Fermi National Accelerator Laboratory

FERMILAB-TM-1970

Fermilab Collider Run 1b Accelerator Performance

V. Bharadwaj, M. Halling, P. Lucas, E. McCrory, S. Mishra, S. Pruss and S. Werkema

Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510

April 1996

Operated by Universities Research Association Inc. under Contract No. DE-AC02-76CHO3000 with the United States Department of Energy

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

FERMILAB-TM-1970

Fermilab Collider Run 1b Accelerator Performance

V. Bharadwaj, M. Halling, P. Lucas, E. McCrory, S. Mishra, S. Pruss, S. Werkema

October 1st, 1995

Introduction

This report summarizes the performance of Run 1b as of the end of July 1995. This run is the conclusion of Fermilab Collider Run 1, which consists of Run 1a (May 1992 - May 1993) and Run 1b (January 1994 - February 1996). Run 1b is characterized by being the first with the new 400 MeV Linac. At this time, the run is not yet complete. Colliding beam physics is scheduled to resume after the summer 1995 shut down and continue until mid-February 1996. We do not, however, believe that the upcoming running period will significantly modify the performance that is being reported. All of the operation to date is at a Tevatron energy of 900 GeV.

This report emphasizes performance numbers and the various improvements made to systems to achieve this performance. It will only discuss the underlying physics to a limited extent. The purpose is to provide a reference document with which to dispassionately compare performance during future runs. The report has the following distinct sections:

1. Run Statistics p. 1
2. Instrumentation and Controls Issues p. 4
3. Proton Source Performance p. 8
4. Antiproton Source Performance p. 11
5. Main Ring Performance p. 28
6. Tevatron Performance p. 33
7. Summary p. 53

1. Run Statistics

The goals set for collider Run 1b were predicated on the benefits expected from the Linac Upgrade. These goals specified the delivery of an integrated weekly luminosity of at least 2.0 pb^{-1} with collider stores having a typical initial luminosity of 1×10^{31} $cm^{-2}\bullet sec^{-1}$. The plots that follow in this section document the struggle to achieve these goals. The first two plots show the initial luminosities (Figure 1) and the proton and antiproton intensities (Figure 2) for each store during the entire three year period of Run 1. Figure 1 shows that it took nearly as long to exceed the Run 1a performance during the Run 1b startup as it had taken to achieve it the first time. Much of the startup difficulty was caused by a rolled Tevatron quadrupole magnet in the B0 low beta insert. The goals set for Run 1b were achieved almost immediately after the realignment of this quadrupole. Only a month later, the entire month of September, 1994, was lost because of an LN_2 supply interruption. It took several months to completely recover from the subsequent unscheduled shutdown and Tevatron warm-up. In March of 1995, following an upgrade of the Main Ring RF bunch coalescing system, initial store luminosities exceeding 2×10^{31} $cm^{-2}\bullet sec^{-1}$ were finally achieved.

In addition to individual store initial conditions, one must consider reliability and integrated luminosity. Figure 3 shows that typically 70% of stores were ended

intentionally. Those stores which were not intentionally terminated were lost due to a variety of different categories of system failure (see Table III in the Tevatron Performance section). Figure 4 shows the weekly integrated luminosity for Run 1. The integrated luminosity generally follows the initial luminosity (Figure 1) with the exception of some poor weeks near the end of the run.

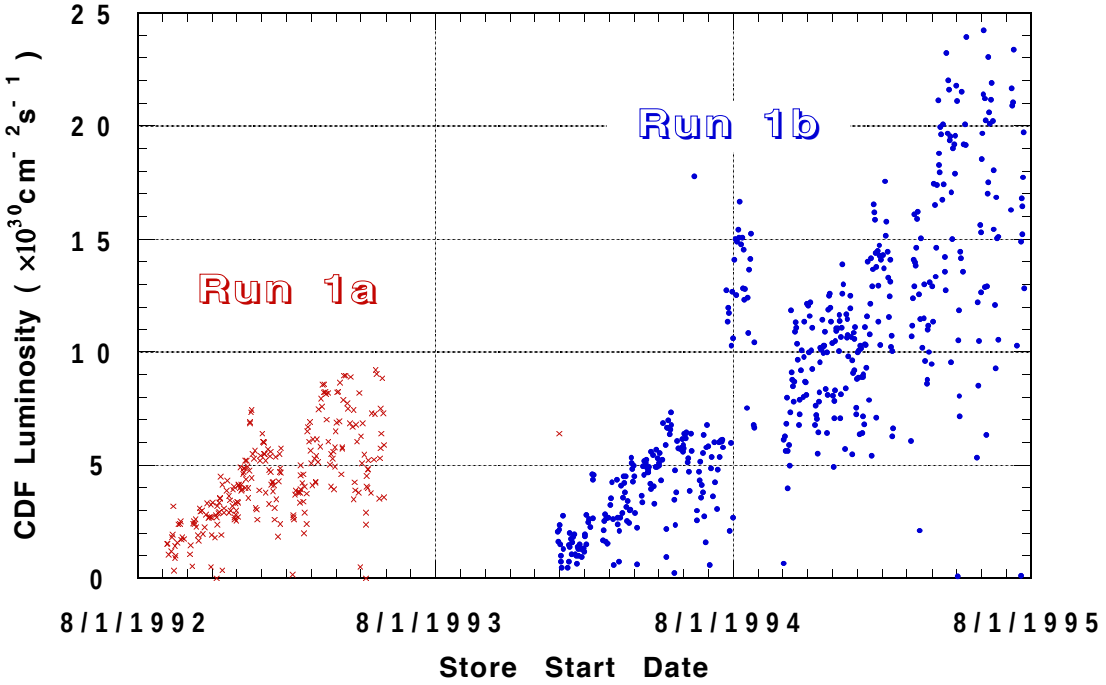

Figure 1. Initial luminosity as reported by CDF for all the stores of Run 1.

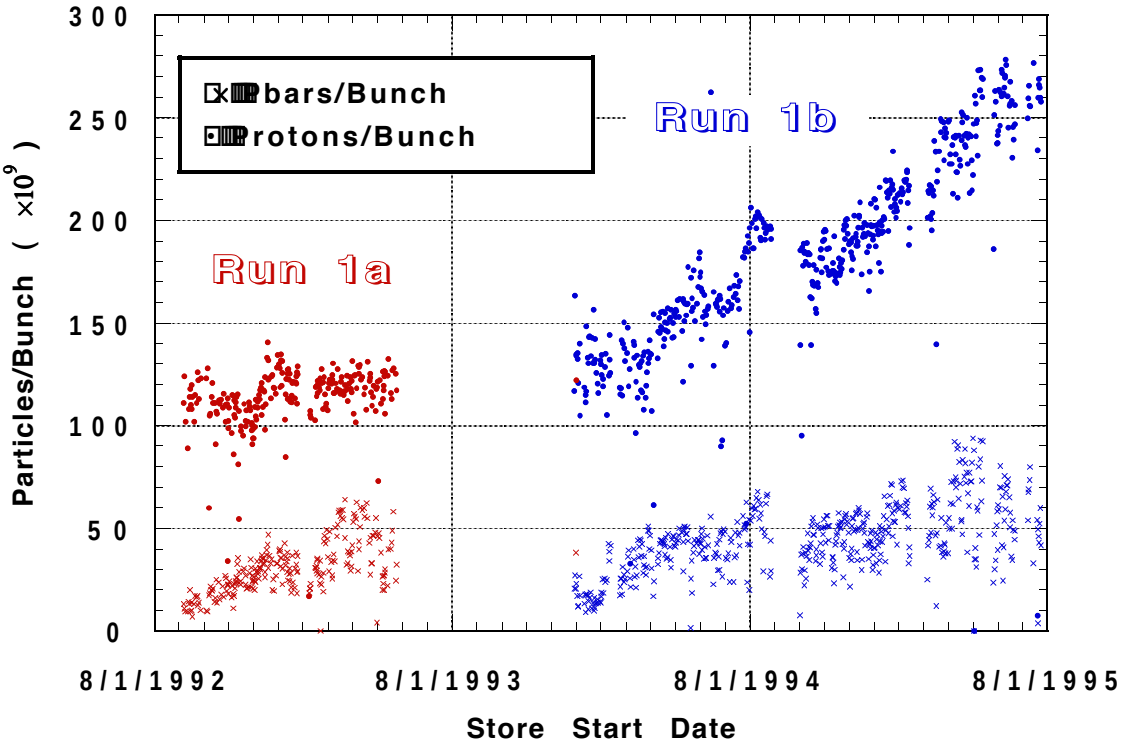

Figure 2. Proton and Antiproton bunch intensities for all the stores of Run 1.

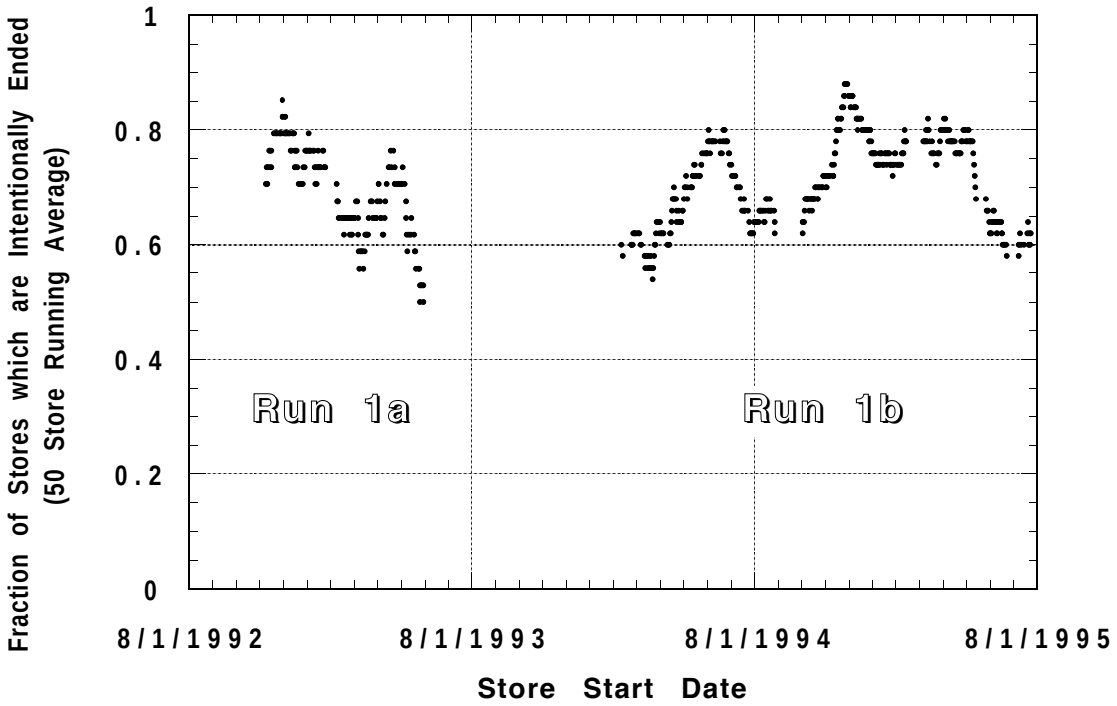

Figure 3. Running average over previous 50 stores of fraction of stores ended intentionally by the run coordinator.

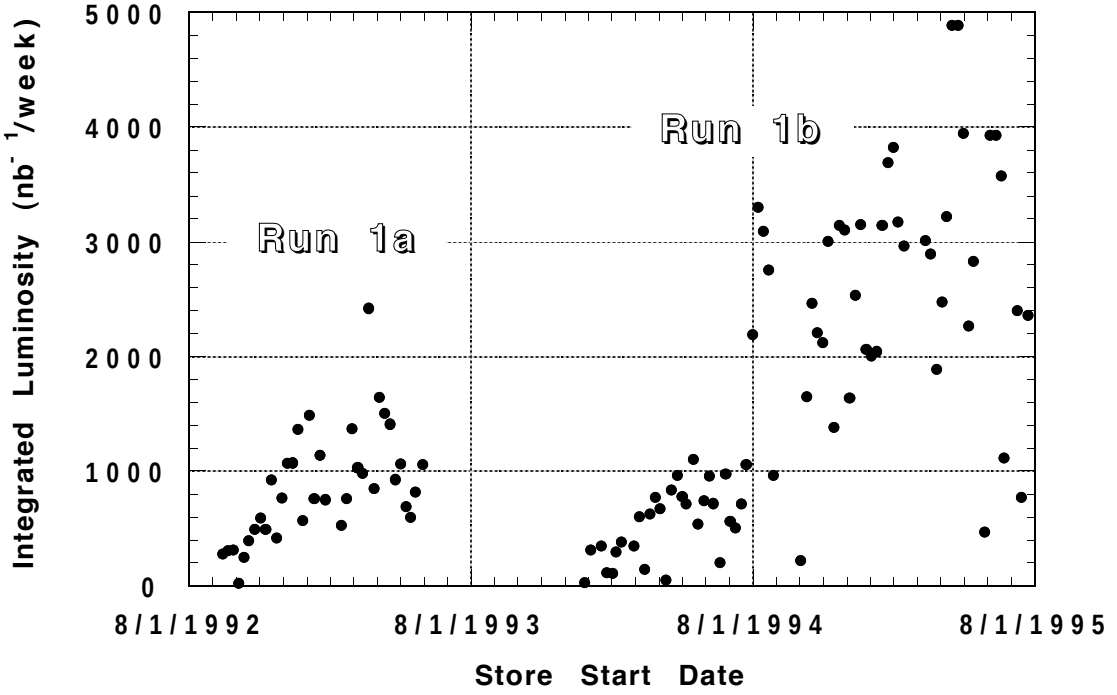

Figure 4. Integrated luminosity/week for all of Run 1 as reported by CDF.

2. Instrumentation and Controls Issues

Discussed here are those pieces of instrumentation used to measure beam positions, profiles and thus transverse emittances, longitudinal emittances, and intensities. There are several devices, particularly in the low energy machines, which have been doing these tasks since before Run 1b, and have continued to function relatively unaltered during it. Among these are emittance probes and wire scanners in the Linac, multiwires in the 400 MeV and 8 GeV lines, the ion profile monitor (IPM) in the Booster, and SEMs in the antiproton source beamlines. Devices in the 400 MeV line were, of course, constructed for Run 1b, but the multiwires operating there are little different from their predecessors. The Main Ring, Tevatron, and Accumulator rings have DC current transformers (DCCTs) which similarly have continued to function as in the past.

There are also a number of devices which were either new for this run or reached maturity during it. There are others for which replacement is being considered and are worthy of comment here.

• *Beam Position Monitors (BPMs)*. The 400 MeV line has BPMs which are the vanguard of the systems of the future. They are able to digitize at a high rate (5 MHz), and come with deep memory buffers for the storage of a number of readings. Similar monitors are installed in Booster, appropriately configured for the Booster revolution time, but have been operational to this point only as prototypes. They are to replace the original monitors shortly. The BPMs of the antiproton source are to be replaced with a modern version, being constructed under an AIP project but without any impact on Run 1b. There are vague plans to upgrade the Main Ring and Tevatron monitors as well, but in the more distant future.

• *Flying Wires.* These devices are used to measure the beam profiles in the Main Ring, Tevatron, and Accumulator; the Accumulator wires having been installed during this run. There are plans to replace major sections of the existing systems with improved versions, addressing problems involving both mechanical movement and mathematical data analysis.

There is a mechanical problem with the wires apparently related to the fact that the apparatus is undergoing an angular acceleration as it passes through the beam. The rotational position is read from the shaft of the apparatus, so that if the assembly were completely rigid a proper profile would be determined despite such an acceleration. However there is apparently a mechanical distortion, in particular of the wire, so that the encoded shaft position does not properly correlate with the disturbance of the beam. The result of this problem is that there exists an 'up-down' effect. Namely as the wires are alternately flown upwards and downwards through the beam, the measured position and width correspondingly alternate between two values. An improvement was made during Run 1b in going from a belt to a direct drive; however some effect remains. The next level of improvement will involve a fly through 540 degrees of arc, with the data taken where no acceleration or deceleration is occurring. The wire will pass through the beam twice and the data from both passes will be used.

During the present run the wires have been the chief source of information on beam profiles and thus transverse emittances. However, the lattice β's at the wires are not well enough known to allow accurate emittance calculations to be made. Day to day variations do seem meaningful, and these are what are used in practice. In the case of horizontal emittances lattice dispersion functions as well as β functions must be known, and in general these values are imprecise and are not used in any significant way. Such statements hold for any transverse profile monitors, not only the wires. The lattice uncertainties affect the results obtained from the instrumentation; they are not related to the instruments themselves.

• *Synchrotron light monitors.* The Tevatron has two synchrotron light monitors, one each for the protons and antiprotons. These devices operate on the principle that light is emitted by the beam as it is accelerated, in particular at the edges of magnets. It is focused onto a photosensitive plate with image intensification, and leads to a depiction of the two dimensional transverse beam intensity profile. The devices were under development throughout the run, and by the end were reliable enough to be used for some diagnostic purposes. These monitors may eventually replace the wires for some emittance measurements, as they give good results in a non-destructive manner. The light output of the beam is significant above about 600 GeV, so the monitors are useful during normal energy stores, but do not provide information at injection.

• *Bunch intensity monitors.* Bunch intensities are measured by two types of devices, Fast Bunch Integrators (FBI) and Sampled Bunch Display (SBD). Although neither device was new for this run, they did mature and come to the point where they were used quite regularly to monitor a variety of effects. The FBI are able to return integrated intensities of each of the several bunches in the Tevatron with a 1 Hz update rate to a standard plotting package. The SBD is a fast digital scope with readout to a LabView Macintosh, where data analysis is performed before the results are transmitted to the control system. Although the SBD activity is slower, the results are more sophisticated than those of the FBI. Namely, the scope is fast enough to make a number of measurements of intensity during passage of a single bunch. Thus it is able to determine longitudinal emittance as well as intensity.

• *Ion profile monitors.* There are such devices in both the Booster and Main Ring. They function by applying an electric field to drift the ions, resulting from the collisions of beam particles with the residual gas in the beam pipe, to a detector. From the transverse distribution of those ions, the corresponding profile of the beam can be inferred. One of these devices has been operational in the Booster for a few years. It has indicated a number of interesting effects and has given insight into transverse emittance growth in that accelerator. The IPM in the Main Ring was installed late in the run, and is in a commissioning stage. As designed it can produce a plot of beam width vs. turn number for an entire cycle.

The excellent operation of the Fermilab ACNET control system has been one of the success stories of Run 1b. ACNET since its inception with the construction of the Tevatron has had a number of strengths which have persisted during this run. Included among these are uniform operation of all accelerators from any console, minimal delays in switching between programs being run, and a strong application programming library and the resulting general uniformity in the user interface and operational aspects of the (hundreds of) various programs. Built on this solid base, a major series of upgrades was instituted in the late 1980s, and a program of creating collider specific software utilizing the workstation consoles of this upgrade program was undertaken preparatory to Run 1a. Many of these additions were still on their learning curves during that earlier run, but are now in a mature state and operating well.

As with any such large system there are problem areas and places where improvements might be realized. The following items represent an attempt to summarize both the present situation and the major areas of concern.

• *System reliability.* As to software, the number of downtime entries pertaining to failures of computers and individual software processes has reached a very low and presumably quite acceptable level. The CAMAC front end computers of the controls upgrade were in the process of being phased in during Run 1a and did cause problems at that time. They have now settled into a mode of very reliable operation, while processing amounts of information inconceivable in the days of their predecessors.

Hardware operation is generally reliable, in particular with the error rate on the real-time clock system having been reduced to negligible proportions. One problem which does exist and is in the process of being investigated is that of CAMAC crate power failures and their frequent root cause, that of overheating. Occasional stores and stacks are lost due to these problems. A major fraction can be traced to failures of cooling fans, and the installation of 'smart fans' in vulnerable crates is being investigated. Such equipment will emit an alarm on fan failure, and will use a certain amount of redundancy to maintain reliable operation until repairs can be effected.

The central computers, some Main Control Room consoles, and the networking electronics connecting them are protected by Uninterruptable Power Supply (UPS) devices. These operated as designed during the recent power outages.

• *Software tools.* The library of routines available to an application programmer, and the applications which utilize this library, represent a strong point and a success. The time required to create a new program is decreased by a large factor from what it was previously, due to the tools available; tools involving both data acquisition and the user interface.

The Sequencer program, in use since the earliest days of Fermilab colliding beams, has become indispensable in performing shots in a reliable and consistent manner. It so automates the process that, when non-standard actions have been taken, human errors brought on by having the automation not available have been common. These errors have decreased of late as the Sequencer has been made easily modifiable and as users have become aware of the pitfalls.

The ability to 'script' a program, to play back at a later time a series of operational actions, has become a standard technique.

• *Collection of massive amounts of data.* During Run 1b the original Datalogger program, residing on a server node, was retired, having been entirely replaced by a series of distributed 'Lumberjack' loggers. The logged data are stored on local disks of the consoles, but are available globally. The number of values regularly being logged has grown to 6500, an impressive figure. An ability has existed for some time to log settings of quantities as well as their readings; it appears that this capability is not being utilized to the level which had been hoped.

The shot data collected by the Sequencer Data Acquisition (SDA) has also grown, in this case by roughly a factor of three over earlier runs. It is often questioned whether these data are of value, but the fact remains that they are collected only because someone has explicitly requested them. There is a popular notion that SDA slows shot set-up to a significant degree; Controls personnel vehemently deny this statement. There has been discussion of moving from the standard file format to use of a relational database for archiving SDA data. Questions of long range plans in this area deserve to be addressed in more detail by the Accelerator Division as a whole.

• *Open Access Front End.* The so called Open Access, or sometimes Modeling, Front End has come into popular use during Run 1b. This computer is able to run models and specific calculations, and to produce complicated derived values, all of the results of which are available to the control system as are normal hardware data. Many collider calculations are performed in this manner. This machine also serves the more prosaic purpose of allowing a new program to be tested without actually referencing accelerator hardware.

• *Outmoded computers.* The control system still utilizes in essential areas three Lockheed MAC and two PDP-11 processors. The risk of using such machines, where hardware and software maintenance abilities have decayed to very low levels, is substantial.

• *Keepers of application programs.* The Accelerator Division has failed to come to grips with the problem of orphaned, i.e. no longer maintained, applications. While some programs have outlived their usefulness and no longer require maintenance, often some of the most important ones are simply abandoned. At present one of the most notable examples is C2, the SDA data display and analysis task. It is little changed from the version of Run 1a, despite the increase in SDA data collected.

• *LabView.* Several LabView instruments are functional and are connected to ACNET. However the level of communication is not complete, and the ad hoc approach to deciding which features to support makes it difficult to understand the future of either the instruments or their controls.

• *Computer system.* As constructed, ACNET, particularly the operators' consoles, is crucially dependent on the byte ordering and floating point format of the computers used, namely VAX stations. It is clear to all concerned that eventually the control system must migrate to some more modern platform. Such a migration will have to be evolutionary and gradual, the new system probably best thought of as the controls of the next major Fermilab project beyond the Main Injector. The ability to operate the accelerators efficiently would be badly degraded by any precipitous changes to this mature control system.

3. Proton Source Performance

The Operation of the Proton Source over the course of Run 1b has exceeded the goals set at the beginning of the run. The beam brightness has steadily increased for most of this period. Of course, the primary source of this improvement has been the presence of the new 400 MeV Linac. Other significant aspects of the improvement have been: (1) The gradual tune-up of the Booster to take advantage of the higher Linac energy; (2) Commissioning of the various Booster damper systems; (3) Improvements in the H-minus Ion Source.

The new Linac and the re-worked Booster were commissioned at the beginning of Run 1b, and it took about 6 weeks to obtain 3.1×10^{12} Particles Per Pulse (PPP) out of the Booster (on Nov. 9, 1993) [1].

The beam current of the Linac is shown in Figure PS-1. Note the steady increase in the delivered beam current over the course of the run. On Aug. 15, 1995, the extraction voltage was turned up in a way that had never before been tried. Subsequent tuning of the correction elements and of the Linac quads has adiabatically increased the current in the Linac to 45 mA. Increasing the Linac intensity is directly correlated with increased Booster extracted current. However, Booster has been unable to run above about 9 turns for antiproton production independent of Linac intensity.

