Electron cooling of 8 GeV antiprotons at Fermilab's Recycler: Results and operational implications *

L.R. Prost[#], D. Broemmelsiek, A. Burov, K. Carlson, C. Gattuso, M. Hu, T. Kroc, J. Leibfritz,
S. Nagaitsev, S. Pruss, G. Saewert, C.W. Schmidt, A. Shemyakin, M. Sutherland, V. Tupikov,
A. Warner
FNAL, Batavia, IL 60510, USA

Abstract

Electron cooling of 8 GeV antiprotons at Fermilab's Recycler storage ring is now routinely used in the collider operation. It requires a 0.1-0.5 A, 4.3 MeV dc electron beam and is designed to increase the longitudinal phase-space density of the circulating antiproton beam.

This paper briefly describes the characteristics of the electron beam that were achieved to successfully cool antiprotons. Then, results from various cooling force measurements along with comparison to a non-magnetized model are presented. Finally, operational aspects of the implementation of electron cooling at the Recycler are discussed, such as adjustments to the cooling rate and the influence of the electron beam on the antiproton beam lifetime.

INTRODUCTION

In the Recycler, the goal of the electron cooler is to reduce the longitudinal phase-space area of the stored antiprotons either to allow for additional incoming transfers or to prepare the bunch for extraction to the Tevatron collider. Soon after the first cooling demonstration, the electron cooler became part of normal operation of the accelerator complex and allowed for the latest advances in the Tevatron luminosity. Meanwhile improvements to the electron beam quality were pursued and the cooling force characterized.

ELECTRON BEAM CHARACTERISTICS

Electron cooling of 8.9-GeV/c antiprotons requires a dc electron beam with kinetic energy of 4.3 MeV and a beam current of 0.1-0.5 A. The main parameters of the cooler are listed in Table 1.

Table 1: Electron cooler main parameters

Parameter	Symbol	Value	Unit
Electron energy	E_b	4.34	MeV
Beam current (for cooling)	I_b	0.1	A
Terminal voltage ripple, rms	δU	250	V
Cooling section (CS) length	L	20	m
Solenoid field in CS	B_{cs}	105	G
Beam radius in CS	r_b	3.3	mm

The dc electron beam is generated by a thermionic-cathode gun, located in the high-voltage (HV) terminal of the electrostatic (Van de Graff-type) accelerator operated in a so-called 'energy recovery' mode [1] and is referred

to as 'recirculation' in this paper.

To provide cooling, the electron beam should recirculate for hours at the nominal energy and a current of hundreds of mAmps. Although this accelerator is in principle capable of sustaining dc beam currents to ground of up to 300 µA, stable operation can only be achieved for very low current losses, typically 2×10^{-5} at 0.5 A (Figure 1). The losses to the acceleration/deceleration tubes are monitored by measuring the resistive divider currents at the top and at the bottom of the column. The beam stability improved greatly when the tube current changes were maintained to less than 1-2 µA, at the detriment of the beam loss to ground. This was achieved through optimization of the optics in the Pelletron, notably in the deceleration tube.
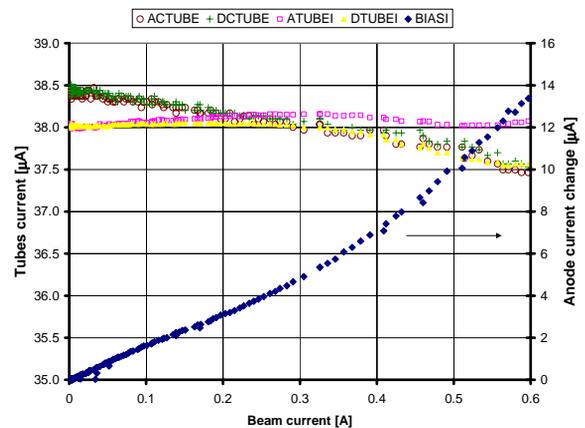

Figure 1: Current losses vs beam current. Blue diamonds are the changes to the anode current, representing the beam current loss. Brown circles and green crosses are the acceleration and deceleration tube resistive divider currents at the bottom. Pink squares and yellow triangles are those at the top.

