Beams-doc-1925-v1
August 2, 2005

TeV BPM Upgrade: Stability and Calibration

Rob Kutschke, CD/EXP

Abstract

This note shows data that illustrate the stability of the TeV BPM measurements with time: the proton measurements are repeatable to better than 40 μm (RMS) over a period of 6 weeks and the anti-proton measurements are repeatable to better than 50 μm (RMS) over the same period. These numbers are upper bounds and the stability of the measurements is probably better than this. This note also shows that the information about the differential attenuation in the analog portions of the system is not useful; improved measurements of the differential attenuation in the cables are needed. Finally this note discusses an anomaly that was reported on VC11; the anomaly is real and I have not yet found a fix.

1 Introduction

The data presented here were taken from 32 HEP shots taken between between June 18, 2005 and July 29, 2005. The data were acquired and analyzed as described in Beams-doc-1863. This document uses the notation defined there.

Figure 1 shows a scatter plot of the size of the proton helix vs the size of the anti-proton helix. The data from all 32 shots is shown at once. In the notation of Beams-doc-1863, this is $p_H - p_C$ vs $\bar{p}_H - p_C$. It sometimes happens that, at a given BPM, the helix opens mostly in the unmeasured coordinate and only slightly in the measured coordinate. In this case, there is little lever arm with which to determine the cancellation coefficients and the anti-proton position is not reliable. Data which might be unreliable in this way are shown in blue; to be specific, the blue points have $|p_H - p_C| < 0.25$ mm.

The data for VF29 are shown in red and appear twice. The points at the lower left were analyzed the same as was all other data points and the position of these points suggests the that the VF29 anti-proton cables were crossed. The second set of red points, labeled "VF29 fixed", show the same data reanalyzed with the cables uncrossed in software. These 32 data points fall within the main body of the data; moreover they fall so close together that they appear as a single large data point. Marv has checked this BPM, found that the anti-proton cables were indeed crossed and has fixed it. In the rest of this document, only the repaired VF29 data will be shown.

The green data points are for VC11 and will be discussed in Section 7.

If one looks carefully at this plot, one can see that each "point" is really a cluster of many points spaced closely together.

Figure 2 shows the data from within the yellow box in Figure 1. In this figure the clustering of the data is more evident. Each cluster corresponds to the 32 data points for a single BPM. The spread of each cluster contains contributions from,

1. The resolution of the BPM system
2. The stability of the BPM system
3. Any real changes to the Tevatron orbits.

The bottom line is that, once the motion of the central orbit is removed, the whole system is remarkably stable. See the discussion of Figure 6 for a more quantitative picture.

The red triangle in Figure 2 shows the data for VC11 for a particular run that will be discussed in Section 7.

2 Outliers

Figure 3 shows the same data as Figure 1 but on an expanded scale and with with the blue and green data points suppressed; those data points are unreliable and sometimes have very large outliers. The scale on this figure is large enough to include all other outliers. The red box shows the region included in Figure 1. The only points outside of the red box are the handful at the far left, from HF36. In this plot, all data points from HF36 are colored red. Note that most of the HF36 points fall within the main body of the data but that a few lie exactly at (0,0); those points have I and Q values of 0 on all channels. I have not yet investigated why HF36 has this odd behavior.

3 Self Consistency

An interesting measure of the self consistency the whole system, Tevatron plus BPMs, can be obtained by projecting the data in Figure 1 onto the diagonal dashed line. This line will be referred to as the "self-consistency axis". Following the notation of Beams-doc-1863, the projection is given by,

$$\Delta^+ \quad = \quad (p_H - p_C) + (\bar{p}_H - p_C) \tag{1}$$

The top plot in Figure 4 shows a histogram of Δ^+ for all data from the first shot in the data set. The lower plot in Figure 4 shows the same for the second shot. For each of these histograms, one can compute the mean and RMS. This process was repeated to give the mean and RMS for all 32 shots in the data set.

The RMS width of these histograms contains contributions from,

1. Imperfect cancellation of the proton contamination on the anti-proton cables. Note that the proton helix is measured without anti-protons in the machine.

2. Biases in the BPM system, which come from ignoring higher order corrections and orthogonal plane corrections.

3. Measurement resolution in the BPM system.

4. True differences in the helix sizes.

5. A change in the central orbit during anti-proton injection.

It is believed that the last two effects are very small. The measurement resolution is known to be at the level of 10 μm. Beams-doc-1893 suggests that the magnitude of the biases within the BPM system are at the level of a few hundred microns. This leaves imperfect cancellation as the most likely explanation for the observed RMS width. If we presume that this is indeed the main contribution, then the RMS of these histograms is a direct measure of the bias due to imperfect cancellation.

The upper plot in Figure 5 shows a time series of the mean of the 32 Δ^+ histograms. There are no significant excursions from zero. The error bars are given by RMS/$\sqrt{N}$, where N is the number of measurements that contribute to the mean, usually about 215. The number varies because, from shot to shot, a handful of BPMs wander in or out of the requirement that $|p_H - p_C| > 0.25$ mm. If the effects that cause this bias are uncorrelated from BPM to BPM then, in a large number of BPMs, Δ^+ should average to zero. This is indeed observed.

The middle plot in Figure 5 shows a time series of the RMS of the Δ^+ histograms. The data is stable around 0.67 mm. Note that the first two data points here are slightly lower than the RMS values printed in the previous figure. That is because the values shown here only include data from the black points in Figure 1, while the values printed on the previous page include all data. The error bars are given by RMS/$\sqrt{2N}$.

The bottom histogram shows N, the number of BPMs used to compute the mean and RMS for each shot. The data from the 13th shot, on July 7, had many BPMS which reported zero for all data. There was a smaller glitch on July 21, point 23. I have not looked into these data.

The conclusion of this section is that the data are remarkably stable. There are not enough statistics to get a good number on the overall reliability of the system. All that we can really say is that a substantial fraction of the system was up 31/32 times, which corresponds to more 91% availability at the 90% confidence level.

4 Stability

These data can be looked at another way. Consider one of the clusters of data in Figure 2. Each of these clusters shows the variability at one BPM over the

32 shots. For each BPM we can compute the RMS spread of the 32 points on both the proton and anti-proton axes.

The top plot in Figure 6 shows a histogram of the RMS width of the proton helix at a given BPM for each of 219 BPMs. All BPMs shown in black in Figure 1 are included. The mean value of this histogram is some sort of overall stability metric and it has a value of about 38 μm. As mentioned above, the stability metric includes contributions from the BPM resolution, the BPM stability, and true changes to the helix size. Previous studies have shown that the best resolution obtained by the BPM system is about 7 μm but that some BPMs have resolutions as poor as about 20 μm; the explanation is that poorer resolutions are caused by the "15 Hz" beam motion. So most of stability metric does not come from the intrinsic resolution of the BPMs. We can interpret the 38 μm as an upper bound either on the 6 week stability of the BPM system or on the true changes to the helix over this time.