Figure PS-1, Linac Beam Current for Run 1B

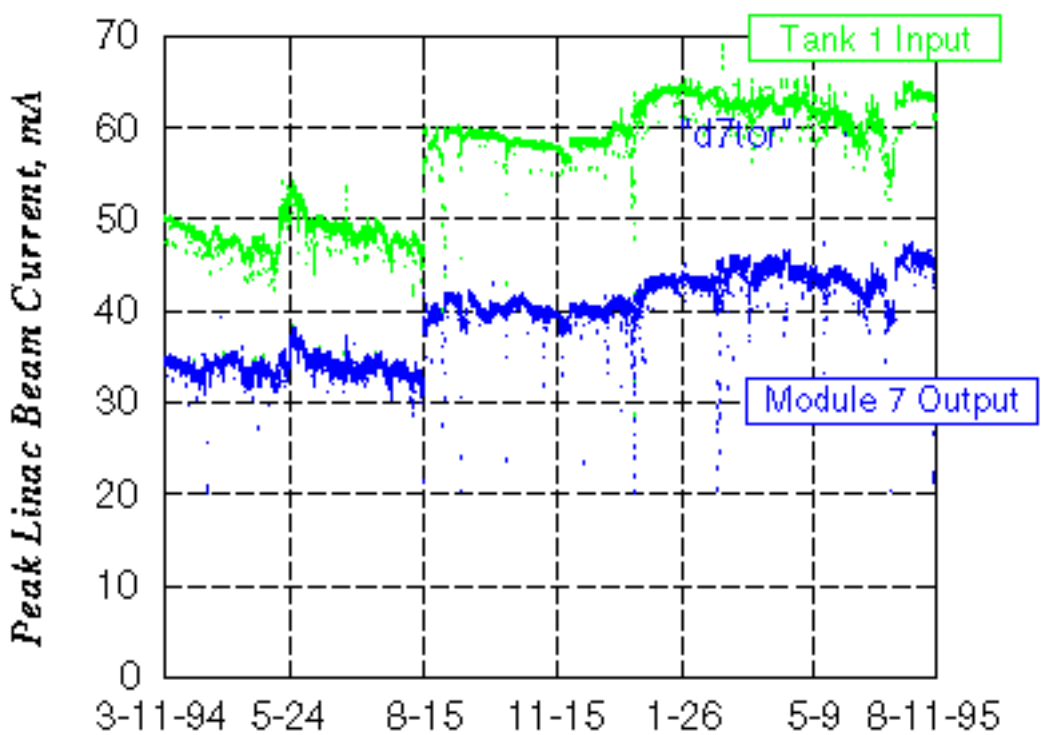

There were many improvements to the Booster after initial commissioning. (A paper on this subject [2] should be consulted for further information. The summary of that paper follows.) On Feb. 15, 1994 a transverse mode damper system was implemented in the

Booster to improve the extracted current. This damper suppresses the betatron-induced 180 kHz horizontal instability due to multi-bunch "mode 7" which probably arises from a resistive-wall impedance. This Booster mode damper permitted the achievement of 4×10^{12} PPP at Booster extraction. However, the 8 GeV emittance was too large to be useful for antiproton production. By the end of the run, another fixed-delay damper was in place operating on the last one-third of the 15 Hz cycle. This system allowed the extraction of high intensity (4×10^{12} PPP) at emittances small enough for antiproton production cycles.

A systematic study and re-alignment of the Booster magnets was completed in February of 1995. These magnet moves increased the aperture of the Booster to approximately 26(H)×12(V) π mm•mrad. The moves having the greatest impact were those of the septa at Long-3 and at Long-13. After these moves, the correction elements throughout the Booster were hardly being used, and the activation of the RF cavities was decreased measurably. There remain questions regarding the accuracy of these aperture measurements, although the general character of the measurements was certainly sound.

Several key features of operations are listed here:

- The lifetime of the H-minus ion sources has continued to be approximately six months.

- Of the seven 12-MW klystrons in the Linac, there were no failures.

- Of the three 0.2-MW klystrons in the Linac, two were changed out, but not for operational reasons (for practice and to refit some revised equipment).

- Long-term (software) compensation of the high-energy Linac cavity gradients has been necessary for best performance.

- In order to operate Booster above 3.5×10^{12} protons per pulse (PPP), the longitudinal (analog) dampers are mandatory.

- In June, 1995, the GMPS (powering the Booster magnets) regulation was changed from voltage to current.

At this time, the performance of the Linac and of the Booster have been measured to be as follows:

Linac

- Beam Current: 45.6±1.0 mA (L:D7TOR, sampled once per hour, averaged over the month of July, 1995).

- Transverse Emittance: Estimated at 7 π mm mr at 95%, both planes.

- Longitudinal Phase Space: Bunch Length = 14 psec rms (805 MHz bunches, July 1995) Momentum Spread = 0.3% (95%, measured from spectrometer).

- Beam "chop" times: HEP On = 1937 μsec
  HEP Off = 1977; the Booster chopper off pulls off low-energy chopper
  NTF On = 1910, Off = 1967

- Operational Reliability: 97.8% overall
  98.6% Sept. 94-June 95

Booster

- BEAM CURRENT: 4×10^{12} PPP (average, 8-10 turns)

- Transfer Efficiencies: 400 MeV: ~99%
  Overall: 71% @ 8 turns; >90% @ 2 turns

- Transverse Emittance: H<15 π mm mr, 95%, V<13 π; @ 2 turns: 10 π.

- Longitudinal Phase Space: Momentum Spread = 0.07% ± 0.1%
  (Bunch length measurements not available)

- Injection Aperture: H: 26 π mm mr; V: 12 π mm mr (See text).

- Begin taking Linac beam, 10 μsec after beginning of Linac beam

- Operational Reliability: ~98%

References

[1] The Commissioning and Initial Operation of the Fermilab 400 MeV Linac, E. McCrory, 1994 International Linac Conference (Tsukuba, Japan), p 36.

[2] The Fermilab Booster After the 400 MeV Upgrade, S. Shukla, 1994 DPF (Albuquerque) Proceedings, World Scientific, p 1485.

4. Antiproton Source Performance

Introduction

The most important limitation to the luminosity integrated by the Tevatron collider is the ability to produce and deliver to the collider abundant quantities of antiprotons. The length of time required to accumulate a usable number of antiprotons is a significant factor in the management of collider operations. In this section we discuss the issues relevant to antiproton stacking and issues relating to the transfer of antiprotons to the collider (unstacking).

I. Stacking

The best reproducible antiproton stacking rate achieved during Run 1b was 7.2 mA/hr for stacks less than 50 mA (1 mA = 10^{10} antiprotons). For stacks larger than 50 mA, the "best stacking rate" falls off with increasing stack size at a rate of approximately 1 mA/hr for every 50 mA increase in stack size (see Figure 1). The "best stacking rate" is the rate that is achieved when everything is working well and the antiproton production target is being delivered beam at the maximum intensity and rate (i.e. NTF is not running and there are only antiproton production cycles - $29's - in the time line). A reasonably accurate, albeit completely empirical, parameterization of the best stacking rate as a function of stack size is given by:

$$R(I) = R_0 \,\mathrm{sech}\left(\frac{I}{I_m}\right) \tag{1}$$

where R is the stacking rate, I is the stack size, R_o is the stacking rate extrapolated to zero stack size, and I_m is the stack size at which the stack rate is 64.8% of R_o. The best stacking of Run 1b is characterized, in this parameterization, by an R_o of 7.4 mA/hr and an I_m of 192 mA. This represents an 85% increase in performance relative to that at the beginning of Run 1b.

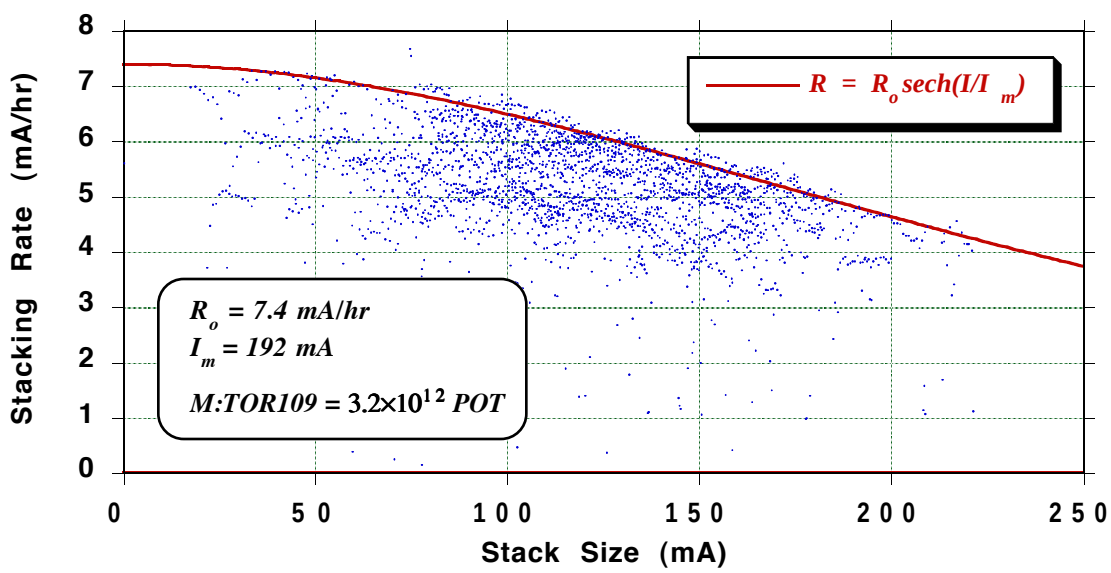

Figure 1. Stacking rate versus stack size for the month of April 1995. One point is plotted for every super cycle during stacking. This month contained the best stacking of Run 1b.
POT = Number of protons on target per pulse.

The beginning of run stacking performance was established in December of 1993 and January 1994. At that time $\bar{p}$ stacking was characterized by an R_o of 4.0 mA/hr and an I_m of 155 mA (Figure 2). The increase in R_o during the run is largely due to an increase in the

number of protons on target attributable to the Linac upgrade and subsequent Booster improvements [1]. The proton intensity on target at the beginning of the run averaged 1.8×10^{12} protons/cycle†, while at the end of the run it was 3.2×10^{12} protons/cycle - a nearly 80% gain. The increase in the value of I_m is due to improvements in the Accumulator which will be discussed later.

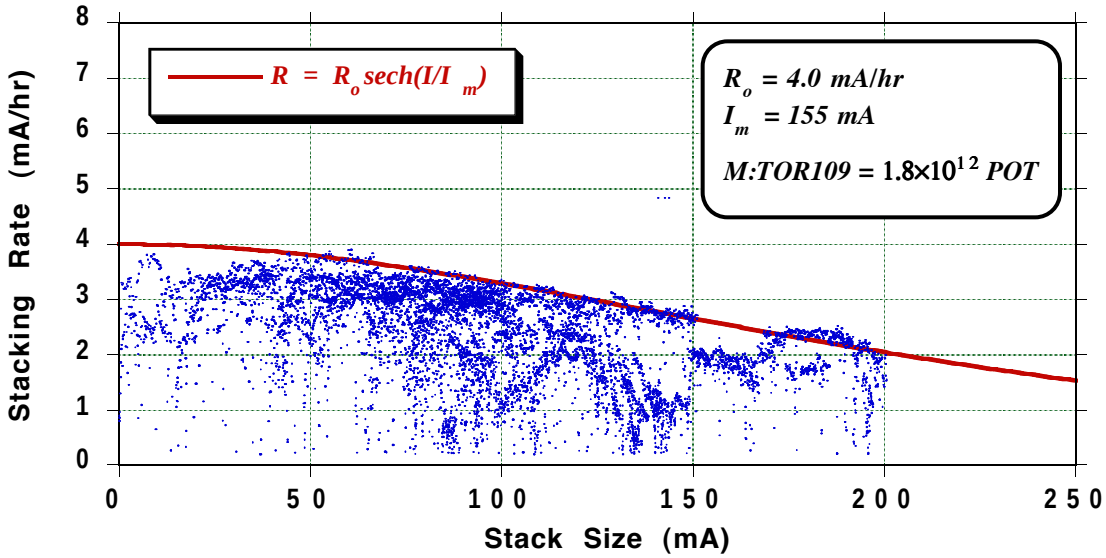

Figure 2. Stacking rate versus stack size for the month of December 1993.

A. Typical end of run stacking performance

The rate of antiproton production is the product of the yields and efficiencies associated with every step in the antiproton production chain. Table I gives a summary of several quantities which characterize the stacking performance of the antiproton source.

Table I
Antiproton Source Stacking Yields and Efficiencies

Protons on target (POT)	3.2×10^{12}/cycle
Antiproton yield into the Debuncher	21×10^{-6} $\bar{p}$'s/POT
Debuncher to Accumulator transfer efficiency	80%
Stacktail cooling efficiency	90%
Overall antiproton production efficiency	15×10^{-6} $\bar{p}$'s/POT

A description of each of the items in Table I is given below.

Protons on Target

The number of protons on target is measured by a toroid in the AP1 beamline which is located upstream of the target. The ACNET designation of this toroid is M:TOR109.

Debuncher Yield

The yield into the Debuncher is the number of antiprotons injected into the Debuncher, as measured by a longitudinal Schottky pickup in the Debuncher (D:FFTTOT), divided by the number of protons on target from M:TOR109. The D:FFTTOT measurement is triggered 350 msec after injection of beam into the Debuncher.

† This number is corrected so that it reflects an 11% calibration change in M:TOR109 that was made on 20 January 1995.

Debuncher to Accumulator transfer efficiency

The Debuncher to Accumulator transfer efficiency is the fraction of the $\bar{p}$'s injected into the Debuncher (D:FFTTOT) that make it into the Accumulator. The number of $\bar{p}$'s injected into the Accumulator is measured by a longitudinal Schottky pickup in the Accumulator which is triggered prior to moving the newly injected beam to the stacking orbit. The result of the Accumulator injected beam measurement is stored in ACNET parameter A:FFTTOT. The Debuncher to Accumulator transfer efficiency is the ratio A:FFTTOT/D:FFTTOT.

Stacktail cooling efficiency

The stacktail cooling efficiency is defined as the fraction of $\bar{p}$'s injected into the Accumulator which make it to the core of the Accumulator beam momentum distribution and stay there. This efficiency, ε, is given by:

$$\varepsilon = \frac{R}{A:FFTTOT} T_{cycle} \tag{2}$$

where R is the stacking rate (A:STCKRT) and T_{cycle} is the time interval between Main Ring stacking cycle resets.

Overall antiproton production efficiency

The $\bar{p}$ production efficiency is the ratio of the number of $\bar{p}$'s stacked in a given time interval to the number of protons striking the production target in that same time interval.

B. Stacking parameters

There are many tunable parameters in the antiproton source complex that bear directly on stacking. For purposes of documentation, the typical values of a few of the more important parameters are given in Table II.

Table II
Antiproton Stacking Machine Parameters

Proton beam spot size on target (sigma)	0.2 mm
Lithium Lens gradient	750 Tesla/m
Debuncher Bunch Rotation RF Voltage	5 MV
Debuncher Stochastic Cooling Power	1100 Watts/plane
Stacktail Momentum Cooling Power	600 Watts
Accumulator Core Cooling Power	30 - 50 Watts
Stacking RF voltage	70 kV
Stacking cycle time	2.4 sec (0 - 50 mA) to 3.8 sec (> 180 mA)

C. Limits to stacking rate

We discuss here various factors which limit the rate at which antiprotons can be accumulated and the improvements that have expanded the limits during the course of Run 1b.

1. Protons on target

Throughout the course of Run 1b there have been frequent step increases in the intensity of the 120 GeV proton beam delivered to the $\bar{p}$ production target. The $\bar{p}$ stacking rate has generally followed this increase. We have, however, reached the point at which an increase in the number of protons on target will effect less than the same percentage increase in $\bar{p}$ stacking rate. Prolonged running with proton intensities in excess of about 3×10^{12} per cycle has been observed to cause damage to the target [2]. This necessitates the use of a larger beam spot size on the target which in turn lowers the antiproton collection efficiency. Significant upgrades in the target station are being planned to accommodate the much larger

proton intensity and shorter cycle time envisioned when the Main Injector becomes operational.

2. Debuncher stochastic cooling

The efficiency with which $\bar{p}$'s are transferred from the Debuncher to the Accumulator depends critically on the transverse beam size after cooling in the Debuncher (see Figure 3). There is as much as a 10% loss of beam during transfer which would be avoided if the $\bar{p}$ emittances prior to transfer were made smaller. At the present time the Debuncher stochastic cooling system is power limited; that is, the system must be operated at a gain setting which is less than the optimum gain to maintain the output power below a level at which damage to the kicker hardware occurs. Thus, the final emittance achieved prior to transfer to the Accumulator is determined entirely by the emittance of the beam injected at the beginning of the cycle. The only way, during Run 1b, to further decrease the Debuncher beam size prior to extraction is to increase the stacking cycle time, allowing more time for cooling. Increasing the cycle time, particularly for large stacks, also increases the stacktail cooling efficiency in the Accumulator. However, only small increases in cycle time are beneficial. The increase realized in Debuncher to Accumulator transfer efficiency and stacktail efficiency will quickly be overcome by the loss of flux from the target due to fewer proton pulses on target per unit time, thus compromising the overall stacking rate. Figure 3 illustrates the effect of cycle time on Debuncher to Accumulator transfer efficiency and Debuncher emittances.

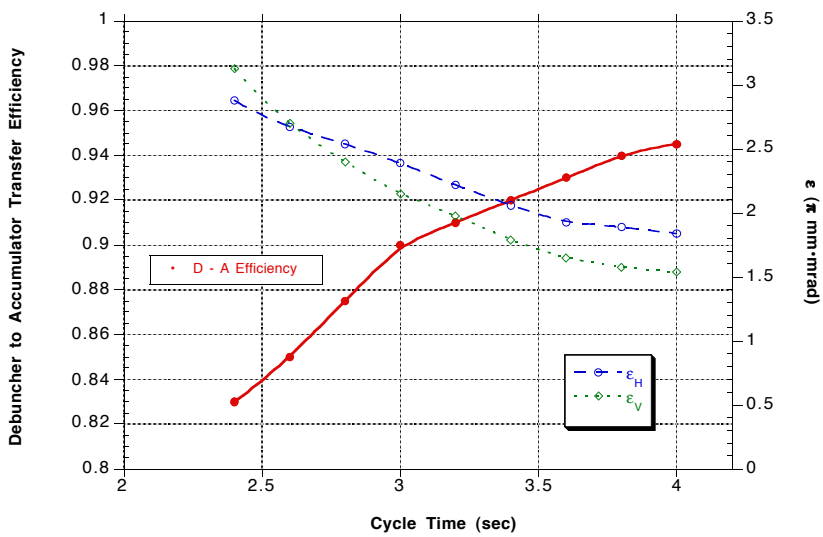

Figure 3. Debuncher to Accumulator transfer efficiency and Debuncher emittances as a function of stacking cycle time. The emittances are measured using a SEM grid in the D-A transfer line (SEM 806). The emittance values plotted in this figure are not normalized. This figure is derived from data taken from reference [3].

3. Stacktail momentum cooling

During stacking the Accumulator functions as follows: Beam from the Debuncher is injected, using a shuttered kicker, onto an orbit at the high energy side of the momentum aperture of the Accumulator. The newly injected beam is decelerated with a 53 MHz (h = 84) RF system (ARF1) to the stacking orbit - an orbit near the center of the momentum aperture. The beam is then stochastically cooled into the existing $\bar{p}$ stack with the stacktail momentum cooling system.

There was a significant amount of effort during Run 1b devoted to understanding and improving the Accumulator stacktail momentum cooling system. There are two issues which have influenced how the stacktail system is operated. First, the gain must be high enough (or the stacking cycle time long enough) to allow the stacktail momentum cooling to

move freshly deposited beam off of the stacking orbit prior to the arrival of the next pulse. Any beam remaining on the stacking orbit will be RF phase displaced backwards, to a higher energy, by ARF1 during the next cycle. Subsequent stacking cycles will eventually cause this beam to be accelerated into the injection kicker shutter. The second issue is transverse and longitudinal heating of the beam at the core of the $\bar{p}$ momentum distribution by the stacktail momentum cooling. In general, the higher the gain setting of the stacktail momentum cooling, the greater the heating of the core will be. A considerable amount of progress was made in mitigating both of these issues during the course of Run 1b.

The rate at which freshly deposited $\bar{p}$'s move into the core is determined by the stacktail momentum cooling gain profile and the shape of the $\bar{p}$ momentum distribution in accordance with the Fokker-Planck equation [4]. The stacktail momentum cooling system consists of a high energy and a low energy pickup connected via independent networks of gain and delay to a common set of kickers. Ideally the high and low energy legs function independently. The high energy leg serves to move newly deposited $\bar{p}$'s off of the stacking orbit while the low energy leg functions primarily to move beam from the stacktail into the core. When the system functions in this manner the total power put onto the beam by the stacktail momentum kickers at any time in the stacking cycle is minimized. The gain profile of the combined system is the vector sum of the high and low energy legs. The overall stacktail momentum cooling gain profile can be tuned by manipulating the gains and delays of the individual legs.

In its design configuration, the high energy pickup was located 16 MeV above the central orbit and the low energy pickup was located 1 MeV below the central orbit. In this configuration there is a significant overlap in the response of the two pickups with the result that the high energy and low energy legs act largely to cancel each other.

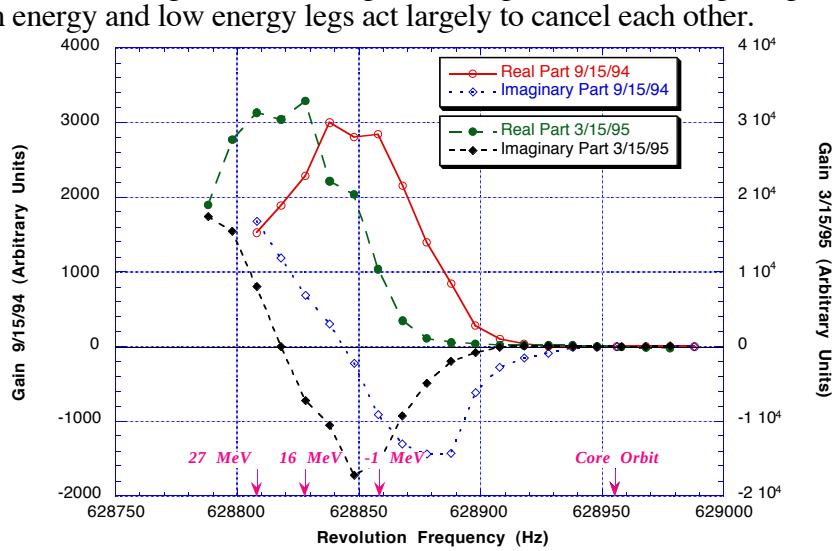

Figure 4. Real and imaginary parts of the stacktail cooling gain profile before and after the hardware changes of Run 1b. The quantity measured here is the response of the stacktail cooling system to a delta function of beam at each revolution frequency. It is the real part of the gain that moves the beam toward the core. The approximate factor of 10 difference in the gain scales between the 9/25/94 measurements and the 3/15/95 measurements is due, in part, to differences in the way the data was normalized.

In November of 1994 the high energy pickup was moved to 27 MeV above the central orbit. The effect of this move was to separate the two pickups by a distance greater than the width of the individual pickup response. This move permitted moving the peak in the system response farther from the core allowing more stacktail system gain for the same amount of core heating [5].

Even with increased separation of the high and low energy pickups there remained a significant response in the -1 MeV leg to beam on the stacking orbit. This undesirable response was greatly attenuated by re-commissioning the -23 MeV compensation leg of the stacktail momentum cooling system. The -23 MeV leg consists of a pickup, located 23 MeV below the momentum of the central orbit, connected via gain and delay into the -1 MeV leg. This leg was originally intended to correct the response of the -1 MeV leg to beam near the core orbit. However, the -23 MeV leg is currently phased to correct the response of the -1 MeV leg to beam near the stacking orbit [6].

