

OPERATION OF CEBAF WITH HEAVY BEAMLOADING

A. Hutton, Thomas Jefferson National Accelerator Facility, 12000 Jefferson Avenue,
Newport News, VA 23606, USA

Abstract

CEBAF is a 4 GeV, 200 μ A five-pass recirculating superconducting electron accelerator that has been operating for nuclear physics research at full energy since November 95. The beam current has been increased to over 180 μ A at 4 GeV with the maximum current in the linac over 900 μ A. The superconducting cavities operate in a regime where the beam-induced voltage is comparable to the accelerating gradient. The operational limits and the issues required to maintain stable operation of the 1497 MHz superconducting cavities will be discussed, together with the implications for the other accelerator systems. There are three experimental Halls which can run simultaneously with three interleaved 499 MHz bunch trains and RF separators. Operation with simultaneous beams to two Halls is now routine, and simultaneous three beam operation has been demonstrated. The maximum design current per bunch train (120 μ A) has been achieved. Hall B eventually requires beam currents as low as 1 nA (200 pA has been delivered) simultaneous with delivery of up to 200 μ A to the other Halls. The required beam current ratio of 10,000 has been achieved; development of 1 nA beam position monitors continues.

1 NOMINAL DESIGN PARAMETERS

CEBAF was designed to accelerate electrons to 4 GeV by recirculating five times through two, 1497 MHz superconducting linacs, each making 400 MeV per pass, with the nominal accelerating gradient of 5 MeV/m. The design maximum current is 200 μ A CW (maximum current per linac of 1 mA CW). The current can be split arbitrarily among three interleaved 499 MHz bunch trains, destined for the three experimental Halls A, B & C. A bunch train may be peeled off to any one Hall after each pass using RF separators and septa, while all Halls can receive the maximum energy. Hall B has a large solid angle detector and requires very low currents – down to 1 nA CW in the presence of beams of up to 200 μ A CW in the other halls. Most experiments will need polarized beams, initially 35% polarization and later up to 80%. Emittance and energy spread criteria of the beams are very demanding (design goal of 2×10^{-9} m for the 4σ transverse emittance at energies ≥ 1 GeV and an energy spread $> 2.5 \times 10^{-5}$). The emittance goal is met routinely, while the energy spread goal has been obtained, and is approached routinely. Low background conditions are vital, so careful attention must be paid to beam halo.

2 INJECTOR

2.1 Thermionic Gun

The beam from the 100 keV thermionic is sent through a 499 MHz chopper which rotates the beam in a circle of ~ 1.5 cm radius. Three independent slits are used to individually control the current for the three Halls. The three beams are then recombined, bunched and accelerated to 500 keV, before reaching the first superconducting cavities. This region of the machine is the most susceptible to space charge deformation and has been the extensively modeled with PARMELA. The simulations have been carefully benchmarked against measurements and have been used to establish focusing parameters that minimize space-charge effects. The maximum current that has been accelerated in a single 499 MHz bunch train is 220 μ A, comfortably above the required 120 μ A.

Over the last year, extensive tuning has been carried out to set up the three simultaneous beams. The major problems encountered were in obtaining the same path length through the choppers for each of the three beams. Without this condition, the bunching for the three beams is different and leads to background at the experiments.

2.2 Three-Beam Operation

The independent currents are set using the chopper slits in the Injector while the energies, or more exactly, the number of passes, is controlled by a set of 499 MHz RF separators, phase locked to the 499 MHz chopper RF system. There is one of these separators in each beamline downstream of the second linac. This enables one beam to be extracted on each of the intermediate passes and all three beams to be separated after five passes. The separators are set up using an air-cored corrector to deflect a single beam and then varying the phase and amplitude of the separator power to exactly re-create the same extraction orbit. This works well; the only remaining hardware is an independent phase reference for setting up an alarm signal for the machine protection system.

Downstream of the RF separators are DC septum magnets to complete the beam separation. These were damaged by the pulsed beam used for set-up and are being updated to improve the cooling.

Operation with two beams is now routine and three beam operation has been demonstrated. The first operation with three simultaneous beams is scheduled for May 15, 1997.

2.3 Polarized Gun

The polarized gun currently installed in the accelerator uses a bulk gallium arsenide cathode illuminated by a 1497 MHz RF gain-switched diode laser operated at 780 nm [1]. The laser makes 510 mW total power or 170 mW per Hall. The combination of laser and cathode can produce a maximum of 1070 $\mu\text{A}/\text{Hall}$ per 1% quantum efficiency. The quantum efficiency of the cathode declines initially at a rate dependent on the vacuum conditions, but stabilizes just below 1% and is fairly stable at this level.

