Journal of Physics: Conference Series **299** (2011) 012015 doi:10.1088/1742-6596/299/1/012015

CEBAF Accelerator Achievements

Y. C. Chao, M. Drury, C. Hovater, A. Hutton, G. A. Krafft, M. Poelker, C. Reece, and M. Tiefenback

Jefferson Lab, Newport News, Virginia 23606, USA

Abstract. In the past decade, nuclear physics users of Jefferson Lab's Continuous Electron Beam Accelerator Facility (CEBAF) have benefited from accelerator physics advances and machine improvements. As of early 2011, CEBAF operates routinely at 6 GeV, with a 12 GeV upgrade underway. This article reports highlights of CEBAF's scientific and technological evolution in the areas of cryomodule refurbishment, RF control, polarized source development, beam transport for parity experiments, magnets and hysteresis handling, beam breakup, and helium refrigerator operational optimization.

1. Overview

Reference [1] describes Jefferson Lab's Continuous Electron Beam Accelerator Facility (CEBAF) circa 2000, when operation for experiments was well established, but when much remained to be learned about improving and optimizing the accelerator. As of early 2011, CEBAF operates routinely at 6 GeV, with substantially improved reliability, and with a 12 GeV upgrade underway [2, 3, 4]. The following sections report highlights of CEBAF's scientific and technological evolution over the past decade, beginning with the cryomodule refurbishment program (Section 2) that has enable robust five-pass CEBAF operation at 6 GeV, 1.5 times the original design for final energy. A new radiofrequency (RF) control system (Section 3) has been implemented using a digital feedback system that eliminated the components most susceptible to temperature-induced phase drifts and that greatly improved injector stability. A polarized-source development process (Section 4), complemented by the optimization of beam transport for parity experiments (Section 5), has enabled the parity-violation and polarization-transfer experiments that are now yielding essential new information on nucleon quark-gluon structure. Section 6, on magnets and hysteresis handling, describes the control mechanisms that have been devised to satisfy reproducibility and stability requirements of CEBAF magnets. Section 7 reports improvements in beam breakup simulations that have enabled confirmation of the higher-order-mode (HOM) damping specification for the 12 GeV upgrade. Section 8 describes the Ganni Cycle, an optimizing reconfiguration of CEBAF's helium refrigeration system that nearly doubles the lifetime of some refrigerator components and that saves more than $1000 dollars/day in power costs.

2. Cryomodule refurbishment project

Jefferson Lab has recently completed the C50 refurbishment project, so called because it sought to increase the peak effective voltage of the refurbished cryomodules to a maximum of 50 MV. (The original specification was 20 MV, although several provide up to 35 MV.) The ultimate goal was to enable robust five-pass CEBAF operation at 6 GeV, 1.5 times the original requirement.

Ten of the weakest of the accelerator's 42 cryomodules were refurbished. This activity included complete disassembly to enable the reprocessing of the SRF cavities with improved protocols (Figure 1). The chief objective was to extend the operating range of each cavity from as low as 4 MV/m to an average of 12.5 MV/m, the performance needed to yield 50 MV for a cryomodule.

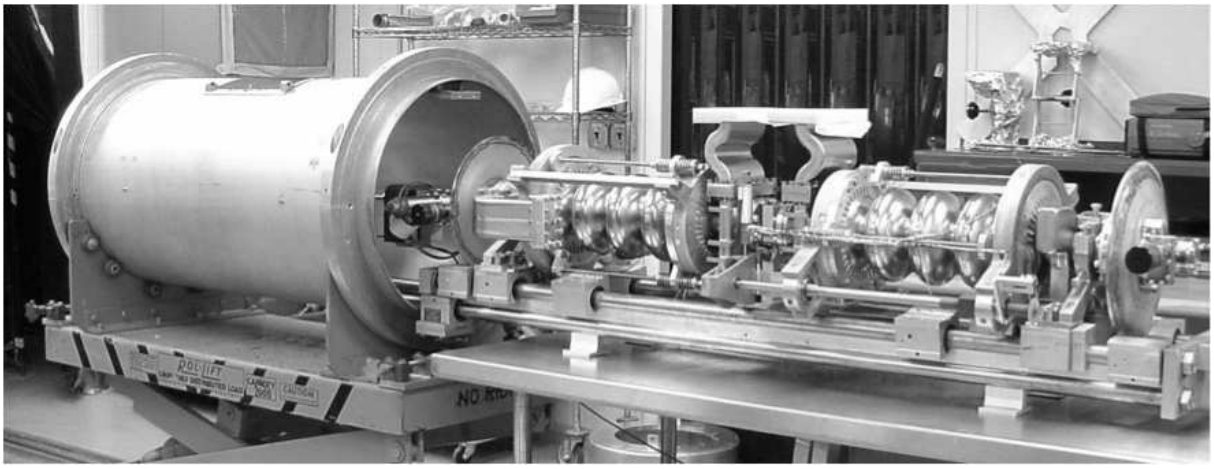

Figure 1. Reassembly of a reprocessed C50 cavity pair into a quarter-cryomodule.

While two of the cryomodules were selected for rework due to vacuum leaks, the principal performance-limiting phenomenon addressed was parasitic electron field emission due to residual particulate contamination on cavity surfaces. Such field emission produces two ill effects: increased cryogenic load and electrostatic charging of the ceramic RF input window which sits at 2 K in the original cryomodule design. This charging in turn induces the occurrence of periodic arc discharges that require both RF and beam interruption. As a result, the available gradient in many cases is governed not by cavity limitations but by arc rate management. Operationally, it is the sum of these trips over hundreds of cavities that has constrained the maximum useful energy of CEBAF. Elimination of the arc rate limitation was a C50 goal.

Two measures were taken in the C50 project to address this issue. First, to eliminate almost all of the offending particulate contamination, the niobium SRF cavities each received a fresh acid etch followed by a high-pressure rinse with ultra-pure water, both within a cleanroom environment. Neither of these procedures was available during the initial construction of CEBAF. Second, the RF window and waveguide attached to each cavity were replaced with a new "dog-leg" configuration that significantly reduces the risk of charging and trips, in the event that field emission is yet present. These measures have proven to be quite effective. For almost all of the cavities, field-emission-induced radiation was eliminated in the operating range of interest. Cavity performance, if limited, was constrained only by quench conditions derived from original fabrication flaws. While window arc-related trips do occasionally occur, they are greatly reduced and occur at much higher operational conditions. None of the C50 cavities have gradient limitations imposed by arc rate management.

During refurbishment prior to receiving their fresh etch, the cavities were subjected to a 600°C bake in a vacuum furnace. Such a bake reduces the hydrogen content dissolved in the niobium bulk. When that concentration is too high, NbH will precipitate on the surface at temperatures around 100 K, given sufficient opportunity. NbH is a normal conducting metal, so its presence can dramatically increase cryogenic losses in what has come to be known as Q-disease. During CEBAF construction this phenomenon was only just being identified and no remedial vacuum furnace was available. To manage the situation, it became necessary to enforce rather fast cooldowns for the cryomodules, especially through the vulnerable hydride

