Antiproton Source Operational Summary Over the E835/E862 Run

James Morgan

April, 1998

Introduction

This document summarizes operation of the Antiproton Source in support of the E835 and E862 experiments during the Fixed Target run of 1996-97. The first effort to run beam took place on July 17, 1996 and the last store was terminated on September 18, 1997. In the following sections I will attempt to quantify the performance of the Antiproton Source over this 14 month period, as well as compare this run to the E760 runs of 1990-92. The E760 data includes both data from the first "engineering" run which ran from early June to early September in 1990, and the physics run which covered the period between July 1991 and January 1992. The E835 data begins after E835 installed their detector in late August and early September 1996. The data presented came primarily from weekly summaries prepared by the Operations Department.

Stacking

Stacking rates at the beginning of the run averaged in the 2.0-2.5E10/Hr range with peak rates of about 3.5E10/Hr. By mid-run the average rates had improved to about 3E10/Hr with peak rates of over 4E10/Hr. This was an apparent drop off from the end of Collider Run Ib when peak stacking rates routinely exceeded 5E10/Hr. For comparison, the highest weekly stacking rate in Collider Run 1b was 5.08E10/Hr as compared to 3.64E10/Hr in the E835 Run. There are several known factors that contributed to a lower stacking rate:

- There were fewer available Main Ring cycles. During Fixed Target operation there is a Main Ring cycle of 5 second duration that is used for Tevatron injection. Also over much of the run no stacking cycles were allowed during fast extraction from the Tevatron resulting in a 3 second dead time during the supercycle. The average stacking rate dropped by about 5-10% due to the reduction in stacking cycles.

- Main Ring intensity on stacking cycles was significantly lower during E835 operation. Typical Main Ring intensity in the latter part of Collider Run 1b was 3.2E12 per pulse compared with 2.2E12 per pulse during E835 operation (a reduction of 31%). The relationship between protons on target and stacking is not linear, the reduction in stacking rate is not as great. With the focus of the Physics program on Fixed Target experiments, the majority of tuning time was spent on those Main Ring cycles destined for the Tevatron. The optimum tune for the multiple-batch Fixed Target cycles differs from the single-batch stacking cycles. The presence of the Tevatron ramp during Fixed Target also adversely affects Main Ring cycles. The Tevatron is normally in a constant energy state during

Collider stacking. Main Ring compensation for the Tevatron is not absolute which results in tune shifts and movement of the extracted beam on target.

•Toward the latter part of Collider Run 1b there were failures to Debuncher stochastic cooling kicker electrodes apparently due to excessive power levels. The power output of each TWT was limited for the E835 run to prevent further damage. The reduced cooling power led to larger beam emittance and reduced transfer efficiency to the Accumulator. The larger beam size resulted in about a 10% reduction in antiproton flux to the Accumulator.

•The stacktail betatron systems were removed from the tunnel and used as prototype tanks for a future system. The stacktail horizontal system could improve stacking by as much as 10%, although this was for larger stacks than those routinely accumulated during the E835 Run.

The total reduction in stacking rate due to these factors would amount to more than the 28% reduction observed. There was one factor that contributed in a positive way to stacking during the E835 run. That was the fact that stacks were much smaller than during the Collider Run, usually by a factor of 2 or more. There is a drop-off in stacking rate with stack size that occurs when the stack size exceeds about 40E10. The average stacking rate during Collider operation was typically 20-30% less than peak stacking rates observed with smaller stacks. Also, it is worth noting that the decrease in stacking rate with stack size was worse in the E835 run than in the last Collider run. The mechanism leading to the reduction in stacking rate is not well understood but appears to involve signals from the core being picked up but the stacktail system resulting in degraded performance.

