Beams-doc-2115-v2
January 27, 2006

More Details of Injection and Extraction Flashes

Rob Kutschke, CD/CEPA
Steve Foulkes, CD/CEPA

Abstract

This note presents a detailed look at flash data from MI State 5, mixed mode anti-proton production plus NUMI, in which there are seven injections followed by two extractions. The data is presented for two different configurations of the timing card, one in which the flash timing is armed by the BES and synced to the $AA marker and a second in which the flash timing is both armed by and synced to the BES. In this study the turn-by-turn delay (house delay) was scanned from 0 to 1180 half-buckets, in steps of 5 half-buckets. The response of all eight flashes will be shown as a function of this delay. This note also studies the disruption of the NUMI batches following the extraction of the anti-proton production batch.

Contents

1 Introduction

The data presented in this note are the eight 512 turn flash buffers from MDAT state 5, triggered by a master reset of TCLK $23. The data were taken at the horizontal MI BPM 412. The first set of data were taken between about 8:30 AM and 9:00 AM on December 30, 2005. The second set of data were taken on January 11, 2006 between about 1:20 and 1:30 PM.

MDAT state 5 is a mixed mode that sends a slip-stacked batch to the anti-proton production target and five batches to the NUMI target. All batches have 80 bunches. The details are:

- Inject one batch and start it slipping.
- Inject a second batch.
- Capture these two batches into a single batch.
- Inject five more batches. Each batch fills an empty spot in the MI.
- In a single turn, extract the slip stacked batch to the anti-proton target.
- In a single turn, extract the remaining five batches to the NUMI target.

This gives a total seven injections and a flash measurement was made for all seven.

Only one extraction flash was taken and it was triggered by the earlier of the two extraction beam sync events, MIBS $79. The two extraction events are close enough in time that both extractions are seen in the single extraction flash measurement.

For both datasets, the Echotek card was programmed to down-convert at 53 MHz in wide-band mode. In this mode, the integration window of the Echotek is about 44.8 buckets, a little larger than half of a batch.

The gain on the transition card was increased after the first data set was taken and before the second data set was taken.

2 Timing Card Configurations for the Two Data Sets

For the first set of data, the timing card was still configured as described in Beams-doc-2060; that is, the pre-trigger delay was computed by counting $AA markers and the turn-by-turn delay was applied following the $AA marker.[1] For this study the pre-trigger delay was set to zero, which means that the turn-by-turn delay started on the first $AA marker following the trigger event. The turn-by-turn delay is denominated in half-buckets.

For the second set of data, the timing card was configured as follows:

1. On the BES start counting RF cycles immediately.
2. Do one of:
    - If the pre-trigger delay is zero, immediately start counting down the turn-by-turn delay.

[1] In this note, as in Beams-doc-2060, the turn-by-turn delay is defined as the sum of the house delay, the board delay and the channel delay, all denominated in half-buckets. For this study, which looks at only a single BPM, the house delay was varied with the other two contributing delays fixed at zero.

- If the pre-trigger delay is non-zero, count turns by dividing the RF clock by 588. Once the pre-trigger is satisfied, count down the turn-by-turn delay.

3. Once the turn-by-turn delay expires, issue the first sync to the Echotek.

4. Every 588 RF cycles following the first sync, issue another sync, until 512 syncs have been issued.

With this second timing arrangement, all seven of the injected batches appear at the same turn-by-turn delay. For the data shown here the pre-trigger delay was set to zero.

For the extraction flashes, the timing card responds to MIBS $79 but otherwise behaves the same was as it did for injection flashes. For the first data set, the extraction flash is armed by MIBS $79 and the turn-by-turn delay countdown begins after the first $AA marker that follows. For the second data set, on the other hand, the turn-by-turn delay countdown begins immediately following the MIBS $79.

3 Examples of Good Data

Figures 1 and 2 show some typical examples of good data from this study. They are taken from the first data set but are also representative of the second data set. The left hand column shows the sum signal plotted against turn number for a particular delay and flash, while the right column shows the corresponding position data plotted against turn number. The key to decoding the delay and flash from the plot title is explained in the figure caption. The top row of Figure 1 shows an example of flash data triggered when no beam is in the machine; this gives some idea of the noise level, about 30 Echotek Units (EU). In Figure 1, the second row shows an example when the Echotek window is centered on a batch that is injected on the current flash; the first few turns show a sum signal at the noise level and, at turn 27 or 28, the large jump in the sum signal indicates that the beam has arrived. Based on these two histograms a threshold of 100 EU was chosen to discriminate between beam being present or absent. The third row shows an example when the Echotek window is centered on a batch that was injected on a previous flash; the sum signal is above threshold for all 512 turns. The first row in Figure 2 shows an example when the Echotek window captures some of a batch injected on a previous flash and some of a batch injected on the current flash; the sum is above threshold for all 512 turns but shows a large jump at turn 27 or 28. The second row shows an example of an extraction flash when the Echotek window is centered on a batch sent to the NUMI target. The bottom row shows an example of an extraction flash when the Echotek window is centered on the batch sent to the anti-proton production target.

For each set of data, the mean value of the sum signal was computed, excluding turns with a sum signal below a threshold of 100 EU. The mean value, and the number of turns included in the mean, are superimposed on the left hand plots. The upper left plot in Figure 1 has a mean of zero since no points were above threshold.

4 Flashes as a Function of Delay

4.1 The First Data Set — Relative to $AA Marker

Each of the data sets contains eight 512 turn flashes for 236 different house delays. For each of 1888 flashes the mean sum signal was calculated as described in the previous section. Figure 3 shows the results for the first data set, in which the timing was synchronized to the $AA marker. The upper left plot shows the mean sum signal, plotted as a function of house delay, for flash 0, the first injection flash. The upper right plot shows the same information for flash 1 and so on. In these figures, no point is drawn for flashes in which the sum signal was always below threshold. The points are color coded by counting the number and type of threshold crossings, as follows:

Blue The flash starts with 27 or 28 turns below threshold and all remaining turns are above threshold. This pattern occurs when the Echotek window includes only a newly injected batch, as is shown in the second row of Figure 1.

Green The flash has all 512 turns above threshold. This occurs when the Echotek window contains some bunches from previously injected batches; it may also include some bunches from the newly injected batch. Examples are shown in the third row of Figure 1 and the first row of Figure 2.

