate the current settings for running some ion to an arbitrary energy we would first of all interpolate in this library of stored information to determine the main field shape and the trimming-coil increments at the particular field strength. This information would then be plugged into Eq. (2) and the computer would grind out the currents.

Each time an operating point was calculated for some energy we would put the results in a library so that subsequently the operator would simply consult the library when he wished to run at that energy. In final analysis then you have a fairly laborious set of field measurements which you make one time; thereafter, parameter adjustments are determined fairly automatically. We feel that to get really convenient and reproducible variation in the machine energy some such automatic scheme is really essential.

The principal possible bug in this system would be a shift of the field shape with time due to some aging process. We think its fairly sure that the shape of the increments produced by the trimming-coils won't be affected by such shifts. First of all, the trimming coil increments are an effect produced mainly by current and secondly, you don't need to impose severe accuracies since the fields are already a small percentage of the total field. The question of the main field is much trickier. We plan to set the main field by monitoring some standard position with a proton resonance probe so that we don't have to worry about hysteresis effects.

The question is, then, if we have one point with the right absolute value, how far off will other points be? All the data which we've looked at indicate that we'll be ok on this question but in most cases the information hasn't been very accurate and we are still somewhat concerned about this problem of aging.

One or two things I wanted to mention finally about the least squares trimming. We've found it essential to always use considerably more data points than variables in the least squares problem, usually we use 3 or 4 times more data points. If you don't you'll often get an effect called over fitting; the currents will work very hard to get close to every data point and if you then compute values between data points you'll usually find fairly severe wiggles. The second thing is that occasionally from our least squares fits we'll find large oppositely directed currents in adjacent coils which essentially null each other. When this happens we split the problem into two steps. Fit first with a subset of the coils including one of the troublesome pair and then fit the residual error from the first run using the remaining coils. This has always worked fine to remove this sort of trouble.

With these precautions we've been very pleased with the performance of the least squares method for setting trimming-coil currents. We feel it ought to make possible a programmed type of operation where you can really shift energy quite easily, and, can rapidly compute the currents required for operation at any arbitrary energy.

CHAIRMAN KELLY: I think the problem that Blosser brought out here is that you never know whether you are really on the best optimum for all of the coils, and there are probably several fairly good solutions or sets of currents that will produce a satisfactory field. If you try to approach this by twiddling, it depends on whether you are close enough to one of these good positions to be able to find it, and if you find one, is it really a very good one? I think this is the question that would go with any sort of display of an error such as this phase measuring device would have. You might adjust the coil currents, and what Dols had in mind, I think, was displaying the phase at a series of radial points all at once, perhaps on the oscilloscope, so that they might form a straight line when the phase error was zero and then adjust the currents, trying to get them all in line, but that does not mean that is the best setting. It is a possible one. Somehow or other you have to get at this best optimum to start with. I think that is an unresolved problem.

GREEN: There is a very powerful mathematical method for doing this, the Southworth relaxation method, which is fitted for handling problems of the combination of equipotential surfaces and currents. It is a method by which one starts with a grid of measured or estimated values and then iterates, and it requires a number of iterations. It is easily adaptable for a hand calculator, the 650 class, or it can be programmed on a 704. There are, I believe, at NYU some sub-routines, but with this relaxation method you can start out with measured values and then calculate the currents to produce your desired field shapes in quite a straight forward method by relaxing through the grid.

BLOSSER: This would substitute for the least-squares process here?

GREEN: Yes. As an example of what it will do, we designed the shape of our magnet poles with the relaxation method. It is the same problem. If you have a contour that is quite complex and you alter the iron in one place you alter the field throughout the gap. Our contour was calculated with the relaxation method, and the gradient, the first derivative on the median plane, was measured and the calculations agreed to better than 0.1%. It is an extremely powerful method.

One side remark might be interesting. We have gone into extensive tests on aging of steel, for obvious reasons. We find that the coercive force and low-induction permeability tend to age in a hair-raising manner. Coercive forces of many steels will increase nearly a factor of 2. We cannot detect appreciable aging in steel at high induction. It just does not seem to change at high induction, no matter what you do with it or how badly you treat it.

CHAIRMAN KELLY: How high?

GREEN: High is above 5,000 gauss.

BLOSSER: It sounds good.

GREEN: At 100 gauss it is awful.

WELTON: I just want to say that Dr. Green is a little spoiled by the simplicity of the A-G magnets.

GREEN: No, this is going down to three dimensional relaxation.

CHAIRMAN KELLY: I would like to get together with you about this afterward. I think it is something that we may be overly concerned about and we will find that when this machine is in operation it is really not as much of a problem as we think. Ken MacKenzie liked to say that if you asked somebody whether it is possible to operate an automobile and to do all the things one has to do simultaneously, he would say that it is impossible and that it would take at least several people; yet we do this quite automatically at times and so perhaps this coil problem isn't so bad. It is just that we are new at it.