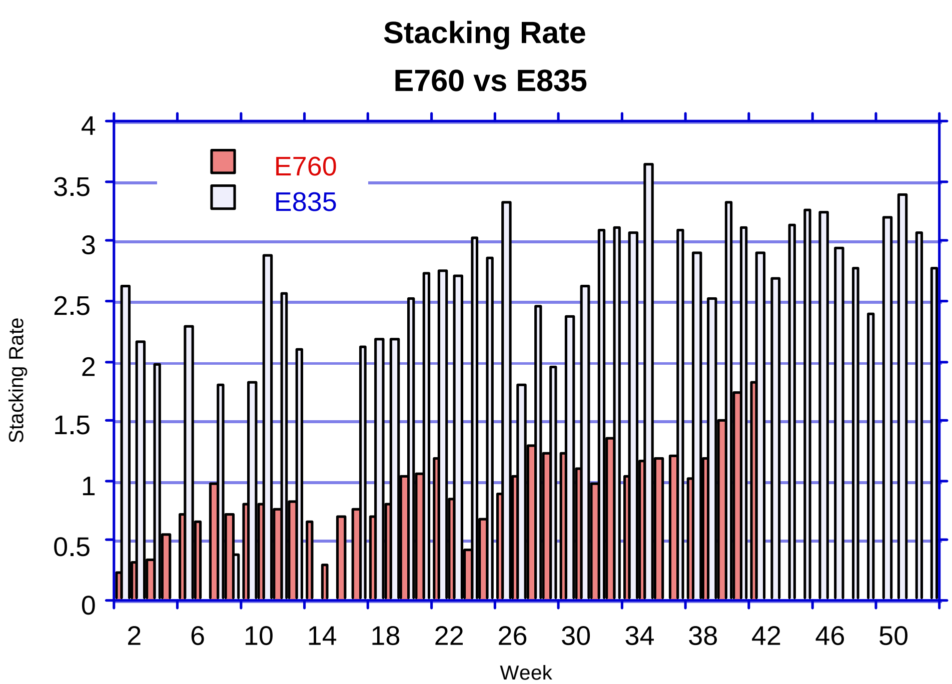

Figure 1

Figure 1 is a weekly plot of stacking rate over both the E760 and E835 runs. The upturn in the E760 stacking rate at the end of the run was due to the installation of the Debuncher momentum system.

<u>Store Hours</u>

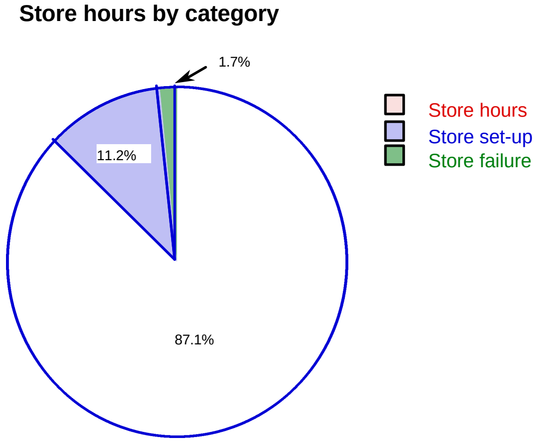

Figure 2

The traditional way to define a store hour was to include any time spent in an experimental store, whether the experiment could use the beam or not. For this run a further breakdown was done on a weekly basis. The categories were store hours, set-up hours and failure hours. Set-up for a store was deemed ended when E-835 turned on their gas jet, thus some inefficiencies within the experiment were included in store set-up. The category of "failure hours" was only made up of accelerator related failures and not problems that the experiment had. Figure 2 is a chart which shows the proportion of hours spent in the three categories.

The mode of experimental operation greatly affects the amount of time spent in set-up for a store. If a stack was decelerated to a single above transition point, a set-up time of less than an hour would routinely yield 48-72 hours of store time. As beam energy changes and below transition decelerations are added, the proportion of set-up time increases. During early decelerations to points below transition, setup times of 4-6 hours were required to yield a store of only 12 hours duration. A comparison of store hours for E835 and E760 is shown in figure 3, this data is the total store hours including set-up and failure. On average, there were more store hours during the E760 run than during the E835 run.

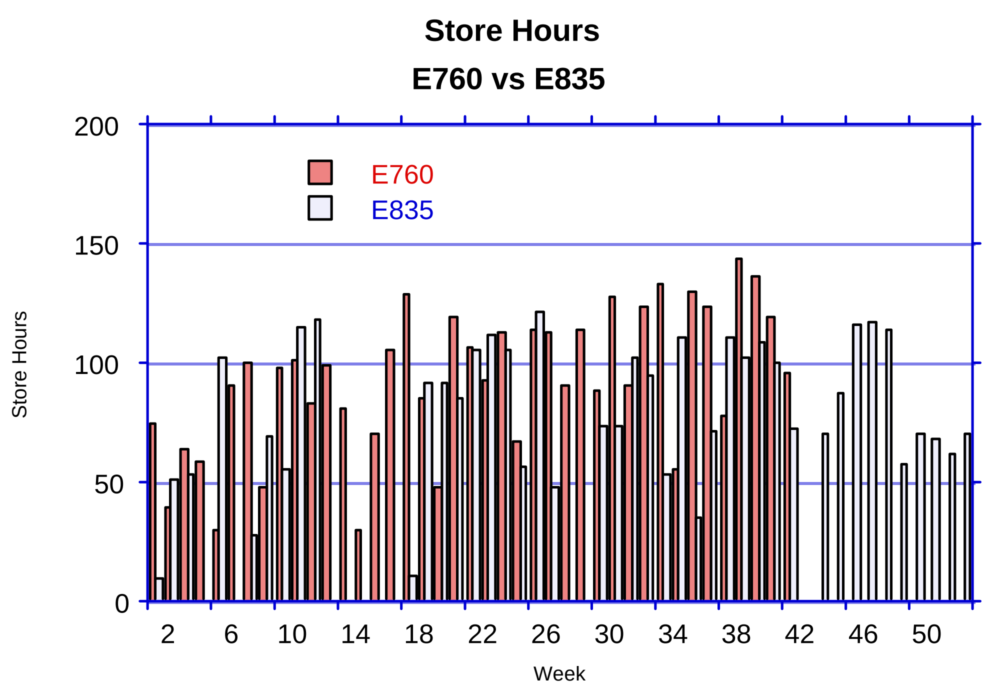

Figure 3

Most of this can be explained by the mode of experimental operation and lower gas jet density used in E760. A large proportion of the stores were at one or two above transition points resulting in little set-up time. The lower gas jet density slowed the rate of stack loss during a store resulting in longer duration stores (at the cost of lower luminosity).