The main figures of merit to assess the quality of the beam are its rms transverse and longitudinal velocity spreads, which, along with the total beam current, ultimately determine the cooling force.

In our case, the longitudinal velocity spread, expressed as the total electron energy spread in the laboratory frame, is estimated to be approximately 250 eV (rms), dominated by the power supply ripple, δU. Multiple-coulomb scattering and electron beam density fluctuations [2] are estimated to contribute to the order of ~100 eV, added in quadrature.

The transverse velocity spread can be expressed as an angle, θ_e, in the laboratory frame and has several origins

listed in Table 2, along with our present estimation of the magnitude of each component. It should be noted that our estimation of the envelope scalloping assumes that the beam is self-similar (linear optics). Because the envelope scalloping measurement is based on the beam boundary at the 10^{-5} level of losses, if non-linearities are larger than anticipated, it could represent an underestimate of the real scalloping for the core particles [3].

Table 2: Electron angles in the cooling section for I_b = 500 mA. Angles are added in quadrature.

Component	Design	Present estimation	Diagnostics
	µrad	µrad	
Temperature	90	70	OTR + pepper pot
Aberrations	90	50 ≤30 (@ 1 mm)	Simulated BPMs
Envelope scalloping	100	120	Movable orifices
Dipole motion caused by magnetic field imperfections	100	40	Magnetic measurements + BPMs
Beam motion	50	40	BPMs
Drift velocity	20	20	Calculated
Total	**200**	**160**	

Nevertheless, as computed in Table 2, the total rms angle meets our design goal of < 0.2 mrad and the electron beam quality achieved successfully cools 400×10^{10} antiprotons to the longitudinal emittance, close to the final goal (Table 4).

COOLING FORCE MEASUREMENTS

Since the first cooling demonstration of 8.9 GeV/c antiprotons in July 2005, measurements of the cooling force (or drag rate) have been carried out a number of times. First we compared cooling rates for various focusing solutions, and converged to an optimum for operation. The electron beam angles presented in the preceding section are the result of such iterations at the electron beam current of 500 mA. Then, we tried to correlate these measurements with the electron beam properties we measured (or estimated) using a non-magnetized cooling force model.

Measurements methodology

For the results presented in this paper, the method we employed to measure the cooling force experimentally is the so-called "voltage jump" method [4].

The typical number of antiprotons is small (1-5×10^{10}) and we start by making a coasting beam (no RF structure) that fills the entire Recycler ring (longitudinally) and has small transverse emittance (typically < 3 π mm mrad, 95%, normalized). Then, the beam is cooled to an equilibrium state with the electron beam. Transverse stochastic cooling may or may not be

on at that time. Finally, the electron beam energy is changed instantaneously by several keV (at this time, the transverse stochastic cooling system is turned off and remains so for the duration of the measurement). The electron cooling force drags the antiproton longitudinal distribution to this new equilibrium momentum which is M_p/m_e times the voltage jump away from the initial equilibrium. An example of a typical sequence of momentum distribution snapshots is shown on Figure 2, where the electron energy jump was +2 keV (i.e. 3.7 MeV for the antiprotons).
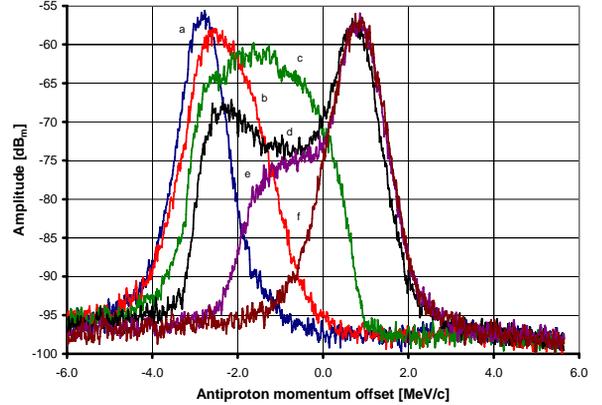