The -23 MeV compensation leg does, in fact, also serve to correct the response of the -1 MeV leg to horizontal beam motion at the core momentum (in accord with its design function). The -1 MeV pickup is sensitive to the horizontal betatron motion of the beam at the core (the pickups are horizontally opposed, thus vertical motion is not detected). The relatively large number of particles at the core give rise to a strong signal at the horizontal betatron sideband frequencies of the core. The -1 MeV leg of the stacktail momentum cooling system responds to this signal by putting power on the stacktail kickers at the betatron sideband frequencies within the stacktail momentum cooling bandwidth. While the kick to the beam is primarily longitudinal, imperfections in the system (which will be enumerated in the next section) couple some of this power into the horizontal plane. The result is horizontal heating of the beam at the core. The -23 MeV pickup, due to its close proximity to the core, is also sensitive to the horizontal motion of the core. It turns out that when the -23 MeV compensation leg is phased to correct the longitudinal response at the stacking orbit it also provides a 10 to 15 dB reduction in the response of the -1 MeV leg to the horizontal motion of the core [7].

In addition to the high energy pickup move and -23 MeV leg re-commissioning, there was a substantial amount of activity directed to measuring and characterizing the performance of stacktail momentum cooling in the Accumulator. These efforts included a series of measurements which allow a complete characterization of the system for purposes of computer modeling. The net effect of all of these improvements was to accommodate the increased $\bar{p}$ flux through the system without further exacerbating the core heating. Figure 4 illustrates the evolution of the stacktail momentum gain profile over the course of Run 1b.

4. Core heating by the stacktail cooling

As was indicated previously, the stacktail momentum cooling heats the beam at the core of the momentum distribution. The amount of core heating increases as the $\bar{p}$ stack grows. Moreover, the effectiveness of the core stochastic cooling systems decreases as the stack size increases. The heating is the result of the core being driven transversely by the stacktail cooling system at the betatron sideband frequencies of the core. The transverse kick comes from mechanical misalignment of the stacktail cooling kicker electrodes, residual dispersion in the Accumulator lattice at the kicker locations, and small difference signals applied to the momentum kicker electrodes due to imperfect hybrids. Core heating from the stacktail momentum cooling system is one of the reasons the stacking rate declines as the stack size increases.

This transverse heating is partially compensated for by horizontal and vertical Δ kicker hardware installed on five of the sixteen stacktail momentum kickers. The Δ kickers allow the application of a variable transverse kick by the kickers on which they are installed by applying a difference signal to the kicker electrodes in addition to the normal sum signal which provides the longitudinal kick. The transverse kick from the Δ kickers is adjusted to approximately cancel the spurious transverse kicks from all of the other kickers.

The Δ kickers work well; however, this correction is not perfect. The Δ kicker correction is optimized for revolution frequencies near the core of the $\bar{p}$ momentum distribution

(revolution frequency of 628955 Hz). There is a region of significant heating at a revolution frequency of 628930 Hz where the Δ kicker correction is much less effective (see Figure 5).

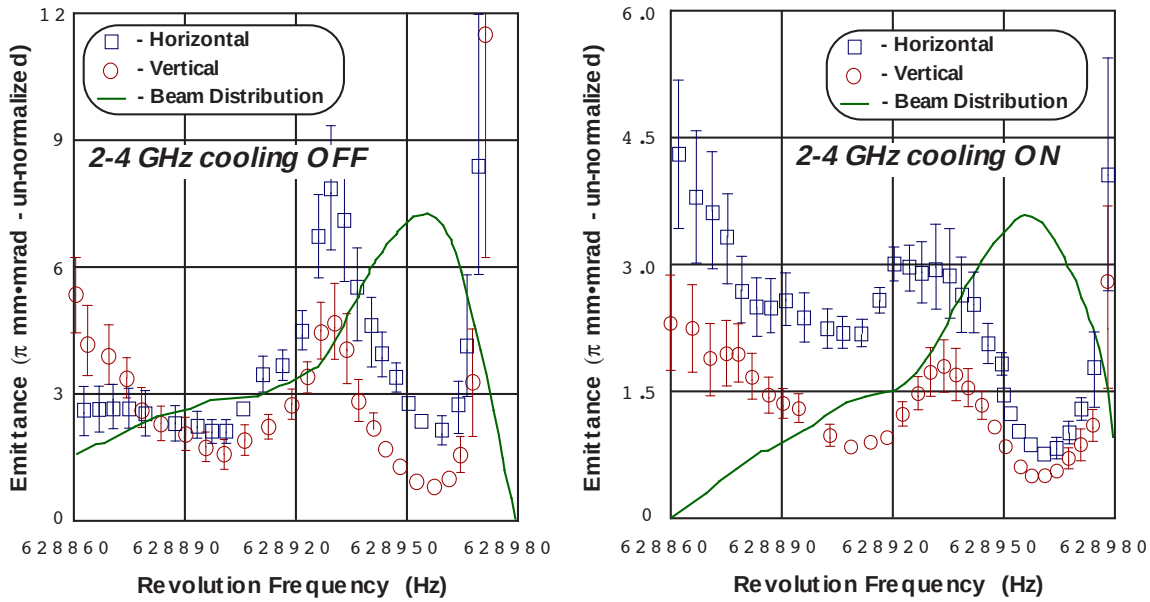

Figure 5. Effect of the 2-4 GHz transverse core cooling system on overcoming the emittance growth due to heating from the stacktail momentum cooling at a revolution frequency of 628930 Hz. The stack size in each case was 135 mA. Note that the vertical scale on the right hand plot is expanded by a factor of two relative to the left-hand plot.

The solution implemented in Run 1b was to install a second core transverse cooling system for each transverse plane. The transverse core cooling configuration at the beginning of the run was a single 4-8 GHz system for each plane. In March of 1995 the existing transverse cooling hardware was modified to include 2-4 GHz pickups and kickers. The lower frequency cooling extends the revolution frequency range over which the cooling electronics is phased to the beam. For a maximum phase error of ±60° (i.e. the error at which the cooling gain is attenuated by $\cos(\pi/3) = .5 = 6$ dB), the range of revolution frequencies (Δf_{rev}) at which effective cooling will occur is given by:

$$\Delta f_{rev} < \frac{f_{rev}^2}{3\alpha f_{\max}} \tag{3}$$

where f_{rev} is the core revolution frequency, f_{max} is the upper end of the cooling band width and α is the fraction of the total circumference of the accelerator which comprises the distance from the cooling pickup to the kicker. For the Accumulator core cooling systems α is $1/3$. For the 4-8 GHz system Δf_{rev} is 50 Hz centered at the core. With only the 4-8 GHz transverse cooling the cooling gain is attenuated by 3dB at 628930 Hz. The 2-4 GHz system extends this down to a revolution frequency of 628905 Hz, which completely contains the region of increased beam heating.

Subsequent to the commissioning of the Accumulator transverse 2-4 GHz cooling the value of I_m, in the stacking parameterization of equation (1), increased from a value of approximately 150 mA† to 192 mA.

5. LCW temperature

The antiproton source magnets, power supplies, RF system high power amplifier tubes, and the stochastic cooling traveling wave tubes are cooled by a low conductivity water

† This value of I_m is determined from the stacking performance of January 1995.

(LCW) cooling system. It has been observed that the antiproton stacking is degraded when LCW temperatures rise. Figure 6 shows a slight decline in average stacking rate when the LCW supply temperature exceeds 85°F with a more pronounced decline when the LCW supply temperature goes above 92°F.

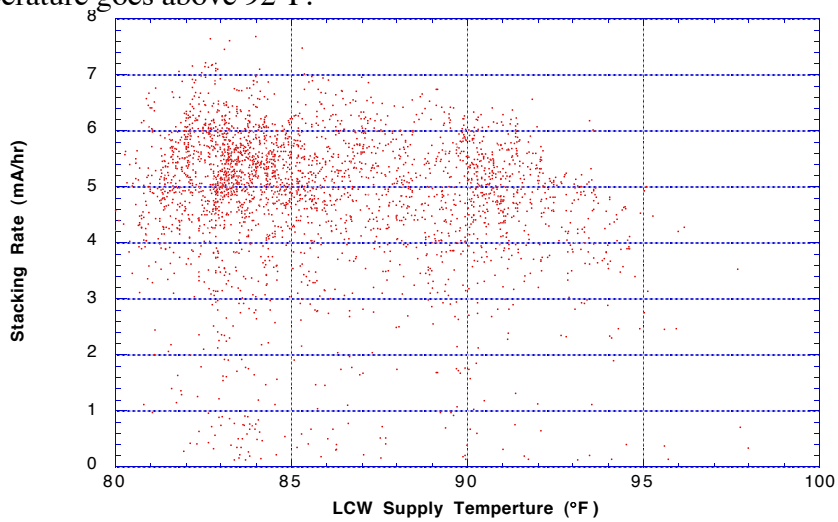

Figure 6. Antiproton stacking rate versus LCW supply temperature for the period from 1 March 1995 through 24 July 1995.

This effect is not yet completely understood. It is well known that the LCW temperature affects the integrated bend field and quadrupole gradient (magnets expand and contract and pole spacings change). However, since field and gradient changes are routinely compensated for with bend bus adjustments and tune corrections, it is likely that something else is adversely affecting the antiproton source when the LCW system gets hot.

II. Unstacking

Unstacking is the process by which antiprotons are extracted from the Accumulator core and transferred to the Main Ring for acceleration and injection into the Tevatron. The beam at the Accumulator core is bunched at twice its revolution frequency by a single bucket (or suppressed bucket) RF system. The bunched beam is accelerated to the extraction orbit of the Accumulator where it is kicked by means of a shuttered kicker into the AP3 beamline for transfer to the Main Ring.

A. Amount of beam unstacked

The goal of this process is the transfer of the maximum possible number of $\bar{p}$'s into the Tevatron. The following three factors limit the number of $\bar{p}$'s which can be transferred:

(1) The maximum RF bucket area which can be used is constrained by the momentum aperture of the Main Ring. The largest bucket area used this run was 1.55 eV•sec.

(2) The amount of beam contained in an RF bucket of a given size is determined by the longitudinal phase space density of the beam being bunched. The maximum achievable longitudinal phase space density is determined by the capacity of the core momentum cooling and the onset of longitudinal instabilities. Figure 7 shows the longitudinal phase space densities achieved as a function of stack size during Run 1b.

(3) The transfer efficiency of the $\bar{p}$'s from the Accumulator core into the Main Ring and beyond is maximized when the transverse emittances of the extracted beam

are maintained less than the admittance of the beam lines and the Main Ring. The minimum achievable $\bar{p}$ transverse emittance is a function of the stack size. This dependence is due to the existence of transverse beam heating mechanisms for which the heating rate becomes greater with an increasing number of beam particles. Therefore, the equilibrium emittances (i.e. when the cooling rate equals the combined heating rate from the various sources of beam heating) increase with increasing beam intensity. Figure 8 shows the dependence of the minimum transverse emittance and momentum spread on $\bar{p}$ intensity.

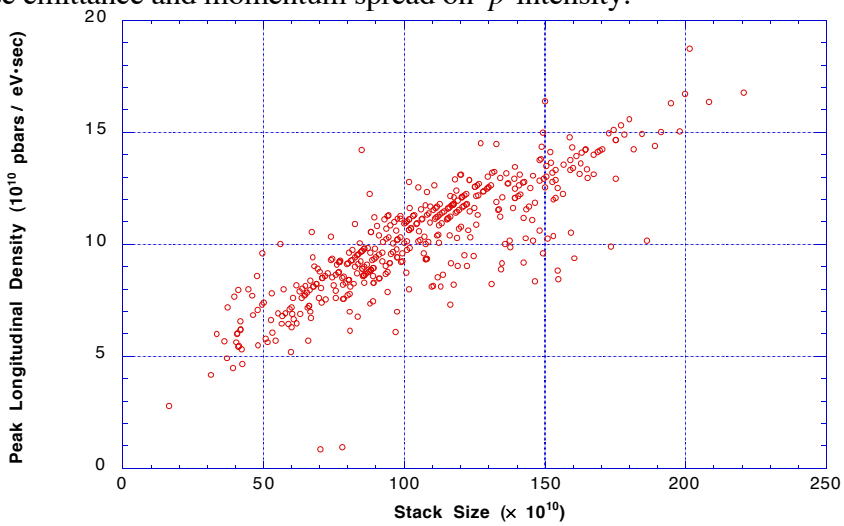

Figure 7. Longitudinal phase space density at the peak of the antiproton momentum distribution just prior to unstacking as a function of stack size.

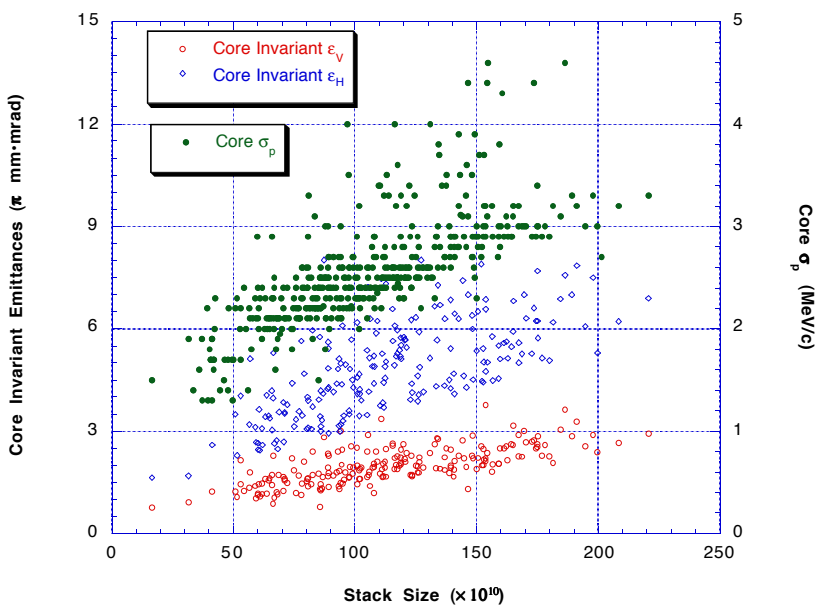

Figure 8. Accumulator core beam size prior to unstacking. In general, the beam is cooled in each dimension to the smallest size possible prior to initiating antiproton transfers.

During the course of Run 1b a variety of RF bucket areas were used for unstacking. The fraction of the antiproton stack extracted as a function of stack size for each size RF bucket is shown in Figure 9. Figure 9 shows that the fraction of the $\bar{p}$ stack which is removed during the course of unstacking decreases with increasing stack size. This is essentially a consequence of the fact that the momentum spread of the beam increases with increasing stack size (see Figure 8).

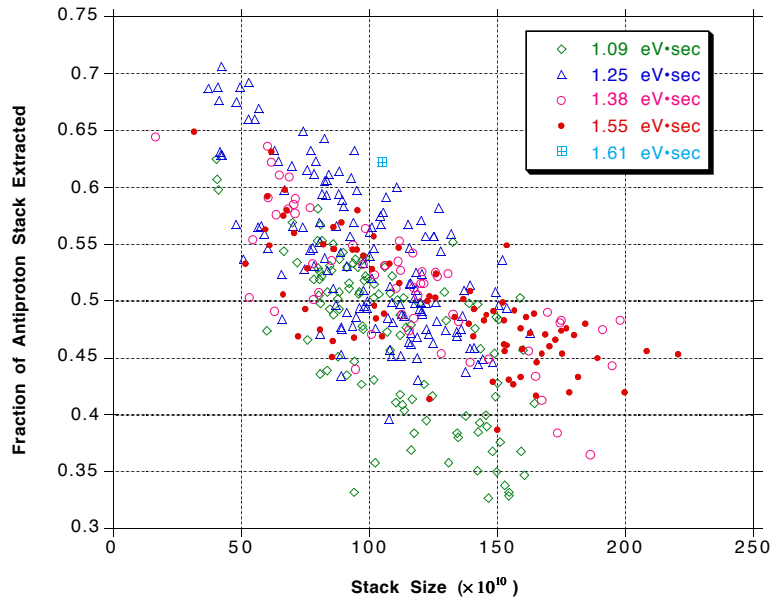

Figure 9. Fraction of Accumulator beam unstacked as a function of stack size for different RF buckets. Note the suppressed zero on the vertical axis.

An approximate understanding of the functional dependence of the fraction of the stack extracted ($\delta N/N$) on the RF bucket area and beam momentum spread can be gained by considering the simple case of a beam with a gaussian momentum distribution bunched at the peak of the distribution. The amount of beam extracted is obtained by integrating the beam distribution over the phase space area enclosed by the RF bucket. The shape of the bucket in longitudinal phase space is described by the function $\Delta p(\phi)$, where ϕ is the RF phase. For a beam distribution of width σ_p, the fraction of the stack extracted by bunching the beam with an RF bucket of area B, is given by:

$$\begin{aligned}
\frac{\delta N}{N} &= \frac{1}{\sqrt{2\pi}\sigma_p}\int_{-\pi}^{\pi}\int_{-\Delta p(\phi)}^{\Delta p(\phi)} e^{-\frac{p^2}{2\sigma_p^2}}dp\,\frac{d\phi}{2\pi h} \\
&= \frac{1}{\pi h}\int_0^{\pi}\mathrm{erf}\left(\frac{\Delta p(\phi)}{\sqrt{2}\sigma_p}\right)d\phi \\
&= \frac{\sqrt{2}}{\pi^{\frac{3}{2}}h\sigma_p}\int_0^{\pi}\Delta p(\phi)\left[1-\frac{1}{6}\left(\frac{\Delta p(\phi)}{\sigma_p}\right)^2+\cdots\right]d\phi \qquad (4) \\
&\cong \frac{1}{\sqrt{2\pi}}\frac{f_{rev}}{c\beta}\frac{B}{\sigma_p}
\end{aligned}$$

For the range of bucket areas (0.5 eV•sec to 1.6 eV•sec) and momentum spreads (1.2 MeV/c to 4.0 MeV/c) relevant to $\bar{p}$ unstacking, the fraction unstacked on a single transfer is very nearly linear in the quantity B/σ_p.

From Figure 8 it is seen that σ_p increases by approximately 40% as the stack size increases from 50×10^{10} to 150×10^{10}. The corresponding decrease in $\delta N/N$ exhibited in Figure 9 is about 30 - 40%, which is in approximate agreement with what is expected from the simple model of equation (4).

One would also expect from equation (4) that an increase in RF bucket area from 1.09 eV•sec to 1.55 eV•sec should yield about a 40% increase in the number of $\bar{p}$'s

unstacked. From Figure 9 it is clear that this was not the case during Run 1b. This discrepancy is not well understood. There are several possible explanations: (1) the unstacking RF voltage calibration is not known and/or changes with time, (2) beam is expelled from the RF bucket by some unknown mechanism on its way to the extraction orbit (about a 10% loss of beam from the bucket is normal) - the transverse dampers are particularly suspect here, (3) interference from the ion clearing RF system, or (4) errors in measuring the beam momentum distribution. Measurements have been made to investigate all four of these possible problems with no clear resolution of the issue.

B. Other unstacking issues

During the course of Run 1b there have been a variety of issues which have adversely affected the quantity and quality of the $\bar{p}$'s delivered to the collider from the antiproton source. These issues are briefly documented here.

1. Transverse emittance growth from the Accumulator core to Main Ring 8 GeV

The most serious unstacking issue is that of $\bar{p}$ emittance preservation. There is a severe growth in the $\bar{p}$ transverse emittance during the transfer of antiprotons from the Accumulator core to the Main Ring. The issue is somewhat obscured by uncertainties in the emittance measurements in the Accumulator and in the Main Ring. It is clear however, that given even the most pessimistic uncertainties in the various emittance measurements, a serious dilution of transverse phase space takes place during $\bar{p}$ transfers. Figure 10 illustrates the problem.

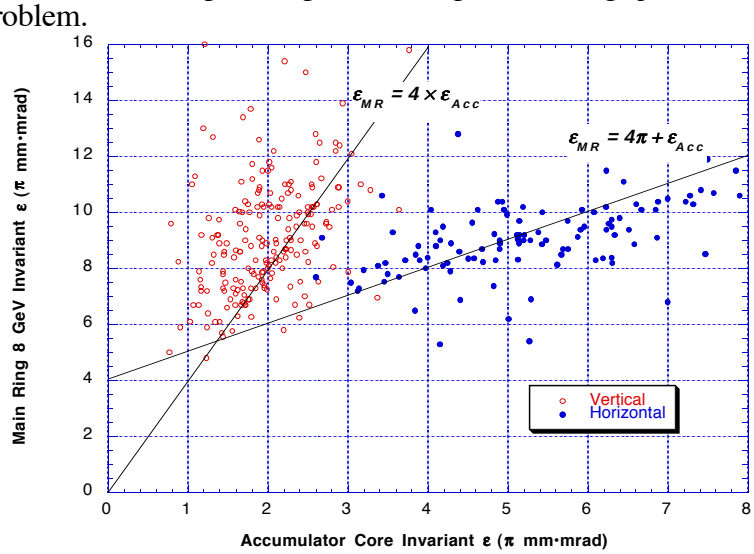

Figure 10. Main Ring flying wire emittances at 8 GeV versus Accumulator core emittances. Main Ring vertical emittance is plotted against Accumulator vertical emittance and Main Ring horizontal emittance is plotted against Accumulator horizontal emittance. The Accumulator emittances are measured using transverse Schottky pickups. Each point in these plots represents the average of all antiproton transfers during the course of a shot (usually 6 transfers). The solid lines indicate that the vertical emittance blows up by at least a factor of 4 while the horizontal emittance growth is 4π mm•mrad independent of initial emittance. Some of this 4π offset may be measurement error.

In Figure 10 the vertical plane exhibits an emittance growth which is proportional to the initial Accumulator core emittance. This growth is at least a 4-fold increase in vertical emittance. Also shown is a 4π horizontal blowup which is independent of the Accumulator core horizontal emittance. The source of the various kinds of emittance growth can be isolated to some extent by using a beamline SEM grid to measure the emittance of the beam just after it is extracted from the Accumulator. The SEM used for this purpose (SEM 900) is approximately 2 meters downstream of the Accumulator and upstream of any beamline

elements. Figure 11 compares the emittances derived from the SEM profile with the Accumulator core emittances.

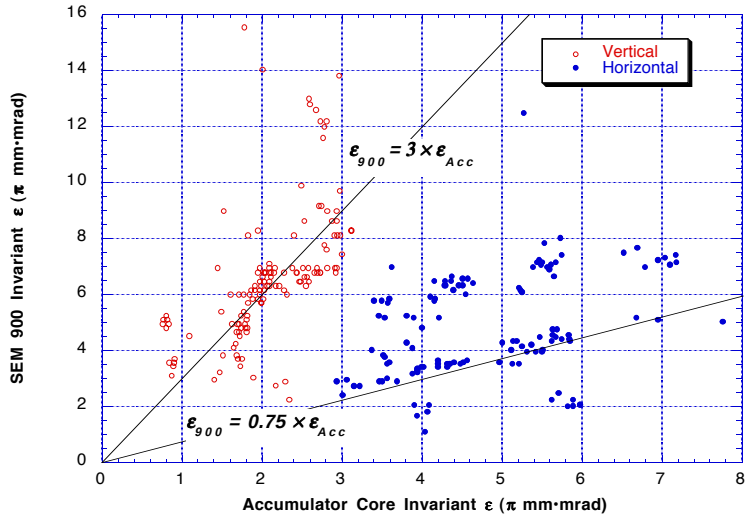

Figure 11. Measured emittances at SEM 900 at the upstream end of the AP3 beamline versus Accumulator core emittances. SEM 900 vertical emittance is plotted against Accumulator vertical emittance and SEM 900 horizontal emittance is plotted against Accumulator horizontal emittance. The data shown here was taken during the period from 17 April 1995 to 23 July 1995. Each point represents a single antiproton transfer.

A comparison of Figures 10 and 11 shows that, if the measurements are taken at face value, most of the vertical emittance growth occurs in the Accumulator. It is also apparent that the horizontal blowup most likely occurs on injection into the Main Ring.

The proportional vertical growth shown in Figure 11 is due, in part, to coupling on the Accumulator extraction orbit. Additionally, since Accumulator transverse emittances are correlated with stack size (see Figure 8), some component of this growth may be a $\bar{p}$ intensity dependent effect.