Reliability

Some machine downtime is unavoidable during accelerator operation, however this run had more major failures than is typical. Initial turn-on was delayed while a chilled water extension to the Main Injector was completed. Problems continued on the first day of beam start-up when heavy storms brought unprecedented amounts of rain to the area. Although the Pbar source recovered quickly, an entrance to the Main Ring tunnel was breached and the resultant flooding held off beam to Pbar for nearly a week. This was by no means the only major interruption of the run, Table 1 summarizes downtime with 1 day or more duration beginning with the first attempts to run beam. The major blocks of downtime add up to about three months, this in addition to the many significant periods of downtime that had a duration of less than a day.

7/17 - 7/22 '96	Flooding rains cause standing water in the M. R. tunnel	5
7/29 '96	Main Ring feeder fault	1
8/16 - 9/10 '96	E835 detector installation	25
10/14 '96	Main Ring magnet replacement	1
10/21 '96	Main Ring feeder fault	1
11/5 '96	Repair laminations on A4B7	1
12/11/96-1/10/97	A20 berm LCW leak followed by holiday shutdown	30
2/26 '97	Replaced D6Q18 due to overheating	1
3/24 - 4/5 '97	E835 forward calorimeter installed	12
4/18 '97	Replaced D2Q3 which had an internal LCW leak	1
5/20 - 5/27 '97	Tevatron component failure	7
6/23 '97	Recovery from power outage	1
7/7 - 7/13 '97	A20 prototype cooling kicker tank installation	6
7/15 '97	Main Ring feeder fault	1

Table 1 Major downtime
(Green planned, Red unplanned)

The most lengthy period of downtime deserves a more detailed description. On December 11, 1996 a major LCW leak developed on the Pbar LCW system. System water pressures went to 0 psi which only occurs when leaks are very large. Tours of the Central Utility Building, service buildings and tunnel enclosures resulted in no evidence of a large leak. There was the potential for weeks of downtime to be spent digging up the LCW headers looking for the leak. An expert in leak detection was called in to provide an alternative leak detection scheme. The location of the leak was found within a few days by injecting hydrogen gas into the LCW system and detecting the gas at ground level. The leak was traced to the underground LCW lines in the vicinity of the A20 berm.

After excavating the area, it was found that a pipe leading to an air bleeder had been severed. Not only did this cause a major loss of water but significant amounts of soil and debris entered the LCW system. Special piping had to be installed in the tunnels to trap the debris before flow was reintroduced to the tunnel headers. Recovery was further complicated by the lab 2-week shutdown that occurred while efforts were being made to clean up the water system. During the down period the a stacktail kicker tank in A20 was removed and replaced with a spool piece. After the holiday shutdown it was noticed that the bellows on one end of the spool piece had collapsed, apparently because the installed spool was too short. A new bellows was installed in its place, but this required starting the vacuum pump-down over again causing further delays. It was not until January 10, 1997 that beam was again established to the antiproton source.

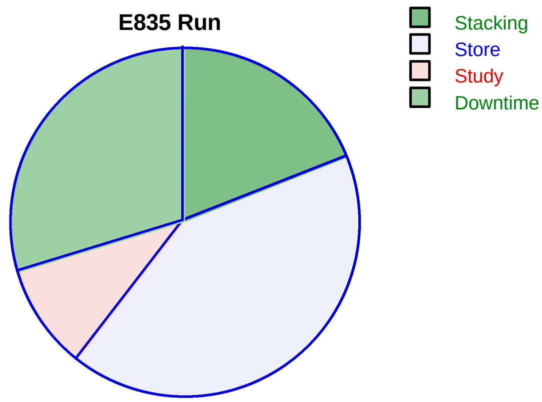

Figure 4

The proportion of stacking, store, study and downtime hours is provided in figure 4. The downtime category includes failures during stores, stacking and studies. Downtime represented nearly 30% of the total hours over the course of the run. This would be considered a bit higher than normal for collider operation, but is definitely higher than expected for an experimental run. When the antiproton source is in a store mode the sources of downtime are reduced as there is no reliance on the other accelerators. During collider operation the antiproton source relies on the other accelerators much of the time.

It is more difficult to compare E835 downtime to that in E760 directly as downtime statistics were not gathered the same way. Adjusting the E835 stacking and store hour statistics to be comparable to the E760 data, a comparison can be made by creating an "other" category that includes downtime and studies. Figures 5 and 6 are charts showing these three categories. Clearly there was a larger proportion of time spent stacking and in stores during the E760 runs. Part of the discrepancy in hours in the "other" category is due to the relatively large number of study hours in the E835 run.