Figure 2: Evolution of the antiproton momentum distribution as measured by a Schottky detector after the energy jump (2 keV). Zero MeV/c offset corresponds to the nominal Recycler energy. (a) Initial distribution. Other traces (from left to right) are taken 2 (b), 5 (c), 18 (d), 96 (e) and 202 (f) minutes after the energy jump. The number of antiprotons was 2.8×10^{10}, the transverse emittance 3-6 π mm mrad (n, 95%) over the duration of the measurement. The electron beam was centered with the antiproton beam and its current was 500 mA.
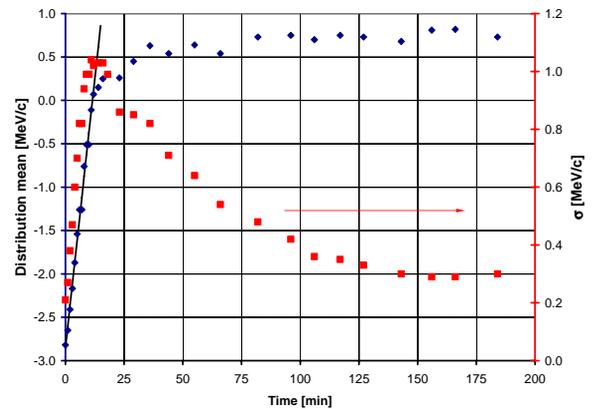

Figure 3: Evolution of the antiproton momentum distribution first (red squares) and second (blue diamonds) moments for the same conditions as for Figure 2. The linear part of the mean's evolution gives a drag rate of 15 MeV/c per hour.

It is clear from the shapes of the distributions in Figure 2 that all particles do not undergo the same drag force. We attribute this behavior to two main effects: the

presence of small diffusion and the fact that the cooling force depends on the betatron amplitude of each antiproton within the distribution – particles with smaller amplitudes are cooled faster. This last effect is combined with the fact that the drag force for large momentum deviations drops considerably and the low-momentum-side tail lags behind the core particles. As a result, the distribution widens even though the initial momentum spread is small. By characterizing the antiproton distribution at a given time by its mean momentum and following its evolution we obtain the corresponding drag rate, which is given by the slope of the linear part of the curve, as illustrated in Figure 3.

Measurements results

As mentioned previously, cooling force measurements were carried out for various focusing settings as a tool for optimization of electron cooling during normal operation. Results presented in this section deal with what we consider our nominal focusing settings, reflecting the latest of all iterations. Once the optics were set, we conducted three types of measurements: 1) cooling force as a function of the electron energy jump, 2) cooling force as a function of the electron beam position and 3) cooling force as a function of the electron beam current.

Figure 4 shows the cooling force as a function of the electron beam energy jump for two electron beam positions: one where the electron beam and the antiproton beam centroid trajectories are the same (so-called 'on-axis' because this also corresponds to our reference orbit) and one where the electron beam is offset by 1.5 mm in the vertical direction with respect to the antiproton beam (parallel shift). Note that in the offset case, the electron beam still overlaps the antiproton beam almost entirely – the antiproton beam radius is ~2.5 mm (95% of particles) and the electron beam radius is ~3.3 mm (~100% of particles). The difference seen on Figure 4 for the cooling force between the two positions is at least partly due to changes to the average electron beam current density and RMS angles that the antiproton beam undergoes.

Additional measurements of the cooling force as a function of the electron beam position were performed for a fixed electron energy jump (2 keV, where the cooling force is close to a maximum). These data have a rough bell shape distribution with the cooling force nearing zero for an offset > 3 mm. Although for this offset amplitude the electron beam still overlaps a large fraction of the antiproton beam, the current density may be locally quite low and the angles quite large [3, Fig. 4]. Unfortunately, direct measurements of the electron beam current density profiles in the cooler were not available at the time, which makes quantitative conclusions difficult since it is not possible to disentangle the relative effects of the angle and current density variations.