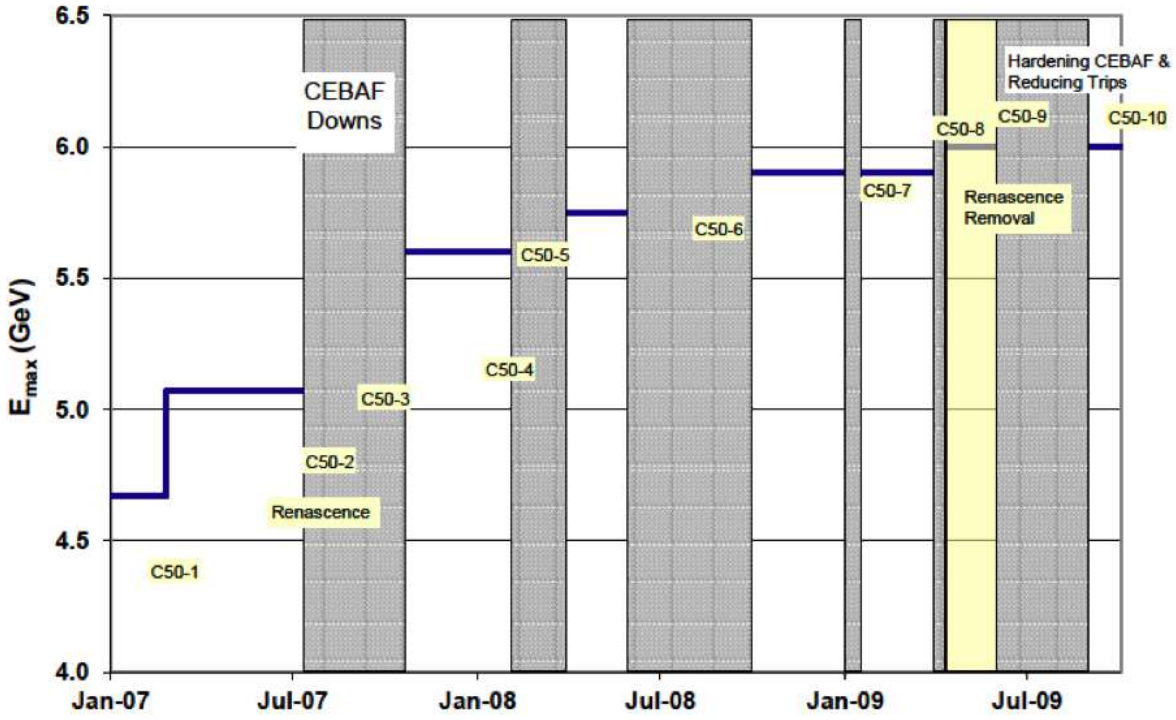

Figure 2. Staggered C50 cryomodule reinstallation and CEBAF's associated performance rise.

precipitation range. With the C50 cavities now baked, we can allow for slower cryomodule cooldown. This in turn reduces the likelihood of mechanical damage, such as the creation of vacuum leaks, during thermal cycles.

Other C50 improvements included replacement of the external polyethylene RF windows with a new ceramic model and modifications to the mechanical cavity tuners to prevent backlash. Damaged or worn parts were also replaced.

The C50 cryomodules were refurbished at a rate of three per year. To maximally support CEBAF operation during this time, cryomodules were rotated in and out during machine downs, all the while maintaining a rough balance between the voltage capacity of the north and south linacs. Figure 2 shows the progression of C50 cryomodule reinstallation and the associated rise in CEBAF's maximum five-pass final energy for experiments. All ten refurbished cryomodules are now once again fully operational in CEBAF.

The raw cavity performance after reprocessing was excellent. Field emission was dramatically improved and the superconducting RF dissipation was essentially textbook performance. After incorporation into finished cryomodules, however, the results were mixed. Of the eighty reprocessed cavities, 75% met or exceeded the goal of 12.5 MV/m. Under operational conditions, the C50 cryomodules have been able to sustain an average voltage of 47 MV, with three of them meeting or exceeding the 50 MV goal. Unfortunately, the installed C50 cavities were unable to meet the improved Q specification. During cavity vertical tests, the reprocessed cavities routinely demonstrated Q's of 1×10^{10} or higher. However, assembly into a cryomodule resulted in significant Q degradation. The RF dissipation losses increased by roughly a factor of two, essentially returning to their previous operational values, less the contribution from field emission. The degradation is now believed to be due to unrecognized magnetized components built into the original cryomodules, such as the helium vessels or the tuners. (The presence of magnetic field during cooldown creates "frozen-in" fluxoids which produce dissipation in

response to RF fields.) The low Q's had been previously attributed to Q-disease. Unfortunately, due to schedule pressure, investigations into the specific causes along with possible fixes were not completed before the end of the C50 project. C50 cryomodules operating at or near 50 MV are generating heat loads for the 2 K cryogenic system in excess of 150 W. Thus, while the C50 project has been quite successful in enabling robust 6 GeV operation, the cost in terms of refrigeration load has been much higher than expected.

3. RF control

A new RF control system has been implemented using a digital feedback system that eliminated the components most susceptible to temperature-induced phase drifts. The new design directly digitizes the intermediate frequency, eliminating all of the analog signal processing. The new system incorporates a field programmable gate array for all of the digital signal processing. The feedback controller uses a proportional-integral controller to maintain cavity field, removing the finite gain offset that was seen in the analog proportional controller. The new controller reduced the phase drift susceptibility from 1°/°C to 0.1°/°C. Starting with the photogun drive laser and up through the buncher cavity, the injector RF was converted to the new design. The new RF control system greatly improved the stability of the injector.

4. Polarized source development

Approximately 80% of the experiments conducted at CEBAF require spin-polarized electron beams, and even those experiments that do not require polarized beam receive it, since all electrons originate from a common photogun. Today's CEBAF polarized source is a load-locked gun (more below) that provides 85% polarization and up to 180 μA delivered to a single hall, with combined current to all three halls in excess of 200 μA. The operating lifetime of the polarized source is limited primarily by the photogun vacuum level, which is difficult to measure because it is so low (low 10^{-12} Torr) and commercial vacuum gauges are not reliable in this pressure range. The photogun can deliver ∼500 C before steps must be taken to restore photocathode yield, or QE. This represents approximately six weeks of beam delivery without interruption under high-current operation, and longer time periods at lower current.

The CEBAF/Jefferson Lab load-locked gun is shown in Figure 3 [5]. It consists of four vacuum chambers: the high-voltage chamber, the photocathode preparation chamber, a "suitcase" chamber used for replacing photocathode samples, and an intermediate chamber that must be evacuated and baked each time the suitcase is attached. The suitcase is normally detached from the photogun and stored elsewhere. This approach helps to reduce the overall footprint of the photogun when in operation. Desirable features incorporated into this design include the ability to store multiple photocathode samples, reliable sample transport from one chamber to the next without dropping, and rapid heating and cooling of samples for fast turn-around at activation.

With a load-locked gun, new photocathode samples are introduced via the suitcase, and then moved from one chamber to the next, with the best vacuum obtained in the high-voltage chamber. In this manner, lengthy bakeouts of the entire photogun are avoided. In addition, since photocathode activation takes place inside a dedicated preparation chamber, cesium is not inadvertently applied to the gun high-voltage electrode as was the case with the previous CEBAF photogun design. Cesium reduces the work function of materials and leads to catastrophic field emission from the cathode electrode. Another benefit of the load-locked photogun is the use of a compact tapered insulator that extends into the gun high-voltage vacuum chamber, which greatly reduces the amount of metal biased at high voltage, and consequently minimizes the likelihood of field emission.

Today's state-of-the-art high-polarization photocathode is the GaAs/GaAsP strained-superlattice structure [6], which consists of a very thin GaAs surface layer (∼5 nm) grown atop

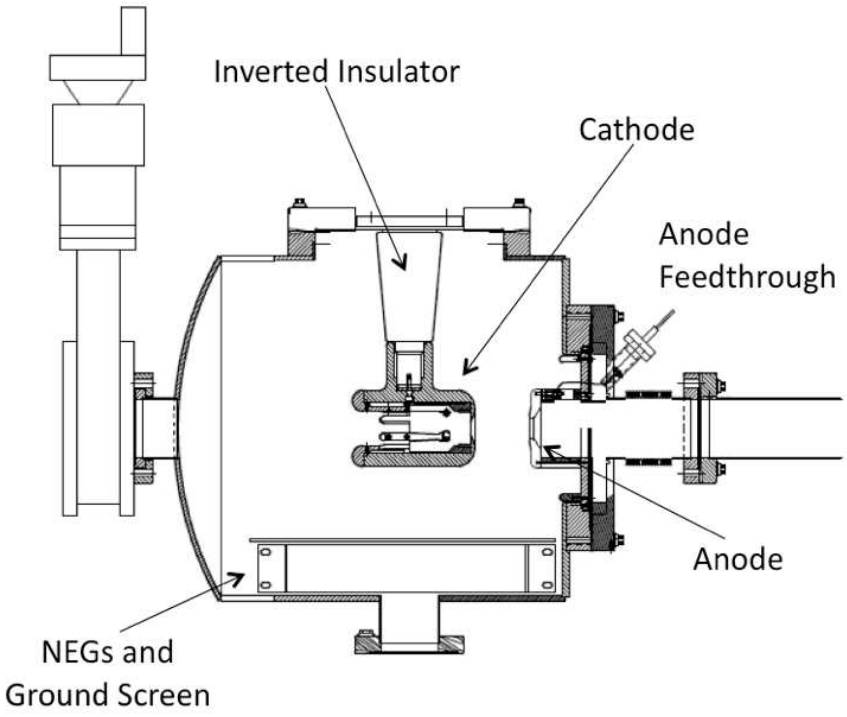

Figure 3. The high-voltage vacuum chamber of the latest polarized electron source, with compact inverted geometry high-voltage insulator.