The unstacked beam spends about 8 sec on the extraction orbit prior to being extracted into the AP3 beam line. Since the extraction orbit lattice is coupled, there should be a time, well before the extraction event, when the horizontal and vertical emittances become equal. If the transverse phase space of the unstacked beam is not diluted in the accumulator, the horizontal and vertical emittances on the extraction orbit prior to extraction should each be equal to the average of the horizontal and vertical emittances at the core. Therefore, if there is no emittance dilution in the accumulator, and using the fact that, at the core, the vertical emittance is approximately $\frac{1}{2}$ the horizontal emittance (see Figure 8) the following relationship between extraction orbit and core orbit emittances is expected:

$$\varepsilon_H(extr) = \frac{3}{4}\varepsilon_H(core)$$

$$\varepsilon_V(extr) = \frac{3}{2}\varepsilon_V(core)$$

The observed vertical growth is about twice what is expected. The assumption that there is no emittance dilution from core to extraction orbit is not valid. The cause of this vertical growth is not known. The SEM 900 horizontal measurements indicate that about half the time the horizontal growth is consistent with the no dilution hypothesis. This behavior is difficult to understand in light of the coupling on the Accumulator extraction orbit. The observation that, frequently, little or none of the vertical growth is coupled into the horizontal plane may be an indication that the blow up occurs very close to the time of extraction (i.e.

within hundreds of beam turns of the extraction event). The relationship between vertical and horizontal emittances at the Accumulator core and extraction orbits and at 8 GeV in the Main Ring are shown in Figure 12.

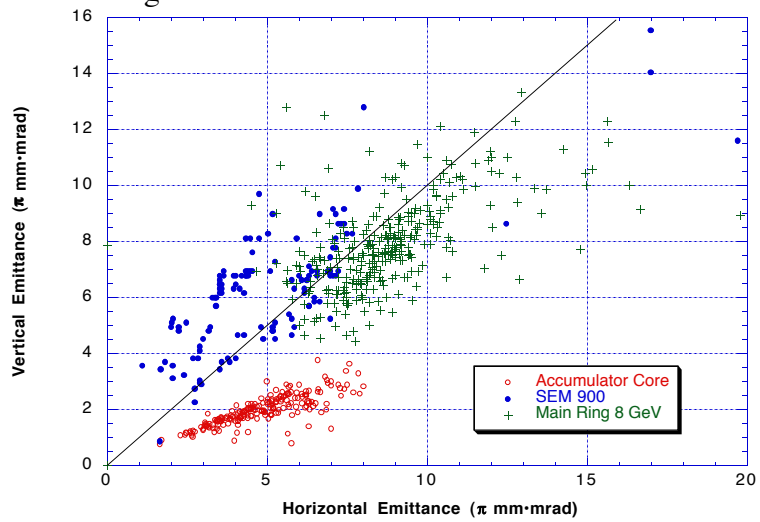

Figure 12. Vertical emittance versus horizontal emittance at the Accumulator core, SEM 900, and Main Ring at 8 GeV. The core and Main Ring data are averages over all transfers in a shot and represent all of Run 1b (i.e. the same data displayed in Figure 10). The SEM 900 data represent individual antiproton transfers during the period 17 April to 23 July 1995 (i.e. the same data displayed in Figure 11).

The horizontal emittance growth on injection into the Main Ring could be the result of an injection steering error or a mismatch in horizontal dispersion. Injection oscillations are normally corrected to an amplitude of less than 1 mm; thus contributing less than approximately $0.3\ \pi$ to the horizontal emittance offset. A large part of the balance of the horizontal growth is likely due to a horizontal dispersion mismatch between the AP1 beamline and the Main Ring. There were several attempts to correct this mismatch with some success; however the mismatch was not completely eliminated due to the difficulty in simultaneously correcting the dispersion and maintaining the β-function match.

Finally, it should be noted that the Main Ring 8 GeV lattice is coupled. It is therefore expected that, since beam circulates several hundreds of turns prior to the emittance measurement, the horizontal and vertical emittances will have the same value. The 8 GeV flying wire measurements show this to be the case to within $\pm 2\ \pi$ (see Figure 12). This suggests that part of the observed horizontal growth may be coupled in from the vertical plane.

A significant amount of effort was expended to understand and correct the $\bar{p}$ emittance growth during transfer to the Main Ring. A more complete account of this work is given in reference [8].

2. Trapped ion instabilities

Trapped ion induced instabilities in the Accumulator severely limited the performance of the antiproton source during the early part of Run 1a. The improvements to hardware and procedures developed during Run 1a were largely successful in controlling these instabilities throughout Run 1b. However, due to the much larger stack sizes accumulated during Run 1b relative to Run 1a (Run 1a largest stack was 150 mA, Run 1b largest stack was 220 mA), the implementation of RF ion clearing had to be adjusted somewhat to avoid rapid growth in the core emittances just prior to unstacking.

RF ion clearing is a technique whereby large amplitude ion motion is destabilized by longitudinally modulating the line charge density of the beam [9]. Use of RF ion clearing permits stable operation of the Accumulator with very large stacks (>130 mA). The RF system used for ion clearing (ARF2) is the same system currently used for unstacking $\bar{p}$'s during a collider fill. The difficulty with this arrangement occurs at those times during the preparation for unstacking $\bar{p}$'s when ARF2 must be switched off to program it for its unstacking sequence. If the stack is large, the beam will begin to oscillate transversely. This oscillation often results in rapid growth in the core emittances. If the emittance growth is large enough, time will have to be taken to re-cool the beam; a process usually requiring about 15 to 20 minutes. Bunching with ARF2 during unstacking will stabilize the beam; however, as soon as the unstacked beam is moved away from the core the transverse oscillation of the beam returns (see Figure 13).

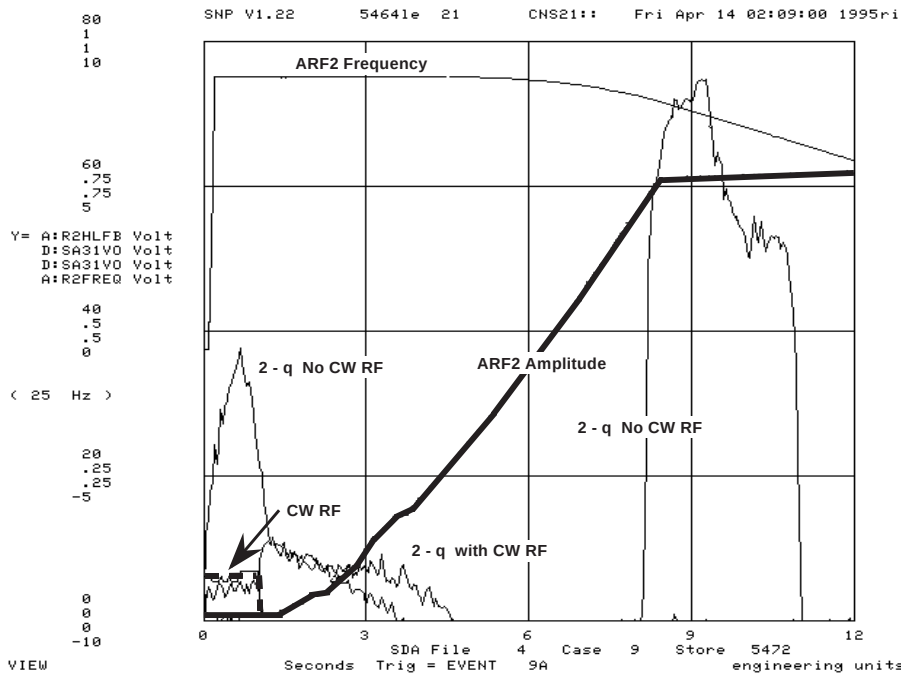

Figure 13. This is a snapshot of several Accumulator parameters during the course of unstacking one antiproton bunch during a collider fill. What is shown here are the first 12 sec of the unstacking sequence (the entire sequence takes 36 sec). Two sequences are superimposed on this plot; one with and one without the CW ion clearing RF. The top curve shows the RF frequency ramp (indicating the revolution frequency of the beam being unstacked). The heavy curve shows the measured voltage on the ARF2 cavity. The heavy dotted curve shows the ARF2 voltage readback with the CW ion clearing RF on. The CW RF is turned off at 1 sec and back on again at 7 sec. The remaining curve on the plot is the amplitude of the beam coherent transverse oscillation at one of its resonant frequencies (the 2 - q resonance). Without the CW RF, there is an initial episode of oscillation for the first second of the cycle - provoked by switching ARF2 from an ion clearing mode to an unstacking mode. There is a second burst of very large amplitude oscillation when the unstacked beam has moved approximately the width of the RF bucket away from the core. When the CW RF is applied the initial oscillation is greatly attenuated and the second period of oscillation is completely absent.

The implementation of RF ion clearing was modified to avoid the conflicts associated with using ARF2 for both unstacking and ion clearing. The new scheme has become known as CW RF ion clearing. CW RF ion clearing is simply the use of a signal generator in place of the ARF2 low level electronics during those times when ARF2 is also being used for unstacking. The CW RF signal generator is tuned to twice the revolution frequency of the beam at the core of the $\bar{p}$ distribution and drives the ARF2 cavities through the final power amplifiers in the ARF2 system. In the sequence shown in Figure 13, the CW RF is

switched off 1 sec after the initiation of the unstacking sequence, and is switched on again 7 sec after the start of the sequence until the end of the sequence. The timing of the CW RF was varied extensively during the remainder of Run 1b.

CW RF ion clearing successfully inhibited coherent oscillation of the beam during unstacking; however it is possible that CW RF interferes with unstacking and may be responsible for at least part of the discrepancy between the actual and expected unstacking yield identified above. Measurements performed to test this suspicion have been inconclusive.

III. Reliability

Any disruption in the operation of the antiproton source has a direct impact on the integrated luminosity of the collider. Lengthy downtime due to equipment failure frequently precludes antiproton stacking thereby reducing the number of antiprotons which are available for transfer to the collider. Failures which cause the loss of the antiproton stack are especially egregious due to the long time it takes to accumulate a sufficient number of $\bar{p}$'s for a new collider store (about 10 to 12 hours). The operations department maintains records of all system downtime. Also, a concerted effort is made to determine the cause of each stack loss. We present here a summary of the antiproton source reliability data collected during Run 1b.

A. Lost Stacks

During Run 1b a total of 70 antiproton stacks were lost with an average interval between lost stacks of 8.2 days. These stack losses caused an accumulated loss of 4.181×10^{13} $\bar{p}$'s. The causes for these stack losses fall into several categories as indicated in Figure 14.

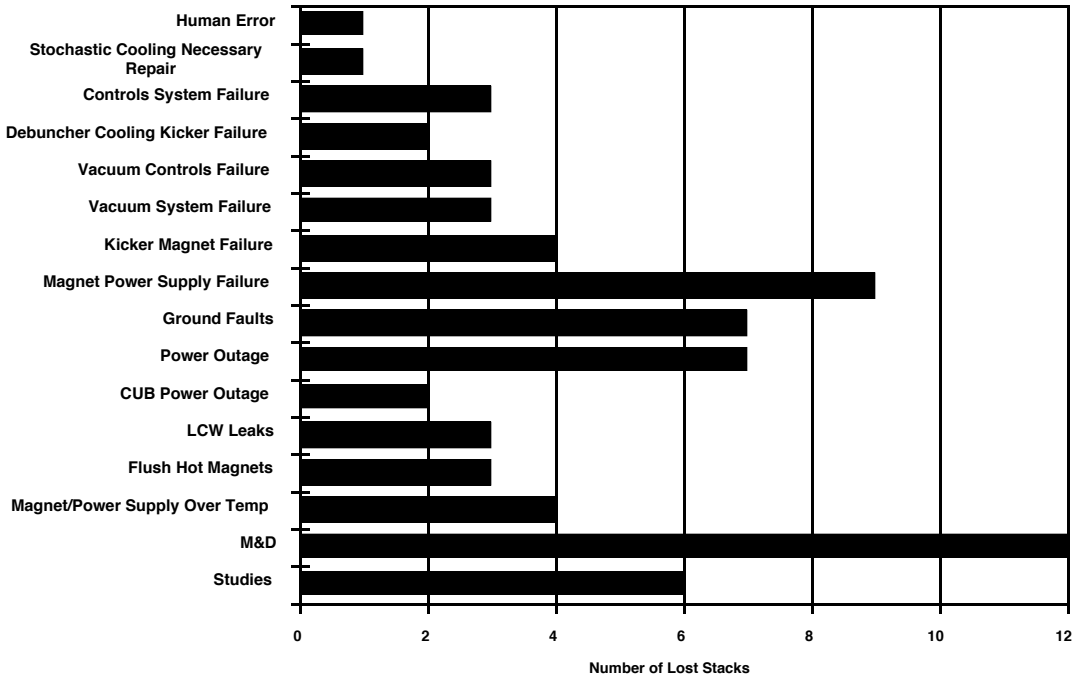

Figure 14. Causes of lost antiproton stacks.

In order to appreciate the value of what has been lost here, it is useful to note that the total antiproton accumulation for Run 1b was 31.680×10^{13} $\bar{p}$'s of which 26.318×10^{13} $\bar{p}$'s were unstacked for use in the collider. Thus, lost stacks consumed 13.2% of the total $\bar{p}$ yield for

the run while 83.1% was used for the collider. The remaining 3.7% of the Run 1b $\bar{p}$ production consists of failures which caused only a partial loss of the stack as well as special tests (e.g. test shots) which consumed some but not all of the stack.

B. Downtime

There were a total of 917 hours of antiproton source downtime during Run 1b, which constitutes approximately 7.2% of the total available running time. Down time is recorded for the antiproton source whenever the normal scheduled operations are interrupted.

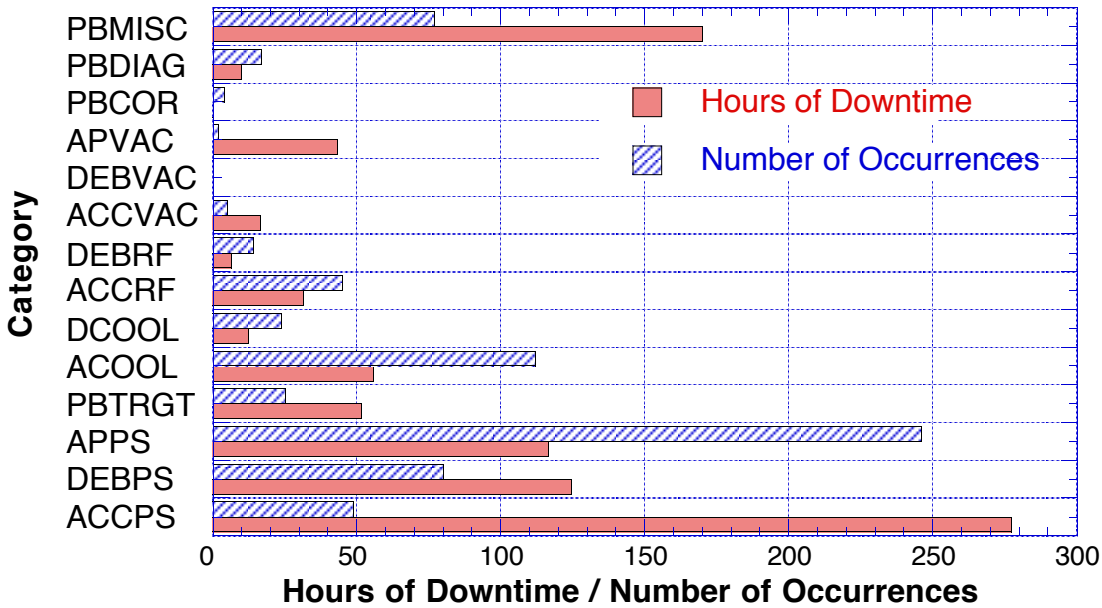

Figure 15. Summary of the antiproton source downtime for Run 1b.

Table III
Antiproton Source Downtime Categories

PBMISC	Everything that doesn't fit into another category
PBDIAG	Diagnostic systems
PBCOR	Antiproton source correction elements
APVAC	AP1, AP2, or AP3 beamline vacuum system
DEBVAC	Debuncher vacuum system
ACCVAC	Accumulator vacuum system
DEBRF	Debuncher RF systems
ACCRF	Accumulator RF systems
DCOOL	Debuncher stochastic cooling systems
ACOOL	Accumulator stochastic cooling systems
PBTRGT	Antiproton source target station systems
APPS	Beamline power supplies and magnets
DEBPS	Debuncher power supplies and magnets
ACCPS	Accumulator power supplies and magnets

Scheduled operations for the antiproton source include stacking, shot setup, accelerator studies, and shutdowns for maintenance and development. On each occurrence of downtime, the duration of the downtime is logged and the general category into which the failure which caused the downtime is recorded. A summary of the downtime for Run 1b is given in Figure 15. The definitions of the various categories are given in Table III.

One common source of downtime during Run 1b was magnet over-temperature trips. When a magnet overheats an interlock trips the power supply to the bus supplying that magnet. The interlock resets when the magnet has sufficiently cooled to safely restore power to the bus. This cool down typically takes 15 - 30 minutes. This type of failure spans several of the categories shown in Figure 15. The total time lost due to magnets overheating during Run 1b was 97 hours, or about 10.5% of the total downtime for the antiproton source. The underlying cause of the large number of overheating magnets in the antiproton source is the build up of copper oxide in the cooling channels of the magnets. This build up restricts the flow of cooling water and lowers the heat transfer coefficient. The most effective treatment of this problem to date has been to flush the magnet cooling channels with a weak acid solution. In addition an effort has been made to minimize the amount of oxygen in solution in the cooling water system.

IV. References

1. S. Shukla, Proceedings of the 1994 Meeting of the APS Division of Particles and Fields, University of New Mexico, 1485 (1994).

2. F.M. Bieniosek, K. Anderson, and K. Fullett, Requirements for a Beam Sweeping System for the Fermilab Antiproton Source Target, Submitted to the proceedings of the 1995 Particle Accelerator Conference, Dallas, TX, May (1995).

3. D. Olivieri, M. Church, and J. Morgan, A Dynamic Momentum Compaction Factor Lattice in the Fermilab Debuncher Ring, Submitted to the proceedings of the 1995 Particle Accelerator Conference, Dallas, TX, May (1995).

4. Fermilab, Design Report Tevatron I Project, section 5.9, September (1984).

5. R. Pasquinelli, private communication.

6. D. Vander Meulen, private communication.

7. K. Bertche, private communication.

8. F. M. Bieniosek, Emittance growth in recent antiproton transfers from the Accumulator to the Main Ring, Antiproton source departmental memo #558, September (1995).

9. S. Werkema, Control of Trapped Ion Instabilities in the Fermilab Antiproton Accumulator, Submitted to the Proceedings of the 1995 Particle Accelerator Conference, Dallas, TX, May (1995).

5. Main Ring Performance

The Main Ring is currently being used in three different modes 1) proton acceleration from 8 GeV to 120 GeV for antiproton production, 2) proton acceleration from 8 GeV to 150 GeV with coalescing and cogging for Tevatron injection and 3) antiproton acceleration from 8 GeV to 150 GeV with coalescing and cogging for Tevatron injection. In the first two modes the Main Ring accepts beam from the Booster and in the third mode from the antiproton source. In all of these operating modes the performance during Run 1b exceeds that of previous runs. This section describes the performance of the Main Ring during Run 1b. Upgrades and tuning steps that helped improve the transfer efficiencies through Main Ring are discussed.

In general, for all modes of operation, the beam intensity is limited by the admittance. In the transverse dimension the beam fills the aperture. A second general area of difficulty is the poor reproducibility of performance from one shot setup to another.

In the antiproton production mode the Booster is run with an intensity of up to 5×10^{10} protons/bunch, in 84 bunches. The Main Ring delivers up to 4×10^{10} protons/bunch to the antiproton production target. The efficiency of the antiproton production cycle has been consistently maintained between 70-80%. Earlier in the Run 1b, an accelerator physics simulation and studies [1] suggested the use of harmonic correctors (trim quadrupoles, sextupoles and octupoles) placed in the lattice to improve the aperture at injection. The use of harmonic correctors, changes in tune and chromaticity, small changes in RF bucket size and use of the Booster damping system to reduce the longitudinal emittance have all resulted in higher transfer efficiency. The maximum total intensity of 3.45×10^{12} protons on target (9 turns) was ultimately achieved during antiproton production.

Fig. 1 shows the efficiency of proton transfer through the Main Ring for different operations. About 10% of particles are lost in the first few thousand turns at 8 GeV. This loss is primarily a consequence of the small injection aperture. The small aperture is caused by a combination of a) a lattice function mismatch between the injection beam line and the Main Ring, b) poor magnetic field quality at low excitations, and c) perturbations to the ring that have been required for the integration of overpasses and new injection and extraction systems. The second considerable loss of particles in the Main Ring, about 10%, is at transition energy. At transition, the growing horizontal beam size at high dispersion points gets scraped further because of the limited aperture.

Fig. 1 also shows the transfer efficiencies for the other two cycles of the Main Ring operation, which are used to fill the Tevatron for collider operations. These two cycles also have particle loss from injection through transition energy. The total loss is about 20%. A few percent (<3%) of particles are lost between 20 GeV and 150 GeV. The proton loss during cogging and coalescing is about 25% due to the limited momentum aperture at flat-top. *The average transfer efficiency through the Main Ring during antiproton production cycle is at best 80%., whereas the average transfer efficiencies of coalesced proton and antiproton bunches are 50% and 70% respectively.* Most of the difference between the proton and antiproton transfer efficiency is due to the difference in their respective coalescing efficiencies. The antiproton transfer efficiency is larger due to the smaller emittance of the antiproton beam from the accumulator.

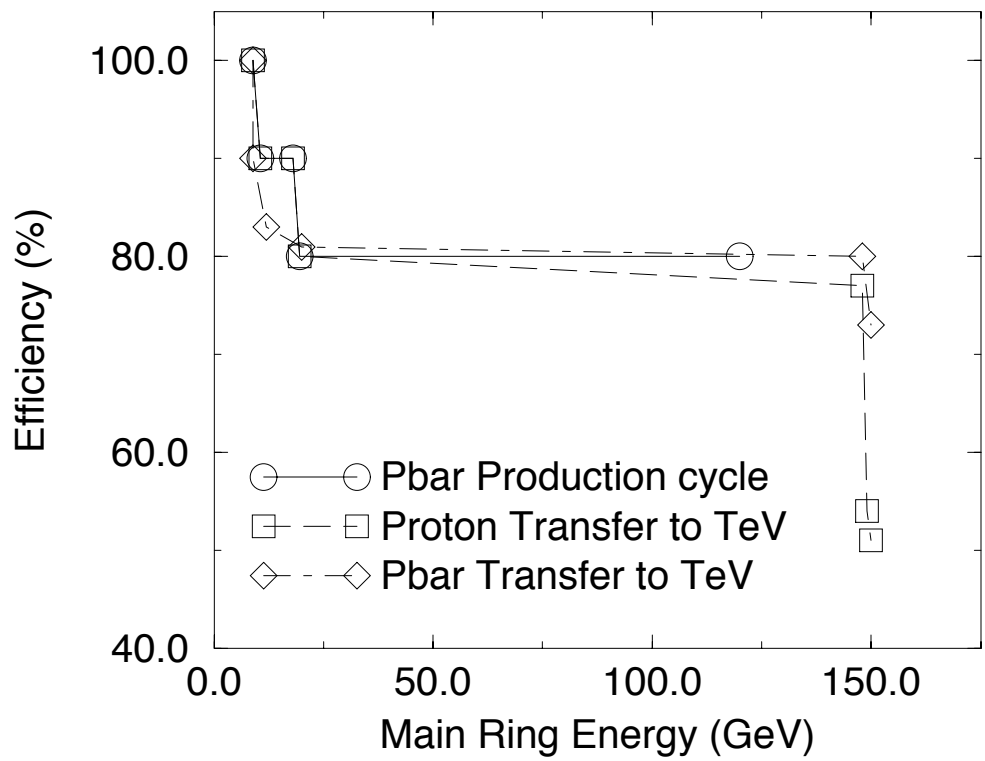

Fig. 1 Main Ring Transfer efficiencies for different particle transfer cycles.