For the last set of measurements, the electron energy jump (2 keV) and the beam position (on-axis) were kept constant, while the electron beam current was varied from 100 mA to 500 mA. Unexpectedly, the measured cooling force was found to be almost constant, at 20 ± 2 MeV/c

per hour over this range of currents. One explanation might be that the combination of the variations of the beam envelope and current density as a function of the beam current is such that the increase of the cooling force that would result from the increase of the current density with the beam current is almost exactly compensated by an increase of the local angular spread.
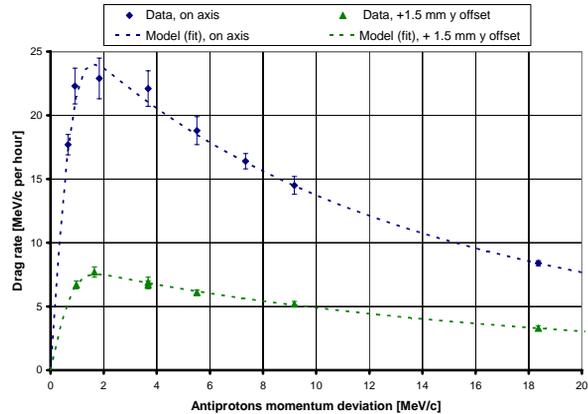

Figure 4: Cooling force (drag rate) as a function of the antiproton momentum deviation. Points are data with error bars representing the statistical error (1σ) of the measurement procedure. Dashed lines are a fit to the data using a non-magnetized model. Blue diamonds correspond to the centered ('on-axis') electron beam. Green triangles correspond to the electron beam being offset by +1.5 mm vertically with respect to the antiproton beam. The electron beam current was 100 mA, the antiproton transverse emittance 2-6 π mm mrad.

Comparison with a non-magnetized cooling force model

An attempt was made to compare our data to the force model where the electrons are not magnetized [5], which should best describe the kinematics of the electron-antiproton scattering in our cooler. This model, in its most convenient format, gives the functional dependence of the cooling force with the antiproton momentum deviation from its equilibrium value.

Employing the same equations with the same assumptions as in Ref. [6] (i.e. Maxwellian velocity distribution for the electrons, zero antiproton emittance), we fit our data with this model, initially using three independent parameters that characterize the electron beam. Fits are plotted on Figure 4 (dashed lines). The three fitting parameters are the lab frame RMS energy spread, δE, RMS angular spread, θ_e, and the local current density, J_{CS}. Results of the fitting procedure on the fitting parameters are summarized in Table 3 for three data sets (corresponding to three electron beam offsets). It should be noted that although all three parameters were used in the fitting procedure for the "on-axis" data set, once δE was determined, it was kept unchanged for the fitting of the "offset" data sets. Although the quality of the fits is rather good, J_{CS}, θ_e and δE found in this manner over

estimate by a factor of 1.2 to 2 our best estimates of what these should be (on-axis case).

Table 3: Fitting parameters results. Other parameters for the model are: I_b = 100 mA, Coulomb log is 10.

	Electron beam vertical offset, mm		
	0	1.5	2
J_{CS}, A cm^{-2}	1.2	0.7	0.3
θ_e, mrad	0.19	0.25	0.25
δE, eV	370	370	370

ELECTRON COOLING IN OPERATION

Electron cooling is used routinely for storing and cooling of antiprotons in the Recycler ring. At the early stages of the electron cooling project, the cooling scenario [7] was to continuously cool with the electron beam the main antiproton stack while cooling freshly injected batches separately with a gated stochastic cooling system before merging them to the main bunch. Although this scenario may reappear in the near future, this has not been the cooling mode of operation so far.

Cooling procedure

Currently, gated stochastic cooling is not being implemented and electron cooling is only used *when needed*, allowing for periods of beam studies. A typical evolution of the longitudinal emittance between transfers/extractions is plotted on Figure 5 (blue crosses).
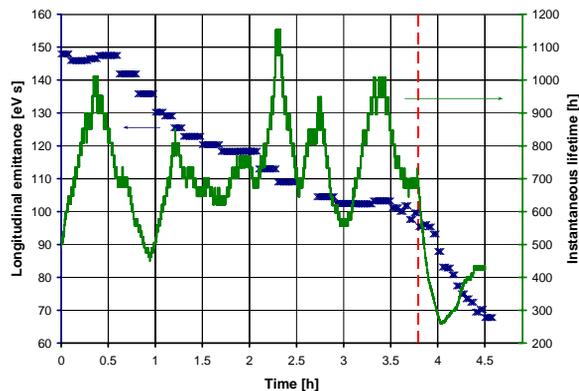