10 to 20 pairs of thin, alternating layers of GaAsP and GaAs. By growing very thin GaAs layers, the strain can be maintained, which improves polarization, and electrons in subsurface layers efficiently tunnel through the GaAsP layers. And by using many thin layers of GaAs/GaAsP, the QE can be considerably higher than obtained from a single (thicker) layer of strained GaAs. The net result is polarization ∼85% and QE ∼1%.

Originally, we expected to use a DC laser to generate a DC electron beam which would be given the appropriate RF structure required for acceleration using the photoinjector RF chopper system. However, this would have been very inefficient, with most of the extracted electrons dumped on a copper plate within the photoinjector RF chopper. Instead, CEBAF drive lasers emit short pulses of light, to directly produce an electron beam with the appropriate RF structure required for acceleration. This approach, termed *synchronous photoinjection*, has now been adopted by other accelerator facilities. CEBAF drive lasers rely on the accelerator-friendly technique gain-switching to generate RF-pulsed laser light, and exploit recent advances from the fiber-based telecommunications industry [7]. Light from a fiber-coupled, gain-switched diode laser is delivered to an ErYb-doped fiber amplifier to obtain ∼5 W average power at 1560 nm. This light is converted to a useful wavelength via second harmonic generation to obtain over 2 W average power at 780 nm with ∼40 ps pulses and pulse repetition rate of 499 MHz.

4.1. Spin manipulation and polarimetry

Polarized-beam experiments require a specific orientation of the electron spin direction at the target, typically parallel to the direction of beam motion. Moreover, the spin direction must

flip sign at some specified frequency. Spin flipping is accomplished by reversing the polarity of voltage applied to an electro-optical element called a Pockels cell, which is located on the drive laser table at the photoinjector. Typically, the polarization direction of the electron beam flips at 30 Hz, and more recently a technique was developed to flip polarization much faster, up to ∼1 kHz. Fast flipping was implemented to help reduce the effects of target boiling for the Qweak experiment and has significantly improved the experimenters' ability to make a precise measurement of an extremely small physics asymmetry.

Electrons leave the photocathode with spin direction pointing parallel/antiparallel to the direction of beam motion, depending on the helicity of the laser circular polarization (right or left circular) created by the Pockels cell. But the spin direction precesses in the horizontal plane as the beam passes through the arcs and transport lines to the halls, and this net spin precession must be cancelled out by orienting the spin direction at the injector by the opposite amount using a spin manipulator.

The original Z-style spin manipulator was overly complicated, and quickly replaced with a Wien filter spin manipulator [8]. A Wien filter is a device with static electric and magnetic fields perpendicular to each other and to the velocity of charged particles passing through it. Unit charged particles with a velocity of $\beta c = E/B$ are undeflected in passing through the Wien filter, while the spin is rotated in the plane of the electric field. Most recently, a second Wien filter was added to the injector beamline, but with field forces oriented in the vertical plane. Two Wien filters and intervening solenoids are used to flip the sign of the beam polarization weekly, to assist parity violation experiments cancel out unwanted helicity-correlated spot size variations.

There are five polarimeters at CEBAF: one Compton polarimeter, three Møller polarimeters, and the 5 MeV Mott polarimeter at the CEBAF photoinjector. The term *spin dance* was coined to describe the process of sending identical beam to multiple polarimeters, while using the photoinjector Wien-filter spin manipulator to vary the direction of the electron spin. Each polarimeter receives polarization with the same magnitude, but the portion of the polarization pointing in the longitudinal direction follows a sinusoidal trend as the spin direction is changed using the Wien filter (see Fig. 4). The spin dance provided the first precise comparison of the analyzing powers of Compton, Møller, and Mott polarimeters [9]. It revealed disagreements among the polarimeters that were associated with errors in the analyzing power of some polarimeters and an incomplete appreciation for all systematic effects.

To understand the different curves in Figure 4, it must be noted that beams delivered to the three experimental halls follow different paths after extraction from the accelerator, and are often derived from different numbers of recirculation passes, so the total spin precession from the polarized source to each experimental hall is in general quite different. The Wien angle that provides longitudinal polarization at one hall does not necessarily provide longitudinal polarization at neighboring halls. In fact, there are approximately 400 two-hall energy combinations that give maximum polarization in two halls [10] and only two beam energies (2.1 GeV and 4.2 GeV) can provide maximum longitudinal polarization to all three halls simultaneously. Recently it was realized that by redistributing the energy gain of the two linacs, while keeping total energy gain constant, there is significantly greater flexibility in maximizing the degree of longitudinal polarization at multiple halls [11].

Finally, it's worth mentioning another beneficial outcome of the spin dance, namely, evaluation of the value of the spin orientation exiting the injector that provides maximum longitudinal polarization in each experimental hall leads to an independent and very precise (better than 10^{-4}) absolute measurement of the final electron beam energy.

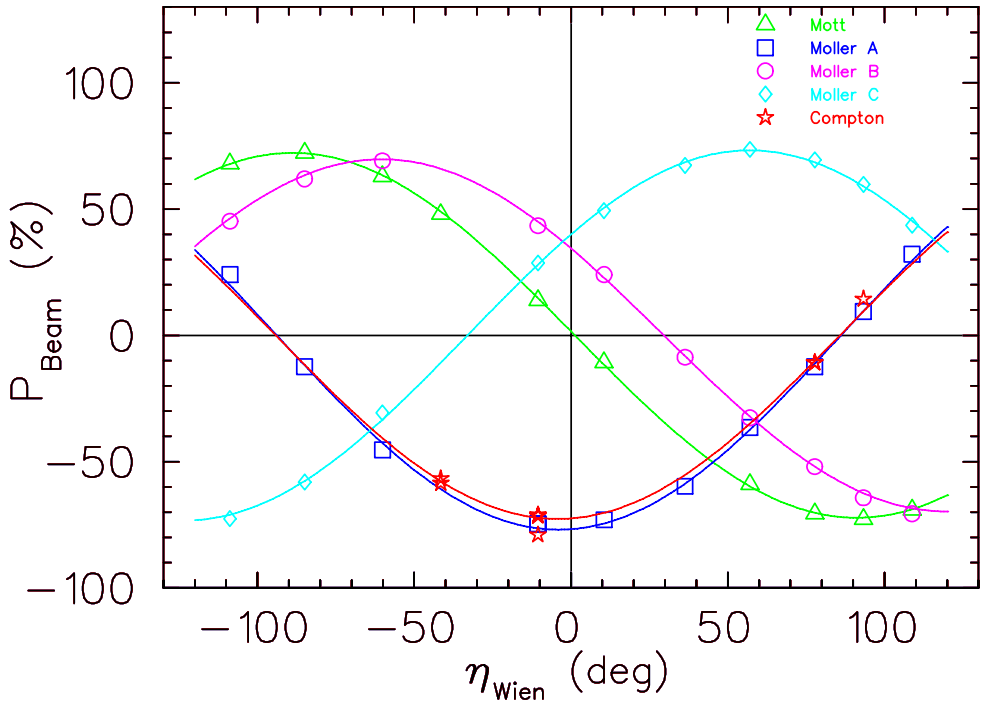

Figure 4. The results of a spin dance, where the longitudinal component of the beam polarization is measured as a function of Wien spin manipulator angle, at multiple polarimeters.

5. Optimizing injector-to-target beam transport for parity violation experiments

Here we discuss efforts to identify anomalous beam envelope mismatches (transport anomalies) across different energy sections of CEBAF and the correction techniques we developed that resulted in an unprecedented low level of helicity-correlated orbit variation, partly contributing to the first Jefferson Lab parity violation experiment to achieve precision better than 100 parts per billion (ppb).

