

How to Cope with the Linear Coupling in RHIC

S. Ohnuma

University of Houston
Houston, TX 77204-5506
September 1, 1995

1. Introduction

As the name “linear coupling” makes it clear, the problem discussed in this report is “linear” so that it should be “unambiguous”. Operational nuisance associated with it has long been recognized by those who work in control rooms. The traditional remedy has been to minimize the difference of two tune values, horizontal and vertical, using two families of skew quadrupoles. Since tune is a global parameter, this may be called a “global correction”.

In 1991, Richard Talman proposed a new decoupling scheme which aims to reduce the amount of linear coupling everywhere in a ring.¹ This was actively pursued by a group of people at the SSC Laboratory but, to the best of my knowledge, it has not been tested on a working machine. In early 1994, a very poor luminosity performance of the Tevatron at B0 (CDF) was found to be caused by a roll of one quadrupole in an insertion triplet at B0, and this seems to have revived an interest in Talman’s decoupling scheme in the RHIC group. The purpose of this note is to question, albeit in a limited manner, the necessity of using this decoupling scheme as a standard operational procedure of RHIC.

In section 2, a story of what happened to Tevatron at Fermilab last year is given in order to understand its implication to the future operation of RHIC. This is a composite picture of several (sometime conflicting) versions I have gathered at Fermilab during June-July of this year. As such, it may contain factual errors for which I take responsibility. In section 3, my concerns with the use of Talman's decoupling scheme are listed, although some of them have been answered by Talman and Fulvia Pilat in private conversations. A simple decoupling scheme with orbit bumps covering one insertion triplet at a time is explained in section 4 with a concrete example to show its practicality. Section 5 contains a variety of topics related to the comparison of Talman's scheme with the minimization of tune split. A "miniature" model of RHIC* has been used for this purpose and this implies a severe limitation of the reliability of results presented in this section. The sole purpose here is to see if one method could replace the other and, if so, to what extent. The quantitative statements given here must be re-examined using more comprehensive codes such as TEAPOT and MAD. The possible importance of off-momentum decoupling has been emphasized by Steve Peggs and others with analytical as well as numerical investigations. In section 6, recent results obtained by RHIC Accelerator Physics Group have been examined to understand what one can do to reduce this effect. A possible way to achieve the reduction of off-momentum coupling with magnet sorting is presented there also. Finally, my personal view of the overall priorities for RHIC during its commissioning period is presented in section 7. The emphasis here is the qualifying word "personal" in that the statements given in this section are entirely based on my past experience.

* All numerical results presented in this note assume RHIC with $\beta^* = 1\text{m}$ at 6 o'clock and 8 o'clock, and 10m at all other interaction points.

2. What happened during Collider Run IB at Fermilab

After a long shutdown, the Collider Run IB of Tevatron at Fermilab started near the end of 1993 and continued its operation until July 1994. During the first part of this period, the measured luminosity at CDF(B0) was 75% of the luminosity at D0 detector.* In addition, the longitudinal distribution of luminosity was not symmetric at B0. Later, somewhat accidentally, it was found out that one quadrupole of the upstream triplet at B0, the one farthest from the interaction point, was rolled by 8 mrad. After the realignment, the luminosity was back to its normal value and the distribution at B0 became symmetric as expected.

Looking from outside, one can ask many questions, “Why didn’t you do this and that...”. Although there is no definitive version of what really happened, the following is a composite of several versions I collected at Fermilab during my visit in June-July of this year. It is easy to criticize them far away from the Tevatron control room, especially knowing what we know now. Instead of indulging in that, I suggest that we all try to learn valuable lessons from their unlucky experience.

