BNL-99178-2013-TECH

C-A/AP/23;BNL-99178-2013-IR

The RHIC Beam Dump Commissioning - Status Report

L. Ahrens

September 2000

Collider Accelerator Department

Brookhaven National Laboratory

U.S. Department of Energy

USDOE Office of Science (SC)

DISCLAIMER

C-A/AP/23
September 2000

The RHIC Beam Dump Commissioning - Status Report

L. Ahrens
BNL

Collider-Accelerator Department
Brookhaven National Laboratory
Upton, NY 11973

The RHIC Beam Dump Commissioning – Status Report
29 June 00 Leif Ahrens

1) Introduction – the points of this note
2) Results Summary
3) Experimental setup – some background information, the tools actually used, some lessons learned, etc.

1) Introduction

The RHIC beam dump is required to be reliably removing the beam from the RHIC rings, both at “store”, i.e. the top energy and anywhere along the magnet ramp. The beam wants 1) to be removed from the ring, 2) to not hit any of the experiments, and 3) even to not significantly hit any superconducting magnets. The first and second are apparently straightforward. At highest energies and intensities, if I understand correctly the reports, the third may not even be possible.

The commissioning – and here that means the "with beam" work – has concentrated on first establishing the timing for the dump trigger, and then making a measurement of the kick given by the dump kicker to the beam. Our goal is to have the dumps "on line" at the design strengths for the FY00 run, and then to study the loss patterns associated with beam dumps as the beam intensity increases.

This note reports the results from the one dedicated dump commissioning study that occurred on 4 Jun 00 during the day shift. The participants were: for Operations Lee Hammond, for the dump kicker system Jian-Lin Mi, and commissioner Leif Ahrens.

2) Results

Timing the blue dump kicker-firing trigger: We find experimentally that a setting of 240 counts on the coarse pulse delay (with 1 count on the fine delay) results in the bunch coming just before the start of the current (or magnetic field) pulse in the kicker. Both the bunch and the kicker trigger are “locked” relative to the "rev tick" on the blue beam synch link. We then reduce this coarse delay by 20 counts to 220. This moves the kicker pulse earlier causing the bunch (which is the first bunch after the abort gap) to come 1.4 us farther up the waveform, which is 400 ns more than should be necessary – which we will revisit when we have a smaller abort gap. This would need to be snuggled down to about 226 counts if the abort gap shrinks to 1 usec. A revisit of the action – looking at the oscillations with a full ring as we kick – would be appropriate. If turn-by-turn of the bunches just before and after the abort gap, for the last few turns of beam were part of the abort postmortem – this would give us information relevant to this.

Timing the yellow dump kicker-firing trigger: We find (somewhat to our surprise) that once again a coarse delay setting of 240 – with the fine control at 1 – allows the bunch to just beat the start of the kicker rise. So this delay also gets set to 220 counts, following the same reasoning as above.

Presumably the observation that blue and yellow want the same delay relative to the appropriate rev tick has to do with 10 o'clock (where the dump timing and electronics live) being just opposite 4 o'clock (where the beam synch clocks are generated). The beam travel path is the same for blue as for yellow, and the rev ticks in each ring have the same fixed relation with the first bunch at 4 o'clock in that ring.

The strength of the kick is measured at injection rigidity, for the center module of the Blue kicker over the low end of its voltage range, using the beam position monitors (bpm) located 9 meters and 31 meters downstream of the kicker. The result is the following. To obtain the design 1.6 mrad at the "dip" in the kicker waveform, one would need about 13 kV in this single module. The highest voltage applied during the study was 5 kV. If we just quietly do the extrapolation linearly, and let one module go to five modules, and let gamma 10.25 go to 70 or 109. **W e need 17.8 +/- 1.4 kV at gamma 70 Gold, and 27.6 +/- 2 kV at 100 GeV/nuc Gold. (gamma = 108.4).** These numbers are about 16 % below the book expectations. They depend directly on the calibration of the two bpms used. The error estimates come primarily from the disagreement between these two bpms, with the assumption that they are accurate on average.