An important and demanding class of experiments conducted at CEBAF is illistrated by the parity violation (PV) experiment [12] devoted to the study of extremely small asymmetries in elastic electron-nucleon scattering generated by electromagnetic and weak interactions. The electron beam is polarized in either positive or negative handedness (helicity state), and the asymmetry, A_{PV}, is measured in terms of left- and right-handed helicity state cross sections, $\sigma_{L,R}$:

$$A_{PV} = (\sigma_R - \sigma_L)/(\sigma_R + \sigma_L) \approx Q^2/M_Z^2 \approx 10^{-6} - 10^{-4}, \quad (1)$$

where Q^2 is the 4-momentum transfer and M_Z the mass of the Z-boson. It is noted that the typical magnitude of such asymmetry is in the range of 1 to 100 parts per million (ppm). To confidently measure such a small asymmetry, the systematic variations of the electron beam must be controlled to a few percent of the asymmetry itself, or to the sub-100 ppb level. This translates into very tight specifications on beam quality, especially the position, angle or intensity of the beam at the target. These are referred to as helicity-correlated (HC) parameters. Table 1 gives an overview of the intrinsic physics asymmetries and requirements on HC beam properties of various PV experiments at Jefferson Lab. The tight constraints on the HC orbit translate into exacting demands on the performance of the accelerator, with over 6 km of beam transport and four decades of momentum gain from the injector to the target location. The typical beam quality measurements performed at CEBAF for non-parity experiments, while sufficient for loss-

Table 1. Helicity-correlated beam specifications for parity violation experiments at CEBAF.

Experiment	Physics Asym. (ppm)	Intensity (ppm)	Position on Target (nm)	Angle on Target (nRad)
HAPPEX-I	13	1.0	10	10
G0	2-50	1.0	20	2
HAPPEX-He	8	0.6	3	3
HAPPEX-II	1.3	0.6	2	2
QWeak	0.3	0.1	2	10
PREX	<1	0.1	1	1

free delivery on target and efficient event generation, are not enough for controlling HC orbit to the sub-100 ppb level. The following describes the mechanism of runaway HC orbit variation caused by transport anomaly and the methods implemented to combat it.

5.1. Helicity-correlated orbit and transport anomaly

Origin of Helicity-Correlated Orbits and Their Consequence — In simple terms, PV experiments rely on comparing scattering cross sections between electron bunch trains in positive and negative helicity states, which in turn are alternatingly generated at 30 Hz from the CEBAF photoinjector through a switching system in the laser path before the laser impinges on the photocathode to produce electron beam. The switching system, an electro-optic element called a Pockels cell, unavoidably introduces systematic steering on the laser spot that develops into fluctuation in the electron beam orbit correlated to the helicity state. The true asymmetry A_{PV} is extracted from raw experimental data by applying experimentally determined correction factors. The measurement of correction factors, however, can be contaminated by HC orbit variations, resulting in deficient correction, and thus compromised precision in A_{PV}.

Consequence of Transport Anomaly — In an ideal world, the phenomenon known as adiabatic damping would ensure reduction of the HC orbit variation by a factor of almost 100 from the injector at 100 keV to 3 GeV, a typical energy for parity experiments. In reality, however, deviations from ideal beam transport can not only prevent the damping from occuring, but can, in extreme cases, actually make the HC orbit grow. This transport anomaly can potentially have detrimental effects on measuring the asymmetry A_{PV}. Transport anomaly can be categorized into cross-plane coupling (coupling) and near-singular transport due to grossly mismatched optics (mismatch). These two effects can compound each other, with consequences not possible by either one alone. Despite having detrimental effects on measuring the asymmetry, their small absolute magnitude may be barely discernible in terms of beam loss or beam size on target, and thus they are very difficult to detect or track. A more subtle effect is the evolution of the mismatch between the beam ellipse and beam orbit (such as the HC orbit itself), even if they start in full congruence [13]. This means in a situation with compounded mismatch and coupling, we can quickly lose the ability to simultaneously correct the beam shape and the orbit as they fall out of coherence.

Methods for Eliminating Transport Anomaly — In 2004 the decision was made to characterize the transport throughout the five-pass CEBAF. A series of beam-based measurements followed. The technique relied on making intentional helicity-correlated orbits by moving the laser beam at the photocathode using a mirror attached to a piezoelectric transducer (PZT). Analysis of the outcome soon pointed to the low-energy front end as the main source of strong beam phase-space blowup. This was primarily attributed to near singular transport across the SRF cavities (cryounit and cryomodules) because effective modeling of beam transport through these cavities

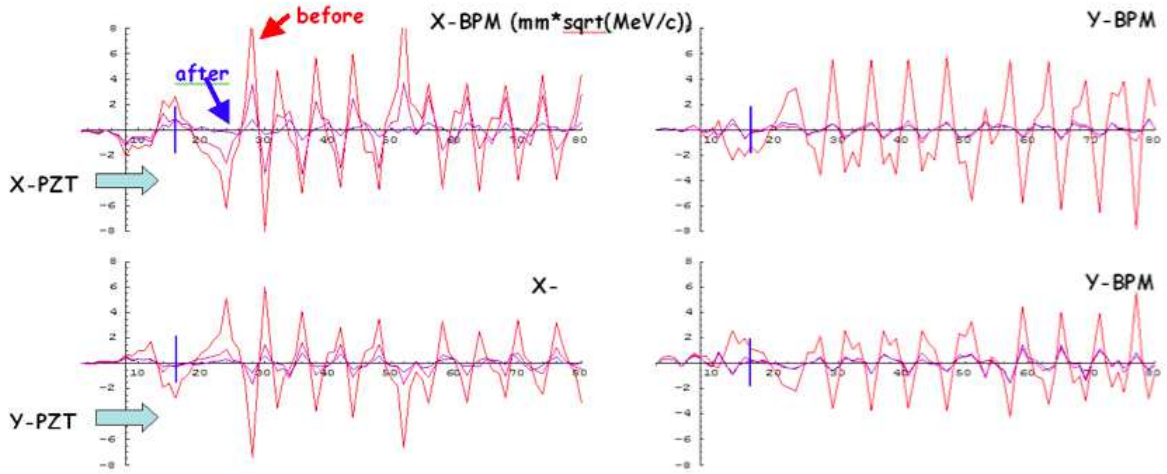

Figure 5. The PZT response before and after injector matching from 100 keV to 800 MeV.

at such low energy was very difficult. This was likely exacerbated by coupling present in the HOM couplers of these cavities, leading to phase-space blowup described above. Coincidence between jumps in PZT amplitude and locations of the cryo-cavities was unmistakable. Similar blowup signatures were also observed in the area between the end of the injector (60 MeV) and the first two linacs, although at a much lower level. With the above observations, a strategy was formulated in two stages to address transport anomaly and reclaim expected adiabatic damping: first, for 100 keV to 60 MeV, with accurately measured transport and sufficient correction elements, a model-based solution is possible, and second, for 60 MeV to 3 GeV, lacking an accurate long-range model, an empirical approach appears more practical.

100 keV to 60 MeV — Four-dimensional (X, X', Y, Y') transfer matrices were measured using difference-orbit methods across all injector cryo-cavities. The result displayed high degrees of near-singularity in 4D transport, confirming initial suspicion. Furthermore all measured transfer matrices satisfy the linear 4D symplectic condition, implying that proper transport can be restored with linear elements only. An optimization program was developed to use existing quadrupoles and skew quadrupoles to eliminate XY coupling and minimize singularity from 100 keV to 60 MeV while complying with multiple constraints dictated by operation and hardware considerations. The solution, once implemented, showed significant improvement in all transport characteristics. The solution eliminated XY coupling (off-diagonal determinant as percent of total), and brought singularity (SVD condition number) in the original transport down to near-ideal level.

60 MeV to 3 GeV — The empirical approach adopted for correcting mismatch from 60 MeV to 3 GeV is betatron matching under the guidance of the X and Y PZT signals, as they define a subspace in the 4D phase space shown to be representative of the characteristics of true HC orbits. This has proved effective except when PZT orbits are not sufficiently congruent with the beam spot, in which case a compromise must be empirically reached.

Outcome of Injector Matching — Figure 5 shows the BPM response, from 100 keV to 800 MeV, to the X and Y PZTs before and after anomalous transport in the Injector was corrected. Figure 6 shows similar signals at 3 GeV in experimental Hall A (HAPPEX). The HC orbit reduction from cathode to target was significant. This translated into unprecedented HC orbit reported by HAPPEX 2005 at 3 GeV and by G0 at a much lower energy of 650 MeV.