Magenta The flash starts with any number of turns above threshold, followed by all remaining turns below threshold. This is typical of an extraction flash, as is shown in the bottom two rows of Figure 2.

Red All points not falling into the above categories; that is there is more than one threshold crossing in the 512 turns. In all cases presented in this note, this occurs only when the sum signal is close to the threshold and the noise causes multiple threshold crossings. This can occur when the Echotek window includes only small fraction of a batch.

This figure contains a number of obvious outliers and a number of data points, those in red, that cannot be classified by the classification scheme described above. The outliers and most of the red points are due to errors in the data acquisition that have since been fixed: these data do not come from MI state 5 but from some other MI state. Moreover the bad data can be identified without reference to the outliers. A second defect is also present but less obvious: there are some repeated data points. This too is due to an error in the data acquisition that has since been fixed. Figure 4 shows these same data, excluding the known bad data. All of the outliers are gone as are most of the

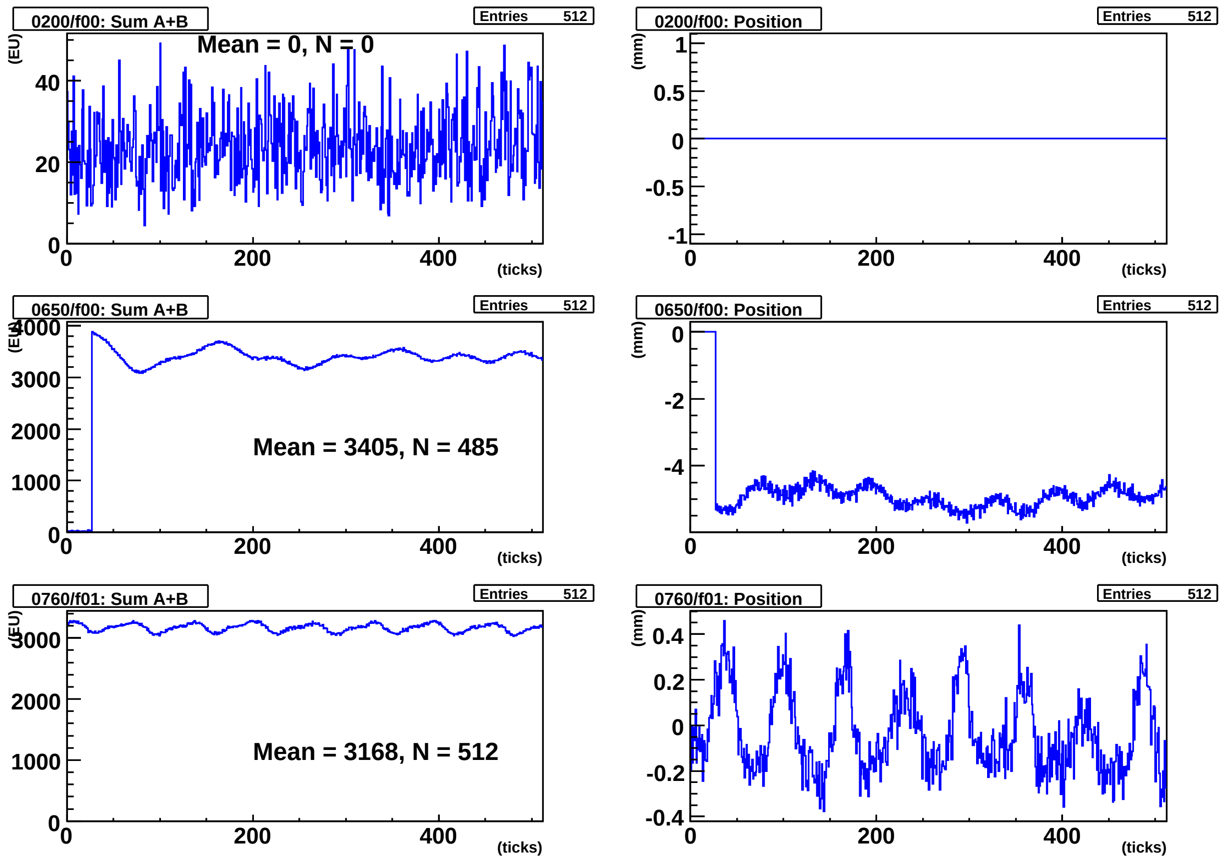

Figure 1: Each row shows an example of the sum and position data from one flash; the left column shows the sum and the right column the position. The position set to zero if the sum was below a threshold of 100 EU. The flashes are identified by the title in the top left corner of each plot; the title has the format $nnnn$/f0m, where $nnnn$ is the turn-by-turn delay in half-buckets and m is the flash number in the range 0 to 7. The mean value of the sum signal, excluding bins below threshold, is superimposed on each of the left hand plots; the number of points included in the mean, N, is also superimposed. The plots are discussed further in the text.