1. First of all, whenever a major shutdown occurs, no single person knows everything that has happened in the tunnel. It has been claimed by someone that Helen Edwards was the only exception to this. Who would be “Helen” of RHIC in the near future?
2. What caused this roll which is truly abnormal in its magnitude? Currently, there are two explanations: shielding added near the triplet in question, and a quench somewhere else which loosened the quadrupole. There seems to be no consensus on this point.
3. Didn’t they notice immediately that the setting of correction skew quadrupoles is different from the one before the shutdown? Yes, they did. But it was “explained” as due to rolled dipoles and quadrupoles done during the shutdown. It is not clear why some quadrupoles had to be rolled but dipoles were rolled to correct the vertical closed orbit which was becoming too much for the vertical

* In the parlance of RHIC, one may say that B0 is at 2 o’clock and D0 at 6 o’clock if A0 is at 12 o’clock. Protons go clockwise in Tevatron.

correctors at the highest energy. Intentional roll of quadrupoles is not new at Fermilab. In the Main Ring, there is something called “Edwards–Stiening Roll” at stations #25 and #43 in each period. The amount of each roll is (I believe) some ten mrad and it has been in existence ever since its introduction in 1973.

4. CDF at B0 was understandably very unhappy but D0 people were quite satisfied with their luminosity. They didn’t want to be interrupted. Apparently, the trouble was well confined to B0. CDF kept taking data in spite of the unhappiness.
5. To be sure, there were some efforts to correct the problem using a pair of skew quadrupoles upstream and downstream of B0. They were, however, not tunable independently and the correction was not effective. Besides, as far as the minimization of tune split is concerned, skew quadrupoles at B0 and at D0 are equivalent in their effect so that there was no way to see which one should be used.
6. Nobody had any convincing argument that luminosity would be affected so much by linear coupling. Now they (and we) know better, of course.²
7. How was it eventually discovered? Rather accidentally. For some reasons not associated with the luminosity problem at B0, the area B1, which includes B0 insertion, had to be warmed-up. Around that time, Craig Moore suspected a mismatch in beta function causing an emittance growth while Jim Holt suspected a transverse (not rotational) misalignment of Q3 (which is next to Q2). Because of these suspicions, they decided to look at the alignment of triplets at B0. In addition to the 8 mrad roll of Q2, they discovered transverse misalignments in Q2 and Q3.

What should be the lessons for us from this incident? Different people would naturally answer differently. Dick Talman believed that this confirms his belief that linear coupling must be controlled locally as well as globally, and that there should be a joint project of beam study somewhat in the manner of E778 at Fermilab. I felt (and still feel) that the incident was quite abnormal (8 mrad!) and one should not overreact. For RHIC, what

RAP members think should be far more important than what I think or, for that matter, what Dick Talman thinks.

3. Some Questions on the Use of Talman's Decoupling Scheme*

When a beam is pinged in a transverse direction in the presence of linear coupling, and its transverse position (x and y) observed by a dual plane of BPM placed somewhere in the ring, one expects

$$\begin{aligned}x(n) &= g \cos \psi_A(n) + h e_D (\psi_D + \epsilon_D), \\y(n) &= h \cos \psi_D(n) + g e_A \cos(\psi_A + \epsilon_A)\end{aligned}$$

for the n -th turn. With this sign convention, ψ_A changes by $(-2\pi Q_A)$ per turn and ψ_D by $(-2\pi Q_D)$ per turn. Talman's decoupling scheme assumes that one can extract a pair of parameters (e_A, ϵ_A) or (e_D, ϵ_D) at all the dual plane BPMs distributed around the ring. He defines a quantity called "badness function"

$$F \equiv \sum e_A^2(d) (\beta_x/\beta_y)_d \text{ summed over all the dual plane BPMs,}$$

which is a measure of the strength of linear coupling in the ring. With correction skew quadrupoles, one then tries to minimize the value of this function. Note that if the beam is pinged in one of the eigendirections, for example A , there will be no contribution coming from eigendirection D so that $h = 0$. In his original report, Talman seems to have assumed that the angle between horizontal direction and eigendirection A is small and the contamination of the signal due to $h \neq 0$ is negligible. Now he (and others) claims that the extraction of e_A and ϵ_A has been tried successfully at CERN. In view of this proof of existence, I should not question its validity. Nevertheless, I am still not completely comfortable because of its complexity.