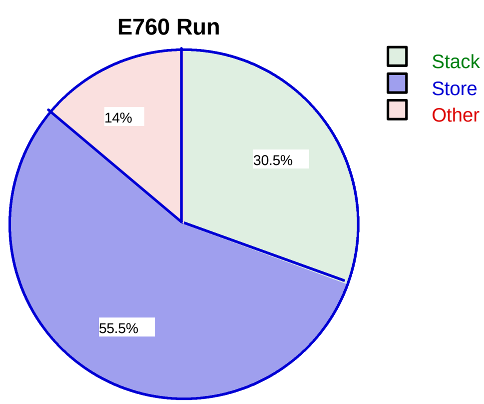

Figure 5

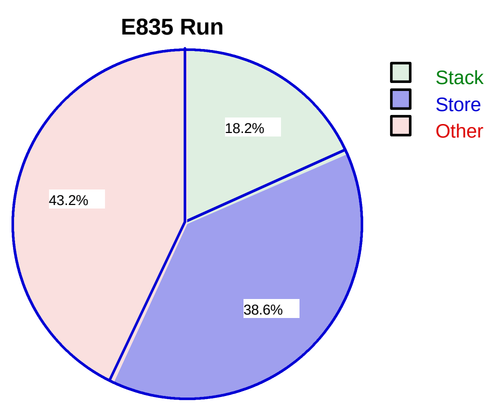

Figure 6

The loss of a stack is another factor that affects machine reliability. When a stack is lost unexpectedly, downtime is only charged for the time required to recover and return to stacking.

actually lost is the hours that have been spent accumulating the stack and the hours it takes to replace it. Table 2 is a list of the stacks lost during the E835 run with a short explanation of the

or other delays follow.

Over the run, antiprotons lost due to dumped s
average stacking rate for the run of 2.79E10/Hr, the lost stacks represent 246 hours of stacking

Date		Stack size (E10)
09/21/96	LCW leak on IQ33	18.40
	A:IB overcurrent trip	36.84
	Access to flush D2Q10 and D2Q11	5.08
	A:QDF stopped regulating	30.07
	Large LCW leak, underground line to bleeder near A20	22.80
	A:LQ output goes to 0	28.84
	D:QD trips, access to flush D1Q10	7.90
	Cycle buses due to Accumulator tunes	3.54
	A:LQ output goes to 0	26.13
	A:LQ output goes to 0	3.55
	A:QDF magnet overtemperature trip	29.24
	Repair LCW leak on IQ1	25.60
	Unexplained partial stack loss (A2B3 suspected)	14.58
	A10 and A60 devices trip	30.14
	Lost during deceleration, controls problem	22.45
	Technician causes Accumulator vacuum valves to shut	14.23
	Bad ramp tables used to decelerate	25.33
	Instability during beam energy change	12.48
	Access to fix LCW leak on IQ1	15.82
	Access to fix LCW leak on IQ1 (again)	16.78
	Access to check vacuum near A:IG603	10.43
	Unknown loss at end of store (A2B3 suspected)	5.43
	E835 gas jet turbo trip	5.64
	Core blew up due to core 2-4 dp system problems	15.49
	Access to repair core 4-8 dp system	16.01
	Site-wide power glitch	43.27
	Unexplained partial loss of stack (A2B3 suspected)	13.01
	Access to flush overheating D3B15	5.38
	Pbar feeder 24 fault caused by Main Injector work	9.55
	Access to repair D50 chilled water leak	20.45
	Partial stack loss, instabilities near transition	14.83
	Power glitch when feeder 48 faulted	40.33
	Partial stack loss, instabilities near transition	17.03
	Partial stack loss, instabilities near transition	7.54
	Instabilities near transition	25.00
	Partial stack loss, instabilities near transition	8.87
	Partial stack loss, instabilities near transition	22.90
	Partial stack loss, instabilities near transition	16.65

time. A similar amount of time is then spent accumulating the pbars that were lost.

Examining the causes for the stack losses, the trend early in the run is the same as at the end of Collider Run 1b. At that time problems with the LCW system and Accumulator power supplies caused the majority of stack losses. Six of the last eight stack losses in Run 1b were due to these problems and eleven out of the first twelve in the E835 run. This situation improved considerably as the run progressed, particularly with the Accumulator power supplies. Some of the main Accumulator power supplies were modified towards the end of Run 1b to be less sensitive to the loss of the power supply enable. In the original circuits the loss of this enable for even one line cycle would cause the output current to droop. After power supply problems of this sort early in the run, all of the main supplies were modified in this way. There were five stack losses in the first four months of the run due to this type of failure, then none for the final eight months. Table 3 provides a summary of stack losses by run, with separate categories for power supply and LCW problems.

Most of the LCW failures fall into two broad categories, leaks and overheating magnets. As mentioned earlier, the most noteworthy LCW failure was the large leak that occurred in December 1996. LCW related problems were considered serious enough so that a group was formed to understand and correct the root causes. LCW related stack losses were relatively rare in the early years of pbar operation, but have progressively worsened with time. In particular the problem with overheating magnets is a relatively recent phenomenon. The first problems with overheating magnets occurred at the end of Collider Run 1a and became more serious over the course of Collider Run 1b. The majority of the overheating magnets were SQ (small quadrupole) series magnets that ran at relatively high currents. It was recognized after the first few occurrences that CuO was forming in the LCW system and building up on the cooling passage walls of the magnet coils. The CuO was created when dissolved oxygen reacted with the copper LCW headers and coils. Efforts to remove dissolved oxygen from the LCW system with filters was inadequate. Prior to the E835 run the only solution to an overheating magnet, albeit a temporary one, was to backflush the magnet with LCW or flush the magnet with a weak acid solution. This process is time consuming and puts a burden on the ion removal system used on the LCW. A more comprehensive oxygen removal plan was implemented including displacing dissolved oxygen with bubbled nitrogen and an oxygen removal system that greatly improves the reduction in oxygen levels after a serious failure (such as a large leak). There were also improvements to the particle filtering in the LCW system that provided a non-destructive means for removing the CuO. The number of overheating magnets dropped significantly after these measures were implemented.