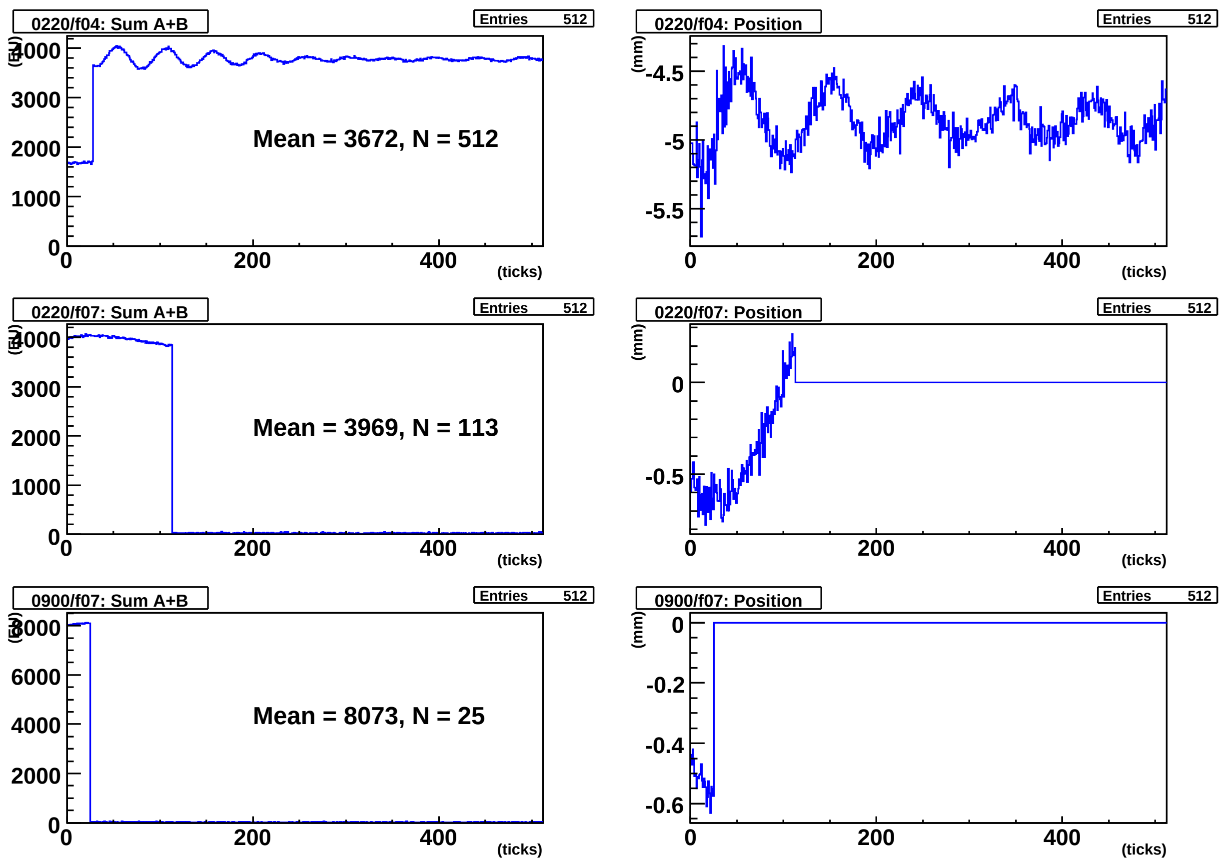

Figure 2: Continuation of Figure 1. In the upper left hand plots all points were included in the mean; no special care was taken to deal with the jump in the data near turn 27.

red data points. The red points that remain have a sum signal close to the threshold and the noise causes multiple threshold crossings. Visual inspection of the turn-by-turn sum signal for these cases shows that the data is good.

The upper row in Figure 3 shows the data for flashes 0 and 1. The main features of these data are the same as those discussed in Figure 2 of Beams-doc-2060: one batch is injected on the first flash; on the second flash the first batch has slipped to a longer delay and the second batch is injected immediately following it. The color coding of these data also matches that features discussed in Beams-doc-2060. In flash 1, the first batch has a smaller sum signal than it did in flash 0. This is probably not a loss of protons but rather an artifact of the instrumentation. In flash 1, the first batch is slipping so it is circulating away from the center frequency of the Echotek digital filter; so the lower amplitude of the first batch in flash 1 is likely due to the roll off of the digital filter.

In the second row, the left hand plot shows the sum signal for flash 2. The green points between delays of about 770 and 1010 half-buckets are the first two batches slipped on top of each other. The group of blue points represents the newly injected batch; it wraps around the right edge of the plot.[2] Comparing this plot to the upper right plot shows that the slip-stacked batch is located in time half way between the two batches from which it was created. One might be tempted to use these first three plots to compute the efficiency of slip-stacking. The bunches, however, are wider after slip-stacking than they are before; because the sum signal gets smaller as bunches are broadened, this causes an artifact that looks like beam loss during slip-stacking.

In the second row, the right hand plot shows the sum signal for flash 3. The slip-stacked batch is present in the same position as in flash 2; similarly for the batch injected at flash 2. Remember that the trigger for these data is synchronous with the $AA marker. The newly injected batch appears as the blue data points following the green data points. A similar story can be told for the plots in the third row and for the left hand plot in the bottom row. In all cases the newly injected batch is injected at the end of the train of previous batches.

Another feature can be seen in the left hand plot in the bottom row. Excluding the slip-stacked batch, the data are scattered about two values, 3866 and 4600. The projection of these data onto the sum signal axis is shown in Figure 5, in which the bi-stable behavior is clear. When all injections are complete, the protons destined for NUMI have the following pattern: 80 buckets filled, 4 buckets empty, 80 buckets filled, and so on. The integration window of the Echotek is 44.8 buckets and the data points are separated by 2.5 buckets. When the integration window is filled with bunches, the upper of the two values is observed. When the integration window includes the gap between batches, the lower of the two values is observed. The ratio of these two values, about 84%, is a little smaller than than the expected value of 40.8/44.9=91%.

[2] One full turn of the MI is 1176 half-buckets; so the rightmost point, 1180 half-buckets, is the same as having a point near the left edge, at 4 half-buckets.

The bottom right plot in Figure 4 shows the mean sum signal as function of turn-by-turn delay for the extraction flashes. Except for one or two turns with low intensity, all measurements have been correctly identified as extraction flashes. Because the timing of all measurements on this page is relative to the $AA marker, the various batches appear the same times for the extraction flash as they do for the various injection flashes.

4.2 The Second Data Set — Relative to BES or MIBS $79

Figure 6 shows the same information as the previous figure except that the data comes from the second data set, which has timing relative to the trigger event (BES or MIBS $79), not the $AA marker. The same general features described for the first data set are also present in this dataset. The main difference is that the newly injected batch always arrives at the same delay: the change in the timing card was explicitly designed to give this result. In this mode a single delay can be specified to correctly set the timing for all 7 injection flashes. A consequence of this choice is that batches previously in the machine move to smaller delays on each flash.

Notice too that there is a timing shift between the last injection flash and the extraction flash. This occurs because the extraction flash is timed relative to the MIBS $79 while all others are timed relative to the BES.

This figure also has a few red data points. For all but one he explanation is the same as before: they occur when the sum signal is close to the threshold. The one exception occurs for flash04 with a delay of 970 half-buckets. For the first 27 turns in this data set, the Echotek window includes only a small part of the batch injected on the previous flash; the sum signal happens to be close to the threshold and there are several threshold crossings. On the next turn the injection occurs and the Echotek window includes a lot of beam from the newly injected batch; at this point the sum signal jumps to about 9000 EU, well above threshold.