Would it be "correct" to include all the dual plane BPMs, in arcs as well as in insertions, on the equal footing in minimizing the badness function? If the decoupling is more

* Someone has commented that this section is "full of holes". After a moment of reflection on this, I decided to keep the section in its original "Swiss-cheese" shape. In doing so, my wish is that nobody would doubt where the prejudice lies.

important in insertions than in arcs (which should be arguably true), the inclusion of arc dual plane BPMs for this purpose might dilute the importance of the information coming from the BPMs in insertions. In this connection, it seems quite reasonable that RHIC is adopting the solution to correct coupling in insertions using the existing (72 altogether) dual plane BPMs, and to use the minimization of tune split for correction of other coupling sources in arcs.³ The latter of course does not require any use of BPMs.

If one insists on using BPMs in arcs, it is necessary to guess the missing value of either x or y from the measurements at two adjacent BPMs (unless dual plane BPMs are installed in arcs as well as in insertions, an unlikely scenario). Since such a guess must necessarily be based on the ideal model free of any errors, certain amount of errors is unavoidable. A simple calculation shows that, for 1,000 random samples, the average error is 0.13mm with $\sigma = 0.10\text{mm}$ and the maximum of 1,000 samples as large as 0.5mm. Parameters assumed for this calculation are:

1. quadrupole misalignment (horizontal and vertical): 0.25mm(rms)
2. quadrupole roll: 0.5mrad(rms)
3. dipole roll: 1mrad(rms)
4. fractional error in the integrated bend field of dipole: 0.05%(rms)
5. Nonlinear field: chromaticity sextupole only.

All random errors are cut off at 3σ . A similar calculation has also been done for one example in the insertion region.

outer arc, vertical BPMs at Q4 and Q6

Guess y at Q5 where BPM is in the horizontal direction only.

Again for 1,000 random samples, the average error = 0.09mm with $\sigma = 0.07\text{mm}$ and the maximum = 0.4mm.

I have no idea if these errors are at all significant or entirely negligible in the application of Talman's decoupling scheme. Nevertheless, it seems prudent to check this point carefully before using the guessed values of beam transverse position.

As for the correction of closed orbit, it is assumed that all the linear parameters such as β , α , and phase advance between correctors and BPMs are known since they are needed to calculate the transfer functions. It is an approximation but seems to be harmless for the

correction of closed orbit as long as the correction is done “step-by-step”. Is this true also for the decoupling? After all, these linear parameters have a meaning in the eigendirections only and not in the usual horizontal–vertical directions when the coupling is significant. It would be interesting to check this with TEAPOT or any other similar codes, paying special attention to simulate the actual operation in the control room as honestly as possible. For example, the lattice to be decoupled should not be used to calculate linear parameters needed for the decoupling algorithm. Rather, these parameters should come from the ideal lattice free of any coupling.

4. A Simple Decoupling Scheme with Orbit Bump

In the talk I gave last year, I mentioned the difficulty in tuning two skew quadrupoles (SQ07B3X and 08B3X, for example) because of the unfavorable phase relations between them and all BPMs inside a localized closed orbit bump in IR. This pessimistic view was also repeated in a report I wrote last year.⁴ This is not correct.

If one looks at Table 2 of ref. 3, page 3, it is clear that one can make a localized horizontal orbit bump covering the downstream triplet only by a suitable combination of a) displacement only at IP, and b) angular change only at IP. One example is given below.

Use (Q6, Q2) correctors upstream, (Q3, Q5) correctors downstream.