3) Experimental Setup

timing

The timing of the trigger for each dump is supposed to come such that the "abort gap" in the train of pulses in that RHIC ring will coincide with the rising edge of the dump kicker magnetic field. We assume that the kick the beam gets passing through the kickers at a particular moment is proportional to the current measured at the magnet at that moment. From the measured current pulse, the "rise time" for the kickers is about 1 usec, corresponding to the time available with 56 of 60 equally spaced buckets occupied. Rise time in this context means the time required for the field to rise from a value giving a kick which keeps the bunch in the ring aperture to a value high enough to kick the entire beam past the edge of the sweet spot on the dump – nominally a 1.6 mrad kick. Unlike most kicker PFNs, these explicitly strives not to have a flat top, but rather to swing from a 1.6 mrad kick up to 2.6 mrad, then back to 1.6 and then up one more time (as the 56 bunches come through) before slowly returning to zero.

We work with a single bunch in RHIC, and at injection energy (gamma 10.25), and start in the Blue ring. We initiate the sequence that fires the kicker only a handful of turns after the bunch is injected, by tying the dump event (ev-bdump) to ev-bfebbunch, (the event most closely associated with blue injection) with 100 usec of additional delay. We measure the fact that the beam has been kicked - and perhaps removed – using the standard turn-by-turn orbit display hardware and software (Blue Orbit application in injection mode). The dump trigger electronics (V125 module) takes the dump event and process it into a trigger for the dump power supply synchronized with the beam bunch in the ring. The electronics include a variable delay relative to the "revolution tick" (also called “rev tick” in the following) on the Blue Beam Synch Link. The delay, there is both a coarse and a fine control, is an ado parameter, available on the V125 pet page. Fourteen ticks of the coarse delay move the kicker by about 1 usec.

By varying this delay, and injecting and dumping a bunch over and over again, one can move the bunch systematically with respect to the kicker. This can be "seen" out in 1010 if one of the signals from the bpm after the kicker is routed onto the scope displaying the kicker waveforms. This was done as a preliminary part to this study, but was not used in the work described here. The connection to the bpm is not done in a way that allows the connection to remain in place. The utility of the bpm for position measurement is destroyed while the connection is in. A more robust connection would be possible with fancy couplers.

If the bunch passes the kicker magnet just as the current begins to rise, one observes that the bunch survives past the dump, but oscillates rather significantly going around the ring. On the next turn, the (now larger) kicker field removes the bunch completely. Moving the kicker pulse (now with a fresh injection of beam) just a bit later results in no oscillations on this turn, but still complete removal on the next. Moving the other way with the timing causes increased oscillations, and a rather small timing change in this direction results in the beam being removed altogether on this first turn. The measurement is rather sensitive provided the leading edge of the kicker pulse is sharp – as the current transformer claims. In the Blue ring with this setup, and a coarse delay less than 220 ticks, the beam would disappear on turn 13. (The beam came in on turn 5. These are "turns" as set in the Orbit program.) With a setting of 240 ticks, the bunch clears the rising edge of the kicker – we see no extra motion on this turn, (and lose the beam on turn 14). At 235 there is clear oscillation of the first turn.

The bunch we were using is the "first" bunch injected, and will maintain its position with respect to the rev tick on the beam synch link as more bunches – even 55 more - are added. Additional bunches are added adjacent to but after this bunch. The abort gap occurs just before this bunch. So we want this bunch to come just as the kicker field gets up to the necessary value on the rising edge of the kicker. Given that we know the timing to put the bunch just before the kicker begins to rise, and that the rise takes 1 us (from the current waveform), we want to move the kicker early by 1 use or about 14 ticks from our measured setting. This gets us to less than 226 ticks but more than 221 ticks. We set this to 220 ticks for now, because the abort gap is wider than 1 usec now. Moving the first bunch farther up the waveform is conservative, until we have to worry about hitting beam on the other side of the gap.