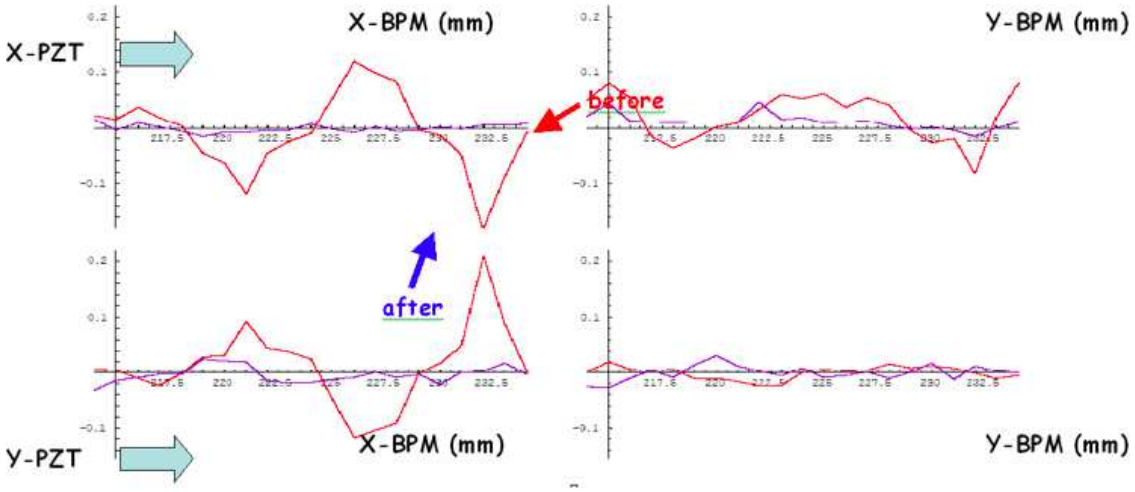

Figure 6. The PZT response before and after injector matching with 3 GeV into Hall A.

6. Magnets and hysteresis handling

The CEBAF accelerator magnets, all normal conducting, are typically static for extended periods, often for weeks with no change of set point. Most magnets in the system have iron flux cores, reducing the total amp-turn requirements to establish the desired magnetic field. For such magnets, providing the desired magnetic field requires control not only of the current in the magnet at the time of interest, but also of the history of the current. This section describes the control mechanisms which have been devised to satisfy reproducibility and stability requirements of CEBAF magnets. It also describes some vulnerabilities of earlier protocols and provisions which have been made to identify soft failures in the magnet power supplies. We do not attempt here to discuss issues such as ripple control, but specialize to hysteresis handling and fault detection.

6.1. Operational philosophy

The magnets are set up during tune-up operations, and hysteresis cycled so that the magnetic configuration can be reliably restored at later times. Stability and reproducibility of the magnetic fields are the principal concerns for magnet operation. The accelerator changes subtly over time scales of a week due to such effects as ground motion driven changes in the machine circumference and changes in the accelerating gradient profile along the linacs as various accelerating cavities are taken out of or restored to service. These influence the detailed optics of the accelerator, and are compensated at setup (and such other times as required) by manual tuning operations which establish certain global transfer function properties. The quadrupoles of the accelerator are almost all independently powered, with independent power supply calibration errors which change when the power supplies are replaced. The procedures used in accelerator operation are, of necessity, tolerant of these variables.

6.2. Magnet power supply characteristics

The dipole bending magnets are powered in groups on common power supplies. The subset of dipoles on a given power supply bus which were required to be independently adjustable were sized to provide slightly greater bending field integral at any given current than their companion magnets. Individual control was provided through a parallel electronic load, a controllable shunt to reduce the individual magnet current without affecting the bus current. The primary dipole

power supplies are physically large units called box supplies. The box supplies are controlled through an 18-bit DAC in the power supply with full-scale currents ranging from 270 to 640 A. The associated set-point granularity is at the mA level. This provides resolution comparable to the intrinsic energy spread of the beam, and is sufficient and convenient for operational purposes.

Almost all of the quadrupole and steering corrector magnets in the CEBAF accelerator are controlled by the trim magnet system. This system is implemented as 32-channel blocks of rack-mounted electronic loads fed from common bus sources at ± 35 V in the rack. Each of the channels provides ground-referenced power up to ± 30 V and (typically) up to ± 10 A. Specialized versions of the trim cards are used for low-current injector magnets (± 1 A) and in a few locations for marginally higher currents (± 12 A). The trim system uses 16-bit DACs with consequent current resolution of 0.3 mA for the usual ± 10 A full range. The sampled read-back current from the trim cards also has 16-bit resolution. While bench calibration of the trim cards is attempted only at the ± 5 mA level, the bit noise on the read-back current is used to trigger a noise counter which is very helpful in identifying marginal stability failures even before the card performance violates calibration thresholds.

6.3. Hysteresis handling

For many magnets in the injector and linacs, as well as a few in scattered service elsewhere, iron cores were not beneficial. These magnets have hysteresis handling disabled and are controlled by simply setting the coil current. However, for the iron-bearing magnets in the system, the history of the current must be controlled in order to ensure that the magnetic fields seen by the beam are as desired. Field integral versus current curves are shown in Figure 7 for unipolar-operated, 3 m length BA dipole magnets and in Figure 8 for bipolar-operated, 0.3 m length QA quadrupoles. The figures' left-hand sides show the field integrals versus current, and the right-hand sides show the nonlinear, hysteretic residue after subtracting the dominant linear contribution in the magnet current. (The current proportionality constant was selected for graphical purposes rather than for having any special physical meaning.) These curves illustrate the need to control the extremes of current during hysteresis cycling.

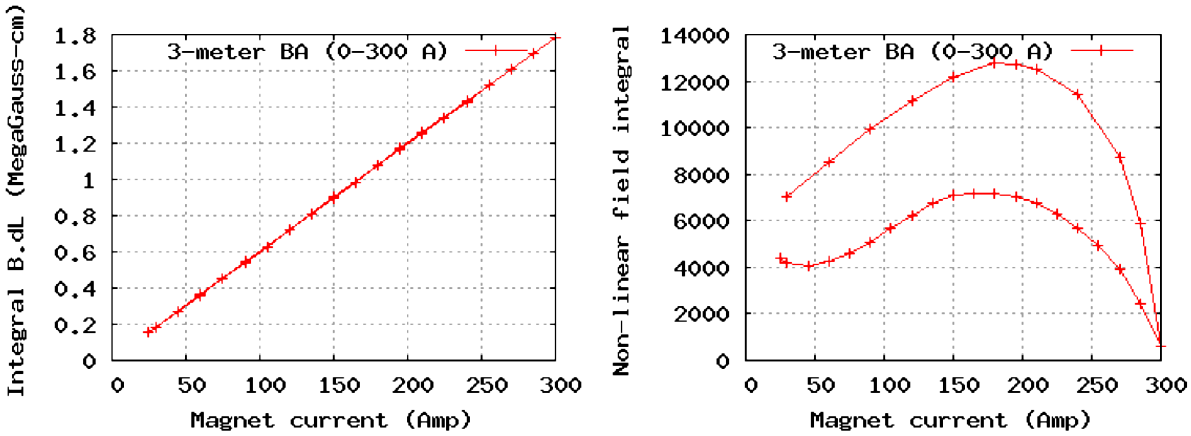

Figure 7. On the left, the field integral (Gauss-cm) through a 3 m dipole magnet as a function of current, and on the right, the field integral with dominant linear portion subtracted, showing saturation and hysteresis.

During commissioning in the early 1990s, the quadrupole magnet controls were configured to use the disconnected parts of the hysteresis curve for decreasing magnitude of current. Positive current set points were approached from the positive extreme current, negative set points

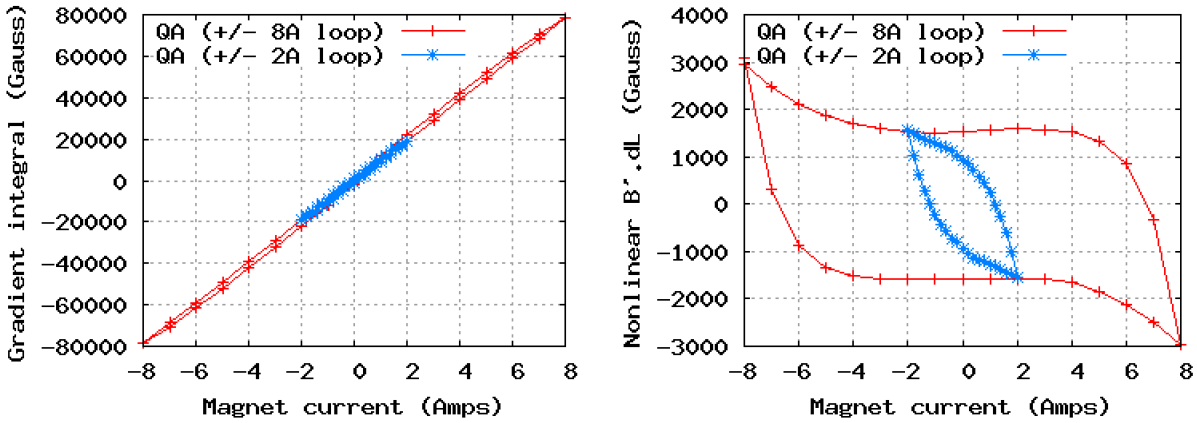

Figure 8. On the left, the gradient integral for a 30 cm QA quadrupole with ± 10 A and ± 3 A hysteresis cycles, and on the right, the nonlinear residue after subtracting the (same) linear component from each.

from the negative extreme. This had the side effect of rendering very low field integral set points unreachable "on-loop". Disabling hysteresis control for these magnets caused confusion concerning the magnetic field configuration. To avoid this difficulty, all bipolar magnets were redesignated to use the portion of the curve starting at the positive extreme and moving in the direction of algebraically decreasing current. This portion of the hysteresis curve remains in use at Jefferson Lab for the (bipolar) trim magnets.