4.3 Number of Turns as a Function of Delay

Consider, for example, the middle row of Figure 1, which shows the turn-by-turn values of the sum and position for an injection flash with a delay of 650 half-buckets. In this particular example there were 485 turns with the sum signal above threshold. The number of turns above threshold was computed similarly for all flashes that were classified as injection flashes, that is for all blue points in Figures 4 and 6. For the first data set, the one with timing relative to the $AA marker, the number of turns so computed is plotted as a function of turn-by-turn delay in the top plot of Figure 7. The points are color coded to indicate from which flash the point was taken. As in Figure 4 the corrupted data is excluded. The data in this plot have two values, either 484 or 485. It is understandable that there are two values: for each flash there is a race condition between the arrival at the VME crate of the beam signal, the BES and the $AA marker. For some values of the delay, the beam signal will arrive just before the $AA marker

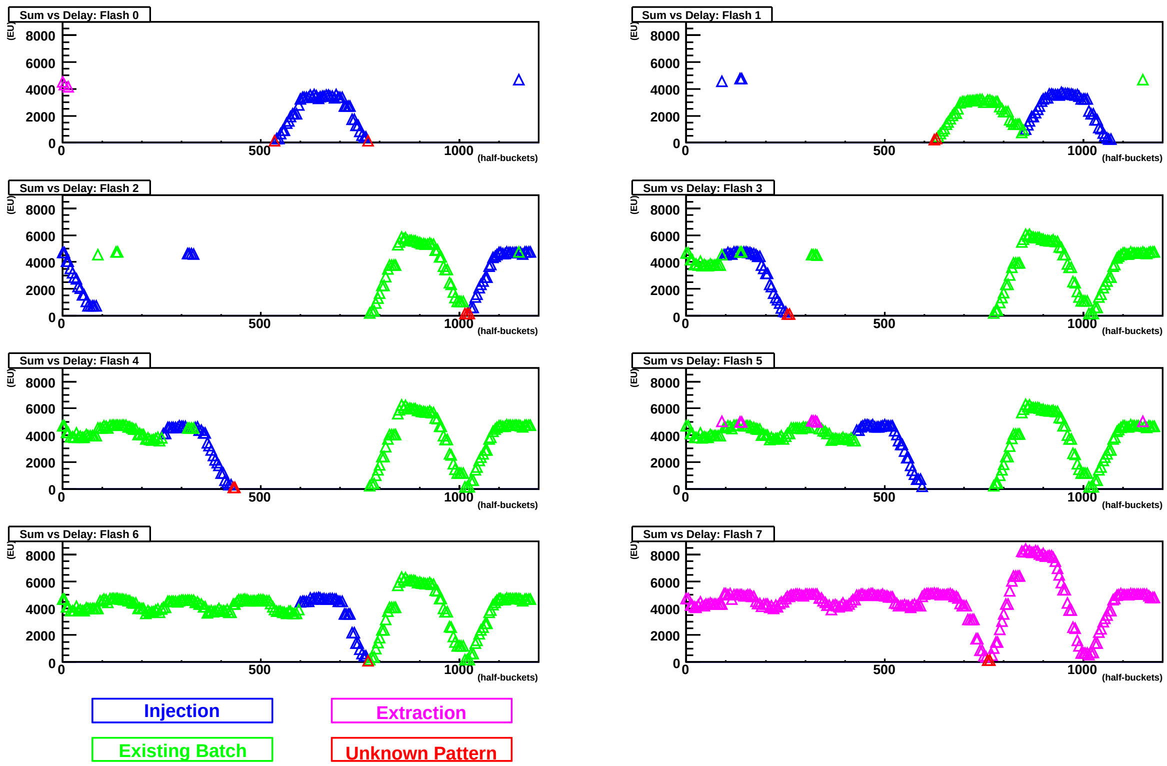

Figure 3: The data points show the mean of the sum signal, as a function of house delay; if a flash has no turns above threshold, no symbol is plotted. The data is shown separately for each of the eight flashes. When the mean was computed only the turns with a sum signal above threshold were included; the threshold was 100 EU. The meaning of the color code is discussed in the text. All eight plots have the same scale.

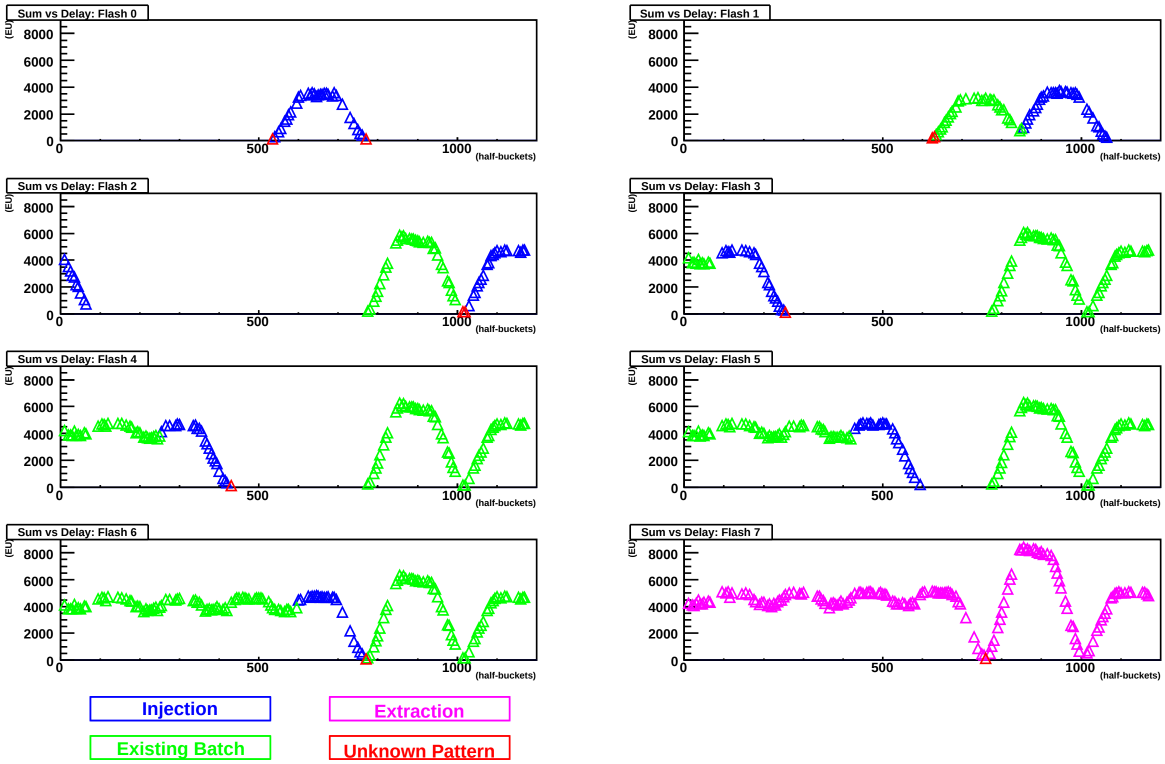

Figure 4: Repeat of Figure 3 with known bad points suppressed.