The bottom plot in Figure 6 shows the results of the same exercise but for the anti-proton position. It has a stability metric (mean value) of about 48 μm. The difference, in quadrature, 48 and 38 μm is about 29 μm. Why is the mean higher for the anti-proton histogram? The most likely answer is that this measures the quality of the cancellation coefficients. Although other explanations are possible.

5 Absolute Positions

Figure 7 is provided for completeness; we don't learn a lot from it. It shows the measured proton position vs the measured anti-proton position for all BPMs for all 32 shots. This differs from Figure 1 in because this shows absolute positions, not helix sizes. The main features of this plot are:

1. We see that, at a few BPMS, the helical orbits are very far from (0,0). (At the last minute I realized that the survey offsets are not in this plot.)

2. Each small cluster in Figure 1 has been smeared out along the main diagonal. This is due to changes in the central orbit during the 6 weeks.

6 Differential Attenuation

We should be able to improve the position measurement by calibrating out the effect of differential attenuation in the analog signal path, as seen by the A and B signals. The path includes,

1. the cables from the tunnel to the houses,

2. the filter board,

3. the Echotek board (the matching transformer between the connector and the digitizer).

We have measurements for all of these quantities and they have been entered into a database that can be read both online and offline.

The self-consistency plots can be used as a metric to see if making these corrections is useful. It is a rather coarse metric but it's what we have.

In the association figures I sometimes used the phrase "transimission coefficients" as a synonym for differential attenuation. That's not really right but I am not going to remake the figures.

The top plot in Figure 8 shows the standard self-consistency plot, all of the black points in Figure 1 projected onto the self-consistency axis. The middle plot shows the same data, after correcting for the differential attenuation in the filter and Echotek boards. This correction is small and makes an insignificant change to the self-consistency histogram. This small change is understood because the differential attenuation in these systems is both small and well measured.

The conclusion is that this metric is not sufficiently sensitive to test these corrections.

An aside: when this calibration is applied it modifies the proton helix size by about 1 μm but it modifies the anti-proton helix size by about 50 μm. I have not yet understood why these are so different.

The bottom plot in Figure 8 shows the same data as the first two but after correcting for differential attenuation in all three components. This makes things much worse. Inspection of the cable attenuation data showed that the anti-proton cables were measured with a precision between 5 and 10%. This is much to crude to be useful. Calibration database software is in place so that, when more precise measurements are available, they can be entered into the database and made available both to online and offline applications. If more precise measurements are made, we can revist these corrections.

7 VC11

Around July 19 Randy Thurman-Keup (keup@fnal.gov) reported that at VC11 the anti-proton position moved around a great deal during a store. VC11 is interesting because it is close to one of the synchrotron light monitors. I don't have his original plot in a form that I know how to include in this note but I can reproduce the effect. His original figure is included as a extra pdf file in the DocDb entry from this document. Figure 9 shows proton and anti-proton positions from the lumberjack data logger for VC11 for three stores. In all three stores, the proton position is stable to less than 1 mm but the anti-proton position wanders by as much as 5 mm; even in the best store, the middle of the three, it wanders by more than 2 mm. In Randy's original plot, the anti-proton position wandered by more than 10 mm while the proton position wandered by about 0.5 mm.

The reason for this wandering is that VC11 is one of the BPMs for which the helix opens almost completely in the orthogonal coordinate; as can be seen in Figure 2, the size of the vertical step of the proton helix is between 0.1 and 0.3 mm. So the cancellation coefficients are poorly determined. The red triangle

in Figure 2 marks the point for the middle store in Figure 9. I have spot checked a few BPMs that have large proton helix sizes and, for those BPMs, the anti-proton position wanders around during a store an amount that is similar to the amount that the protons wander around. I have not made a quantitative comparison.

Here is the history of the cancellation coefficients for online positions. For each store, Mike Martens' program records the raw IQ values for all BPMs just before the helix opens, just after the helix opens and just after anti-proton injection is complete. In the following I will refer to the first two of these points as the "calibration points". From this he computes the cancellation coefficients and writes them to an ACNET device. When this ACNET device is written, the front end computers start to use the new cancellation coefficients. This program of Mike's is the same one that recorded the data for the studies reported earlier in this note.

For the plot in this note, I data-logged the IQ data for VC11 and computed the cancellation coefficients myself. I recomputed them for each store, as is done online. I verified that, given the same input data, Mike and I produce the same cancellation coefficients and the same anti-proton positions. The IQ data is data-logged at 1 Hz, which means that, for each store, there are of order 5,000 pairs of points that I could choose as calibration points. I tried about 10 of these combinations for each store and did not see any significant change in the anti-proton position, relative to that in Figure 9. This study probably could be extended to see if there are "magic" points that are better calibration points than others; however I don't think that this has a high probability of payoff.

In Figure 9, the middle store has the least variation in its anti-proton position. So I tried the following: I recomputed the anti-proton positions for the first and third store using the cancellation coefficients from the middle store. As can be seen in Figure 10, this did not help.

The black points in Figure 10 are the same proton position data as appears in the previous figure. The red points show the anti-proton positions for all three shots computed using the cancellation coefficients determined during the first shot. The green points show the anti-proton positions for all three shots computed using the cancellation coefficients determined during the second shot. And the blue points show the anti-proton positions for all three shots computed using the cancellation coefficients determined during the third shot. None of these variations produces a systematic change in the stability of the anti-proton position.

The conclusion is that we need a better scheme for finding the cancellation coefficients for VC11. One thought is to constrain the anti-proton position to be constant during the full period of the anti-proton injection. There are several ways to combine this with the constraint that the anti-proton A and B signals individually be zero when the are only protons in the machine. I have done the math but have not yet started to code this. Another thought is to systematically study the cancellation coefficients at other BPMS, including knowledge of the orthogonal coordinate.

In September we can evaluate if it is worth pursing these or other strategies.

8 Summary and Conclusions

See the abstract.

The outstanding issues are:

1. Can we do anything to improve the antiproton measurements at VC11.
2. When I apply the corrections for differential attenuation in the filter and Echotek boards, the effect on the proton helix size is much smaller than the effect on the pbar helix size. Why?
3. Can we find a better metric to evaluate the correctness/usefullness of these corrections? Or are stuck with dead reconning and test benches?
4. Why are there some stores in which the number of BPMs used to compute the mean and RMS of Δ^+ drops? Is this a real reliability issue?