Internal magnet LCW leaks have been increasing in frequency over the past few years. Such failures were rare over the first ten years of antiproton source operation (the first internal leak occurred in 1990), but have occurred more than a dozen times since (five times this run). The leak rate is generally low and LCW lost from the system is not a primary concern. Of more concern is that magnets with internal water leaks tend to eventually ground-fault, probably due to a breakdown in the epoxy surrounding the coils. Examining the magnets that failed indicated that most leaks occur at splice joints in the coils. These joints either were drilled off-center in the coil, leaving a thin wall that could erode through abrasive action or had inadequate solder, possibly due to chemical action. Since most magnets have multiple splice joints, there are probably a number of additional magnets that will develop internal leaks. None of the magnets replaced during the E835 run ground-faulted, the magnet changes were planned in advance and often done in parallel with scheduled downtime.

Two other magnets, SQE007 (IQ33) and SQC304 (D6Q18), were replaced due to chronic overheating. The replacement of IQ33 was motivated by the severe discoloration of the epoxy surrounding the coils. After years of overheating problems there was concern that the magnet would fail. An unplanned replacement would have been time consuming due to its

find a cause for this magnet's chronic overheating problems. Unfortunately no other SQE spares were available so they were reluctant to destroy the coils on the only viable spare. D6Q18 was

were spares available. This magnet was indeed opened up and a curl of copper was found wedged at one of the splice joints. SQ magnets that are built in the future will not have splice

LCW leaks external to the magnets occurred periodically during the run as they have

IQ1 located in a high radiation area near the target vault. The plastic tubing used to connect different coils sets in the small quadrupoles tends to become brittle in the presence of radiation.

quadrupoles near the target vault with metal tubing to “rad harden” them.

From the

the fewest stack losses. Collider stack losses are expected to be higher than in experimental mode. Every stack dumped in Collider is considered a stack loss, whether planned or not. Stacks

	Weeks	Stacks lost	(per week)	Power supply problems	LCW related problems number lost (per week)
	15	13		1 (0.067)	0 (0)
	49	32		4 (0.081)	2 (0.063)
E760 Run 1	*12*		*0.83*		*2 (0.167)*
	30		*0.33*		*1 (0.033)*
	42		0.48		3 (.071)
	55		1.00		8 (0.145)
	98		1.04		21 (0.206)
	48		0.79 [0.65]		9 (0.187)

Table 3

prerequisite to returning to stacking. Referring back to Table 3, E835 operation resulted in much more frequent stack losses than the E760 run. The differences are not as dramatic if the two

unstable lattice point are ignored. In historical context, it appears that the number of stack losses in the E835 run was consistent with past runs while the E760 run stands out as a period with

During the E835 run there was an unprecedented number of stack losses where the cause was not clearly understood. This was particularly surprising in light of the improvements in the run there were numerous occasions during which the horizontal orbit appeared to rapidly shift, often resulting in beam loss. Orbit differences primarily pointed to the A2B3 and A1B3 locations. Power supply current readbacks for all Accumulator supplies were found to be constant during these episodes of beam loss. The shunt on the A1B3 magnet (A:BS103) was replaced on the theory that the shunt might change behavior without a readback change. The stack losses continued after the new shunt was installed. The A1B3 and A2B3 magnets, which are both the modified B-1 style, measured low inductance late in the run. It is believed that these magnets would at times develop turn-to-turn shorts which would cause an orbit change. These magnets are being removed during the shutdown for repair. It is interesting to note that out of only eight modified B-1 magnet locations in the pbar source, there have been four failures (IBV1 has been replaced twice in previous runs).

Now that magnets have been identified as a likely source of orbit changes, problems encountered at the beginning of the run can be understood. While progress on the above transition ramps went fairly smoothly at the start of the run, many iterations were performed on transition crossing and below transition ramps. At the lower energies it was found that the orbit would spontaneously change at times, shrinking the aperture and shifting tunes and chromaticity. Periodically study time would be required for ramp development to correct these problems. Most of the orbit shifts occurred at the lower current settings found below transition. It is unusual for a turn-to-turn short to manifest itself under low field conditions. Magnet Test Facility personnel will carefully examine A2B3, the first bad magnet removed, as time allows.