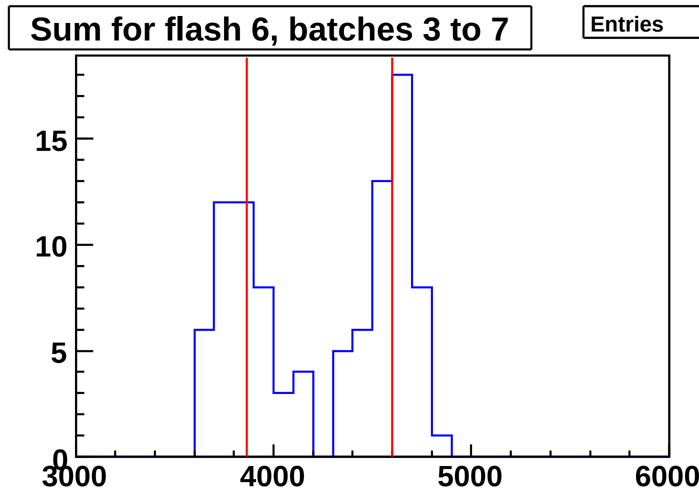

Figure 5: Projection of the data points from NUMI batches from flash 6 in Figure 4 on to the vertical axis. The data are cluster around two values of 3866 and 4600 EU, which are indicated by the vertical red lines in the figure.

and others just after. Moreover the relative timing of the BES and the $AA marker changes with each flash; this explains why there can be both black and green data points for delays between about 1000 and 1100 half-buckets. It also explains why there can be magenta and cyan points in the same delay range as the red points.

What I do not understand, however, is the large red triangles: these are points from flash 0 that appear at turn 484, not at turn 485 as do the majority of the red data points. I can understand how there could be a single step from 485 to 484 in the red data points but I don't understand how they can jump back and forth between the two values.

The corresponding plot for the second data set, with timing relative to the BES, was not made. For these data, all injection flashes have 485 turns above threshold.

The middle plot in Figure 7 shows, as a function of delay, the number of turns above threshold for data from flash 7 from the first data set; these correspond to the magenta points in the lower right plot of Figure 4. The blue scale on the left is for the blue data points and the red scale on the right is for the red data points. In this figure one can see that the batch destined for the anti-proton production target always leaves the MI after 25 turns. However the batches destined for NUMI sometimes leave the machine after 113 turns and sometimes after 114 turns. Again, I can understand why there are two values but I do not understand why the data jumps back and forth between the two values. I had expected one step up and one step down (or vice-versa).

The lower plot shows the same information as the middle plot for but the

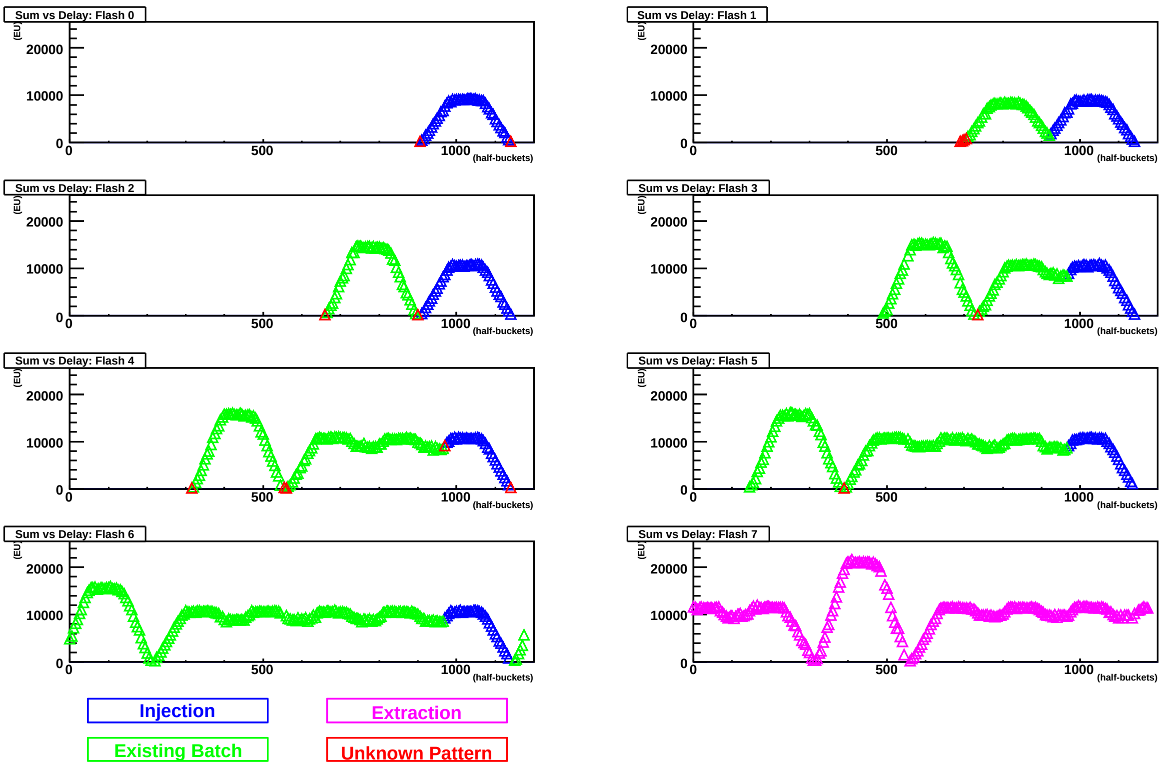

Figure 6: The same information as Figure 3 but for the second data set, which has all timing to the BES, not the $AA marker. With this timing the newly injected batch always appears at the same value of the delay.

second data set, for which the timing is relative to MIBS $79. It has the same difficult to understand behavior as the middle plot.