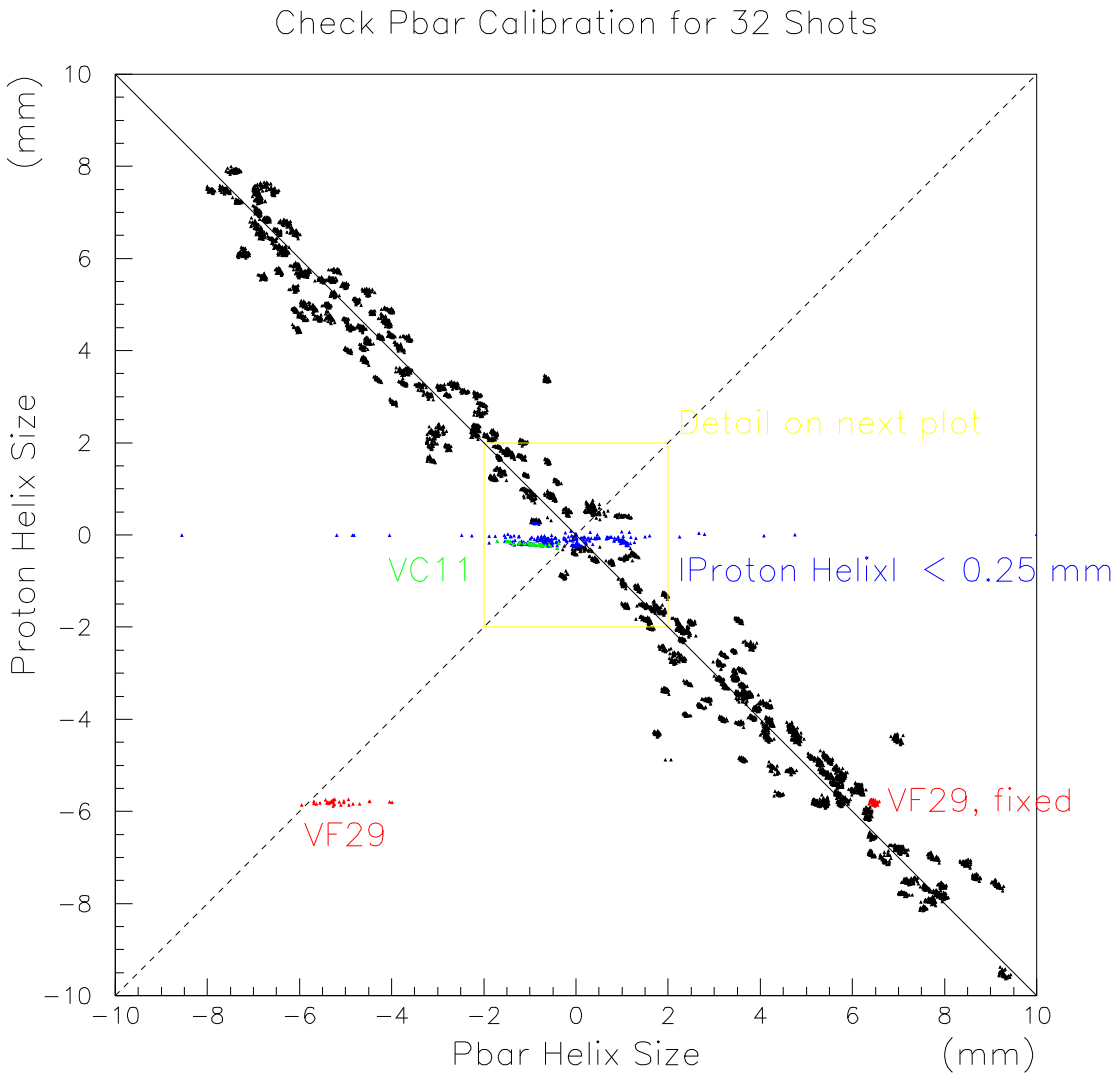

Figure 1: Scatter plot of the size of the proton helix vs the size of the Pbar helix, at 150 GeV. These data include 32 shots. The plot looks better on a screen than on most laser printer output. Both horizontal and vertical BPMs are shown. The black data points show normal BPMs. The blue data points show those data for which the step of the protons, at the opening of the proton helix is less than 0.25 mm. The red data points show the data for VF29, and the green for VC11; these are discussed in the text. The data within the yellow box is repeated on the next page on a larger scale.