The dipoles also initially used the field versus current relation from the upper part of the hysteresis curve, relaxing from the extreme current. Hysteresis cycling was done by 1) raising the bus power supply to peak current, 2) then shunting the desired current around the magnet, and 3) reducing the bus current to the desired operating point. As the beam energy was raised toward 6 GeV, some of the shunted magnets were more saturated than their partners, making supplementary power supplies necessary. Retaining the initial protocol would have resulted in adding supplementary current at the peak bus current, which would alter the hysteretic contribution to the magnetic field. Operating the magnets by setting them on the rising current curve was a better fit for a unified protocol using both shunts and supplementary power. The dipole magnets are now cycled from peak to minimum current. The supplementary power supplies and shunt controls are activated at the hysteresis minimum, and the bus current is then raised to the desired set point. The shunts, being fundamentally passive devices, will never reverse the magnet currents. They only begin shunting current after the bus voltage begins to rise, and each magnet experiences its intended peak and minimum values of current.

The nonzero lower hysteresis bound for the (unipolar) dipoles is driven by a hardware constraint: the power supplies trip off when run too close to zero. For the BA dipoles, for set points in the upward-concave region of the lower curve in Figure 7 (currents less than about 30 A), varying the lower extreme current alters the field from the expected value. During operation at very low beam energy, this has caused confusion over energy calibration at the parts per thousand level.

The rate of change of current during hysteresis cycling is also important. For the solid-core dipoles, there are eddy current effects which should be limited. We have settled on dI/dt values of approximately 5 A/sec, implemented in the power supply hardware, to provide reproducible fields. The maximum and minimum set points are issued by the control system, which waits for the power supply to achieve the set current before issuing the next current set point. The cycle from maximum current to minimum current is executed twice, and then with all shunt

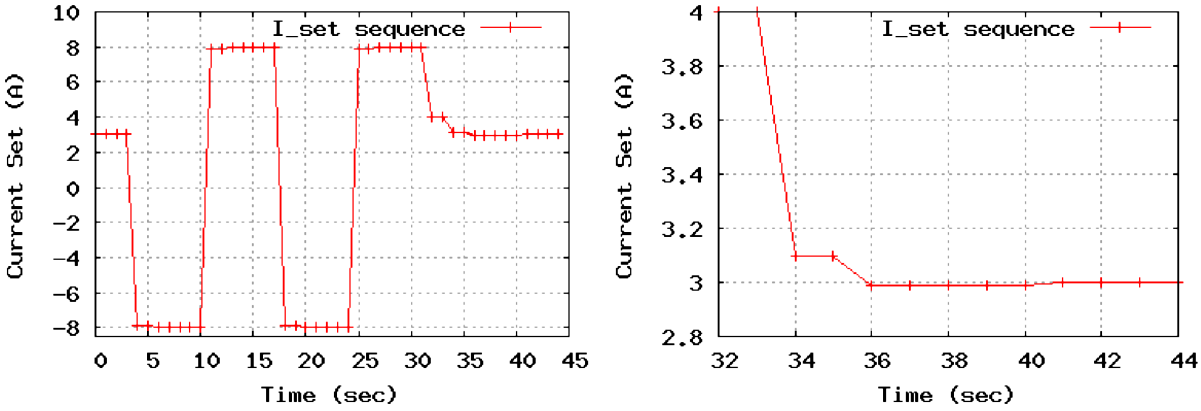

Figure 9. The left figure shows the time course of hysteresis for QA magnet (± 8 A) and the right figure shows the deliberate undershoot by 10 mA.

and supplementary power enabled, the ramp to final set point is done.

CEBAF quadrupoles are laminated, rather than solid, iron. While testing and improving our hysteresis protocols, we uncovered surprisingly large variability due to latency in execution on the local I/O controllers both in EPICS, the Experimental Physics and Industrial Control System (http://www.aps.anl.gov/epics/EPICS), and in the previous TACL systems. We obtain measurably more reproducible magnet behavior by relying upon internal power supply time constants to control fine-grained dI/dt, rather than by trying to specify the I(t) curve explicitly. The detailed set-point sequence now in use avoids vulnerability to latency effects and communication errors by issuing an arguably minimal number of set points, and avoids vulnerability to the variable characteristics of power supply and magnet combinations by controlling the size of the steps issued.

The time course specified for the trim magnet hysteresis protocol is shown in Figure 9. These illustrate the protocol applied to a QA magnet with ± 8 A current extremes. The current is taken from the negative extreme to the positive extreme twice, preventing the underdamped power supply from exceeding the ± 8 A bounds by using initial targets of ± 7.9 A. After the final positive extreme is attained, the current is stepped to 1 A above and then to 100 mA above the final current, and then to the final value. These steps are selected to avoid undershoot of the final current value and consequent variation in the remanent field contribution to the magnets.

The remanent field in the laminated iron quadrupoles relaxes by a small but detectable amount during the few hours after cycling. This relaxation shifts the field of all of the quadrupoles slightly more negative (vertically focusing) by about 50 Gauss of gradient integral. The 100 mA undershoot in current was selected to target this property of the QA quadrupoles. Reducing the current by 100 mA reduces the gradient integral by 100 Gauss, after which restoring the current to nominal set point restores only 50 Gauss. This tactic has been demonstrated to improve the stability of the large-scale optics of the CEBAF accelerator, rendering systematic drifts undetectable.

6.4. Fault detection

During hysteresis cycling, all magnet power supplies are verified to achieve each current set point to within the relevant calibration standards for each type. During operation, the measured currents are also required to remain within a calibration window. However, fluctuations of the magnet current even within the calibration window can cause unacceptable fluctuation of the beam position in the accelerator or at the experimental target. The 16-bit resolution of measured

current in the trim magnets is much finer than the calibration window (0.3 mA versus ± 5 mA for 10 A cards). A cumulative noise alarm counter was implemented to allow identification of magnets showing sample-to-sample variation in current. The noise threshold is selectable, but presently generates a count for steps of three bits. The noise alarm technique was also applied to the measured voltage, anticipating that magnet faults may occur even when the current is held constant. These counters have been instrumental in identifying incipient wiring faults, resistive connections, and card failures and preventing accelerator down time.

7. Beam breakup

In superconducting recirculating linacs, multipass beam breakup (BBU) instability can create a significant limitation on average current [14, 15]. In this instability, roughly equivalent to multibunch instabilities in storage rings, inadequately damped transverse deflecting higher-order modes (HOMs) are excited by beam displacement on subsequent beam passes causing the beam to become deflected transversely and eventually lost, should the HOM excitation become large enough. The instability grows exponentially in time as long as the beam average current exceeds a threshold current, given by the condition that the energy deposited by the displaced beam can no longer be dissipated by the usual HOM damping in the HOM absorbers.

The CEBAF accelerator was constructed with cavities developed for storage ring applications at Cornell University. These cavities had been designed for very high HOM damping to mitigate ring multibunch instability, and applying these cavities in CEBAF was natural, given the previous difficulties with multibunch BBU. Early in the CEBAF project, it was shown both theoretically [16] and experimentally [17] that the actual threshold current for multipass BBU was at least two orders of magnitude beyond the maximum operating current of the accelerator. It was known at the time that CEBAF was overdesigned in HOM damping. Because a recirculating linac at the CEBAF scale and cost had not been built previously, and because at the time the theoretical codes were not deemed to be very reliable, completely eliminating this potential source of trouble to the project was thought to be wise.