Figure 7 is a plot of pbars stacked over the E835 run. This plot is provided to illustrate the frequent interruptions in operation caused by failures, study periods and detector modifications. The E760 run was preceded by a lengthy period of ramp development followed by nearly dedicated experimental operation for the entire run. The E835 run from the beginning was guaranteed to have numerous interruptions even in the absence of failures. The experiment was anxious to begin taking data early in the run and pushed to begin the run with minimal time spent on developing the deceleration ramps. The lack of time allotted at the beginning of the run to develop the ramps was paid for later in reduced efficiency. This was combined with trying to hit a "moving target" caused by the shifting orbits.

The measurement of beam energy came into question over the course of the run due to problems with the orbit length calculation. The experiment desires a high precision measurement of the beam energy which in turn requires an accurate orbit length measurement. On several occasions a Beam Position Monitor (BPM) would not work properly and cause a shift in the calculated orbit length. It was also found that there had been several BPM's with saturated Analog to Digital Converters (ADC's). As the orbit became distorted, presumably caused by the two weak dipoles, the ADC readback had become ranged at several locations. The fix for subsequent stores was to improve the orbit and increase the gain (lowering the resolution) on the ADC's. It was much more difficult to correct the energy measurements on earlier stores where the beam position at these locations were unknown.

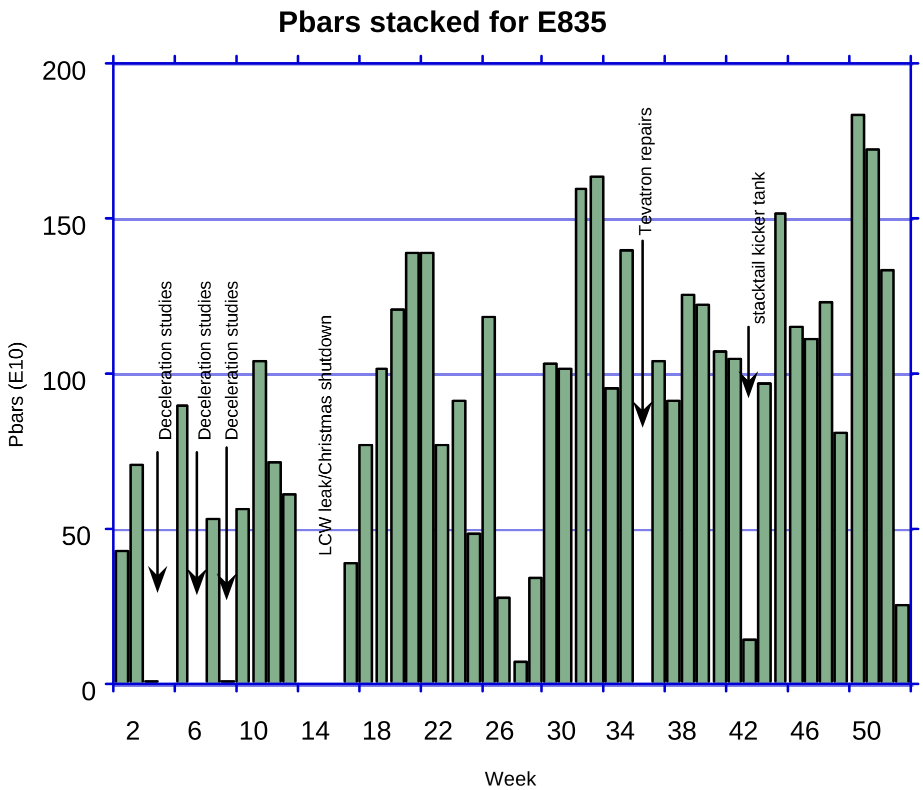

Figure 7

Decelerations

routine decelerations and beam energy changes to the Operators in the Main Control Room (MCR). Although Operators participated in decelerations and beam energy changes beginning

Department handle most of the decelerations and energy changes. Because so few of the stores were considered "routine", the Operations Specialist for the Antiproton Source (myself) assisted

There were several factors that combined to force physicists to supervise the majority of the beam manipulations. The most serious problem was the lack of consistent machine

made under identical circumstances. It is now believed that the two malfunctioning dipoles, A1B3 and A2B3, caused shifting orbits from day to day resulting in a reduced aperture. The

deceleration efficiency. Also, the operating schedule dictated by the experiment involved frequent changes in the target energy for stores. There was little opportunity for the Operators to

Control of the deceleration and beam energy change processes reached a high state of

orchestrated most of the required steps. A deceleration sequencer provided the specialized commands for loading and executing the appropriate ramp tables. The collider sequencer already had more generic commands available that could be used in conjunction with the deceleration sequencer. Despite the complexity of the controls required for deceleration, problems were infrequent and quickly corrected.

Early in the run a commitment was made to switch from a paper log book to an electronic log book for documenting machine information. There were several advantages to an electronics log: Easy to read, entire graphics images saved into the log, back-ups of the log to protect against lost data and searching capabilities. Metacard was an electronics log that was in use by the Booster Department at the time so it was adopted for use by the Antiproton Source Department. A programmer was assigned to improve the functionality of Metacard and address any complaints lodged by the users. Despite all of the advantages of the electronics log in principal, in practice the log was somewhat difficult to use. Even after numerous refinements the log was sluggish in initializing and responding, prone to crashing with accompanied loss of the current entry and full of odd idiosyncrasies that steepened the learning curve. Most felt by the end of the run that an the idea of an electronics log was good, but Metacard was not a good enough product for our needs.