To investigate this further, the top four plots in Figure 8 show the sum signal as a function of turn number for four representative delays taken from the second data set, two for which there are 113 turns above threshold and two for which there are 114 turns above threshold. The bottom four plots in this figure show the same four histograms, zoomed in to show the first few turns without beam in detail. There is nothing unusual with these data. So this remains a mystery.

5 Disruption of the NUMI Batches by the Anti-proton Production Kicker

In MI state 5, once injection is complete, the circumference of the MI is essentially full. This can be seen in the two lower plots of Figures 4 and 6. After injection the beam is accelerated to 120 GeV. Then the slip-stacked batch is extracted to the anti-proton production target. The rise and fall times of the extraction kicker magnet are slow enough that it is possible for this kicker to disrupt the NUMI batches that remain in the machine.

In the first data set, when the turn-by-turn delay is set to 1015 half-buckets, the Echotek window contains the first few bunches of the NUMI batch that immediately follows the slip-stacked batch. The Echotek window also contains the gap between these two batches. In the upper left plot of Figure 9, the sum signal for flash 7, with a delay of 1015 half-buckets, is plotted as a function of turn number. The beam leaves the machine after 112 turns. The sum signal has a small value since most of the Echotek window includes the gap between the batches. The middle left plot shows the position data for this flash. One can see that the position is fairly stable for the first 25 turns, after which large oscillations develop. The start of the large oscillations coincides with the extraction of the slip-stacked batch. To quantify this effect, the Fourier transform was computed for the position data between turns 26 and 112. The magnitude of the Fourier transform is shown as the bottom left plot in Figure 9; a large amplitude is seen at the betatron frequency. In order to characterize the magnitude of the betatron oscillation, the highest bin within a ± 4 bin window of the betatron frequency was found. The magnitude of the Fourier transform in this bin in 1.32 mm. In the following this will be refered to as the betatron amplitude.

In the first data set, when the turn-by-turn delay is set to 500 half-buckets, the Echotek window contains beam from NUMI batches that are far from the slip-stacked batch. This beam should be minimally disrupted by the extraction of the slip-stacked batch. The right hand column of Figure 9 contains the same information as the left hand column but for a delay of 500 half-buckets. In this data set the position data is much more stable and there is no obvious betatron line. Using the algorithm described in the previous paragraph, the betatron amplitude was measured to be 0.03 mm.

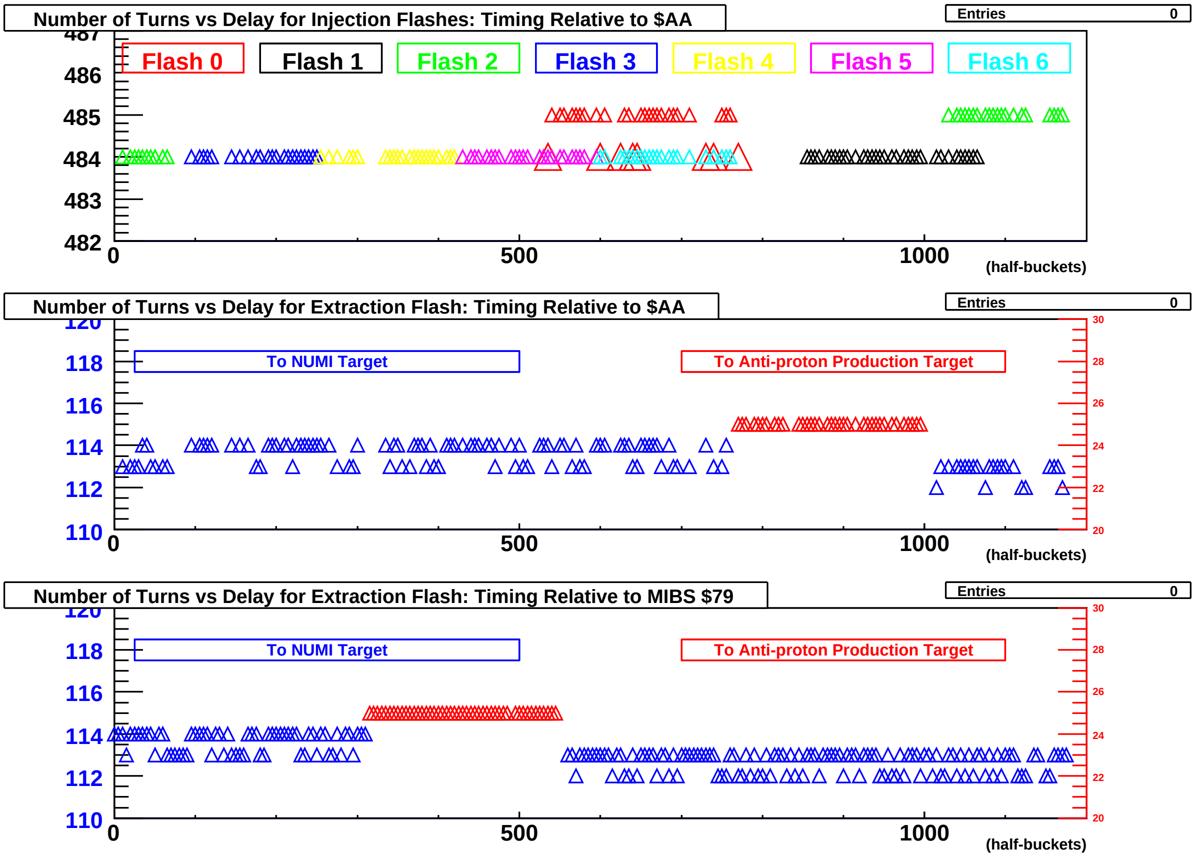

Figure 7: All plots show the number of turns with intensity above threshold as a function of turn-by-turn delay. The upper plot show the information for the seven injection flashes from the first data set; the large red triangles are discussed in the text. The middle plot shows the information for the extraction flash from the first data set. The bottom plot shows the data for the extraction flash from the second data set. In both the middle and bottom plots the points are color coded to match either the left or right vertical scales. The data for the injection flashes of the second data set are not shown since all delays all have 485 turns above threshold.