Because of possibilities afforded by the Cornell/CEBAF cavities, other projects were keen to apply these cavities to yield high average current. In particular, the series of high-average-current energy-recovering linacs built as free-electron laser drivers at Jefferson Lab at first used essentially the same cavity design, but later used cavity designs with worse HOM damping incorporated from the beginning as cost-saving design choices. These devices ran at beam average currents up to 10 mA, roughly two orders of magnitude above that at CEBAF. The free-electron lasers recirculated beam at lower energies than at CEBAF, enhancing the possibilities for instability. Multipass BBU instability was indeed predicted [18] observed [19] and eliminated [20] by a variety of means. Indeed, one significant conclusion of the work during this period was that BBU calculations and HOM damping measurements had become good enough that reliable predictions of the instability threshold current could be obtained. Predictions good to much better than 10% were possible on machines of the small scale of the free-electron lasers, consisting of perhaps two dozen accelerating cavities.

For a brief period, multipass BBU was observed in the full CEBAF accelerator [21]. The reasons for this situation are instructive, and provide compelling evidence about the accuracy of present simulation methods. The HOM damping requirements for the new type of cavities to be used in the 12 GeV upgrade project were relaxed by about two orders of magnitude compared to the original CEBAF cavities. The 12 GeV HOM damping requirements were obtained after thorough calculations taking into account the most unfavorable operations scenario supported by the 12 GeV accelerator and by using a statistical analysis including the fact that in a large machine the HOMs will not necessarily be degenerate in frequency [22]. The original 12 GeV HOM damping specifications were promulgated in 2004 [23], and have a safety factor of only 2 compared to two orders of magnitude previously. This situation reflects our greater

confidence in the results of thorough, modern simulation calculations. Since soon after the specification calculations were completed, new cavities installed in the CEBAF accelerator have been nominally required to meet only the 12 GeV requirements, as they should allow for stable operation of the present 6 GeV accelerator.

In 2007, a new cryomodule incorporating cavity prototypes of differently shaped cell profiles — high gradient (HG) and low loss (LL) — was installed in the fourth slot of CEBAF's north linac (NL04) as part of ongoing activity supporting increasing the final energy of the baseline CEBAF accelerator to 6 GeV. The LL cavity shape (see Fig. 10) was chosen as the most suitable for the 12 GeV era because of its better cryogenic efficiency. Soon thereafter, multipass BBU instability was observed, including the interesting observation that the operating current could be tuned precisely enough to the threshold current so the HOM excitation level could be made static in time. The instability manifested itself initially as a vertically extended spot at the end of the linac that eventually becomes large enough that the machine protection system trips the beam off.

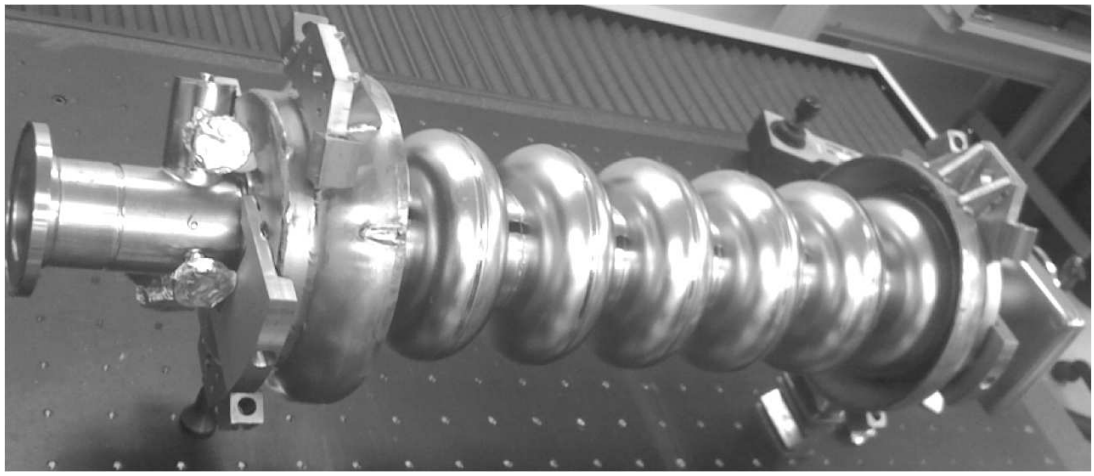

Figure 10. Image of a new high gradient and low loss cavity used to increase the CEBAF base energy to 6 GeV as well as prepare for the 12 GeV upgrade.

Evidence was gathered confirming the effect was multipass BBU. Voltages on beam position monitors were Fourier analyzed when the instability was present indicating transverse movement at an aliased frequency of 1920 MHz. This information, coupled with the fact that HOM frequencies were recorded with a spectrum analyzer from coaxial HOM coupler ports of cavities in the newly installed NL04, led to direct observations of the excited HOMs in the cavities in question when BBU instabilities were triggered. These observations unambiguously confirmed that multipass BBU was present. Particularly, RF measurements detected loaded Q-values of two BBU-corresponding vertically polarized dipole modes in one HG cavity (HG002) residing in the second passband (TM_{110}-like modes) as high as 10^8. The nominal phase advances of the two modes ($4\pi/7$ and $5\pi/7$ respectively) correspond to phase velocities highly synchronized with the electron velocities yielding comparably large geometry-dependent beam coupling factors R/Q. Together with the unexpectedly high Q's, the HOM impedances in fact exceeded the 12 GeV requirement by a factor of 3. Any BBU-driven instability was cured during the period the module was installed in the north linac by judicious changes in the recirculation optics. Also, the experimental observation was made that in some situations of beam recirculation optics, three passes through the accelerator can be more stable (have higher threshold current) than two passes, even though there is greater average current in the former case. In 2009, the module

was eventually removed from the north linac.

Investigations were launched during the period 2007–2009 to understand the lack of HOM damping. With NL04 still in CEBAF, extensive numerical investigations have been carried out to predict the cavity shape of the BBU culprit HG002 merely from frequency and external Q measurements at HOM couplers using today's high-end supercomputing capabilities combined with an elaborated cavity-shape algorithm tool [24]. This effort has forecast that both BBU-causing modes are tilted unfavorably away from the HOM couplers. The findings could not be proven until HG002 became available for field profile measurements after removing NL04 from CEBAF and disassembly in 2009. In fact, the tilted modes could be experimentally verified, though profiles deviated somewhat from the numerical results. For a more thorough understanding, the helium vessel of HG002 was removed to map external cell contours with a coordinate-measuring machine (CMM). This enabled reconstruction of the cavity interior with a parametric CAD model assuming uniform wall thickness [25]. Utilizing this model, HOMs have been found by numerical eigenmode analysis that differed significantly from expected fields in an undeformed cavity. When compared with field profile measurements, the BBU-generating HOMs could eventually be identified unambiguously (see Figure 11). Residual deviations from experimental results can be explained by the fact that the wall thickness is not uniform as assumed and that cells are not fully concentric, as confirmed by CMM measurements. The analysis also revealed that the culprit of the BBU scenario was a single deformed cell that required extensive tuning to flatten the accelerating field. The source of this deformation was traced to shrinkage due to a single extra weld pass on this cell during fabrication. This eventually caused modes to tilt severely, leaving practically no field at the HOM coupler positions, preventing adequate coupling of the HOMs and hence couplers, and leading to the observed elevated Q-values. It was also determined that the corresponding R/Q-factors would be adjusted to a value very close to the value deduced experimentally from the observed BBU thresholds. Another factor conspired to maximize vulnerability of this cavity to BBU; it so happened that the input waveguide window on this cavity had developed a leak, requiring installation of a shorting cap and auxiliary vacuum pump. This configuration not only eliminated opportunity to operate HG002, but also eliminated the input waveguide as a means for damping the TM_{110}-like modes now tilted toward it, thus contributing to raising their external Q values and lowering the BBU threshold.

To summarize, recent studies of multipass BBU theory and experiment show that BBU simulations are now good at predicting the instability threshold for small machines at the sub 10/of HOM damping Q-values and accurate calculation of beam coupling factors are needed for the simulation results to be correct, that the reasonableness and adequacy of the present 12 GeV HOM damping specifications have been confirmed. Such a level of accuracy is far beyond that believed achieved during the original CEBAF project.