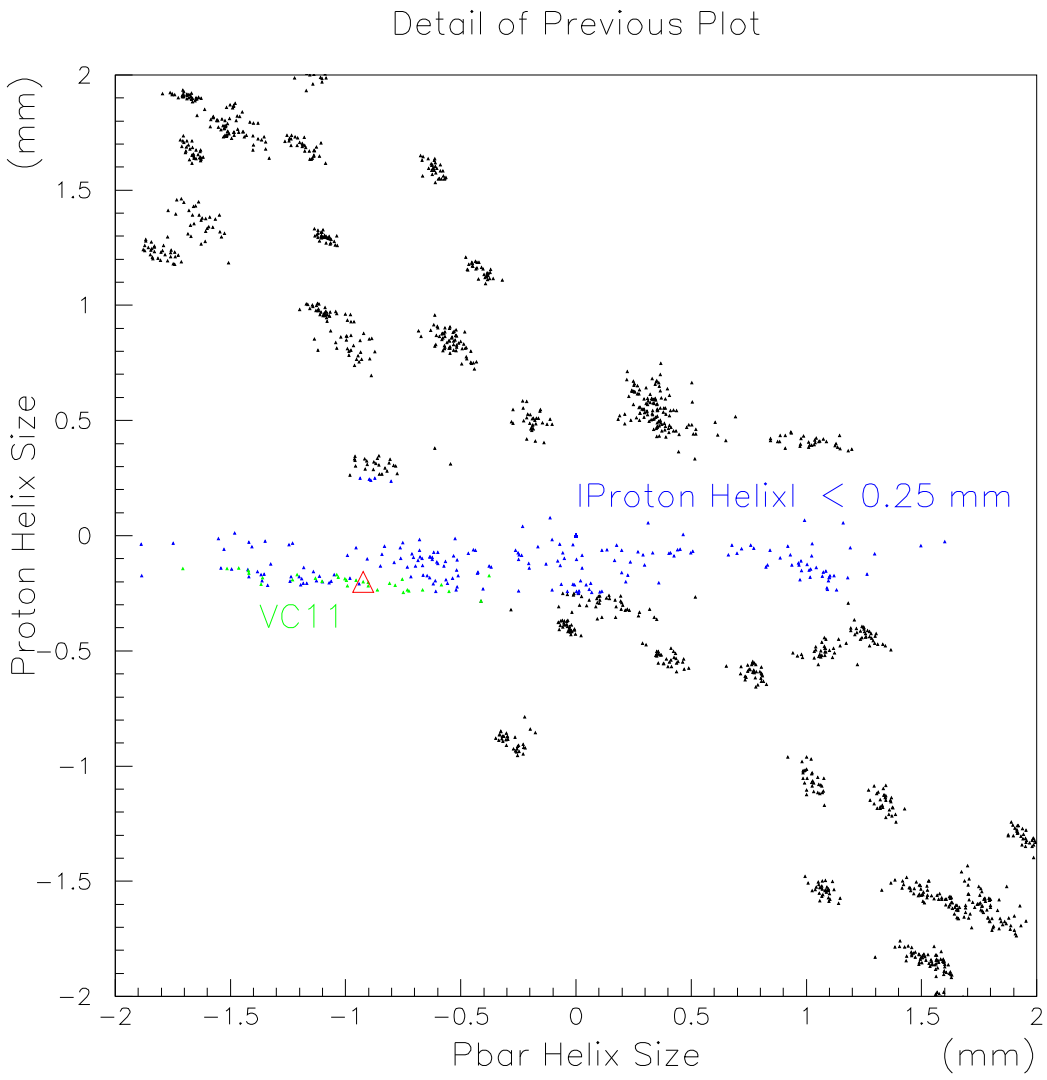

Figure 2: The data from within the yellow box on the previous page. The color coding carries over from the previous page. Note that the data form clusters. These clusters correspond to one BPM over many shots. The red triangle shows a particular shot for VC11; this is discussed in the text.

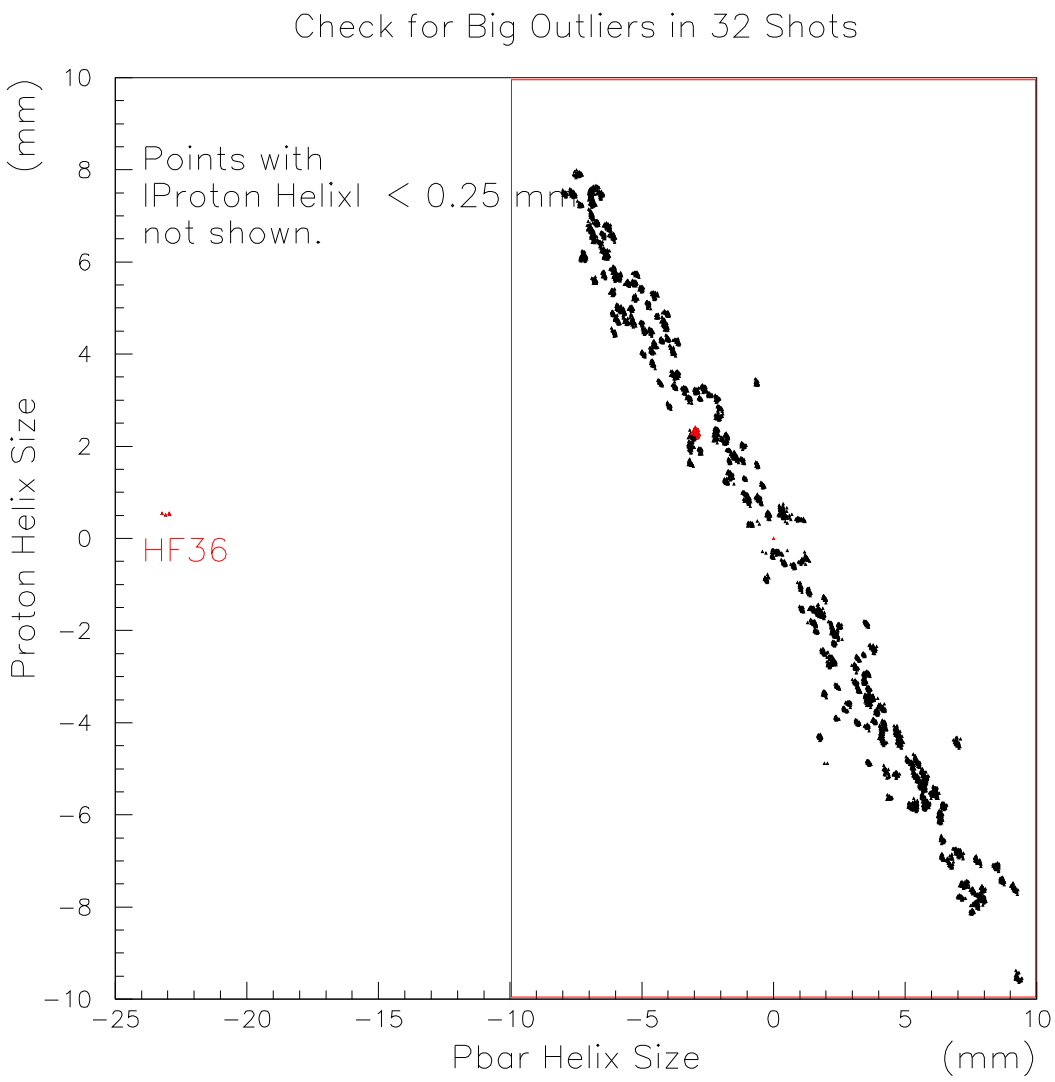

Figure 3: The same data as Figure 1 shown on an expanded scale. The red box indicates the region shown in Figure 1. In this figure, the data points with a proton step less than 0.25 mm have been suppressed but all other data points are shown. The data for VF29 are shown with the crossed cables fixed in software. The only data outside of the main body is a subset of the points for HF36; all data for HF36 are shown in red.