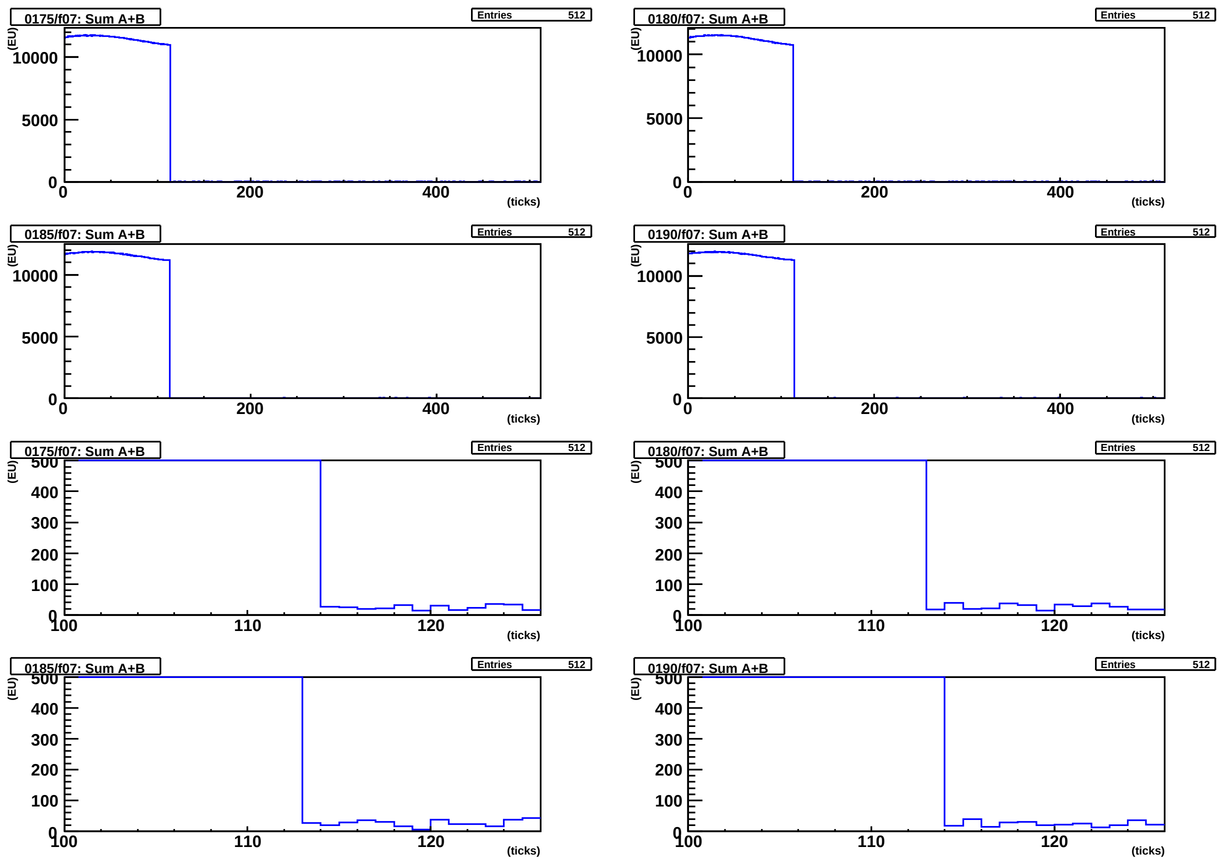

Figure 8: The top four plots show the sum signal vs turn number for the full 512 turns for four representative flashes taken from the second data set, which has timing relative to MIBS $79. The bottom four plots show the same four histograms zoomed in to show details of the first few turns without beam. This figure does not provide any clue why some flashes have 113 turns above threshold and some have 114.

In a previous version of this work, the betatron amplitude was defined as the mean of nine bins in the neighbourhood of the betatron line. In the final preparation of this note, I decided that the contents of the highest bin was a more useful measure.

The bottom plot in Figure 9 shows the betatron amplitude as a function of delay for the first data set. This plot contains no points in the region of the slip-stacked batch because the Fourier transform was computed using data only from turns after the slip-stacked batch was extracted. For reference, the top plot shows the mean sum signal as a function of delay for flash 7 of this same data set. In this figure one can see that the NUMI bunches in the neighbourhood of the slip-stacked batch are indeed disturbed by the extraction of the slip-stacked batch.

In this figure the data points are separated by 2.5 buckets and the integration window of the Echotek is 44.8 buckets, a little more than half a batch. Consider the betatron amplitude for a delay of 755 half-buckets, the last data point before the gap. For this data point, the Echotek window contains the trailing few bunches in the last batch before the slip-stacked batch. The rest of the Echotek window contains the gap between the two batches. As one moves to smaller the delays the betatron amplitude falls, reaching the baseline at a delay of 665 half-buckets.[3] That is, the betatron amplitude reaches baseline when the Echotek window first excludes the few bunches that were included in the point at a delay of 755 half-buckets. The interpretation of this data is that only the trailing bunch or two in this batch are disturbed by the kicker; the slope between delays of 665 and 755 half-buckets is consistent with the width of the Echotek window.

On the other hand, the slope of the data after the gap, and wrapped around to short delays, is much shallower. The interpretation of this slope is that almost all of the first batch following the slip-stacked batch is disturbed by the kicker. The slope is mostly governed by the fall time of the kicker, with only a small contribution from the width of the Echotek window.

[3]The data points with corrupt data have been suppressed so not all points are shown.