Improvements in Bunch Coalescing

The Main Ring improvement having the greatest impact on machine performance during Run 1b was an upgrade to the Bunch Coalescing System[2]. During bunch coalescing several (~11) proton and antiproton bunches are first rotated at the fundamental rf harmonic of h=1113 (53 MHz) to reduce the momentum spread, then are rotated for a quarter of a period at a lower harmonic (h=53 or 2.5 MHz). The bunches are then recaptured in a single 53 MHz bucket. The coalescing upgrade tripled the voltage available for 2.5 MHz to 60 kV by installing 3 new cavities. The result was improved transfer efficiency in the Main Ring for both the protons and antiprotons, permitting the achievement of a luminosity of more than 2.5×10^{31} $cm^{-2}sec^{-1}$. Fig 2 shows the proton coalescing efficiency before and after the upgrade. An average increase of 10% was achieved in proton coalescing resulting in similar increase in proton transfer efficiency.

Prior to this upgrade, it was not possible to recapture all of the protons in a single rf bucket due to the large longitudinal emittance. Consequently, some of the protons were not captured at all or were bunched into neighboring buckets, causing satellite bunches and DC beam. Fig. 3 shows the percentage of protons in satellites before and after the upgrade. It is clear from the figure that after the upgrade the satellites were smaller. The typical satellite bunches after the coalescing upgrade have intensities less than 2% of the central bunch

intensity compared with 8-12% before. This resulted in considerable reduction in background at the two detectors due to parasitic collisions.

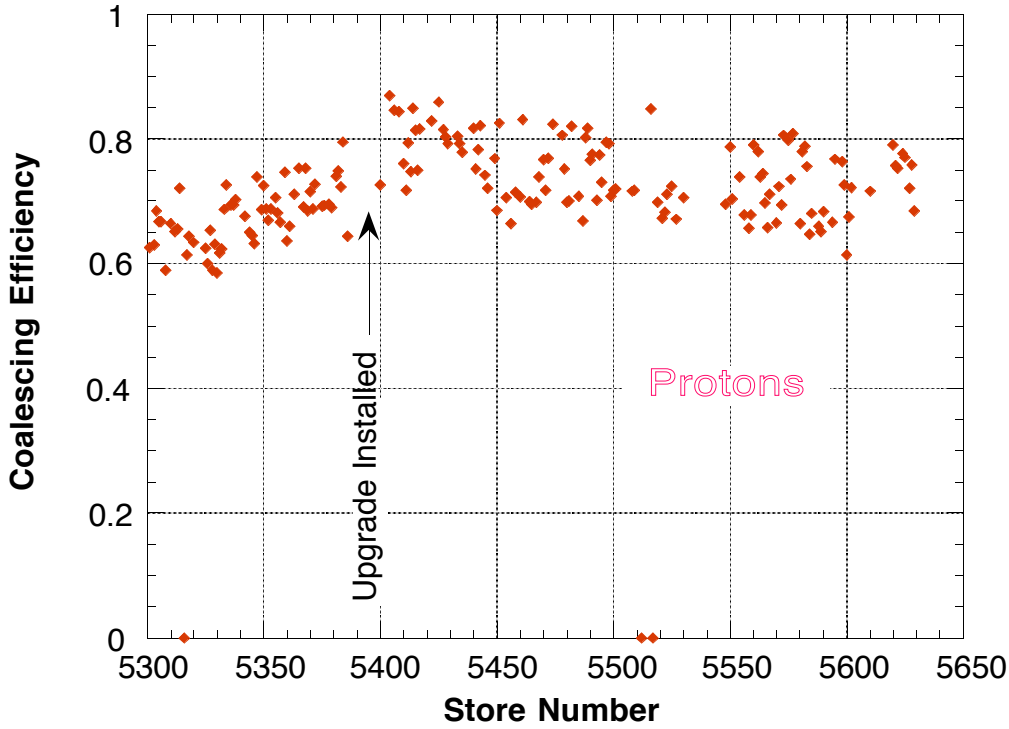

Fig. 2 The Proton Coalescing efficiency before and after the coalescing upgrade

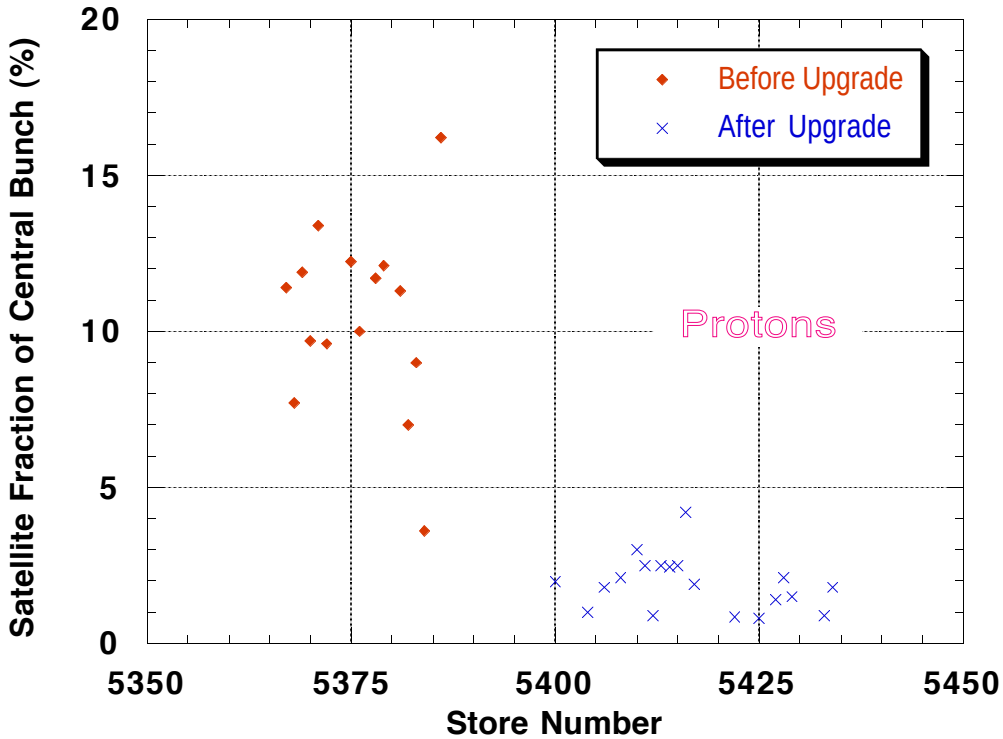

Fig. 3 Proton Satellite percentage before and after the coalescing upgrade.

Similar improvements in the antiproton coalescing efficiency (Fig. 4) and reduction in satellites (Fig. 5) were observed.

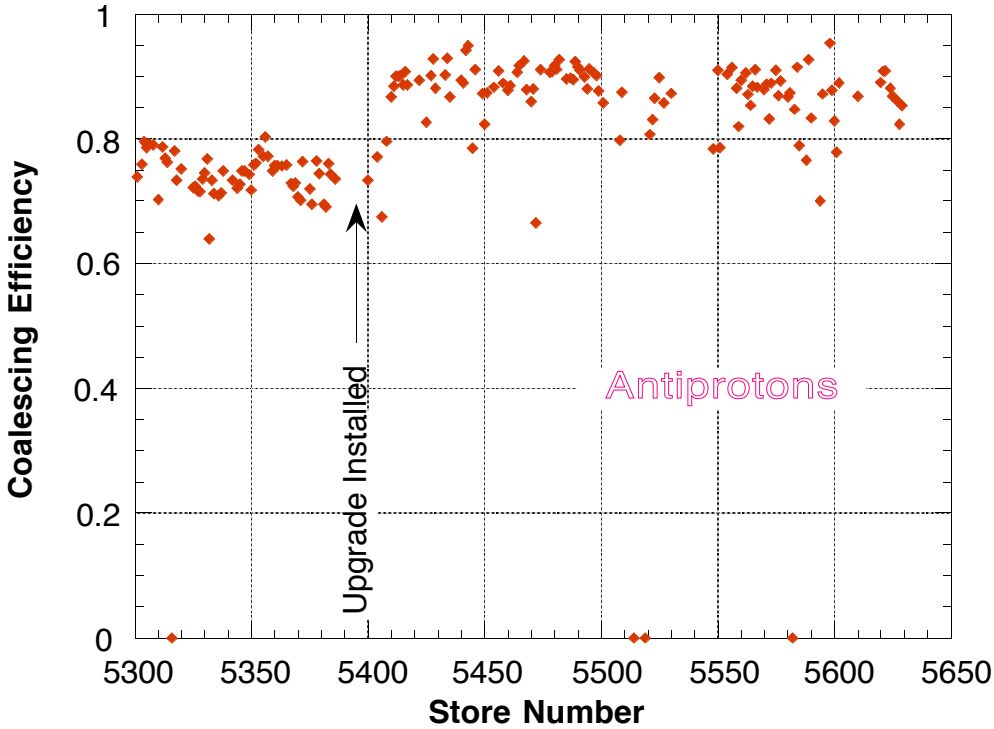

Fig. 4 The antiproton coalescing efficiency before and after the coalescing upgrade

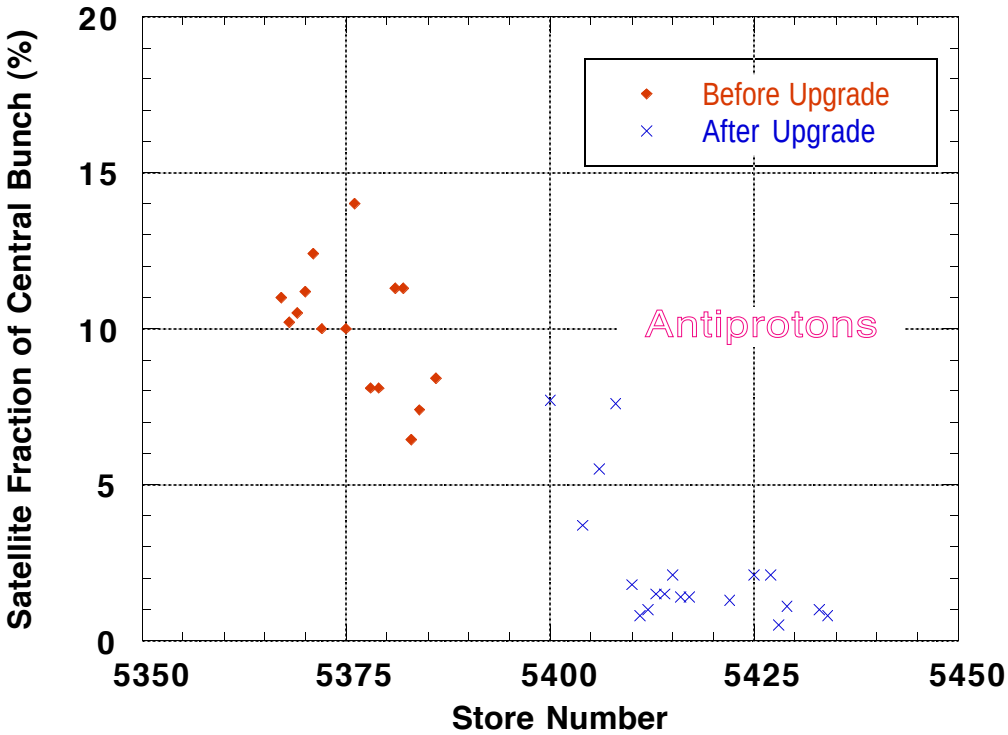

Fig. 5 Antiproton Satellite percentage before and after the Coalescing Upgrade

Other issues

There has been considerable discussion of the issue of antiproton emittance growth during transfers. This issue has not yet been resolved. The protons injected into the Main Ring get scraped by about 10% due to limited transverse aperture. Hence the measured proton emittance is essentially the aperture of the machine. The antiproton emittance appears to blow up by a large factor after extraction from the accumulator core and injection into the Main Ring. This blowup of emittance is not fully understood. This is discussed in more detail in the antiproton unstacking section.

A significant effort was made to understand and improve the match between the transfer beam lines and the Main Ring lattice during Run 1b.

One frequently occurring issue has to do with the operating point of the Main Ring. The horizontal and vertical tunes are coupled and are not constant functions of beam momentum. The preferred operating point at injection is close to a betatron resonance. The tunes are made to cross resonances during the ramp to flat-top. The resultant particle loss is reduced by amplitude dependent detuning using octupoles. The Main Ring also has a beta wave that limits its aperture. While these problems have been extensively studied throughout the operating history of the Main Ring, these issues remain far from being resolved. A considerable amount of dedicated study time as well as extensive detailed simulations will be required if these issues are to be settled.

The operations group is required to tune the Main Ring on a continual basis to keep the transfer efficiency high. This is required due to changing injection conditions. The aperture is not large enough to accommodate any significant change in the injection conditions. This poses limitations on how many particles can be injected and transmitted through the Main Ring. This serious limitation will be eliminated by the construction of the Main Injector.

References

[1] C. S. Mishra, M. J. Yang, S. Assadi, Studies of Main Ring Aperture, Main Ring Internal Note, To be published.

[2] J. Dey et al., Improvements in Bunch Coalescing in the Fermilab Main Ring, Submitted to Particle Accelerator Conference 95, Dallas.

6. Tevatron Performance

Introduction

This report summarizes the typical running conditions of the Tevatron for the later part of Run 1b, when the Tevatron was running well. During this time the Tevatron had some up and down periods; we will attempt to characterize its best reproducible performance during this time.

Since most performance issues in the Tevatron involve transverse emittance growth and luminosity lifetime we will focus on those issues. The flying wire system was improved towards the end of the run, so most transverse emittance data will be taken from that time period. The discussion is divided into several main categories that seem appropriate, listed below.

Proton & Antiproton Injections

The most significant problem encountered in beam transfers from the Main Ring to the Tevatron is transverse emittance growth. The emittance growth on injection is expected to have three dominant contributions, as described in detail in EXP-175 and the Run II Handbook.

1) The vacuum windows in the transfer line are expected to add about 0.5π.
2) Lattice function mis-matches with the main ring and transfer line should result in an emittance growth that scales with main ring emittance. The known dispersion mis-match is expected to give about 0.5π growth.
3) Injection closure errors cause an emittance growth that increases as the square of the injection error. Data concerning this effect are shown below. A typical injection error of 0.25 mm gives a 0.5π growth.

Figure 1 shows a horizontal BPM readback for the first 1024 turns after a proton injection of the most recent shot 5629. A small synchrotron oscillation is visible. This small longitudinal mis-match causes negligible longitudinal emittance growth and so is not considered to be an operational problem. Earlier in the run, however, the longitudinal mis-match was larger as is shown in Figure 2 for shot 5272, and caused a 5-10% loss of $\bar{p}$s on transfer.

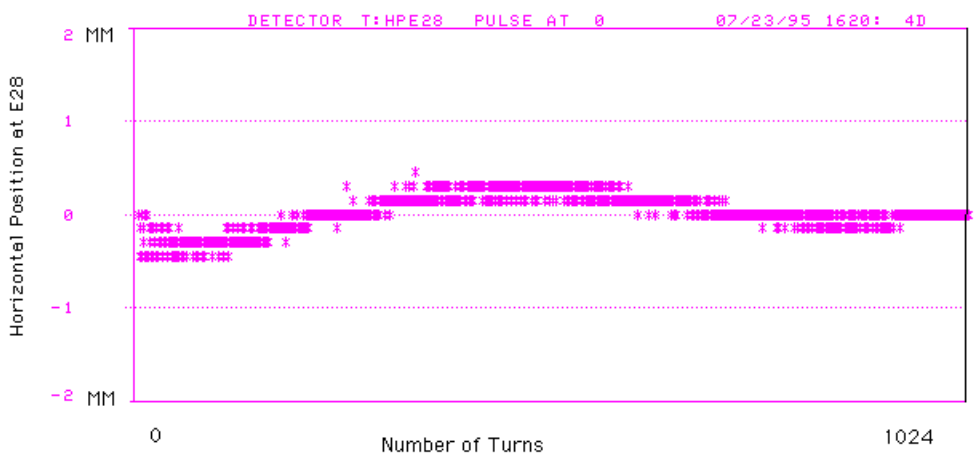

Figure 1 Horizontal TBT for P1 shot 5629

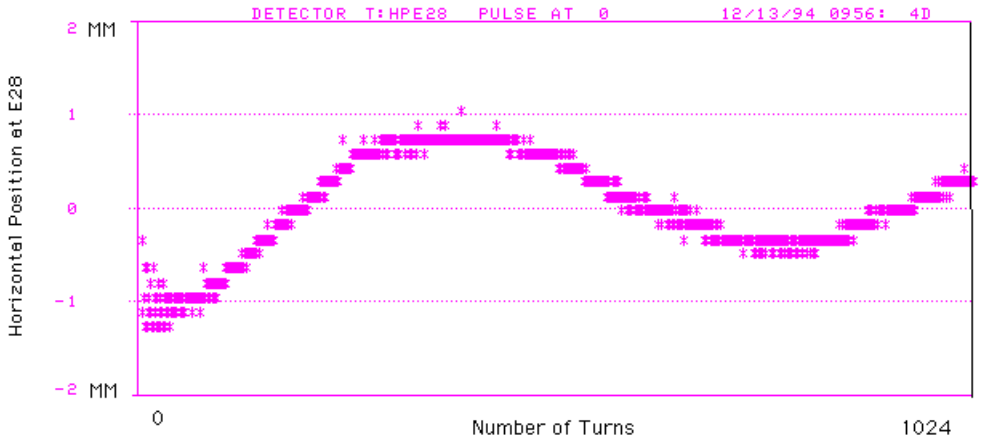

Figure 2 Horizontal TBT for P1 shot 5272

Figure 3 shows the "Bull's Eye plot" (i.e. the magnitude & phase of the injection oscillations for the first 10-15 turns in the Tevatron), as calculated by the new T121 beam line tuner by Jerry Annala and Wim Blokland. The data shown represents all proton and antiproton injections for the 10 most recent shots. Proton injections for these most recent shots were all within a 0.5 mm tolerance. The operational criterion is 0.3 mm. The expected emittance growth for 0.5 mm oscillations is about 1π. $\bar{p}$ transfers, especially A1 & A2, sometimes had larger closure errors as is apparent in Figure 3. During shot setup if A1 & A2 had large errors the operators would adjust the injection steering to close better.

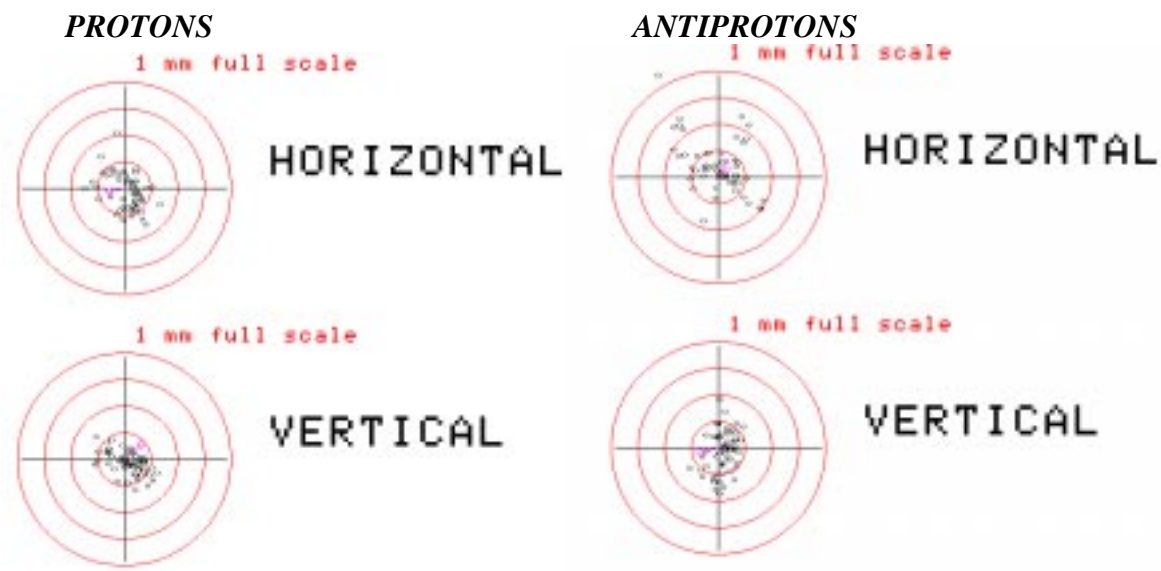

Figure 3 Injection Steering Error Magnitude & Phase for the Most Recent 10 shots (5603-5629)

Emittance Growth in the Tevatron

The $\bar{p}$ vertical emittance growth during recent transfers is shown in Figures 4 & 5 as a function of two quantities, run number and main ring emittance at extraction. Only the vertical emittance growth is considered because this measurement does not require the de-convolution of the longitudinal emittance. Since the tunes are coupled at injection during normal running the "emittance" should be the same in both planes. Both Figures 4 & 5 show a growth of roughly 2π independent of the initial emittance.

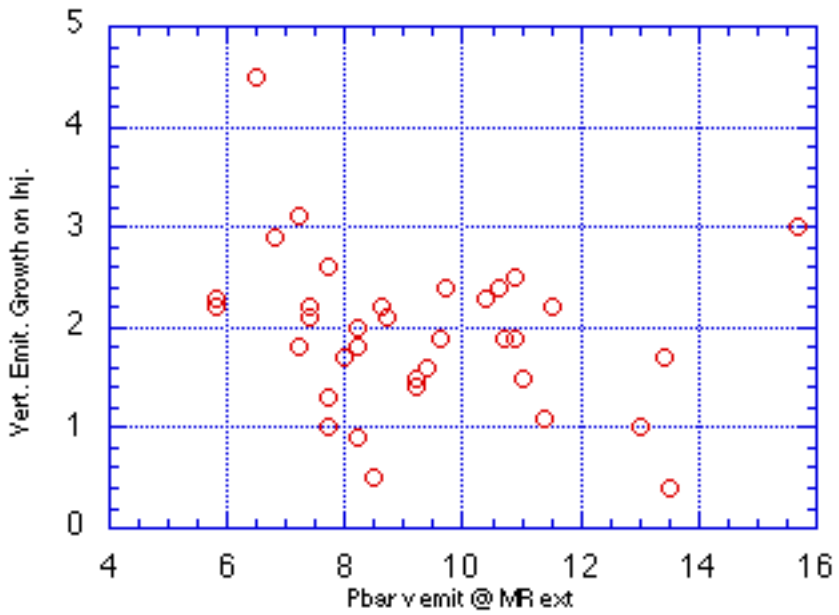

Figure 4 Antiproton Vertical. Emittance Growth on MR to TEV Transfer

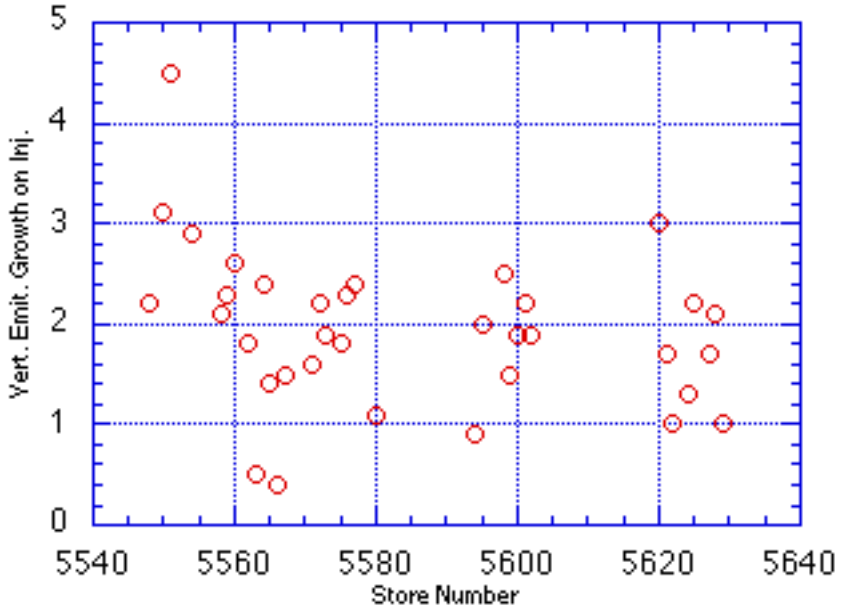

Figure 5 Antiproton Vertical. Emittance Growth on MR to TEV transfer

Table I shows a summary of the transverse emittance for protons and antiprotons at the various steps in the injection and acceleration process in a fashion similar to the online display. Unlike the online display, the values in this table are the average of the last 10 stores that had good flying wire measurements at all steps, i.e. shots 5599, 5600, 5601, 5602, 5620, 5622, 5624, 5625, 5628, and 5629. This table represents a scientific approximation, since the elapsed time between steps in each shot varies.