Resulting horizontal closed orbit deviations are

upstream		downstream	
BPM at Q3	−0.12mm	BPM at Q1	9.43mm
SQ07B3X	−0.17mm	Q2 corrector	8.47mm
Q2 corrector	−0.19mm	SQ08B3X	9.85mm
BPM at Q1	0.32mm	BPM at Q3	10.49mm

Since there is practically no change in phase within a triplet and almost 180 degrees between two triplets, the vertical closed orbit at Q3 BPM (upstream) is

$$\Delta y(Q3) \simeq \cot(\pi Q_y) \Delta y'(\text{source}) \sqrt{\beta_y(Q3)} \sqrt{\beta_y(\text{source})}$$

where β_y (Q3) is 1,250m. If the source of $\Delta y'$ is Q2 (downstream) rolled by θ_R , for example, β_y (source) = 1,100m and

$$\Delta y' = (3.2 \times 10^{-3}) \theta_R$$

With $\theta_R = 1$ mrad, $\Delta y = 3.0$ m. If this is not large enough, one can increase the magnification factor $1/\tan(\pi Q_y)$ by almost 50% by changing Q_y from the nominal 29.18 to 29.12.

Inevitably, there will be noises in the optics, some arising from the nonlinear field and others from the roll of horizontal correctors that are used to create the orbit bump. The latter is expected to contribute to the vertical orbit distortion. With as much as 10 to 20 mrad of roll, however, the resulting vertical deviation will be at most a few percent of the horizontal orbit bump while the signal (the vertical orbit deviation induced by rolled quadrupoles in the triplet) is one-third of the horizontal orbit distortion/1 mrad.

It is obvious from the antisymmetric nature of the insertion optics that the upstream triplet can be covered by a vertical closed orbit bump with correctors at (Q5, Q3) upstream and at (Q2, Q6) downstream.

This scheme is conceptually so simple and understandable compared with the scheme proposed by Talman that it should be tried first. Admittedly, it does not cure the coupling anywhere in the arcs. If a substantial amount of coupling exists in a regular cell, one would notice it when the standard procedure of a local three-magnet bump (for example) is applied.

5. Comparison of Talman's Scheme and the Tune-Split Minimization in a "Miniature" RHIC Model

The experience at Fermilab (see section 2) may have demonstrated that a local decoupling is needed to ensure an optimum performance of storage rings or accelerators. For those of us who still believe that the standard method of tune-split minimization is sufficient for most cases, it would be reassuring to "demonstrate" the validity of our belief. Before going into the presentation of a modest trial for this purpose, it is important to emphasize that, strictly speaking, the tune-split minimization may be "non-global". In

order to be “global”, one should not ask which correction skew quadrupoles to use since the game is to control the global (complex) parameter called the linear coupling constant. As long as there are two (more or less) orthogonal families of skew correctors, one should be able to achieve the goal regardless of where the correctors are placed in the ring. For example, in Tevatron, the tune–split minimization is usually done with correctors near A0 (12 o’clock).

We recognize the correctors near a particular insertion, which is found to be a cause of the coupling (see section 4), should be used, but the method is still the minimization of tune split and not the minimization of Talman’s badness function. In a sense, this is “local” or at least “non–global” but does not depend on the extraction of two parameters ϵ_A and ϵ_{-A} using dual plane BPMs around the ring. Once we admit that not all the skew correctors are equal in their effectiveness even for the minimization of tune split, it is natural to ask whether Talman’s decoupling scheme is superior, and if so, to what extent, compared with the tune split minimization.

With the use of a simple–minded model of RHIC linear lattice with $\beta^* = 1\text{m}$ at 6 o’clock and at 8 o’clock but 10m elsewhere, the following questions will be examined:

1. Scaling of Talman’s badness function in terms of the roll angle of each quadrupole of triplets. That is, how large its value is when one quadrupole is rolled by, say, 1 mrad. Without this scaling, it will be difficult to see how much improvement one is really achieving by a reduction of badness function.
2. With one or more correction skew quadrupoles (SQ05, SQ06, SQ07, and SQ08B3X), how much can one reduce the magnitude of Talman’s badness function? In particular, can a “wrong” corrector be used for this purpose?
3. If the tune split minimization is used to find the optimum currents for skew correctors, how much reduction will result in the magnitude of badness function? Is it enough?
4. For the minimization of tune split, are skew correctors at 6 o’clock (SQ05 and 06B3X) equally effective as skew correctors at 8 o’clock when the source of linear coupling is at 8 o’clock?