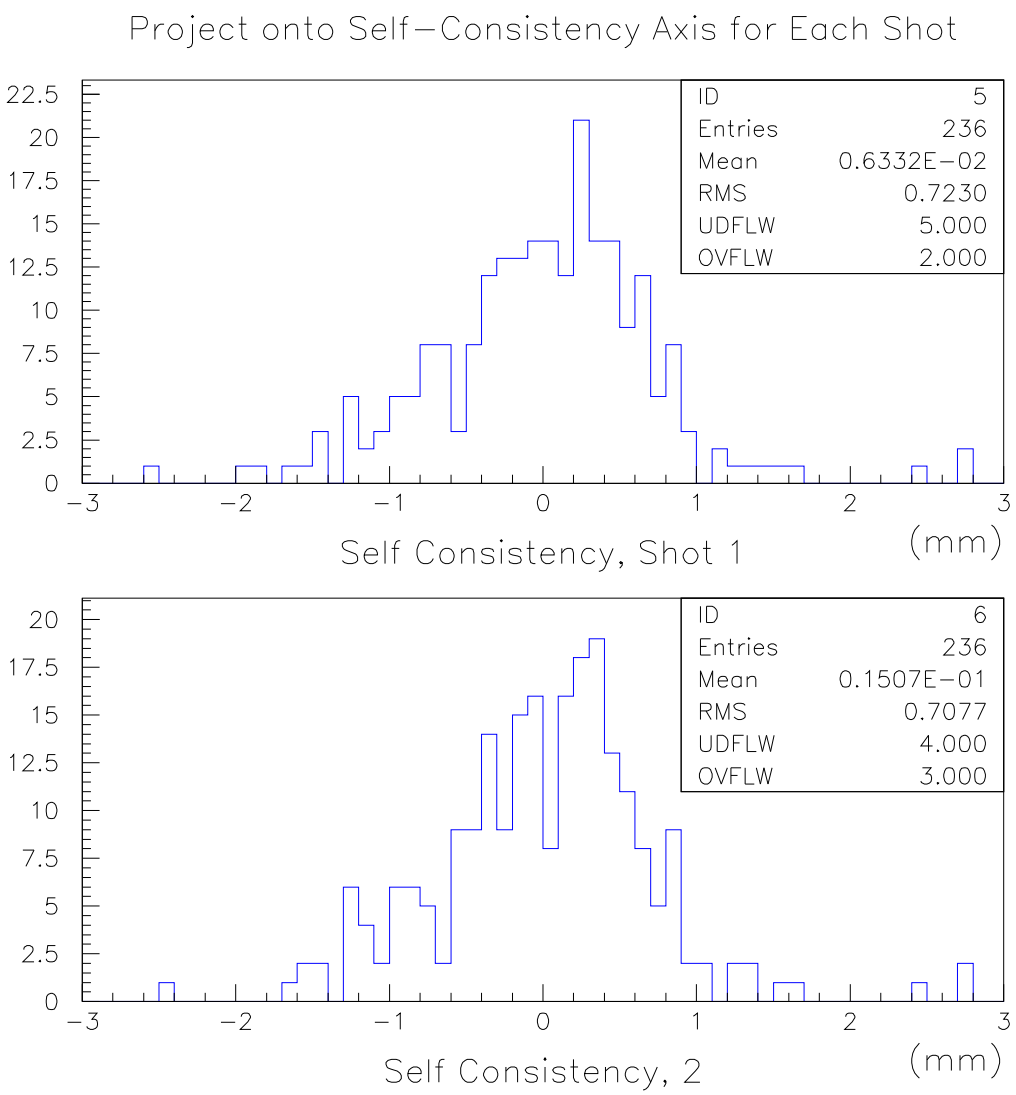

Figure 4: Projection of the data from Figure 1 on the self-consistency axis (the diagonal dashed line in Figure 1). The projection is done separately for each shot. The quantities of interest are the mean, the RMS and the number of entries in these histograms. Timelines of these quantities are shown in the next figure.

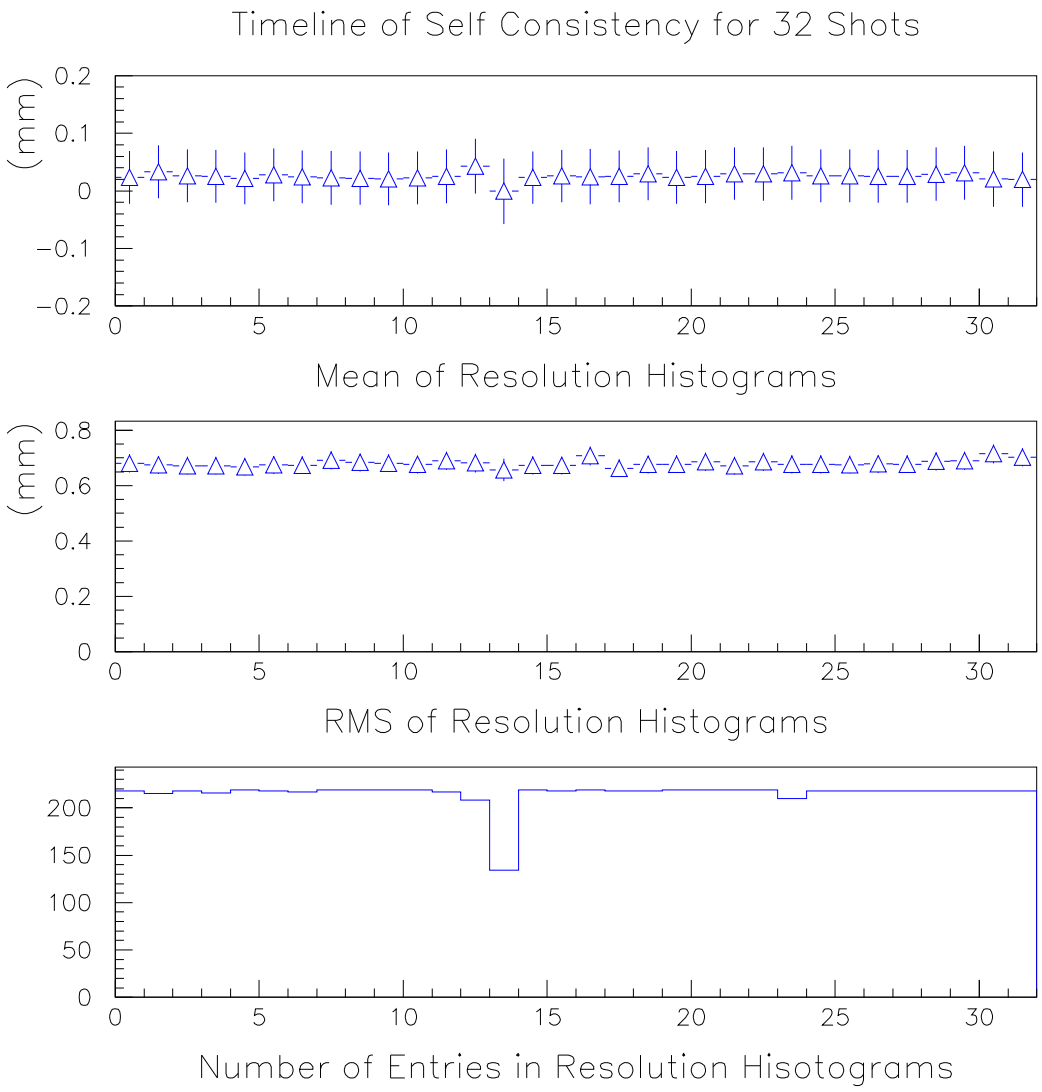

Figure 5: Time series of the self-consistency data for each of the 32 shots. The data in these figures is derived from histograms like those on the previous page. The upper plot shows the mean value of the histogram for each shot. The middle plot shows the RMS value of the histogram for each shot. The bottom plot shows the number of BPMs contributing to the histogram in each shot; for one shot many of the BPMs had sum signals below threshold. This is not yet understood.