TABLE I
Vertical Emittance Measurements averaged for 10 recent shots

Protons

	ENERGY	P1	P2	P3	P4	P5	P6	WT.AVG	
MR inj.	8.938	14.77	14.22	13.01	14.07	13.58	14.76	14.04	pi
MR flattop	150	11.97	11.89	12.01	12.05	12.15	12.01	12.02	pi
MR extraction	150	14.20	14.14	14.29	14.40	14.70	14.36	14.35	pi
TEV P1 inj.	150	16.35	nobeam	nobeam	nobeam	nobeam	nobeam	16.35	pi
TEV P2 inj.	150	16.73	16.36	nobeam	nobeam	nobeam	nobeam	16.55	pi
TEV P3 inj.	150	17.06	16.58	16.63	nobeam	nobeam	nobeam	16.76	pi
TEV P4 inj.	150	17.77	16.94	16.85	16.68	nobeam	nobeam	17.06	pi
TEV P5 inj.	150	19.01	17.66	17.12	16.89	17.23	nobeam	17.58	pi
TEV P6 inj.	150	20.47	19.17	17.97	17.35	17.37	16.70	18.16	pi
TEV open helix	150	23.84	22.45	21.18	20.53	20.55	19.85	21.40	pi
TEV A1 inj.	150	23.84	22.31	21.10	20.46	20.97	19.70	21.39	pi
TEV A2 inj.	150	23.72	22.25	21.05	20.43	21.42	20.19	21.51	pi
TEV A3 inj.	150	24.47	22.40	21.16	20.52	21.59	20.52	21.77	pi
TEV A4 inj.	150	24.69	23.00	21.16	20.58	21.50	20.36	21.88	pi
TEV A5 inj.	150	24.80	23.27	21.85	20.75	21.59	20.49	22.12	pi
TEV A6 inj.	150	24.72	23.29	22.18	21.28	21.59	20.60	22.28	pi
TEV before ramp	150	25.25	23.64	22.74	21.67	21.85	20.83	22.66	pi
TEV on ramp	500	27.38	25.83	25.00	23.96	24.17	23.14	24.90	pi
TEV flattop	900	30.54	28.99	28.03	26.89	27.07	26.06	27.92	pi
TEV BD Step 15	900	25.21	23.92	23.22	22.32	22.43	21.59	23.11	pi
TEV collisions	900	26.52	25.15	24.15	23.29	23.44	22.83	24.22	pi
TEV in store	900	25.91	24.85	23.73	22.89	23.06	22.40	23.80	pi

Antiprotons

	ENERGY	A1	A2	A3	A4	A5	A6	WT.AVG	
MR inj.	8.938	10.62	10.85	10.94	10.40	9.98	10.07	10.50	pi
MR flattop	150	9.70	9.75	9.78	9.36	9.10	8.82	9.44	pi
MR extraction	150	10.70	10.69	10.76	10.23	10.02	9.96	10.42	pi
TEV A1 inj.	150	13.49	nobeam	nobeam	nobeam	nobeam	nobeam	13.49	pi
TEV A2 inj.	150	12.97	12.75	nobeam	nobeam	nobeam	nobeam	12.89	pi
TEV A3 inj.	150	12.92	12.67	13.15	nobeam	nobeam	nobeam	12.93	pi
TEV A4 inj.	150	12.43	12.42	12.57	12.32	nobeam	nobeam	12.44	pi
TEV A5 inj.	150	12.53	12.21	12.09	12.37	12.37	nobeam	12.31	pi
TEV A6 inj.	150	12.29	12.11	12.45	12.08	12.20	11.98	12.19	pi
TEV before ramp	150	12.31	11.98	11.97	11.71	11.56	11.44	11.87	pi
TEV on ramp	500	13.48	12.95	12.74	12.38	12.06	12.10	12.67	pi
TEV flattop	900	14.37	13.88	13.57	13.05	12.82	12.94	13.50	pi
TEV BD Step 15	900	13.55	13.09	12.76	12.58	12.10	12.06	12.75	pi
TEV collisions	900	13.39	13.14	12.59	12.58	11.96	12.16	12.70	pi
TEV in store	900	13.25	12.98	12.14	12.37	11.70	11.92	12.45	pi

The average emittance growth from Main Ring extraction to Tevatron injection for $\bar{p}$s is 2.26π, and for protons 2.3π. Both are consistent with our estimates of the systematic errors of the flying wire system, and are roughly what is expected for the sum of the contributions from the vacuum windows, the known dispersion mis-match, and typical injection oscillations.

Figure 6 shows the average emittance of each bunch as a function of "Time", where "Time" is measured as the number of proton injections. There is a variable length of time between injections, so this way of presenting the data does not give a calibrated measurement of the emittance growth rate, but does allow us to correlate emittance growth to events that are synchronized to the injections like the firing of kickers. It is clear that the vertical emittance of the proton bunches grows as a function of proton injections, and is fit well by a 2nd order polynomial. Growth during this stage that is caused by injection kicker mis-timing would show up as a discrete jump, there is no evidence for this.

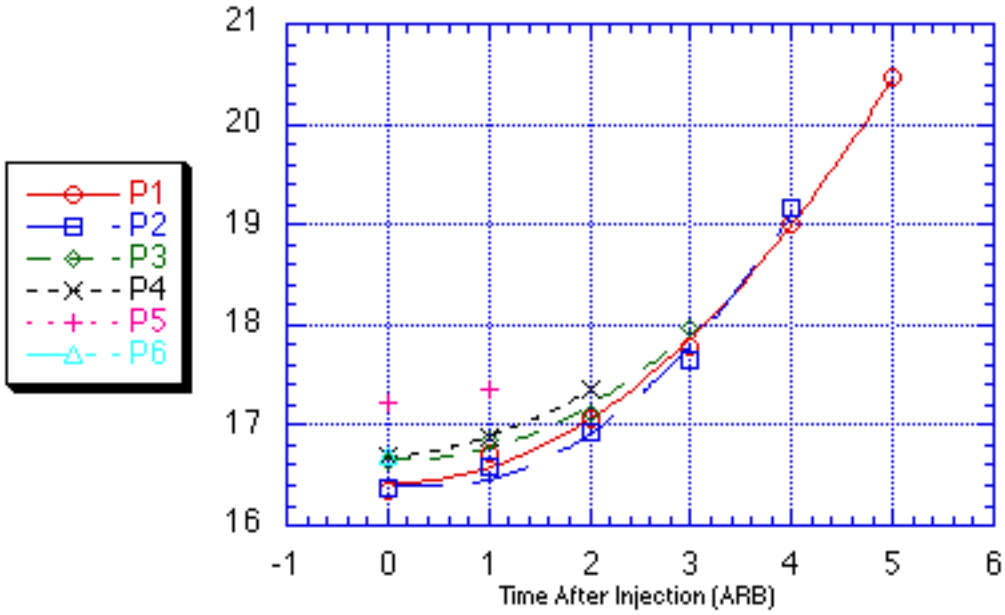

Figure 6 Proton Vertical. Emittance When Later Proton Bunches are Injected

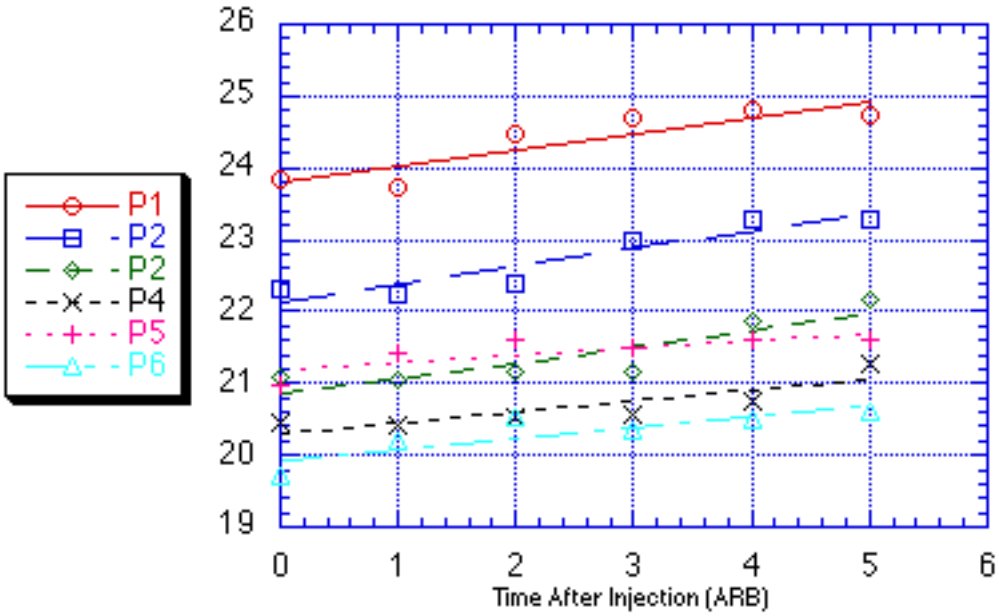

Figure 7 Proton Vertical. Emittance When Later antiproton Bunches are Injected

The proton emittance also grows during antiproton injections, as is shown in Figure 7. It has been pointed out that Figure 7 could easily be fit by a series of discrete jumps caused by the $\bar{p}$ kicker blowing up the nearby proton bunches.

The observed emittance growth could also be caused by the kickers if there is prolonged ringing. Figure 8, taken from a report in preparation by Jerry Annala and Bruce Hanna, hints that there could indeed be such ringing. The figure shows the amplitude of motion observed on a horizontal BPM as a function of time after the kick, and shows that one kick may disturb all the beam in the machine.

The apparent 3π growth of protons while opening the helix is thought to be an instrumental effect, since at times in the past the emittance would return to its previous value if the orbits were returned to the machine center.

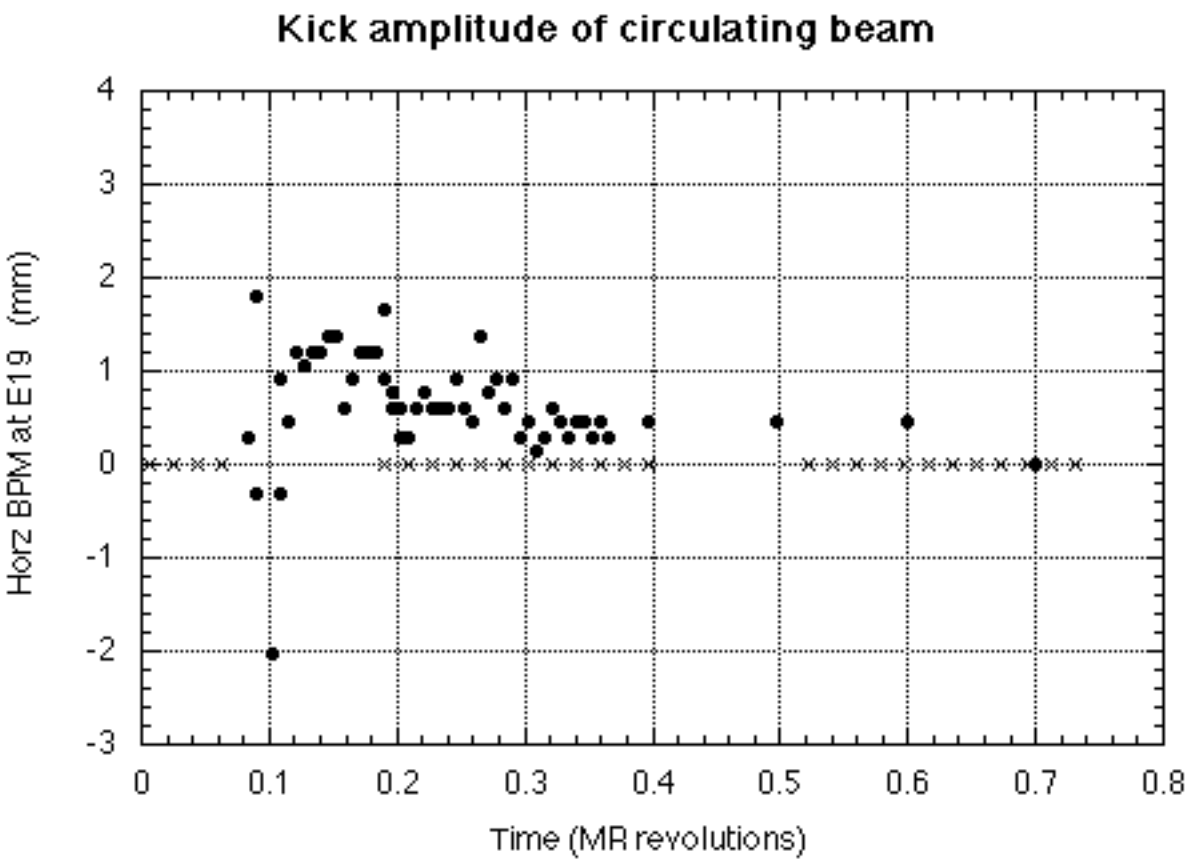

Figure 8 Observed Beam Motion as a Function of Time after Injection Kicker Firing

Another interesting effect is shown in Figure 9, where the $\bar{p}$ transverse emittance is shown as a function of time after injection. The transverse emittance of the $\bar{p}$s clearly shrinks as a function of time. In addition, A1 clearly starts with a larger vertical emittance than the others. This may result from the larger injection errors at the start of $\bar{p}$ injections in both the Main Ring and Tevatron. Since $\bar{p}$s are lost during the injection process, it is likely that this decrease in vertical emittance is real.

There is no clear signature in Table I of emittance growth up the ramp. The apparent emittance growth at "TEV flattop" is clearly an instrumental effect since the emittance shrinks afterwards without significant particle loss (as will be shown below).

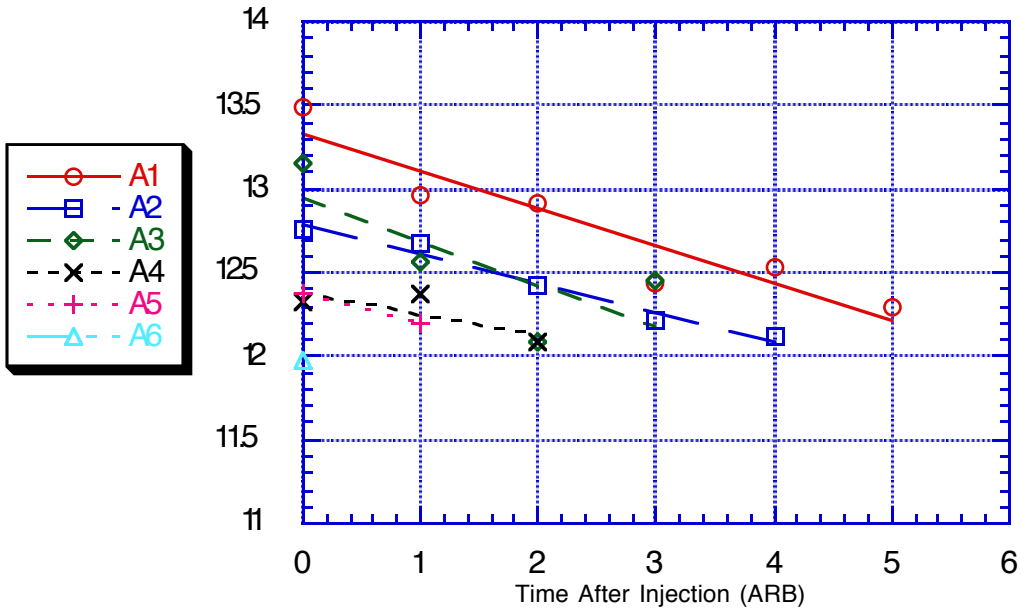

Figure 9 $\bar{p}$ Vertical. Emittance as Function of Time after $\bar{p}$ Injection

In summary, for the last 10 stores there is clear evidence for proton emittance growth while at 150 GeV both on and off the helix. There is also clear evidence that the antiproton emittance shrinks. The data suggest a few π emittance growth of both protons and antiprotons at injection, consistent with our expectations. The data also suggest a few π growth from 150 GeV to collisions and opening the helix; however, systematic errors dominate these measurements.

Longitudinal Emittance Growth

Table II shows proton and antiproton longitudinal emittance measurements by the Tevatron SBD (Sampled Bunch Display). The values in this table are averages for the last 10 shots containing valid measurements (i.e. shots 5602, 5610, 5620, 5621, 5622, 5624, 5625, 5627, 5628 and 5629). It is clear from the results in the table that there is very little longitudinal growth at injection, but there is significant growth up the ramp. It should be noted, however, that the RF bucket is nearly full at injection, so there is very little room for longitudinal emittance dilution. Even a small amount of longitudinal emittance growth at injection would result in dc beam.

TABLE II
Proton and Antiproton Longitudinal Emittance at Various Steps
(Averaged for 10 Shots)

PROTON LONGITUDINAL SBD

	ENERGY	P1	P2	P3	P4	P5	P6	WT.AVG	
MR extraction	150	2.87	2.86	2.86	2.86	2.87	2.56	2.57	eV-sec
TEV injection	150	3.02	3.01	3.00	2.99	2.99	2.68	2.69	eV-sec
TEV open helix	150	2.93	2.93	2.94	2.93	2.94	2.89	2.93	eV-sec
TEV before ramp	150	2.89	2.89	2.89	2.88	5.59	2.79	2.88	eV-sec
TEV flattop	900	This measurement is invalid - taken while still ramping							
TEV squeeze	900	3.74	3.73	3.79	3.74	3.79	3.82	3.76	eV-sec
TEV collisions	900	3.77	3.79	3.84	3.76	3.81	4.31	3.80	eV-sec
TEV in Store	900	3.74	3.77	3.82	3.75	3.79	3.53	3.77	eV-sec

Antiproton LONGITUDINAL SBD

	ENERGY	A1	A2	A3	A4	A5	A6	WT.AVG	
MR extraction	150	2.65	2.63	2.67	2.66	2.74	2.43	2.67	eV-sec
TEV injection	150	2.76	2.70	2.75	2.75	2.82	2.15	2.76	eV-sec
TEV before ramp	150	2.67	2.65	2.71	2.72	2.80	2.83	2.71	eV-sec
TEV flattop	900	This measurement is invalid - taken while still ramping							
TEV squeeze	900	3.29	3.31	3.31	3.36	3.43	3.22	3.34	eV-sec
TEV collisions	900	3.37	3.36	3.42	3.43	3.49	2.11	3.42	eV-sec
TEV in Store	900	3.41	3.41	3.48	3.49	3.51	2.16	3.47	eV-sec

Particle Losses in the Tevatron

The particle lifetimes at 150 GeV are known to be small. Bunched beam lifetimes of 13 hours are considered good. Figures 10 and 11 show the fraction of protons and antiprotons lost during the loading of the Tevatron at 150 GeV. This is defined as (SBD from "TEV before ramp" case - T:STC1MI(SBD) measurement on the first turn in the Tevatron)/T:STC1MI. Typically about 6-7% of the protons are lost while at 150 GeV/c, and about 2-3% of the antiprotons are lost. These losses appear to be constant during the time period shown, from March 6, 1995 through July 23, 1995.

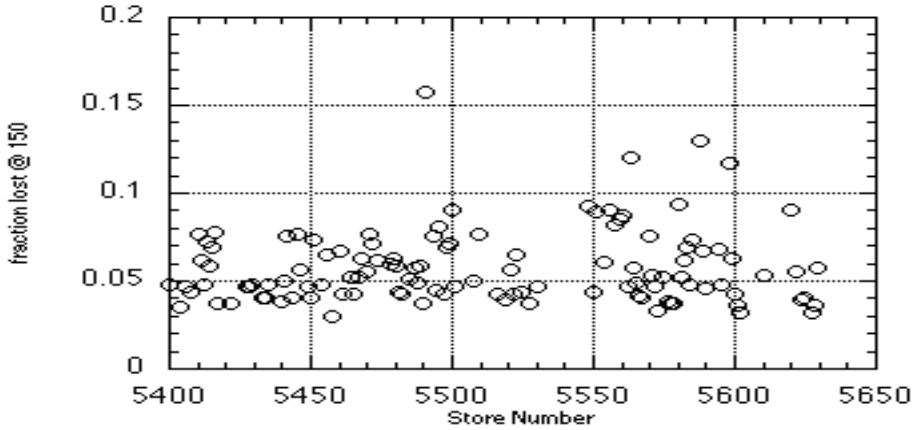

Figure 10 Proton Losses at 150 GeV in Tevatron

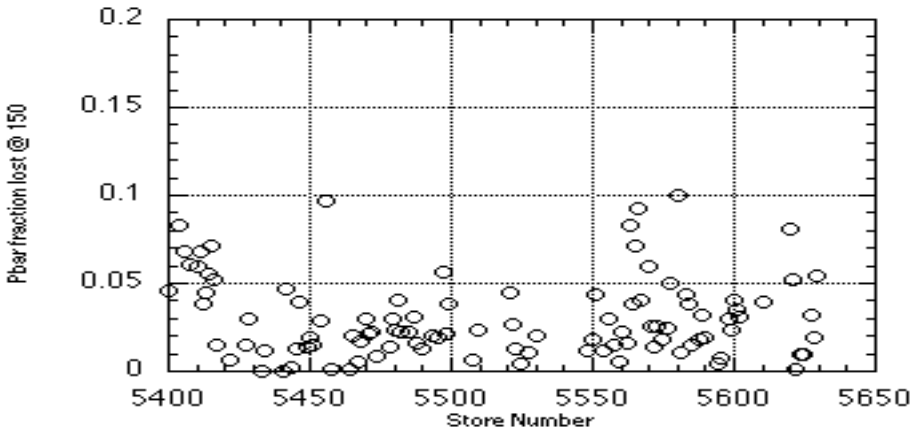

Figure 11 Antiproton Losses at 150 GeV in Tevatron

Figures 12 and 13 show the proton losses up the ramp and at 900 GeV before scraping for the last part of the run. The "fraction lost up the ramp" is defined as (SBD from TEV @ 900 case - SBD from before ramp case)/SBD from before ramp case. This has the flaw mentioned above that the SBD @ 900 case is measured while still ramping, but this shouldn't effect the intensity much. These losses show more variation with running conditions, with noticeable improvements following Run 5500.