The model I have used for RHIC is completely linear and consists of the following elements connected by suitable linear transfer matrices, which are constructed from the linear parameters of MAD output:

1. Four dual plane BPMs at 8 o'clock, two near IP and two at Q3s,
2. Four skew correctors (thin lens) SQ05, 06, 07, and 08B3X,
3. A pair of triplets (thick lens) at 8 o'clock,
4. Upstream of Q4T at 8 o'clock to scan the tune split.

In view of the simpleness of model, I may be simply playing games which may not really represent the real RHIC. It is therefore desirable (maybe even necessary) to check the results presented here using a more elaborate program such as TEAPOT or MAD. Within a limited length of time available, however, this is the best I can do to say something.

5.1. Scaling of (modified) badness function

Simply for the sake of convenience, Talman's badness function is slightly modified in this calculation:

$$F \equiv \sqrt{\frac{1}{4} \cdot \sum e_A^2 (\beta_x / \beta_y)}$$

The absolute magnitude of the function has of course no meaning in itself unless scaled to something real such as the roll angle of a quadrupole in triplets.

The value of modified badness function is plotted as a function of roll angle for (Q1, Q2, Q3) of the upstream triplet at 8 o'clock (Blue Ring). What should be the acceptable value of badness function defined here? It does not seem necessary to go below, say, 0.05 which correspond to less than 0.05 mrad roll of one quadrupole. It is of course important to note that the badness function is defined here for only four dual plane BPMs at 8 o'clock. If all the dual plane BPMs in the ring are taken into account, the situation may differ substantially (though I don't believe so).

5.2. Effectiveness of skew correctors in reducing the value of badness function

It is not surprising to see that skew correctors at 6 o'clock, SQ05 & 06 are not very effective in reducing the value of badness function when the source of coupling is at 8 o'clock. For the maximum available current (50 A) for each corrector, the best one can do is

roll angle = 2 mrad

	no correction	with correctors
at Q3(upstream)	F=0.875	0.719
Q2(upstream)	0.917	0.890
Q1(upstream)	0.803	0.474
Q1(downstream)	0.807	0.476
Q2(downstream)	0.914	0.746
Q3(downstream)	0.850	0.729

If correctors at 8 o'clock, SQ07 & 08 are used, the improvement is much better even with the “wrong” corrector:

roll angle = 2 mrad

	with SQ07	with SQ08	with both
at Q3(upstream)	F=0.030	0.040	0.012
Q2(upstream)	0.022	0.065	0.011
Q1(upstream)	0.029	0.031	0.012
Q1(downstream)	0.032	0.031	0.014
Q2(downstream)	0.056	0.017	0.010
Q3(downstream)	0.037	0.032	0.009

It is not clear from these examples alone (with the simple-minded RHIC model used for this work) whether it is “necessary” to use the “right” corrector (that is, if the rolled quadrupole is in the upstream triplet, use SQ07) in view of the fact that the badness function is down to the level of 0.05 or so even with the “wrong” corrector.

If the roll angle is as large as 4 mrad, there is a clear indication that the “right” corrector is definitely superior in performance compared with the “wrong” corrector. For example, when Q3 (downstream) is rolled by 4 mrad, SQ07 (“wrong” corrector) alone cannot reduce the value of badness function below 0.075 while SQ08 (“right” corrector) can reduce it down to 0.010. In any case, the right corrector alone should be enough if one is satisfied with the value of 0.03 or less.

5.3.

When the source of linear coupling is at 8 o'clock and two correctors SQ07 & 083X are used for the minimization of tune split, one can reduce the tune split to less than the resolution of tune measurement (0.0005?). The question is: what is the resulting reduction in the value of (modified) badness function? The same question can be asked even when only one corrector is used for the minimization of tune split although the smallest tune split one can achieve with one corrector may be as large as 0.001 (good enough?) if the roll angle is taken to be up to 4 mrad.