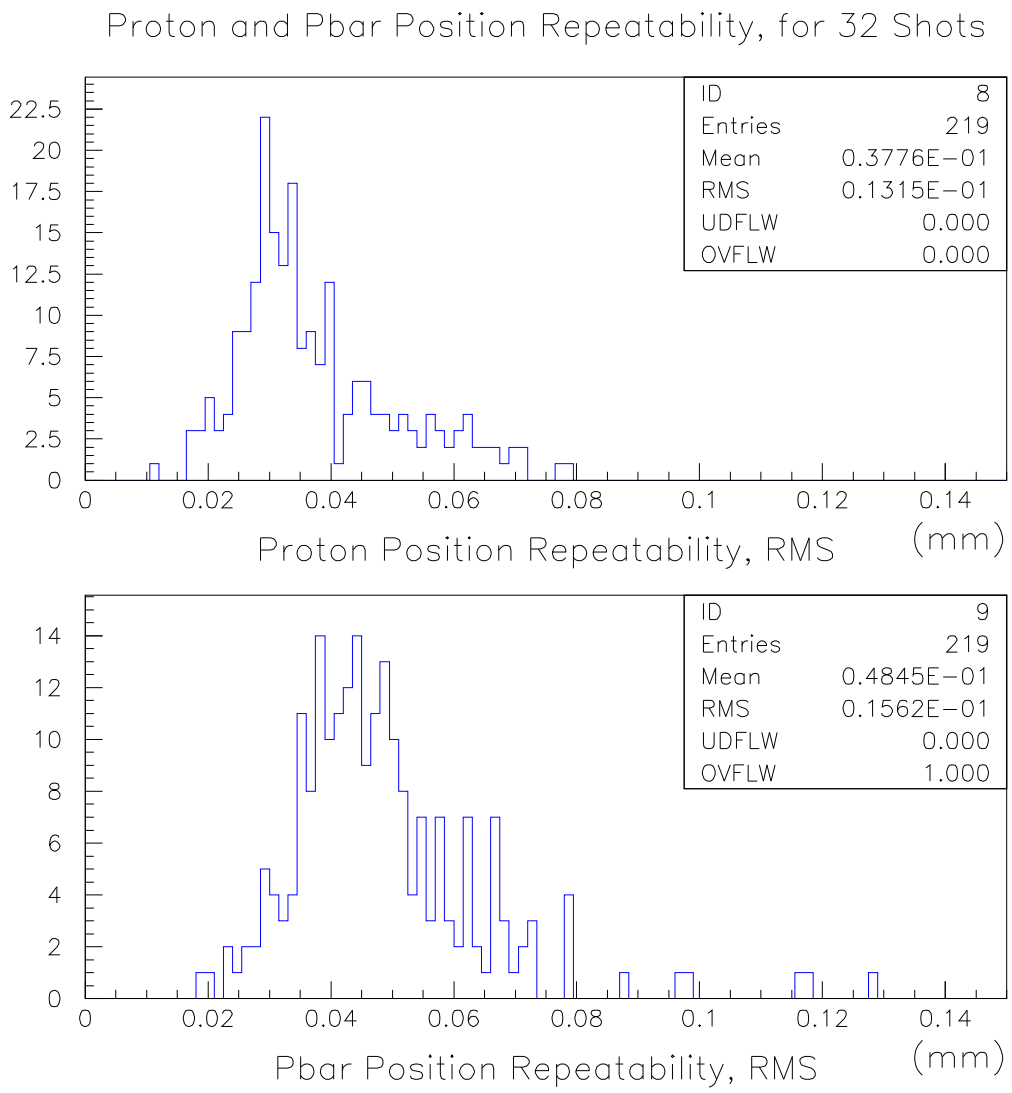

Figure 6: In both plots there is one entry for each of 219 BPMs, those BPMS which have black points in Figure 1. For each of these BPMs, I computed the RMS of the set of 32 measurements of the proton helix size; I also computed the RMS of the set of 32 measurements of the anti-proton helix size. The upper plot is a histogram of the 219 proton helix RMS values; the lower plot is a histogram of the 219 anti-proton helix RMS value.