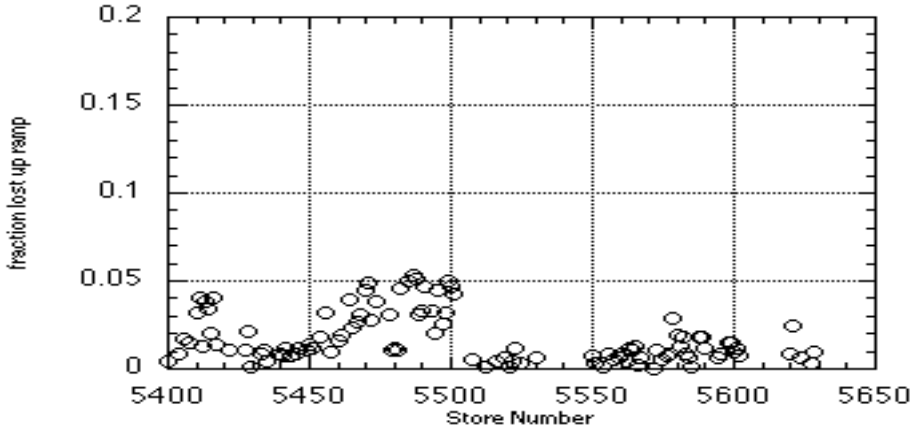

Figure 12 Proton Losses up the Ramp in Tevatron

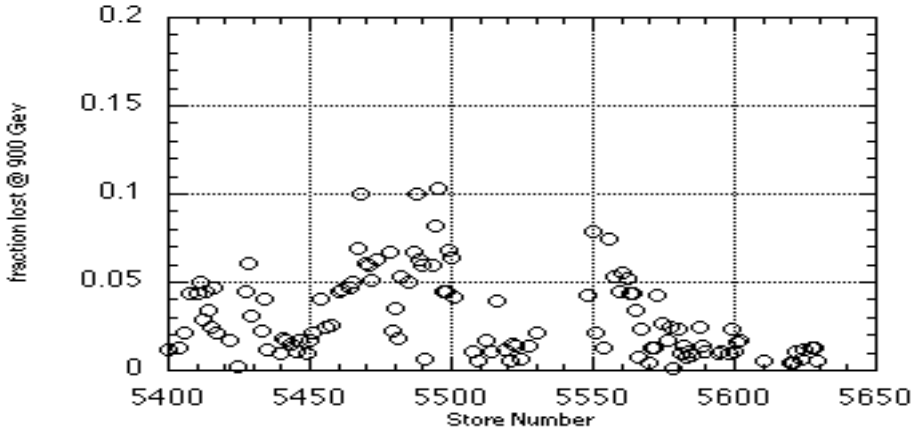

Figure 13 Proton Losses at 900 GeV Before Scraping in Tevatron

Figures 14 and 15 show the antiproton losses up the ramp and at 900 GeV before scraping. Both show a noticeable improvement for stores 5500-5550.

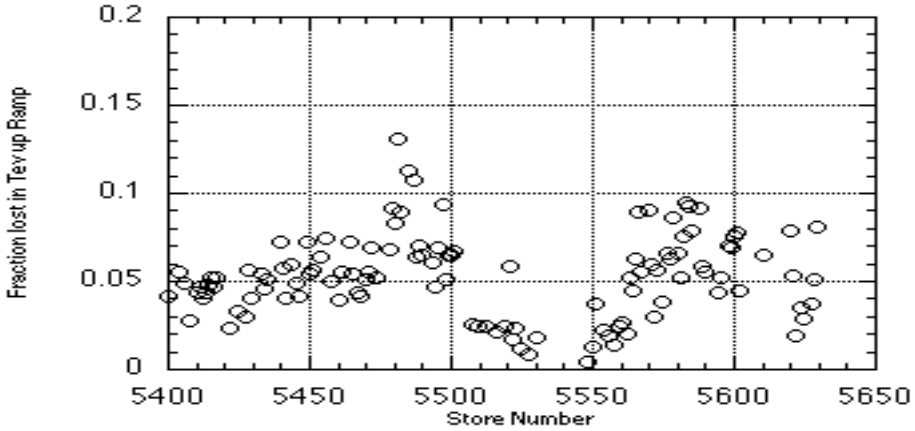

Figure 14 Antiproton Losses up the Ramp in Tevatron

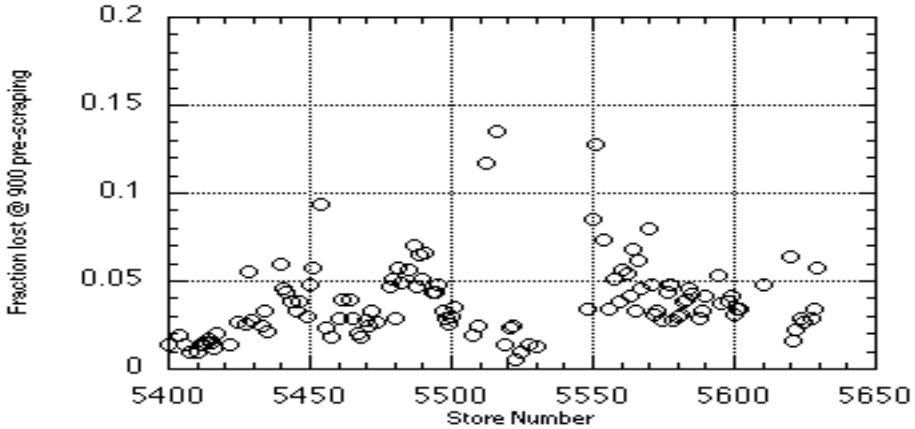

Figure 15 Antiproton Losses at 900 GeV Before Scraping in Tevatron

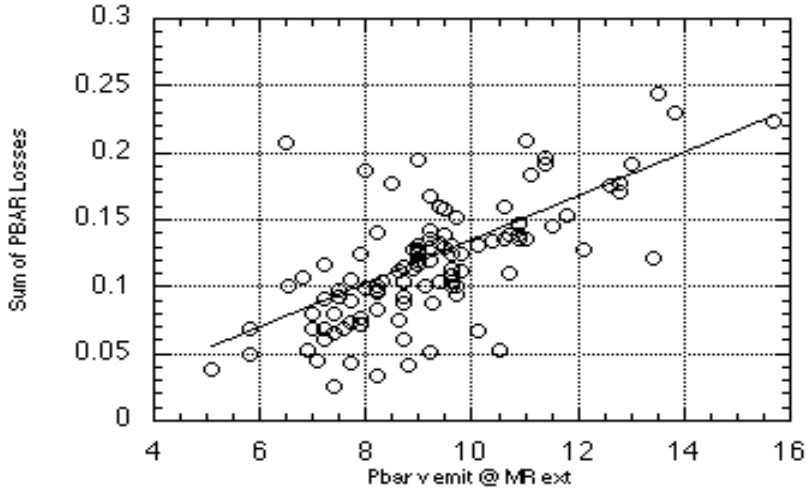

Figure 16 Sum of Antiproton Losses as a Function of MR Extraction Vertical Emittance (Runs 5400-5529)

The transverse emittance of the $\bar{p}$ delivered to the Tevatron varies by a large amount, so the $\bar{p}$ losses could be correlated to transverse size of the beam at Main Ring extraction.

The data in Figure 16 show such a strong correlation. There would certainly also be a correlation with $\bar{p}$ intensity, or initial luminosity, since $\bar{p}$ transverse sizes are known to be correlated to those quantities as well.

Figure 17 shows the subset of the data shown in Figure 16 from stores 5500-5550. It is clear that the $\bar{p}$ losses were smaller than our normal expectations during this time.

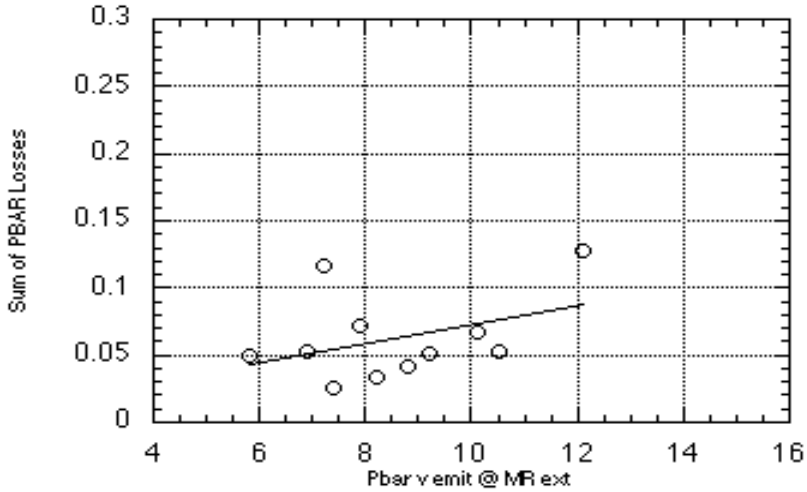

Figure 17 Sum of $\bar{p}$ Losses during the record running period as a function of MR extraction vertical emittance (Runs 5500-5550)

In summary, proton losses can be reduced to about 5% with careful tuning. $\bar{p}$ losses can also be reduced to this level, but that appears to be a more subtle art. Under the best running conditions the dominant particle losses are at 150 GeV in the Tevatron. It would be useful to document what was different about the running conditions for Stores 5500-5550, where the $\bar{p}$ losses up the ramp and through the squeeze were significantly smaller than normal.

Luminosity Lifetime

During the course of a normal collider store, the luminosity decays to roughly 1/3 of its original value before refilling the collider. This decay rate has a large impact on the integrated luminosity. Figure 18 shows data from Store 5416, which was chosen as a typical store without any outstanding problems. The CDF luminosity measurement for this store is shown with the scale on the left, and the luminosity lifetime in hours is shown with the scale on the right. The horizontal axis is time in minutes from the beginning of the store. The rise in luminosity lifetime with time is typical of all stores.

Figure 19 shows the measured (CDF) and calculated luminosity for Store 5416. The calculated luminosity is determined from measurements of transverse and longitudinal emittances of each beam, and proton and antiproton intensities. The calculated luminosity shows discrete jumps as a result of the forward-backward asymmetry of the flying wire emittance measurements. Due to the relatively large uncertainties in the measurement of beam quantities, it is the luminosity measurements made by the collider experiments that are normally used for the official collider run statistics. An excellent description of the experimental luminosity measurement techniques is contained in J. Bantly et. al. TM-1906.

The calculated luminosity shown is a full overlap integral calculation using the measured quantities of each bunch as they collide at B0. The good overall agreement between the measured and calculated luminosity confirms the validity of the beam measurements,

especially the flying wire measurements. However, for most stores the measured luminosity falls faster than expected on the basis of beam measurements. That is, the lifetime of the luminosity calculated from beam measurements over-estimates the measured luminosity lifetime. This anomalous lifetime behavior is not understood.

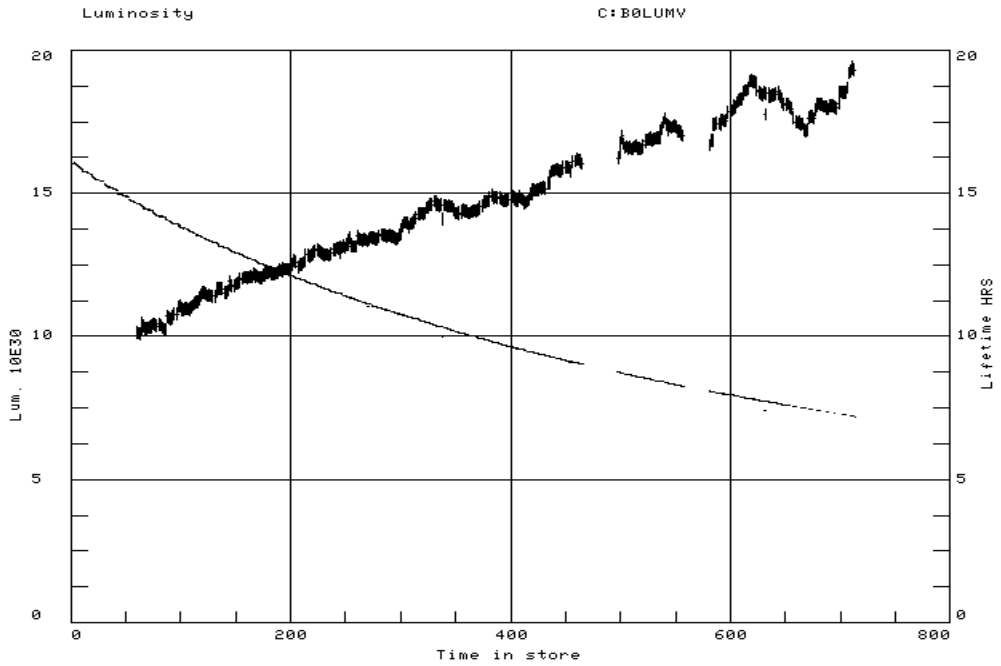

Figure 18 CDF Luminosity & Lifetime for Store 5416

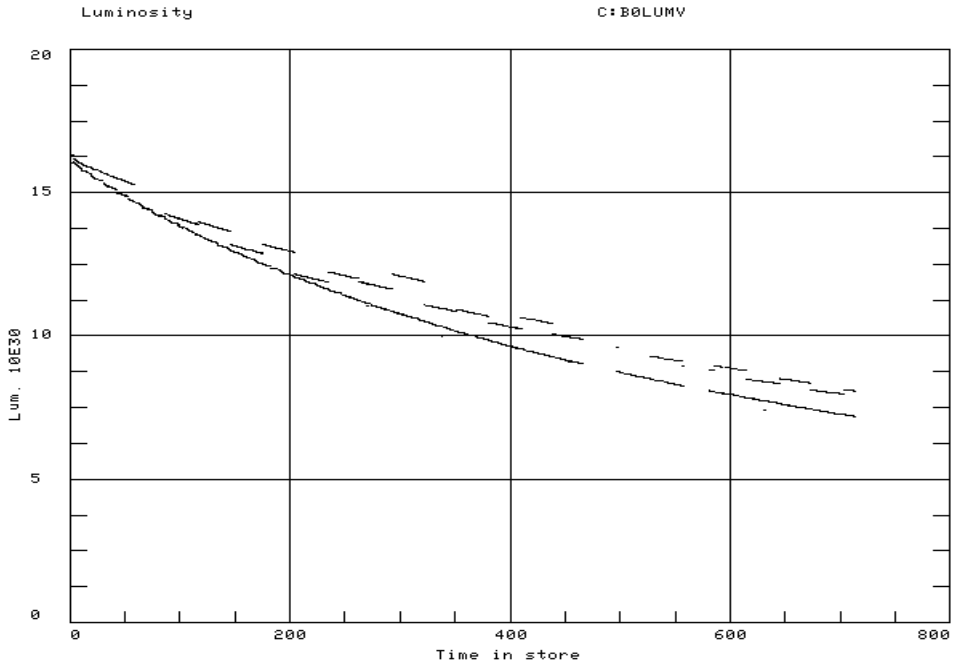

Figure 19 Experimental and Calculated CDF Luminosity for Store 5416

Figure 20 shows the average transverse emittances of the six bunches measured for this store, and the emittance lifetimes. The proton horizontal emittance (calculated with a $\Delta p/p$ correction based on the SBD bunch length measurement) starts at about 22π and rises to about 25π at the end of the store. The lifetime is off scale in some places, but it is roughly 80 hours. The proton vertical emittance starts at about 25π and rises to about 30π at the end of the store, yielding a proton emittance lifetime of about 75 hours. The $\bar{p}$ horizontal emittance starts at about 16π and rises to 20π at the end of the store, giving a $\bar{p}$ horizontal emittance lifetime of about 75 hours. The $\bar{p}$ vertical emittance starts at about 15π and rises to 20π at the end of the store, yielding a $\bar{p}$ vertical emittance lifetime of about 50 hours.

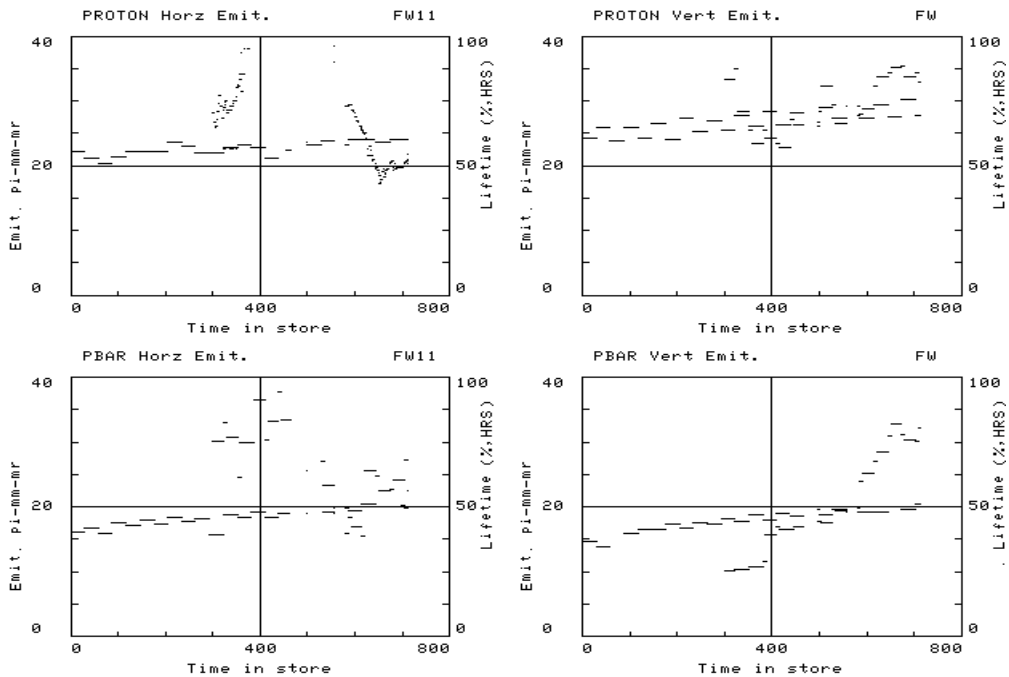

Figure 20 Flying Wire Emittances and Emittance lifetimes for Store 5416

Figure 21 shows the relative contributions of each of the four quantities in Figure 20 to the measured luminosity lifetime. This is the result of a numerical differentiation of the full bunch by bunch overlap integral luminosity formula. The contribution of each individual transverse quantity is 10% or less. The sum of all the transverse emittance growth contributions to the luminosity is about 40%, which is typical of most stores.

Figure 22 shows the measured proton and antiproton bunch lengths and intensities, and the lifetimes of each quantity. In these plots the smooth curves are the measurement (scale on left), and the other curve is the calculated lifetime (scale on right). The lifetimes of all four quantities typically rise during the store in a manner similar to what is shown. In particular, the increase in lifetime of the bunch length agrees with what is predicted from intra-beam scattering.

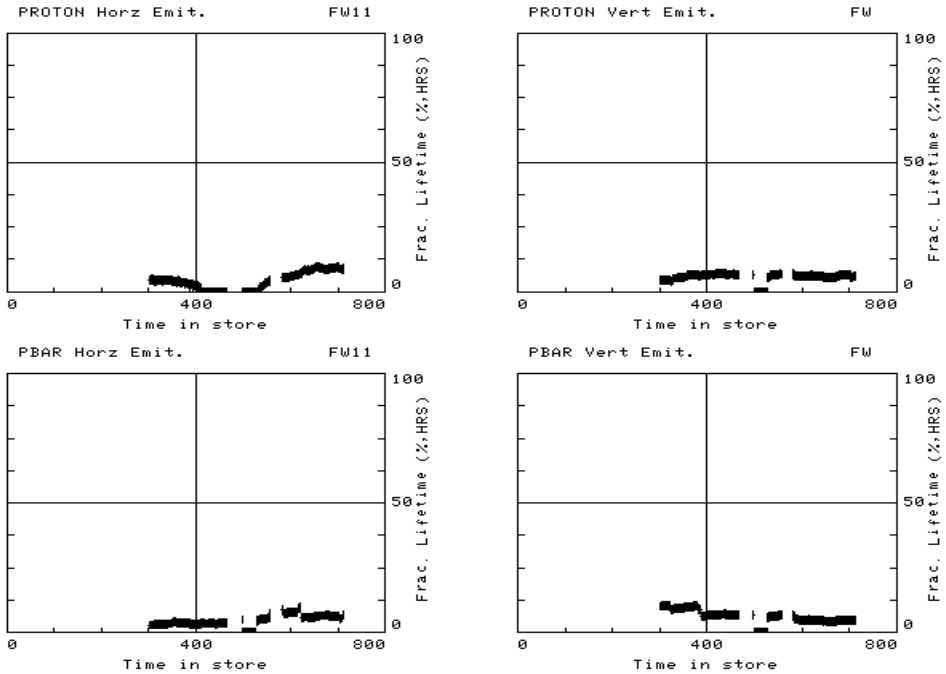

Figure 21 Transverse Growth Contribution to Luminosity Lifetime for Store 5416

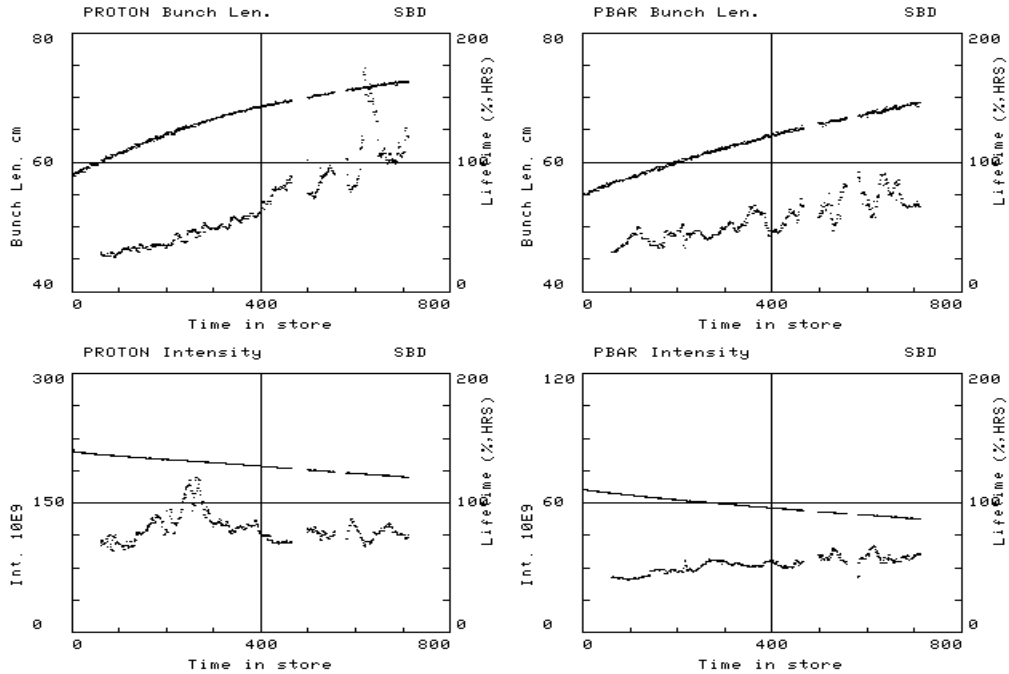

Figure 22 Bunch Length & Intensities with Lifetimes for Store 5416

Figure 23 shows the relative contributions of each of the four quantities in Figure 22 to the measured luminosity lifetime, as described above. The growth in antiproton and proton bunch lengths contributes roughly 10% apiece to the luminosity lifetime, the proton particle lifetime contributes roughly 20%, and the $\overline{p}$ particle lifetime contributes roughly 30%. For store 5416 the integrated luminosity delivered to CDF and D0 totaled 906 inverse nanobarns. For a 75 millibarn total cross section, this consumes 6.75×10^{10} $\overline{p}$s (and

protons). The total number of $\bar{p}$s lost during this store was 8×10^{10} out of the initial 40×10^{10} circulating at the start of the store.

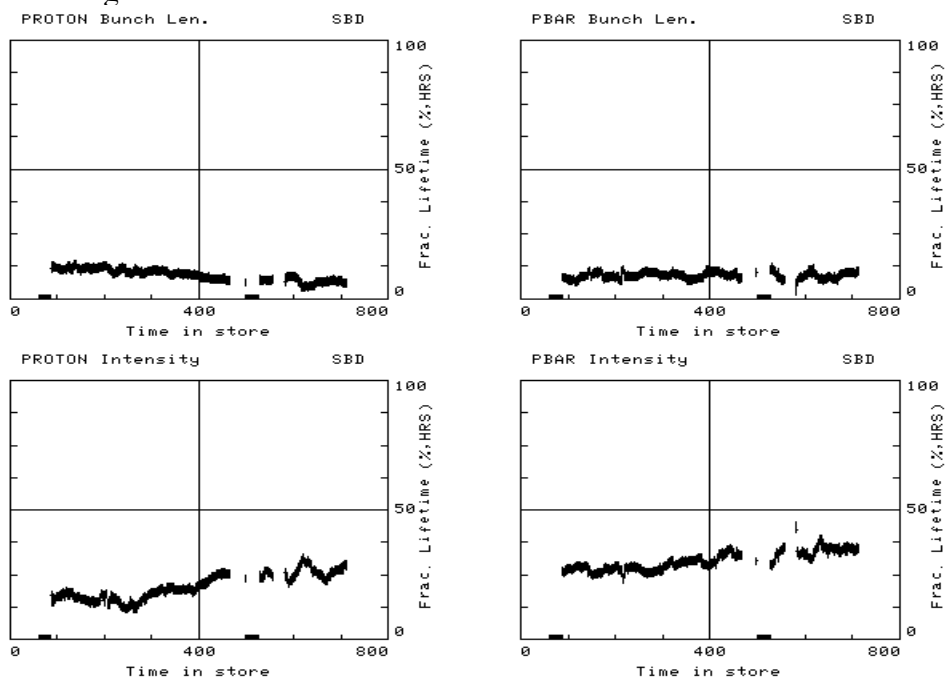

Figure 23 Bunch Length & Intensities Contribution to Luminosity Lifetime for Store 5416

For store 5562 and afterwards the luminosity monitors for CDF and D0 were switched to new devices that do not saturate at high luminosity. This gave us a new opportunity to check if the luminosity lifetime varies with initial luminosity. Figure 24 shows the quantity $\mathcal{L}/(N_p N_{\bar{p}})$ for these stores, where the particle intensities are measured at collisions. This quantity is normally used as an indicator of transverse blowups. To filter out bad stores, Figures 25-27 only show those stores in Figure 24 with $\mathcal{L}/(N_p N_{\bar{p}}) > 11.0$.