Results from the simple-minded model of RHIC lattice are surprising.

roll angle = 2 mrad

	with SQ07	with SQ08	both
at Q3(upstream)	F=0.0329	0.0399	0.0238
Q2(upstream)	0.0222	0.0773	0.0114
Q1(upstream)	0.0286	0.0311	0.0127
Q1(downstream)	0.0431	0.0310	0.0138
Q2(downstream)	0.0555	0.0172	0.0250
Q3(downstream)	0.0602	0.0334	0.0282

roll angle = 4 mrad

	with SQ07	with SQ08	both
at Q3(upstream)	F=0.0570	0.0789	0.0225
Q2(upstream)	*	*	0.0421
Q1(upstream)	0.0550	0.0614	0.0249
Q1(downstream)	0.0655	0.0641	0.0323
Q2(downstream)	*	*	0.0228
Q3(downstream)	0.0753	0.0683	0.0151

* = Minimum tune split larger than 0.001.

The reduction in the value of badness function is of course not as great as the best one can do with Talman's decoupling procedures. Nevertheless, in terms of the equivalent roll angle (see Fig. 1), the reduction seems to be satisfactory especially when the "right" corrector is used. This conclusion must be confirmed with a more elaborate RHIC model such as the one for TEAPOT.

5.4.

Finally, for the minimization of tune split, can we use correctors at 6 o'clock (SQ05 & 06) even when the source of linear coupling is at 8 o'clock? One is tempted to say "yes" in view of the fact that the relevant phase ($\psi_x - \psi_y$) is -1° for SQ05&06 and 349° for SQ07&08, the difference being only 10° .

Figure 2 shows the minimum tune split one can achieve with skew correctors at 6 o'clock when Q1, Q2 and Q3 of triplets at 8 o'clock is rolled one at a time. There is no difference between upstream triplet and downstream triplet. The minimum one can get even at roll angle of 0.5 mrad is definitely larger than a conservative estimate of tune measurement accuracy (0.001).

Is this reasonable? Again this must be confirmed by TEAPOT or MAD.

6. Comments on Off-Momentum Coupling

There has been so much work done on this topic by RAP members that I can add nothing substantial. Indeed, the RHIC Accelerator Advisory Committee was so impressed by the work that it included the following statement in their latest report (April 1995):

“This problem has been studied in more detail and is now resolved. The Committee is impressed with this study and finds the findings convincing.”

Recently, Fritz Dell repeated the previous calculation and clearly identified the major sources of the off-momentum coupling to be the skew sextupole a_2 at the current lead of regular dipoles and the normal sextupole b_2 expected to exist at the current lead of D0 insertion dipoles.⁵

Since one is here interested in the coupling term linear in $(\Delta p/p)$, it should be proportional to

$$(1) \quad X_p \ a_2 \sqrt{\beta_x \beta_y} \exp i(\psi_x - \psi_y)$$

$$(2) \quad Y_p \ b_2 \sqrt{\beta_x \beta_y} \exp i(\psi_x - \psi_y)$$

where X_p and Y_p are the horizontal and vertical dispersions, respectively. X naturally exists in the arcs and predictable from the ideal linear lattice. Small deviation in X_p will come from the fluctuation in b_1 (normal quadrupole component) of regular dipoles but it should be negligible. Y_p , on the other hand, is produced primarily by two unknown quantities, the non-systematic part of a_1 (skew quadrupole component) in regular dipoles and random rolls of quadrupoles.