Exactly the same exercise was carried out in the yellow ring. The events involved just change from blue to yellow. The drill was worked out in Blue, and applied rather quickly in Yellow. We had not thought about whether the answer could be predicted in advance from the Blue result. Last summer we think the answer was no. Anyway, we were a bit surprised to see at the end that we had moved to exactly the same delay for

Yellow as for Blue. I am hoping this is a satisfactory observation rather than a coincidence.

strength

This work has only been done in the blue ring. Indeed only the middle kicker module, one of five identical units, has been measured. The commissioning plan, apriori, was to kick the injected beam sufficiently gently that the bunch would not scrape on the dump but continue to circulate in the ring. The ring bpms would then give a good redundant measurement of the kick strength, which could be translated to motion at the dump knowing the relative beta functions. Along the way we wanted to see if the bpms in the straight section between quad before the kicker and the quad after the dump would give consistent position information. We were pessimistic because we see large (relative to the bunch-induced wiggle) signals on the cables from the close in bpm when the kicker is fired.

This strength study adjusted the timing of the kicker relative to the bunch such that the bunch came at the "dip" in the kicker current waveform. The amplitude of the kick at the dip is about the same as on the leading edge just where we want to start injecting. With this timing, we reduce the kicker voltage as much as we can without losing solid triggering. At 2 kV we fail to trigger. We increase the voltage in .5 kV steps, saving the turn-by-turn bpm data in each setup. At 5 kV we note that no beam survives beyond the dump. Since we were primarily interested in the ring data, we did not go to higher voltages. In retrospect because the close in bpms look very good, more points at higher voltage would have been useful. In fact, given that the response of the bpms located 9 and 31 meters downstream of the middle kicker appear to be linear with the kicker voltage setting, it is that data which we have concentrated on and discuss here. The ring data, over the small voltage range possible (limited by the dump aperture on the one side and the kicker lower voltage on the other) is all very large oscillation stuff - +/- 20 mm – and so may not be so nice as far as bpm linearity is concerned.

Detailed analysis:

We want to do some calibrating. The parameter being varied to change the kicker amplitude is the setting for the waveform generator (WFG) providing the reference to the charging supply. In addition we measure a digital voltage readback in building 1010a for each WFG setting. We have the current waveforms (transformers on each of the PFN outputs) measured using scope traces for all the voltage settings. We assume that these currents are proportional to the field in the magnet and the kick given the beam. In particular, we assume zero current in the transformers means zero field and should imply zero voltage. We define our reference kilovolt from the 1010a meter, but with an offset such that zero voltage gives zero current. The three numbers – reference setting (here multiplied by 100, voltage readback, and current amplitude –show good linearity (figure 1). To get consistent zeros the voltage reading in 1010a must be reduced by 90 Volts (figure 2 gives the situation before this correction). The subsequent analysis of the bpm data is done against this adjusted voltage scale.

For the kicker calibration, we concentrate on the two bpms in the field free straight section downstream of the kicking magnet. Knowing only the beam position shift at these bpms (and the distance between each and the kicker center) gives redundant measurements of the kick angle for a given voltage setting. We take these position shifts for several voltage settings – though over only a small part of the available voltage range. The relevant bpm data is 'difference' trajectory data. If the bpms showed no motion turn-by-turn for no kick, simply taking the difference between the kicked turn and the previous turn would yield the desired data. We are working shortly after injection, and there is significant coherence in the data. There is visible turn-by-turn motion even with no kick applied. (Delaying to the end of the 128 turn data is not a big help in reducing this coherence – our machine has small chromaticity). To minimize the degradation to the data from this variation, we select "by hand" a turn trajectory which nearly overlays the bpms upstream of the kicker on the kick turn. Then we subtract this selected turn's data for the bpms downstream of the kicker. The position motion is converted to a kick angle knowing the distance to the kick (taken to be the center of the magnet being powered). This angle is plotted against the voltage defined above. Figure 3 gives the results for the two bpms.