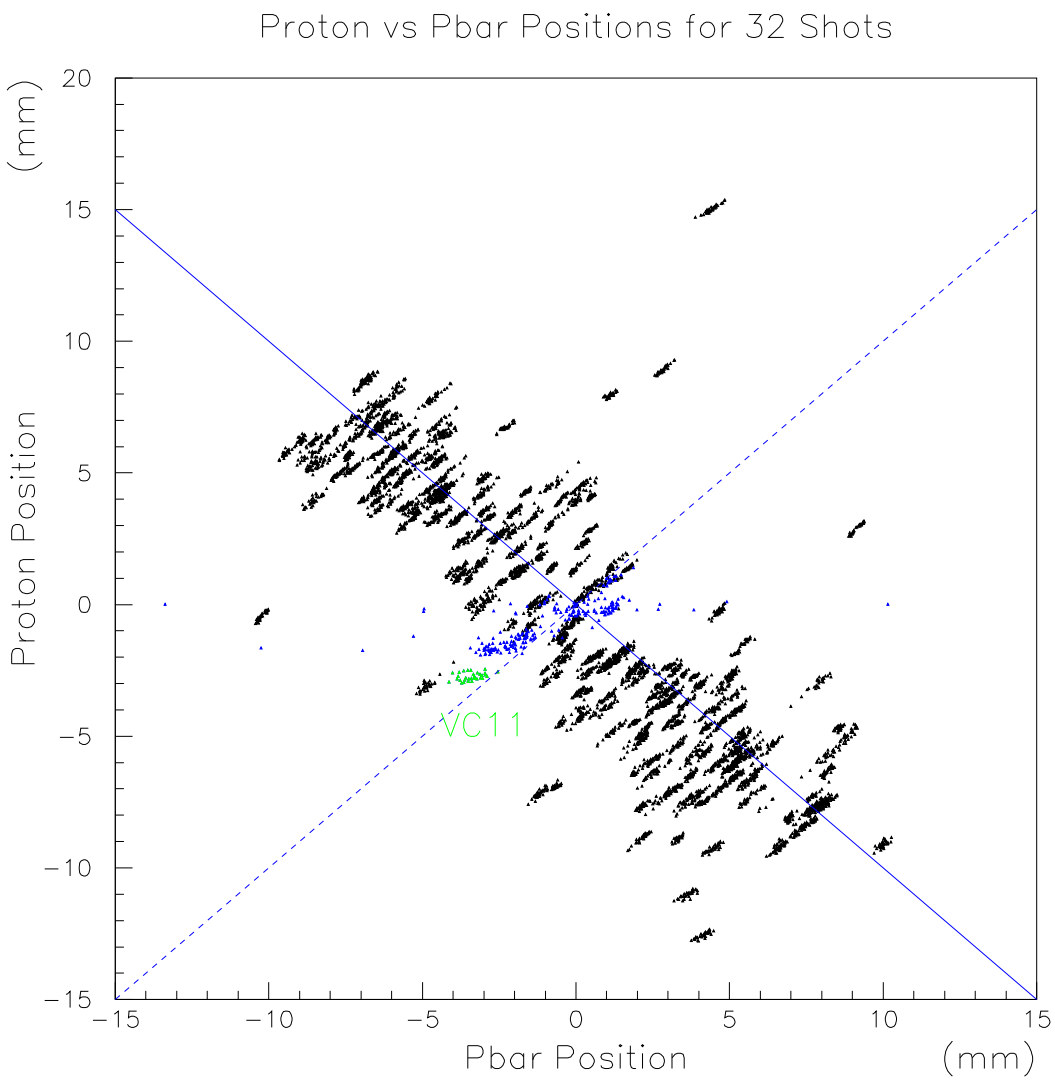

Figure 7: A scatter plot of the proton position vs the anti-proton position, for all 32 shots. One caveat: the positions do not include the survey offsets. The color coding is the same as Figure 1 and only the corrected data for VF29 are shown. In this figure the changes in the central orbit over the course of 6 weeks are visible as the smearing of each cluster along the dashed diagonal axis.

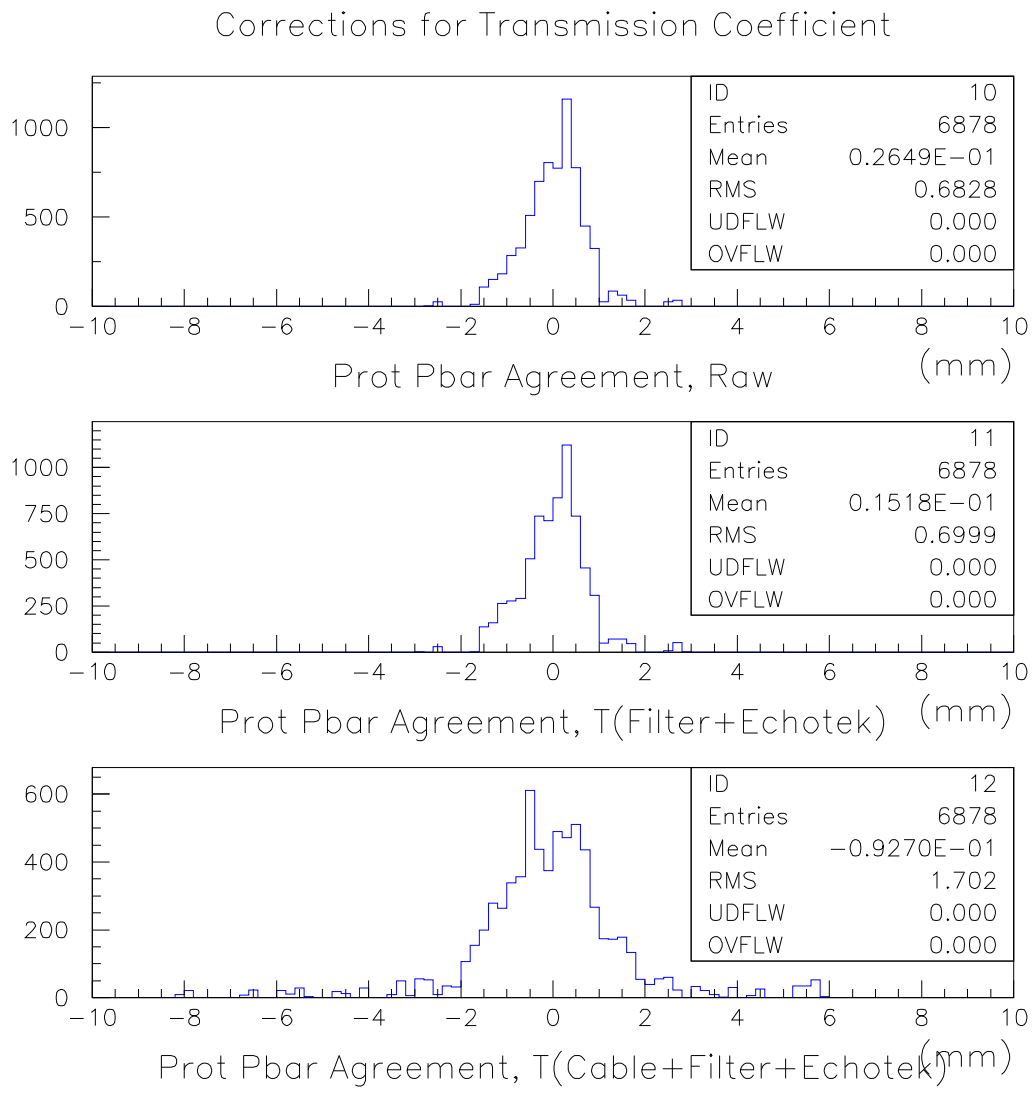

Figure 8: The upper plot shows a projection of a subset of the data from Figure 1 onto the self-consistency axis. It includes data from all shots. The middle plot shows the same data but a correction was applied for the relative attenuation in the filter boards and Echotek boards; this makes little difference because the self-consistency metric is crude. The bottom plot shows the same data but with a correction for the relative attenuation in the cables, filter boards and the Echotek boards. This degrades the self-consistency metric. The explanation is that the relative attenuation in the cables is poorly measured.