Figure 5: Example of the longitudinal emittance and beam lifetime evolution in between transfers and/or extractions. The blue crosses are the longitudinal emittance as measured by a Shottcky detector. The green continuous curve is the instantaneous lifetime ($-N \times (dN/dt)^{-1}$) where N is the number of antiprotons). The vertical dashed red line indicates when the electron beam was brought in closer to the axis i.e. effective electron cooling began. The electron beam was 100 mA, the number of antiprotons 195×10^{10}.

At the beginning, only the stochastic cooling system is in-use (both for the longitudinal and transverse dimensions) and the antiproton bunch occupies almost the entire ring (longitudinally). The electron beam, though on, is shifted away from the antiproton beam. To start

effective cooling with the electron beam, it is brought closer to the axis (parallel shift), thus overlapping the antiproton beam. The cooling strength is then adjusted by varying the offset of the electron beam with respect to the antiproton beam. The rate at which the beam is brought in is dictated by various variables (longitudinal emittance reduction needed, synchronization with the rest of the accelerator chain, lifetime, transverse emittance…) but, typically, it takes an hour to reduce the longitudinal emittance from 100 to 60 eV s, during which the beam is moved from a 5 mm offset to the axis in somewhat linear steps. As pointed out by the cooling force measurements, although the maximum cooling rate varies significantly from 0 to 100 mA, it remains practically unchanged for higher currents, up to 500 mA. Hence, for cooling purpose, the nominal electron beam current is 100 mA. Note that it is not the cooling strength that limits the total cooling time and cooling rates of ~80 eV s/h were achieved for stacks > 200×10^{10}.

The Recycler goals and typical performance are summarized in Table 4.

Table 4: Recycler performance main parameters

	Goal	Typical	Best
Number of stored antiprotons (before extraction), $\times 10^{10}$	500	250	410
Number of bunches extracted	36	36	36
Longitudinal emittance (per bunch), eV s	1.5	1.8	1.6
Transverse emittance (n, 95%), π mm mrad	2-3	4-6	2.5
Extraction efficiency, %	95	90-95	>95

The impact of electron cooling on the Recycler performance is mostly to allow for the storage of larger stacks while keeping the longitudinal emittance small enough for efficient extraction. Note that in the 'best' column, the numbers reported were not achieved concurrently but that the number of extracted bunches is almost always 36 (except for rare exceptions).

Operational issues

Since we started to use the electron beam for cooling, we dealt with three main problems: lifetime degradation of the antiprotons, fast beam loss and transverse emittance growth.

Lifetime degradation is illustrated on Figure 5 where the antiproton lifetime is plotted along with the longitudinal emittance. At first, the longitudinal emittance decreases steadily and the lifetime remains >500 hours (oscillations between 500 and 1000 hours are the result of details of the Recycler operation and the instrumental resolution). When electron cooling is 'turned on', and the cooling rate increases, the lifetime quickly drops. We have not identified which mechanism would lead to this lifetime degradation yet. However, a beam-beam quadrupole instability [8] might be a good candidate.

The beam loss observed for high phase space densities achieved with electron cooling has previously

been attributed to resistive wall dipole instabilities [9] and was almost entirely suppressed by adding a wide band damper [10] system. However, for very large stacks ($>400 \times 10^{10}$), during the 'mining' process [11], the phase space density increases beyond the current damping capability of the damper. Therefore, it has recently been upgraded and is being commissioned. It is expected to increase the instability threshold limit by a factor of 4.