8. Floating-pressure Ganni Cycle helium refrigeration

In the CEBAF Central Helium Liquefier, or CHL, the compression of helium to between 16 and 21 atm means substantial electrical power costs. Operating such a system at maximum capacity, even at the frequent times when maximum capacity isn't needed, is analogous to driving a car by flooring the accelerator while controlling the speed by strong use of the brake. However, by reconfiguring the CHL control structure, the Ganni Cycle enables the efficient scale-back of compressor energy consumption by reducing the operating pressures when full capacity isn't needed. Again figuratively speaking, it's like controlling automobile speed with varying acceleration and without use of the brake to offset over-acceleration. The process is seen in terms of floating pressure because it involves regulating the compressor discharge pressure to match the refrigeration load. The reconfigured control system senses when more refrigeration is needed by monitoring the amount of helium being used for cooling. The pressure floats to

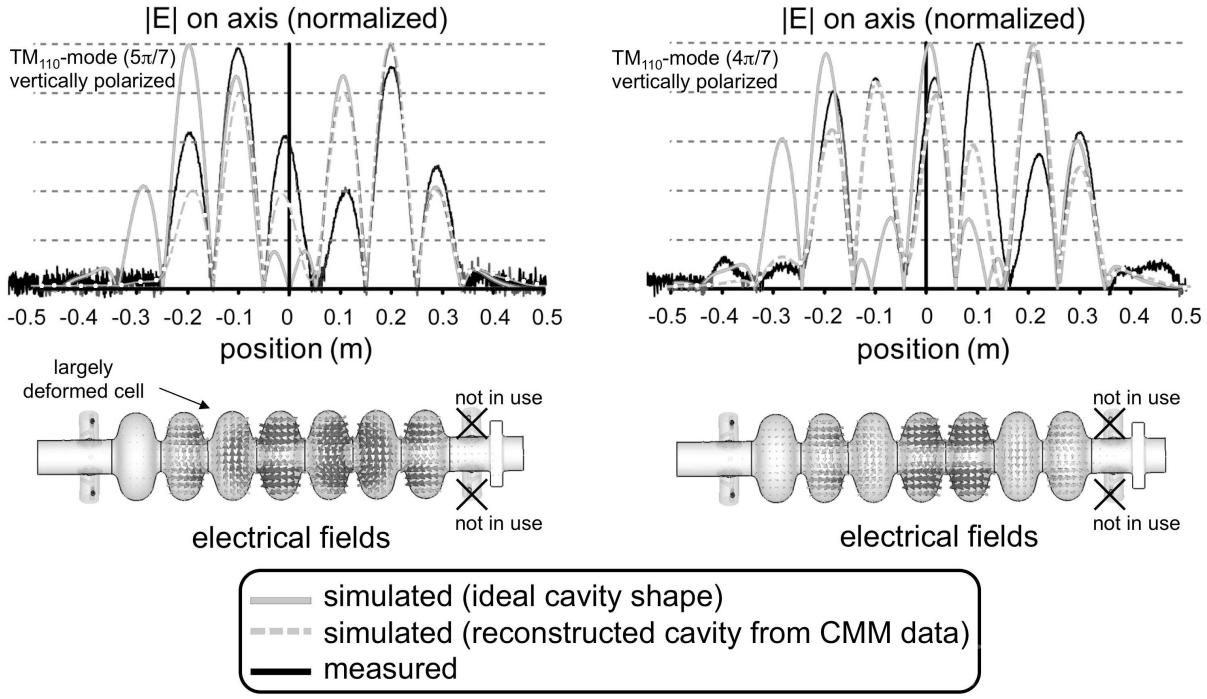

Figure 11. Calculated and measured high-order modes for a prototype high-gain low-loss cavity.

meet the load. If there is a reduction in load, there is a reduction in pressure, in power for compression, and consequently in electricity consumption.

The Ganni Cycle provides a basis for an optimal design at maximum load, turn-down cases and mixed modes, addressing the compressor system as the major input power loss contributor. It provides a solution to implement on as-built systems (existing or new) to improve system efficiency, reliability, availability and load stability under actual loads and help to improve the experimental envelope. It invalidates the philosophy that operating as-built systems at the presumed optimum of the manufacturer's theoretical design is optimal by properly identifying the fundamental process system parameters for control. It is a constant-pressure-ratio process cycle and maintains the compressor efficiency for varying loads. The Ganni Cycle maintains a constant volume flow (and thus velocity) at any point in the system and preserves the expander efficiency and the oil-removal effectiveness during the turn-down cases.

The Ganni Cycle nearly doubles the lifetime of some refrigerator components. Implementation for CEBAF, which mostly involved reconfiguration, did not require extensive capital expenses. For several years, this approach to CHL operation has been saving more than $1000/day in power costs, having lowered the CHL's electrical demand from 6 MW to 4.2 MW. The process has also been adopted at Brookhaven's Relativistic Heavy Ion Collider and at the Spallation Neutron Source at Oak Ridge. It has been licensed by Jefferson Lab to Linde Cryogenics, Division of Linde Process Plants, Inc., and Linde Kryotechnik AG for worldwide commercialization. NASA is using it in a cryogenics plant supporting the preflight, Earth-bound testing of components of the James Webb Space Telescope, Hubble's envisioned replacement, in ultra-cold operational conditions.

References

[1] C. Leemann, D. Douglas, and G. Krafft, Ann. Rev. Nucl. Part. Sci. **51**, 413 (2001).
[2] L. S. Cardman and L. Harwood, Proceedings of PAC07, Albuquerque, New Mexico, USA page 58 (2007).
[3] A. Lung, Proceedings of the XVI International Workshop on Deep-Inelastic Scattering page 230 (2008).
[4] R. McKeown, AIP Conf. Proc. (2011).
[5] P. Adderley, J. Clark, J. Grames, J. Hansknecht, K. Surles-Law *et al.*, Phys. Rev. ST Accel.Beams **13**, 010101 (2010).
[6] T. Nakanishi *et al.*, AIP Conf. Proc. **421**, 300 (1998).
[7] J. Hansknecht and M. Poelker, Phys. Rev. ST Accel. Beams **9**, 063501 (2006).
[8] C. Sinclair *et al.*, Phys. Rev. ST Accel. Beams **10**, 023501 (2007).
[9] J. Grames *et al.*, Phys. Rev. ST Accel. Beams **7**, 042802 (2004).
[10] C. K. Sinclair, Jefferson Lab Technical Note CEBAF-TN-96-032 (1996).
[11] D. Higinbotham, AIP Conf.Proc. **1149**, 751 (2009).
[12] K. Paschke, Eur.Phys.J. **A32**, 549 (2007).
[13] Y.-C. Chao, Proceedings of PAC07, Albuquergue, New Mexico, USA page 780 (2007).
[14] P. Axel *et al.*, IEEE Trans. Nucl. Sci. **24**, 1133 (1977).
[15] C. Lyneis, M. Mcashan, R. Rand, H. Schwettman, T. Smith *et al.*, IEEE Trans. Nucl. Sci. **28**, 3445 (1981).
[16] G. Krafft and J. Bisognano, Proceedings of the 1987 IEEE Particle Accelerator Conference page 1356 (1987).
[17] N. Sereno, L. Cardman, G. Kraft, C. Sinclair, and J. Bisognano, Proceedings of 1993 Particle Accelerator Conference page 3246 (1993).
[18] K. Beard, Jefferson Lab Technical Note JLAB-TN-05-060 (2005).
[19] C. Tennant *et al.*, Phys. Rev. ST Accel. Beams **9**, 074403 (2005).
[20] D. Douglas, K. Jordan, L. Merminga, E. Pozdeev, C. Tennant *et al.*, Phys. Rev. ST Accel. Beams **9**, 064403 (2006).
[21] R. Kazimi, A. P. Freyberger, C. Hovater, G. Krafft, F. Marhauser *et al.*, Proceedings of the 2008 Europeon Paricle Accelerator Conference page 2722 (2008).
[22] I. Shin and B. Yuun, Jefferson Lab Technical Note JLAB-TN-09-004 (2009).
[23] B. Yunn, Jefferson Lab Technical Note JLAB-TN-04-035 (2004).
[24] Z. Li *et al.*, SLAC-PUB-13266 (2008).
[25] F. Marhauser *et al.*, To appear in Proceedings of the XXV Linear Accelerator Conference (2010).