Above transition deceleration efficiencies during E835 compare favorably with similar decelerations during E760. Typical efficiencies were in the 94-98% range for both experimental runs, but the stacks were typically 10-20E10 larger during E835. Below transition decelerations were a different story with disappointing efficiencies during E835. Typical below transition deceleration efficiencies during E760 were in the 90-95% range compared with only 60-75% during E835. Stacks were smaller for E760, mostly less than 22E10 compared to 25-35E10 for E835. When stacks were larger than about 30E10, longitudinal instabilities associated with the transition crossing resulted in excessive emittance growth. The momentum spread of the beam became too large for the stochastic cooling to recover it all. The stack sizes used in E760 were small enough so that the longitudinal emittance growth was manageable. This would still only account for 5-10% of the discrepancy in efficiency between the two runs.

Intermediate cooling steps were required when going to the lowest energies, such as the eta-c. Deceleration efficiency on the earliest efforts were in the 50-60% range and required 5 hours or more to accomplish. There were numerous examples of orbit shifts which reduced efficiency, apparently caused by the two malfunctioning B-1 magnets. On several occasions ramp development shifts were required to return orbits to nominal. With orbit improvements and operational experience the decelerations improved. Typical decelerations to the eta-c during the latter part of the run were typically in the 65-75% range and took 2-3 hours. Because of the problems with longitudinal beam growth while passing through transition stacks were generally limited to 35E10 or less. Beam energy changes starting and ending below transition were normally above 80% as long as the beam was sufficiently cooled.

Considerable study time was spent trying to identify the source of the longitudinal beam growth while crossing transition. Although considerable useful information was gathered, little was learned about the specific cause or ways to reduce the growth. Beam was very unstable on either side of transition as well, efforts to maintain stores near transition to study the χ_0 resonance mostly ended with catastrophic beam loss. To even achieve modest success, the strategy was to spend as little time near transition as possible with as little beam as could be managed. These difficulties led to a number of extremely small stacks reaching the eta-c during early efforts with large set-up times. Later in the run decelerated stack sizes had increased to the 20-25E10 range for stores at the eta-c.

The experimenters frequently required a beam energy that would have less than 50 keV variation over the course of the store. This was very challenging to achieve, particularly maintaining a constant beam energy at each point while performing frequent beam energy changes. At the beginning of the store when the gas jet is first turned on, the beam energy would drop sharply by 300 keV or more in the first hour. The two methods used to compensate for this effect were either to intentionally start the store with the beam energy too high or to make a rapid momentum cooling pickup position change when the jet was turned on. During the course of the store the experiment would at times cause beam energy shifts by making large spontaneous changes to the gas jet density.

Without intervening, the beam energy during a store would drift upwards about 200 keV over the first 10-15 hours of the store. This appeared to be caused by the magnets cooling with the reduced heat load, causing the bfield to rise as the magnets shrunk. This would not directly affect the beam energy, but the shift in position through the momentum cooling pickups would result in an energy shift. To combat the problem, Operators would need to make periodic small changes to the 4-8 momentum pickup position, fortunately controlled by a motorized stand. There were some below transition points where the 4-8 momentum system would not work efficiently and the 2-4 momentum system would be used instead. The pickups for this system are not controllable requiring an orbit bump at the pickup location. The energy range allowed on most of these data points was broad so adjustments were usually not required. If future experiments require more precision at the lower energy regimes, it may be necessary to provide motion control for the 2-4 momentum pickups.