Transverse emittance growth was observed during the mining process where both the physical (i.e. peak current) and phase space density become relatively large. The damper could not suppress it implying that the growth mechanism was of a different nature than the dipole instability due to resistive wall. As for lifetime degradation, a perturbation theory approach of the coherent beam-beam effect in the presence of the electron drift motion in the cooling section shows that a quadrupole instability may develop and lead to emittance growth [8]. In this theory, the growth rates are strongly determined by the machine transverse coupling, as well as the electron and antiproton beam currents but can be mitigated by separating the horizontal and vertical tunes. Thus, it was proposed to change the Recycler operating point in the tune space to allow for larger tune separation. Figure 6 shows the evolution of the transverse emittance after the mining process for the nominal tunes (A), where the tunes separation is 0.004, and for the proposed new tunes (B), where the tunes separation is 0.017.
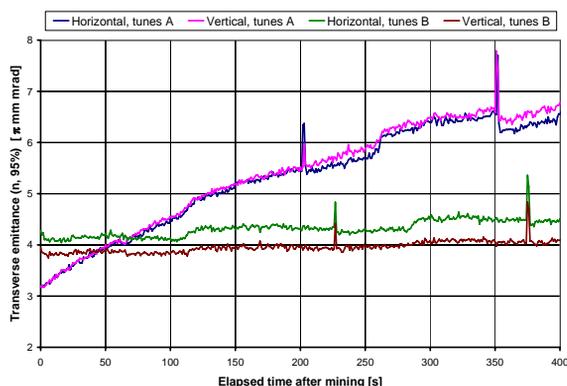

Figure 6: Evolution of the transverse emittance after mining for two operating points. Tunes A (blue and pink curves): 25.414 horizontal, 24.418 vertical. Tunes B (green and brown curves): 25.451 horizontal, 24.468 vertical. In both cases, the number of antiprotons was ~330 × 10^{10}, the bunch length 6.1 μs, the longitudinal emittance ~75 eV s, the electron beam was on-axis and the electron beam current 100 mA. The stochastic cooling systems were turned off prior to mining. Spikes are artifacts that correspond to the time at which bunches are extracted to the Tevatron (4 at a time).

The initial growth rate, defined as the slope of a linear fit over the first 5 minutes, is 36 π mm mrad per hour in the first case and 3 π mm mrad per hour in the second, a factor of ~10 different, consistent with predictions from the quadrupole instability theory. Note that because the transverse stochastic cooling system is kept off during that time, a ~0.5 π mm mrad per hour increase of the emittance is expected in our machine due to diffusion processes.

So, although not completely eliminated, the transverse emittance growth has been reduced to an acceptable level. In addition, it appears that the 'new' tune settings improved the antiprotons lifetime but a definitive conclusion awaits dedicated studies.

CONCLUSION

Fermilab has a unique electron cooling system routinely used for cooling 8.9 GeV/c antiprotons in the Recycler ring. Cooling of stacks up to 400 × 10^{10} to longitudinal emittances of 60-70 eV s was successful.

Lifetime degradation is mitigated by a progressive cooling procedure whereas emittance growth during the mining process was almost completely eliminated by reducing the machine transverse coupling.

The cooling force measurements at 100 mA were compared to a non-magnetized model, for which the electron beam parameters agree with our current estimates to within a factor of 2 or better.

ACKNOWLEDGMENTS

REFERENCES

[1]   A. Shemyakin *et al.*, Nucl. Instr. Methods Phys. Res. A **532** (2004) 403

[2]   A. Burov *et al.*, AIP Conf. Proc. **Vol. 821** (2006) 159

[3]   A. Burov *et al.*, AIP Conf. Proc. **Vol. 821** (2006) 139

[4]   H. Danared *et al.*, Phys. Rev. Lett. **72** (1994) 3775

[5]   Ya.S. Derbenev and A.N. Skrinsky, Particle Accelerators **8** (1977) 1

[6]   S. Nagaitsev *et al.*, Phys. Rev. Lett. **96**, 044801 (2006)

[7]   A. Burov, Nucl. Instr. Methods Phys. Res. A **532** (2004) 291

[8]   A. Burov *et al.*, *these proceedings*

[9]   Laslett, K. Neil and A. Sessler, Rev. Sci. Instr. **36** (1965) 436

[10] A. Burov, FERMILAB-TM-2336-AD (2005)]

[11] 'Mining' is the first step in the extraction process that consists of capturing the core of the antiproton beam in 9 individual square buckets, leaving out the high momentum particles

D. Beam cooling and intra-beam scattering