Finally we make one extension to a more relevant voltage regime. The voltages reported by the meters in building 1010a are measured as the references sent to the WFGs are changed. The data is plotted in figure 4. Here we are back to raw data. If we take the meter as truth, the two systems are behaving differently – at the 4 percent level. A good shutdown exercise is to repeat this keeping track of the amplitudes of the resulting current pulses. The four voltage points from the bpm study are also shown in figure 4. Voltage measurements at the same WFG settings agree well. It appears there is a slight deviation from linearity between the two voltage ranges.

For what is left of the 2000 run, we intend to set the WFGs identical for both rings and varying proportional to the rigidity RTDL frame such that at gamma = 70 the voltage will be about 18 kV. We set the WFG to output 18 at gamma = 70.

Figure 5 gives one other measurement taken during the study while locating the dip in the waveform. We are in the “small kick” configuration. The two bpms in the kicker drift region are analyzed as mentioned above. The timing of the kicker is varied relative to the bunch. This is the “first turn” position change. It is direct evidence that for the shape of the kicker in time – which agrees well with the current pulse. If the positions plotted for the two bpms overlaid here, they would agree perfectly. The close in bpm result is scaled up by the inverse of the relative distance from the kicker center. The downstream bpm (bh4) fails to report when the amplitude approached 20 mm.