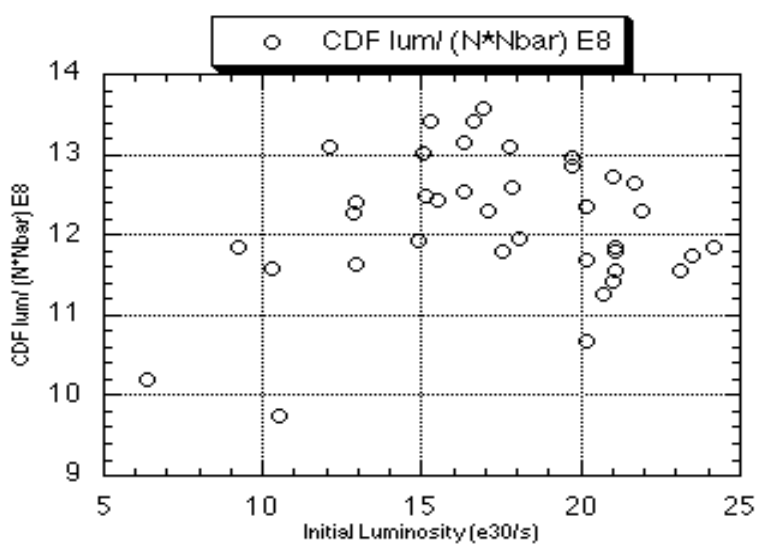

Figure 24 CDF $\mathcal{L}/(N_p N_{\bar{p}})$ for Recent Stores (Stores 5562-5629)

Figures 25-27 show the measured luminosity lifetimes for the selected stores measured at three different times.

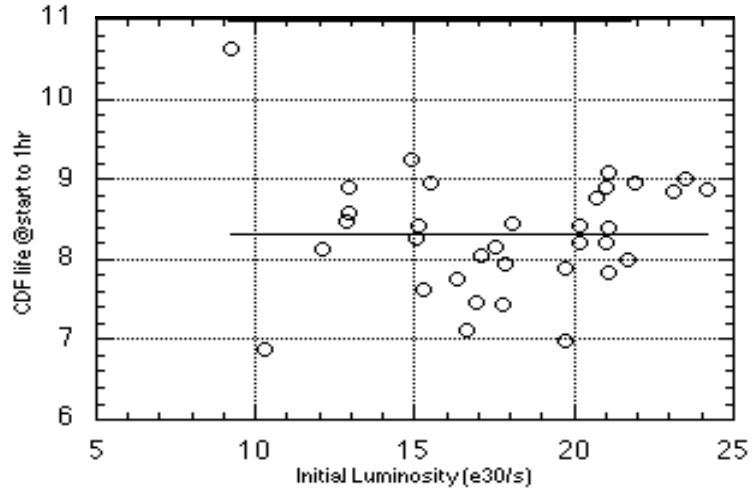

Figure 25 Measured CDF luminosity lifetime in the first hour of the store

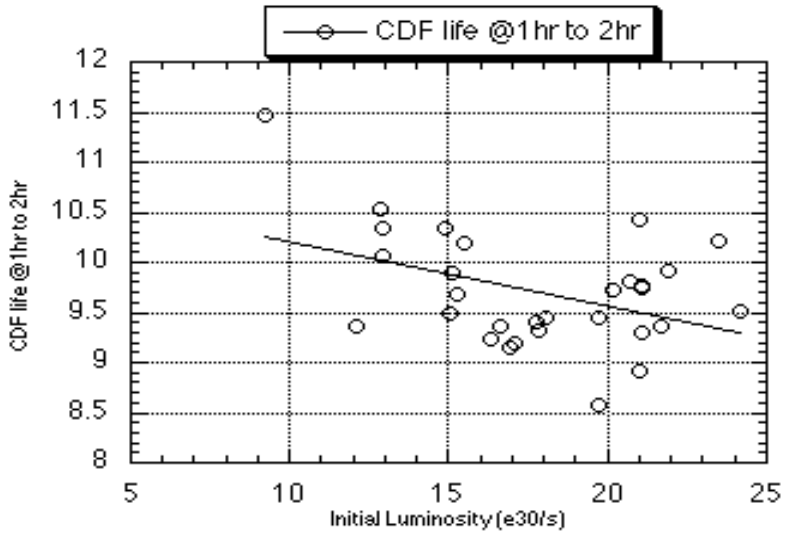

Figure 26 Measured CDF luminosity lifetime in the second hour of the store

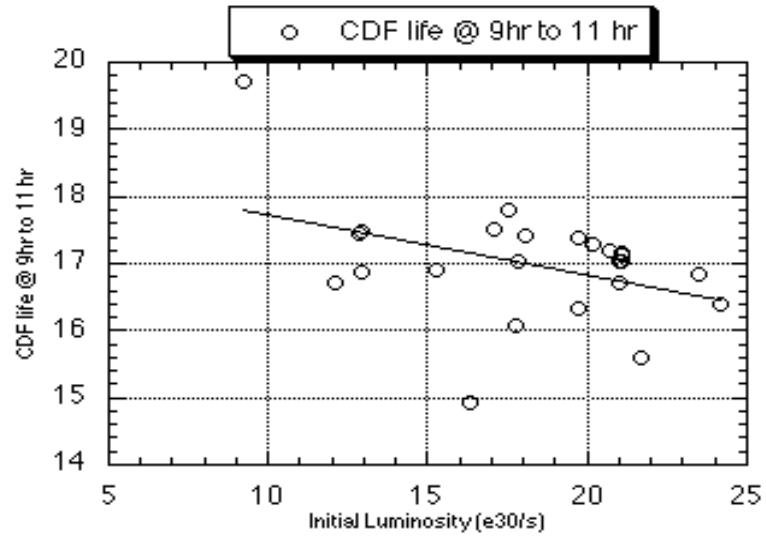

Figure 27 Measured CDF luminosity lifetime at ten hours into the store

Tune Spectra

Figures 28-33 show the transverse Schottky spectra from the control room spectrum analyzers at several points during a typical shot setup. The vertical lines on these spectra mark the proton tunes. In these figures, the upper tune appears more strongly on the horizontal pick up than on the vertical pick up, thus the upper tune is identified as the horizontal and the lower tune as the vertical. However, both frequencies generally appear on the vertical pick ups. This may indicate that the machine is not locally decoupled at the vertical pick-up.

In some of the spectra (particularly the one at collisions) there are hints of additional tune lines. In some cases, these lines are more distinct. It is suspected that they are the $\bar{p}$ tunes. (The $\bar{p}$ tunes are generally expected to differ from the proton tunes both due to beam beam effects and due to their separated orbits). However, the tune lines of even a single beam sometimes do not have a simple shape. It is not at all certain that the $\bar{p}$ tunes are visible in any of these spectra.

When the tunes are adjusted, the primary concern is to keep them off the 4/7ths = 0.5714 and well clear of the 3/5ths, below about 0.595. Although the 7/12th = 0.5833 resonance is nearby, it usually does not strongly affect the beam. An effort is also made to keep the widths of the tune lines reasonable. A very sharp and narrow line, or a very wide line may suggest that the chromaticity is too low or too high, respectively for that plane.

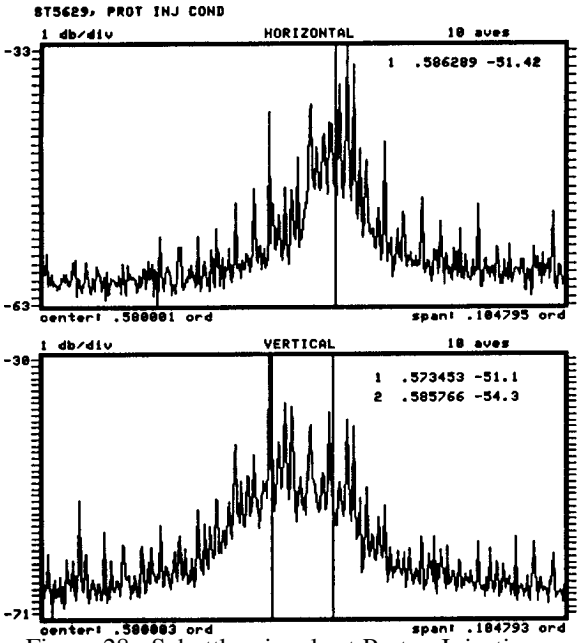

Figure 28. Schottky signals at Proton Injection.

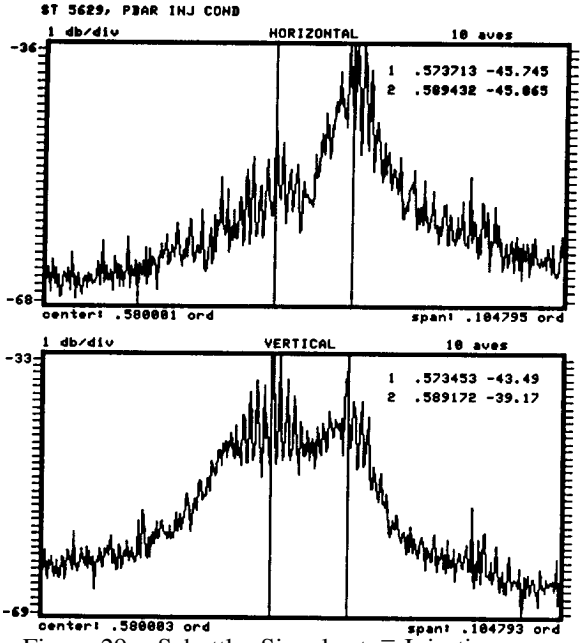

Figure 29. Schottky Signals at $\bar{p}$ Injection.

At 150 GeV (the proton injection, antiproton injection, and ready to ramp spectra) the lower tune is near and/or overlaps the 4/7ths. Despite this, the overall Tevatron performance (particle lifetimes and transverse emittance growth rates) was good. To get good particle lifetimes for the $\bar{p}$ injection and ready to ramp conditions it is often necessary to move the lower tune down to about .565, below the 4/7ths.

At 150 GeV the chromaticity is intentionally kept fairly high, roughly 20 to 30 units, to avoid an instability at the start of the ramp from 150 GeV to 900 GeV. During the first 3 seconds of the ramp the time dependent persistent currents in the main dipoles unwind. Left uncorrected these would change the chromaticities by tens of units. A special

sextupole ramp is used to compensate for these changes. This typically works quite well, for correcting the bulk of the effect. But to provide a little extra margin against small day to day variations in the time dependent persistent currents, (e.g. from differences in the history of the machine) the 150 GeV chromaticities are kept fairly high. However, if the chromaticities at 150 GeV are too high, the beam lifetimes are often seen to suffer.

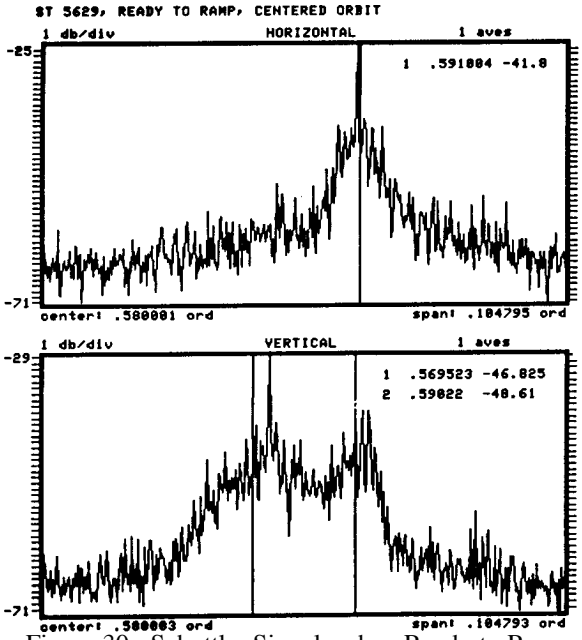

Figure 30. Schottky Signals when Ready to Ramp.

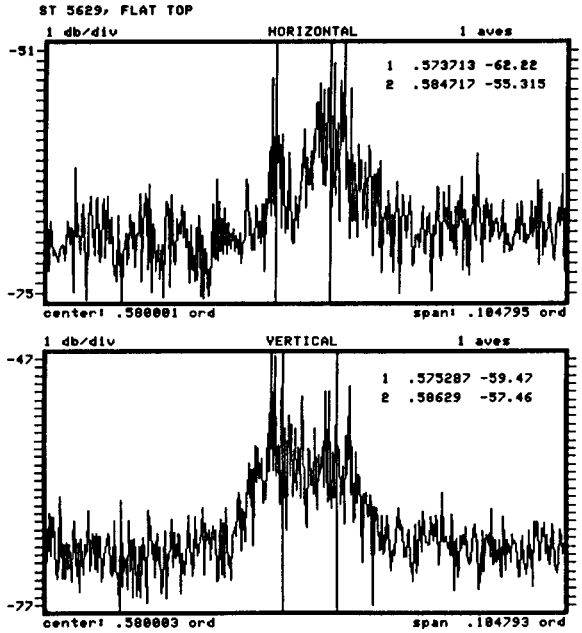

Figure 31. Schottky Signals at Flat Top.

At collisions, the chromaticities are quite small and may even be slightly negative. (They are probably between about -5 and 5 units in both planes.) This seems to be a requirement for good particle and luminosity lifetimes. It has also been observed that when the beams are colliding they can tolerate chromaticities that would make a single beam unstable.

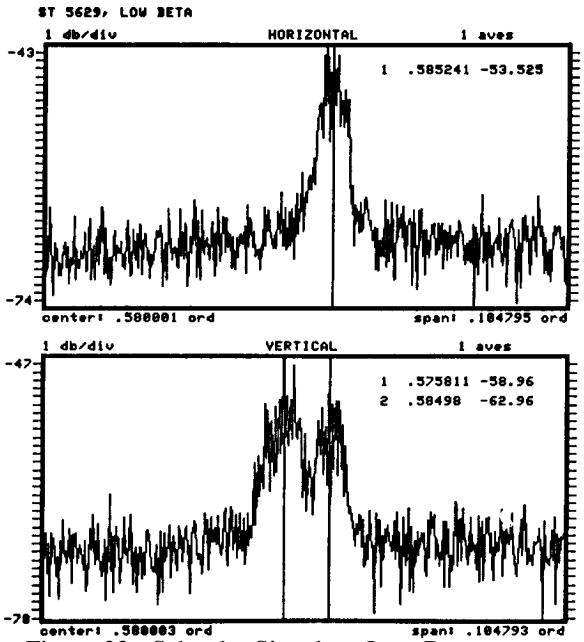

Figure 32. Schottky Signals at Low Beta.

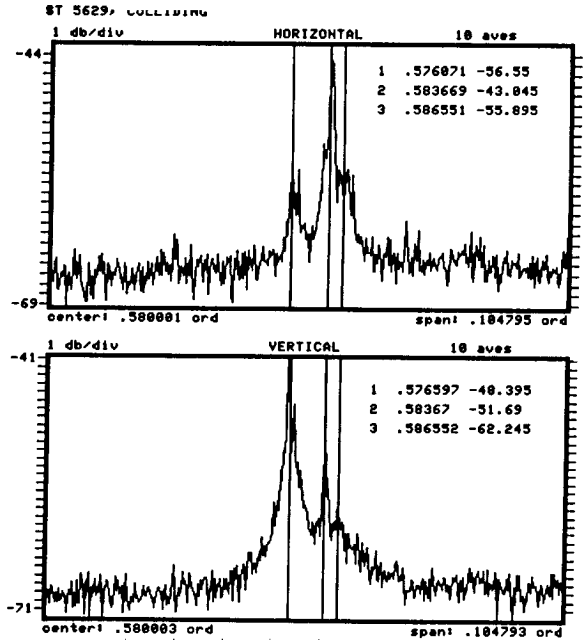

Figure 33. Schottky Signals at Collisions.

Reliability and Shot Setup Time

Table III gives a summary of the reasons for ending all stores in the time period 12/15/93 to 7/24/95. About 71% of the stores were ended intentionally. The most common faults were with controls, the quench protection system, and power supplies at about 4% each.

Table III
Tevatron Collider Reliability Summary

Collider Run 1b
12/15/93 - 07/24/95

Reasons for terminating stores	*Number of Stores Terminated*	*Store Hours*
Intentional	352	5294.74
Controls	20	182.35
Quench Protection System	19	129.57
Miscellaneous	12	89.91
Correction Magnet Systems	6	22.31
Cryogenics	11	64.37
Low Beta Quadrupoles	16	73.01
Utilities	9	48.24
Human Error	9	77.00
Tevatron Power Supplies	19	127.39
Glitches	6	50.30
Tevatron RF	10	105.20
Vacuum	3	25.39
Instrumentation	1	1.73
Experimental Areas	1	13.97
Magnet Failure	0	0.00
Kicker Pre-Fire	2	40.45
Total	***496***	***6345.93***

Total Store Hours = 6345.93 for a mean store length of 12.79 hours
352 stores were ended intentionally with a mean store length of 15.04 hours
144 stores were ended by failure with a mean store length of 7.30 hours
71.0% of the stores were ended intentionally
Current through store #5629

Figure 34 shows a histogram of the shot setup times from 12/15/93 to 07/24/95. Typical shot setups, including the half hour of quiet time, took about 2.5 hours. There are no obvious steps in the shot procedure that dominate the shot setup time, there are simply a large number of steps that require human intervention and feedback.

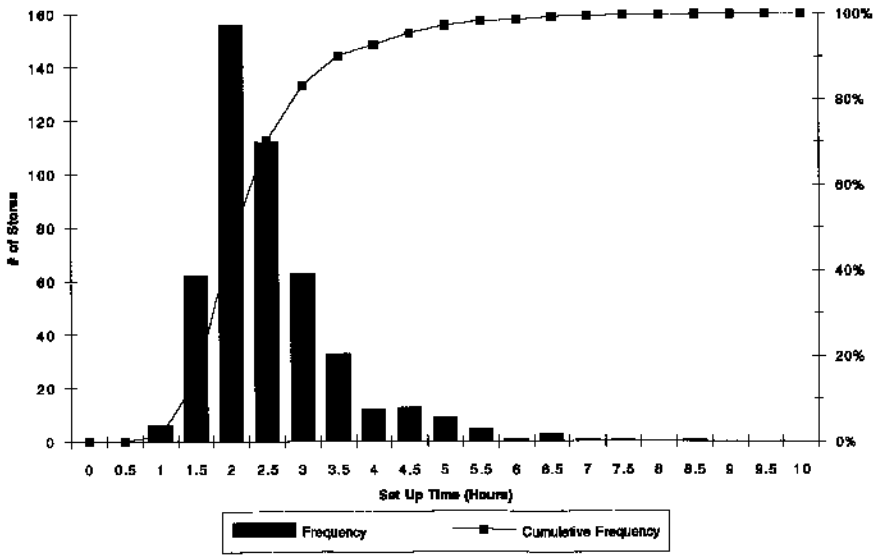

Figure 34 Shot Setup Times

Tevatron Conclusion

For a number of years we have talked about "emittance growth on the ramp" as being a major problem with the Tevatron. The present data taken from "good" stores does not favor any large emittance growth on the ramp, although the data cannot rule it out either. What is clear is that there are significant changes to the transverse emittance while injecting at 150 GeV. There is very clear emittance growth of the protons during the injection process. This does not appear to be entirely a beam-beam effect, since most growth of protons occurs before antiprotons are loaded. There is some suspicion that this growth is caused by ringing of the injection kickers. The $\overline{p}$ bunches show a significant transverse shrinkage, suggesting a physical or dynamic aperture problem that may be related to beam-beam effects.

7. Summary

Fermilab collider Run 1b began in the wake of the very successful Run 1a and with the anticipation of large increases in antiproton production and overall collider integrated luminosity from the successful installation and commissioning of the Linac Upgrade. However, the onset of the expected good running was delayed due to anomalously poor Tevatron performance. The Tevatron difficulties were eventually traced to a rolled quadrupole in the B0 low beta insert. Once this problem was corrected, the performance of the collider complex improved dramatically and rapidly. The accelerator performance goals set for this run were very quickly achieved. It is certainly accurate to say that collider Run 1b was ultimately a success.

The Linac Upgrade project performed as designed and the Booster was able to accept and accelerate the higher intensities generated while maintaining emittances small enough that the beam could be accepted by the Main Ring. A major reason for this success was the active beam damper systems that were developed for the Booster - a testament to getting good people involved in R&D on needed systems very early (Booster damper R&D was started in 1989).

The Main Ring bunch coalescing system was upgraded in March of 1995, resulting in a significant performance improvement. It was unfortunate that this upgrade had not started much sooner. After this upgrade the average $\bar{p}$ "overall" efficiency from accumulator to collisions was 60%.

In the Antiproton source, improvements to the stacktail momentum cooling and the addition of 2-4 GHz transverse core cooling in the Accumulator led to improved stacking with large $\bar{p}$ stacks.

About half way into the run, the Accelerator Division management structure for collider operations was changed from that of a single run coordinator serving a three month term, to that of a semi-permanent run coordination team consisting of a collider coordinator and five run coordinators. Each member of the run coordination team served as the overall run manager for intervals of about one month at a time. This position rotated through all members of the run coordination team for the remainder of the run. It is our belief that this team approach was a good idea.

The "typical best" performance of Run 1b is characterized by initial luminosities of over 2×10^{31} $cm^{-2}\bullet sec^{-1}$, average antiproton stacking rates of 5.5 mA/hour at stacks of 120 - 160 mA, and integrated weekly luminosities of almost 4 inverse picobarns. The best collider store initial luminosity was approximately 3×10^{31} $cm^{-2}\bullet sec^{-1}$. The exact value of the best initial luminosity is uncertain since the highest luminosity stores occurred at a time when there were high rate saturation problems with both the CDF and D0 luminosity monitors at these values. The best weekly integrated luminosity was 4.9 inverse picobarns.

26×10^{13} $\bar{p}$s were unstacked during Run 1b for use in the collider (as of July 1995). Of these, 50% (13×10^{13} $\bar{p}$s) survived to collisions. The total delivered luminosity to both CDF and D0 consumed 1.9×10^{13} $\bar{p}$s (and the same number of protons).

Acknowledgments

The authors would like to thank Frank Bieniosek, Mike Church, Elvin Harms, Jim Morgan, David Olivieri, Ralph Pasquinelli, and Dave Vander Meulen for much of the data relevant to antiproton source operation and for many helpful discussions. Comments, and corrections from Gerry Annala, Saeed Assadi, Peter Bagley, John Marriner, and Mike Martens were incorporated into the Tevatron performance section. Ionis Kourbanis and Dave Capista contributed to the Main Ring section. John Crawford provided the Collider reliability data.

The authors also acknowledge the efforts of the Accelerator Division operations department for their superb skill in operating the Fermilab accelerator complex and for much useful dialog during the course of the run.

The success of Fermilab Collider Run 1 was the result of a cooperative effort that benefited from the ability, hard work, and dedication of technicians, engineers, physicists, and others from every part of the Accelerator Division organization. It is certainly true that each person who was a part of Accelerator Division during Run 1 owns a part of this success.