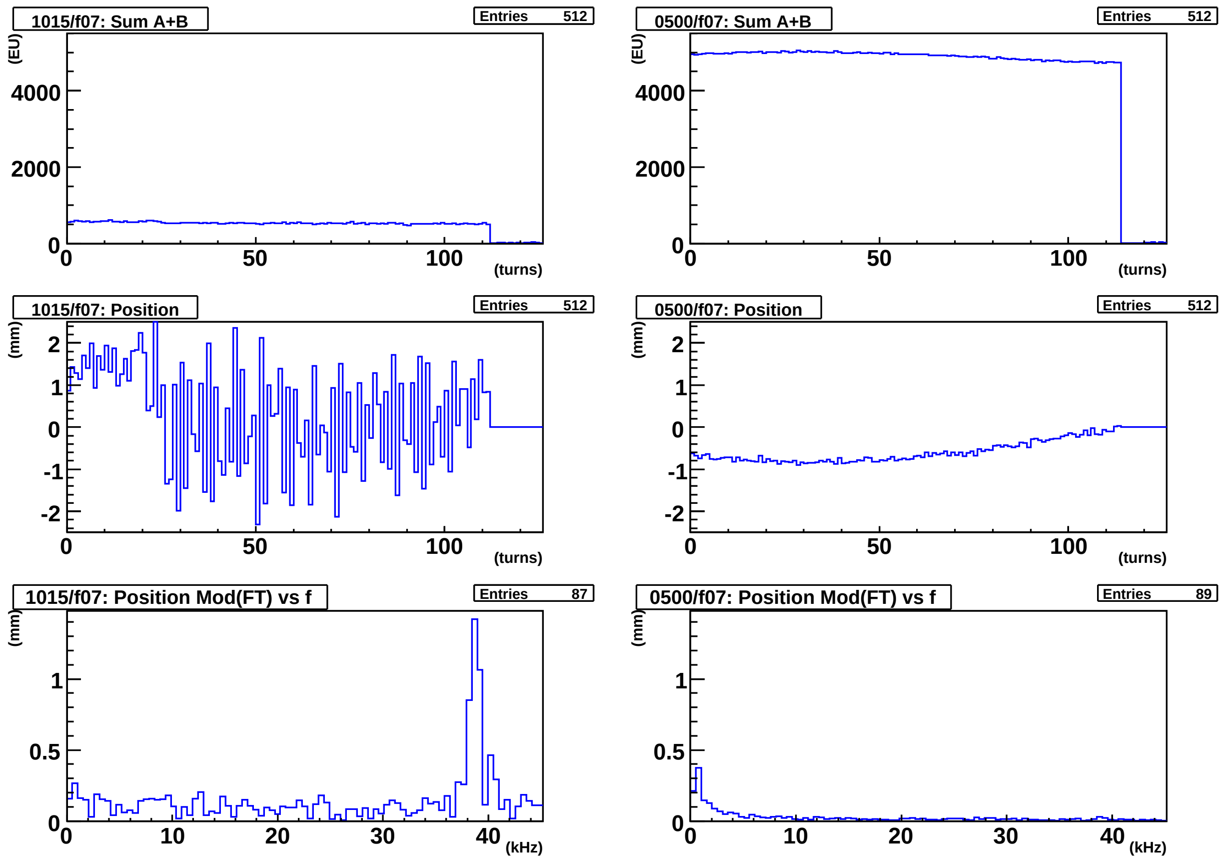

Figure 9: The plots in the left hand column show the sum, position and position FFT for the leading edge of a NUMI batch that was disrupted by the extraction of the slip-stacked batch. Before turn 25, the position is stable but large oscillations begin after turn 25. The bottom plot shows the Fourier transform of the position data, starting at turn 26 and ending with the last valid data point. The right hand column shows the same information for a NUMI batch that is far from the slip-stacked batch and, therefore, should be minimally disrupted by the extraction of the slip-stacked batch. Each left-right pair of plots has the same vertical scale.

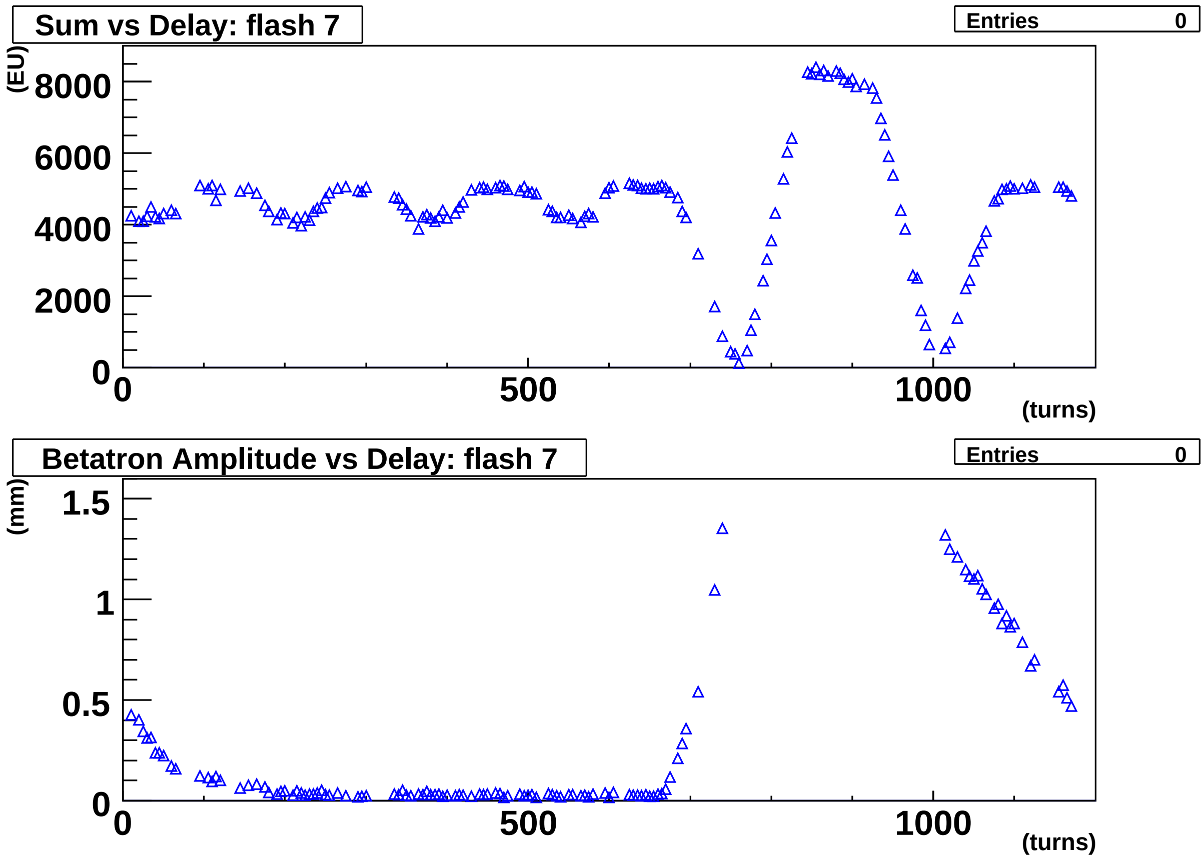

Figure 10: The upper plot is provided for reference; it shows a repeat of the bottom right plot of Figure4, the mean sum signal vs delay for flash 7 of the first data set. The slip-stacked batch is seen between delays of about 650 and 1000. The bottom plot shows the betatron amplitude for the NUMI batches as a function of delay. The betatron amplitude was measured using the position from turns after the slip-stacked batch was extracted. A more precise definition of the betatron amplitude is defined in the text.