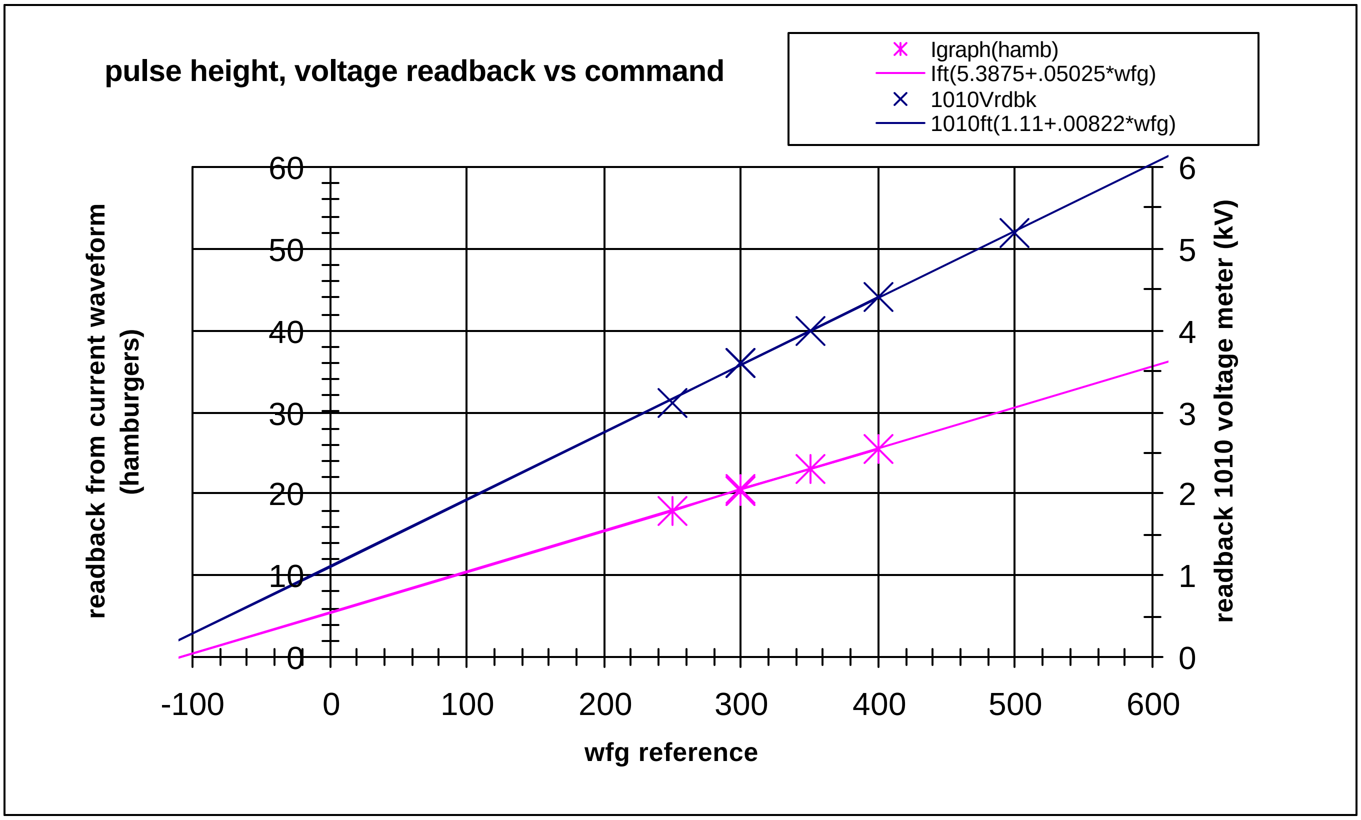

Figure 1 Calibration of the Pulse Forming Network Voltage

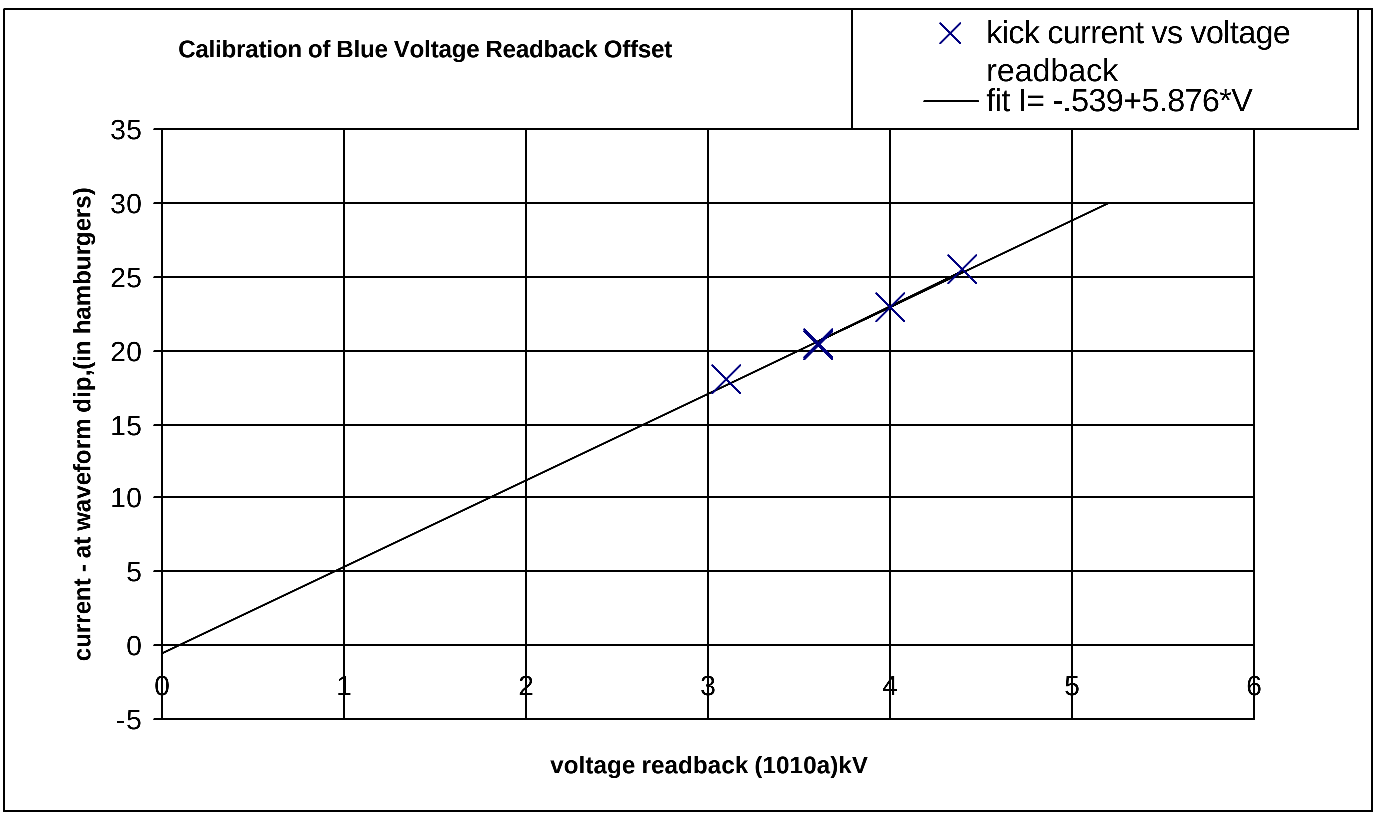