Because of the nature of its origin, there is not much one can do to reduce the term proportional to $X_p \ a_2$. The term proportional to $Y_p \ b_2$ can be reduced, however, if one is allowed to sort the regular dipoles in the following manner. Since the phase advance per regular cell is almost 90 degrees, the first and the fifth dipoles, the second and the sixth dipoles, and so on will cancel each other in their contributions to the vertical dispersion Y_p anywhere outside the pair of these dipoles if a_1 of the pair is identical. Sorting will then be done such that two dipoles of the similar a_1 are installed at the first and the fifth, the second and the sixth, and so on in each arc. Since there are twenty-two regular dipoles

in each arc, one can make ten pairs with two unpaired dipoles, the 19th and the 20th, for example. Presumably, two best looking magnets (that is, those with a_1 closest to the average a_1) would go to these unpaired spots.

There are two minor questions regarding the results made available to me on this subject.

1. Skew sextupole correctors are to be placed at Q3 in each IR replacing the originally planned normal sextupole (b_2) correctors. Because of its phase ($\psi_x - \psi_y$), any such corrector in either 6 o'clock or in 8 o'clock IR should contribute predominantly real component in $Y_p a_2$ term regardless of what value of Y_p happens to be there. This is true in the figures presented by Jörg Kewisch at the RHIC Advisory Committee meeting in April 1995 with one exception. For seed 3, the imaginary contribution from a_2 corrector at 6 o'clock (downstream) is as large as the real contribution and they are both the largest in all cases examined. This is difficult to understand and a bit disquieting.
2. Contributions proportional to $X_p a_2$ coming from the regular dipoles should be almost predictable due to their systematic origin. One expects more or less the same value regardless of seed numbers and the real and the imaginary parts should not be too different, the latter because of the uniform distribution of regular dipoles in $(\psi_x - \psi_y)$ phase space. This (naive?) guess on my part turns out to be not quite correct according to the latest result obtained by Fritz Dell:

For mac95/storage multipoles,

	real part	imaginary part
seed 0	0.0217	-0.3768
seed 1	-0.2682	-0.0413
seed 2	-1.0909	0.3910
seed 3	-1.5805	0.3100

Perhaps it is premature to claim any contradiction from only four random samples, hence the qualifying term "minor" attached to "questions".

7. Concluding Remarks

For many issues in accelerator physics (as in politics, religion, ...), it is not easy to take a balanced view especially as one gets older. Worse still, one's prejudice based on the past experience tends to be fossilized over the years, ignoring the ever-present advance in technologies. The following view is no exception to this sad fact of life. The only valid statement I can make would be that what you think as members of RAP is far more important than what I think as far as RHIC is concerned. Nevertheless, I cannot resist the temptation to conclude this report with the following observation regarding the priorities during the construction and commissioning period.

1. Mechanical and electrical integrity of magnets.
2. Alignment of insertion triplet quadrupoles (relative to each other but not as a unit).
3. Closed orbit. Find the "golden" closed orbit giving the best performance regardless of what BPMs tell.
4. Tunes and chromaticities.
5. Minimization of tune split using two or four families of correction skew quadrupoles.
6. Minimization of the "cross-plane" closed orbit with a local orbit bump covering one (and only one) triplet at a time.
7. If you are still not satisfied with the machine performance, try more sophisticated methods such as Talman's decoupling scheme provided that you know precisely what you are doing.

References

1. Richard Talman, "A Universal Algorithm for Accelerator Correction", AIP Conference Proceedings No. 255, p. 67 (1991).
2. J.A. Holt, M.A. Martens, L. Michelotti, and G. Goderre, "Calculating Luminosity for a Coupled Tevatron Lattice", to be published in the Proceedings for 1995 IEEE/High Energy Accelerator Conference at Dallas, May 1995.

3. F. Pilat, S. Peggs, S. Tepikian, D. Trbojevic, and J. Wei, "The Effect and Correction of Coupling Generated by the RHIC Triplet Quadrupoles", BNL-61165 (to be published in the Proceedings of IEEE/High Energy Accelerator Conference at Dallas, May 1995).
4. S. Ohnuma, "Low Beta Insertions in RHIC – Miscellaneous Issues", RHIC/AP/38, September 21, 1994.
5. G.F. Dell, "Calculation of Skew Chromaticity in RHIC", an informal note, July 1995.