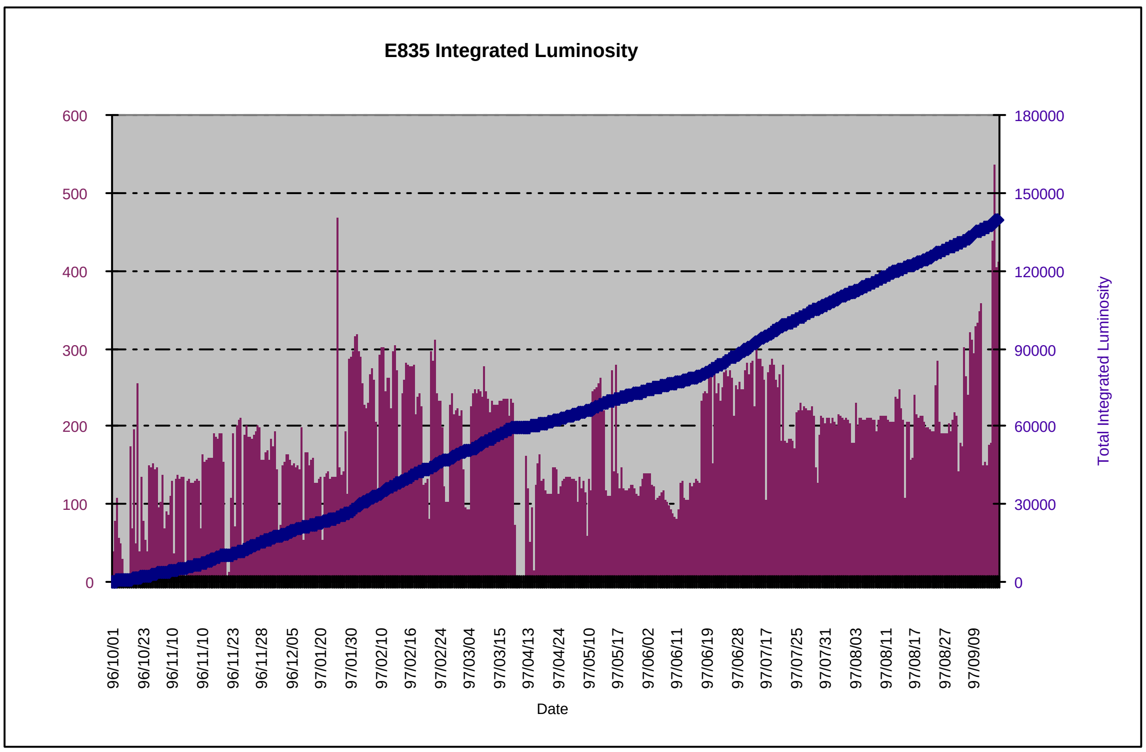

Figure 8

Difficulties with beam energy drift were most common in short duration stores and stores with multiple beam energy changes. Long duration stores at a single energy point tended to be very stable. The experimental program for E835 required shorter stores and more frequent beam energy changes than that used for E760. Manipulating the momentum cooling pickup was the only means successfully used to control drifting beam energy. This required constant monitoring of the beam energy with the burden of the corrections falling on the accelerator physicists and operators. I would recommend that any future experiments be responsible for correcting their own beam energy drift. The experimenters have a better perspective as to what level of drift they can tolerate on a particular store.

Luminosity

The ultimate measure of the operational success of an experimental run is how much luminosity was integrated. An experiment's luminosity not only reflects how well the accelerator is performing, but also includes inefficiencies of the experiment. Figure 8 is a graph of the integrated luminosity for E835 over the course of the run. The luminosity goal for the E835 run was a somewhat optimistic 200 pb^{-1}. The integrated luminosity for the run was actually 143 pb^{-1}, so the goal was not achieved. Table 4 summarizes the run duration and integrated luminosity for the E760 engineering run, the E760 physics run and the E835 run. There was nearly 4 times as much luminosity integrated in the E835 run than in both E760 runs combined. Figure 9 shows the luminosity for each run normalized to the length of the run, obviously there was a substantial improvement in delivered luminosity between E760 and E835.

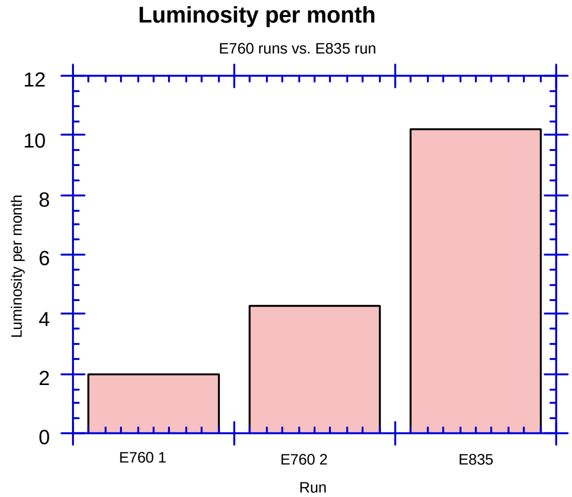

Figure 9

	Run duration (months)	Luminosity (pb^{-1})
E760 (1)	3	6
E760 (2)	7	30
E835	14	143

Table 4

Conclusion

There were few periods of uninterrupted beam to the experiment over the course of the E835 run. A number of large blocks of downtime as well as study periods punctuated the run. The interruptions required "starting over" with a machine that would behave differently and prevented operation from ever becoming routine. When considering the success of the run one must consider the requirements, which included frequent decelerations below transition, high precision energy scans, increased gas jet density, stores near the Accumulator transition energy, Collider Run II studies, a holiday shutdown, coexisting with another experiment and the mid-run installation of a forward calorimeter. Even with a minimum of downtime there would have been fewer weekly store hours for E835 than E760. When you add in the unusually high number of serious failures, and weak magnets that created an ever-shifting closed orbit, it is remarkable that as much luminosity was provided to the experiment as there was. If there is a future experimental run, ramp development should be completed before the experiment begins serious data-taking. It is too difficult playing "catch-up" with ramp corrections over the course of the run. From an operational standpoint it would be advantageous to do all of the data taking in an energy range over consecutive stores. The practice of jumping from one type of deceleration to another often results in longer set-up times. Well defined goals and long range planning by the experiment would result in increased luminosity. Finally a significant amount of downtime over the course of the run was related to degradation in the magnets and the LCW system. It is important to improve the deteriorating pbar infrastructure, especially the magnets themselves, during this shutdown for the Main Injector installation.