Figure 2 Voltage Offset determination

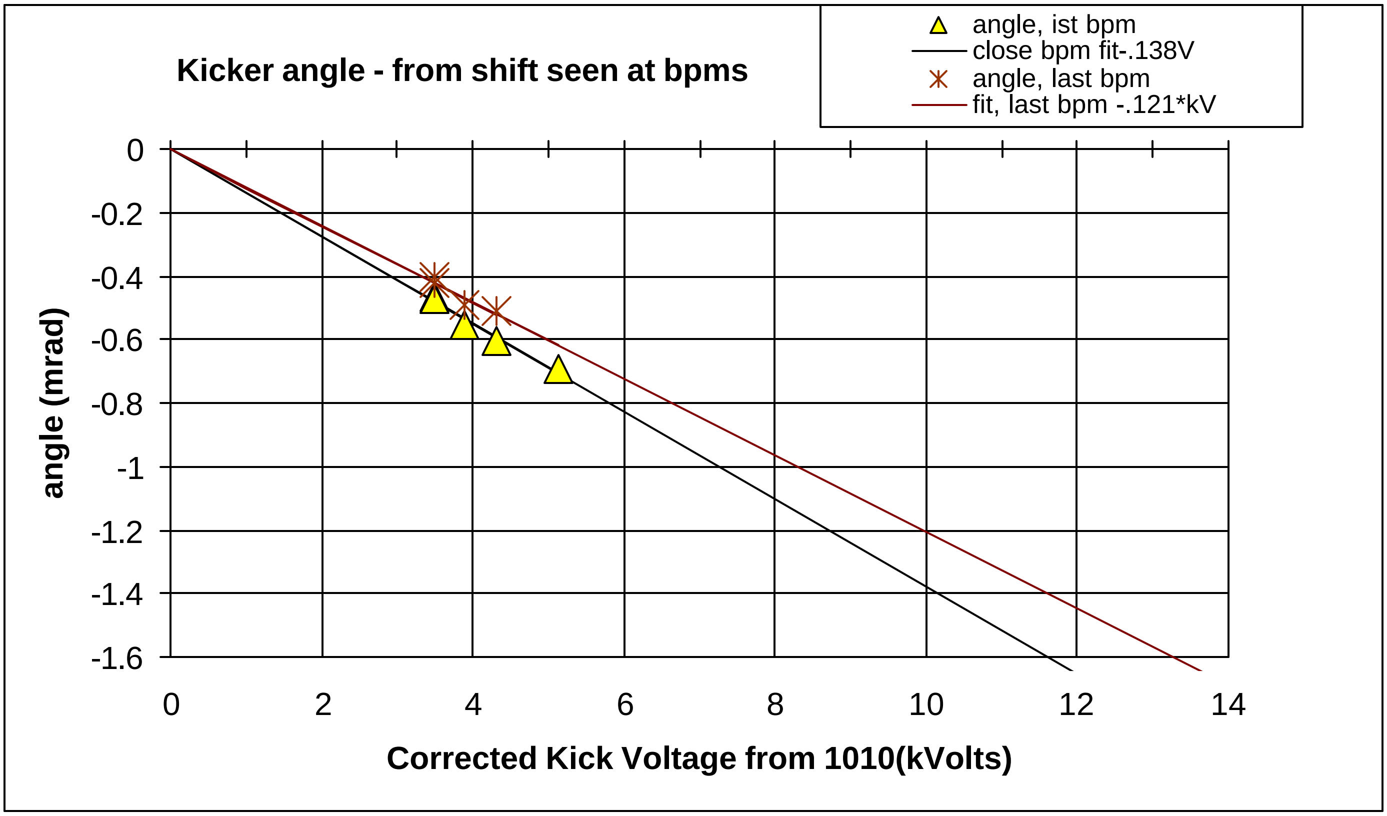

Figure 3 Kicker Angle vs Voltage