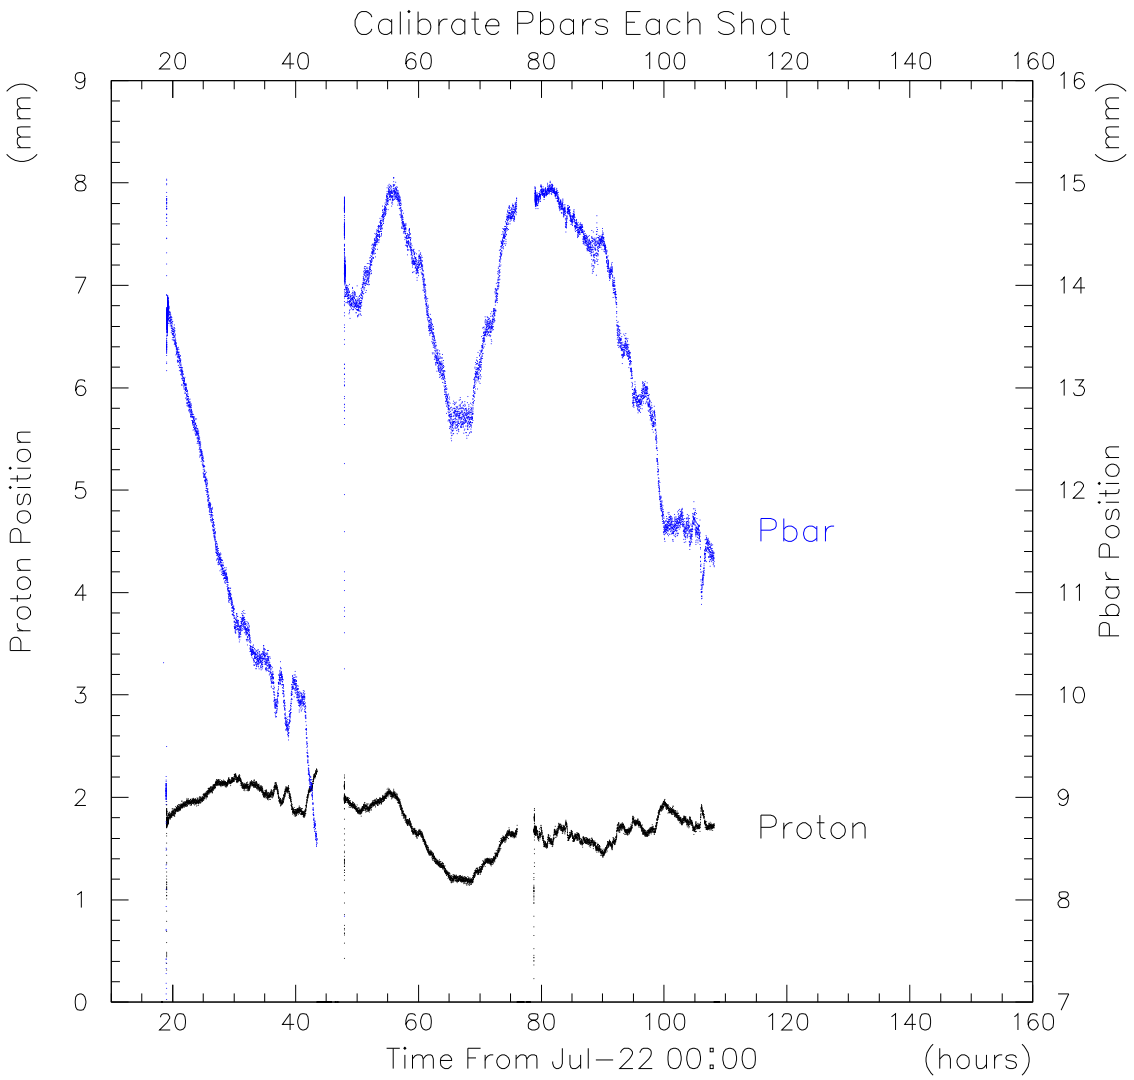

Figure 9: Proton and anti-proton positions from the lumberjack data-logger for VC11 for a period of three stores. The horizontal axis is time, in hours, from midnight on July 22. The full vertical scale is the same for the protons and the anti-protons. The large excursions in the anti-proton position are discussed in the text.

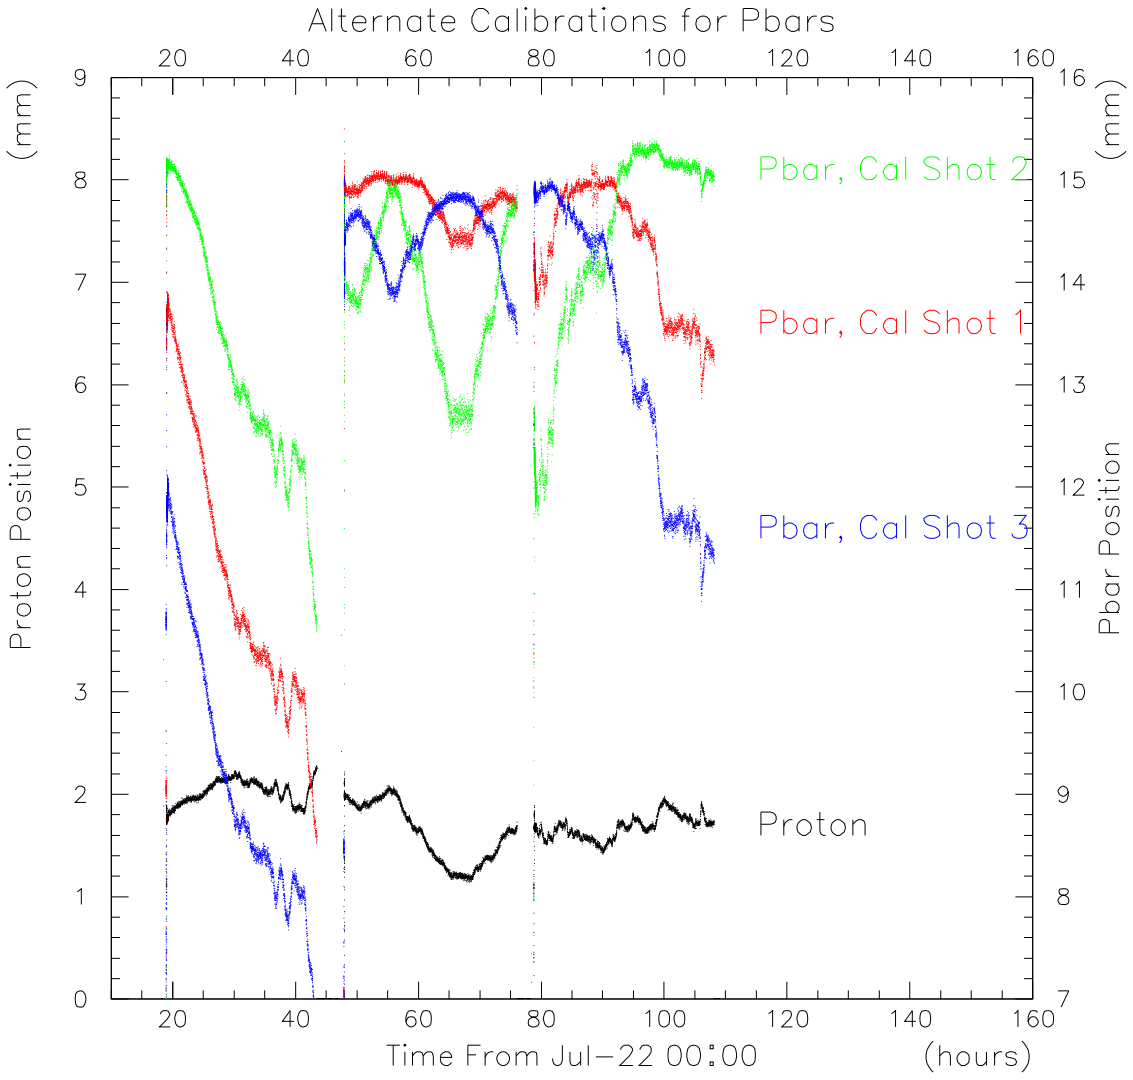

Figure 10: The black points show the same proton position data as in the previous figure. The colored points show different variations for computing the anti-proton position; these variations are discussed in the text but none is truly successful.