Applied to the PFN

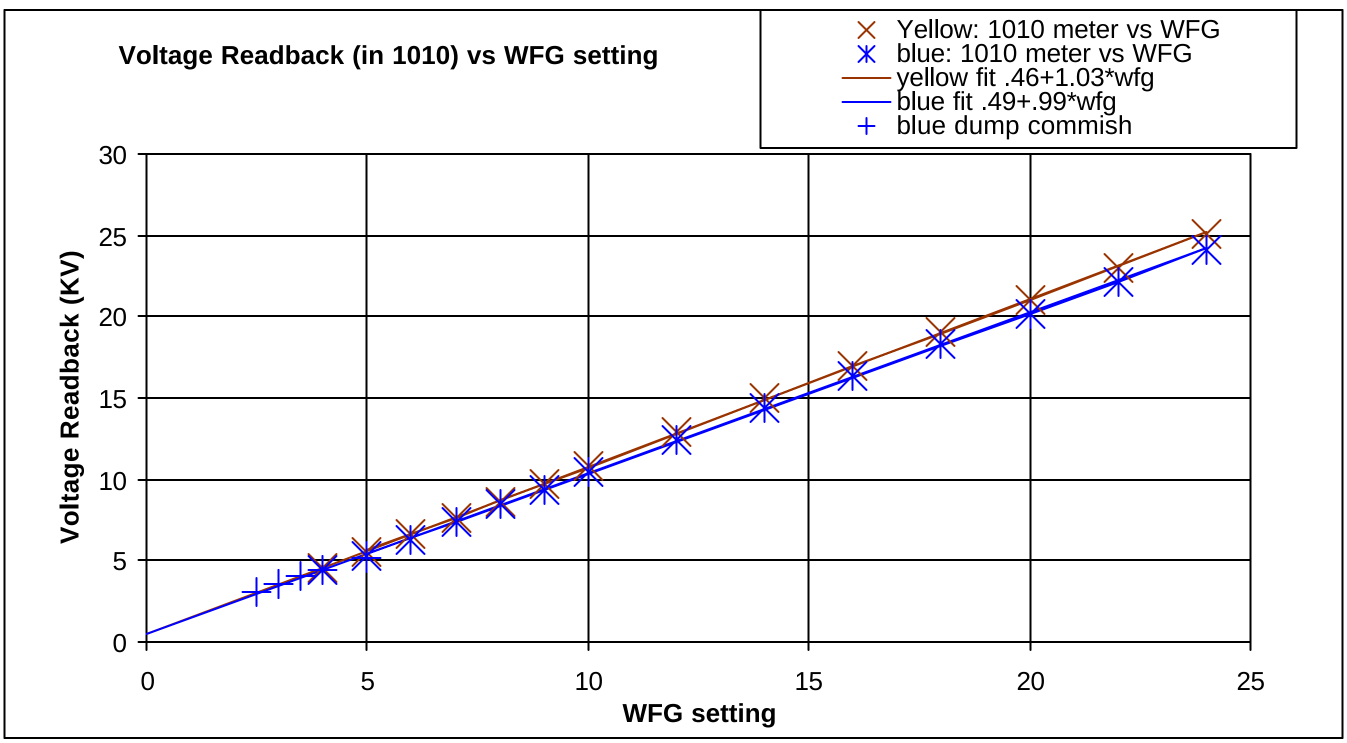

Figure 4 Raw Voltage Readback vs WFG Setting (taken after study)

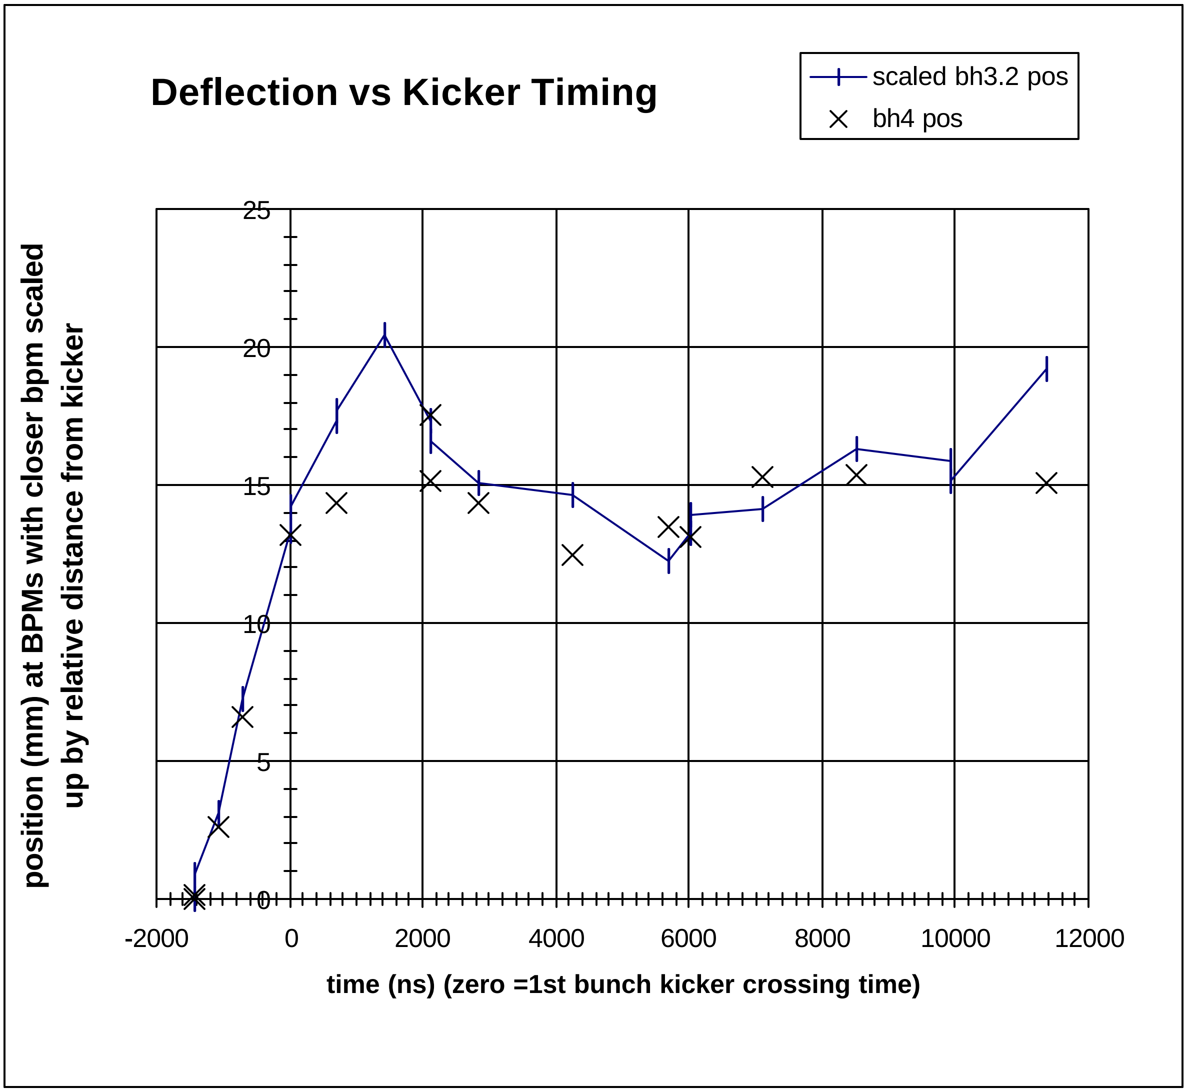

Figure 5 Blue Kicker Timing Scan