The polarized gun was pressed into service on short notice and operated routinely, around the clock for five weeks. Typical currents required for operation were 25 μA CW, and this was obtained with a measured polarization of 36%. The maximum current ever captured in the Injector was 140 μA .

When the Injector is set up correctly, it is possible to change from the polarized to the thermionic gun with no change in RF phases. The rotation of the polarization vector using a Mainz style Z-bend is still somewhat slow, as the solenoids used to perform the rotation also steer the beam. This is currently being improved.

The polarization was measured using a Mott scattering chamber [2] operating at 5 MeV, with four counters at 173°. The counting rate is very low at these extreme backward angles, but we have high currents and the backgrounds are so small that measurements can be made to a precision of better than 1% in less than a minute.

Two techniques were tried to increase cathode lifetime. The optics in the Injector region were carefully optimized to obtain ~100% transmission at low current. Adding a pre-buncher reduces the chopping losses at high currents and, with a higher power amplifier than we currently have, 100% transmission of beams of 120 μA per bunch train should be obtainable. The ability to move the laser spot (300 micron diameter) over an area of the cathode of at least 6 mm diameter, while maintaining the transmission, has been demonstrated successfully.

A technique using atomic hydrogen (produced by RF discharge in low pressure hydrogen) has been extremely successful in cleaning bulk gallium arsenide [3], quantum efficiencies of 10% at 780 nm being regularly obtained. The process has been transferred to MAMI at Mainz where the process was able to recover "dead" cathodes with nearly twice the previous quantum efficiency. This technique is extremely important for high polarization (~80%) strained cathodes where the polarization is a surface, rather than a bulk, process and other cleaning techniques cannot be used on these very thin layers. If this technique can be used for *in situ* cleaning, the use of load-locks for transferring clean cathodes from the laboratory to the accelerator may be avoided.

The lifetime of these cathodes at high currents is the principal focus of further study.

3 SUPERCONDUCTING CAVITY LIMITS

3.1 Characterizing the Cavity Gradient Limits

The accelerating gradient in the linac is limited by many different phenomena, either inherent to the cavities, or due to protection interlocks, or system limitations. The phenomena and operational criteria are as follows:

- 1 Field Emission – the most important effect, limiting the gradient in 86% of the cavities. The initial criterion was 1 W heat load per cavity, and the operational limit is 1 Rad/hr per cavity, established during tests; the main effect is electrostatic charging of the cold RF window leading to arcing – we interlock a photodiode detecting light from the arc. If the arc is real, it is always accompanied by a waveguide vacuum trip.
- 2 Waveguide Vacuum: $>10^{-7}$ torr – usually due to an arc, but can be outgassing from excessive reflected power.
- 3 Quench: 1 MeV/m below the gradient at which quench observed in tests – with this limit, a quench has never been seen during accelerator operations.
- 4 Beamline Vacuum: $>3 \times 10^{-9}$ torr – usually due to thermal profile shifts which release adsorbed hydrogen.
- 5 Q_0 (intrinsic cavity quality factor) $<1.0 \times 10^9$ – affects the cryogenic plant and can be ignored for single cavities.
- 6 Window temperature: the infrared pyrometer tripped – this is usually due to bad waveguide RF seals, which can cause warm window failures.

Measurements are under way to characterize the applicable limit for each of the 330 superconducting cavities installed in CEBAF. This involves the original cavity acceptance data, and measurements in the tunnel at various beam currents to obtain information about the integrated systems. This data base is currently about 80% complete.

3.2 Helium Processing

The average gradient of the installed superconducting cavities was 7.4 MeV/m (predicted from the cryomodule commissioning data). Every cryomodule (which contains eight cavities) exceeded the design specification of 5 MeV/m. This year, two different approaches were taken to improve the operating performance of the cavities, both carried out *in situ* [4]. The aim was to reduce field emission and arcs using helium processing, and to condition the waveguide in the thermal transition region between the warm and cold RF windows using waveguide vacuum processing.

Helium processing is a well known technique for cleaning cavities in laboratory conditions. A gradient of a few MeV/m is applied to the cavity and low pressure helium gas is bled in to bring the pressure up to $\sim 4 \times 10^{-4}$ torr. These conditions create active field emission, but not an RF discharge. The fields were slowly increased to the gradient limit due to other effects and maintained at this level for as long as the schedule permitted. Most of the improvement was obtained in the first half hour, followed by diminishing returns.

One interesting problem arose in all of the cavities. If the cavity trips off, there is a hard multipacting limit at -0.7 MeV/m which prevents the RF from being switched back on. This was addressed by using a variable frequency phase lock loop RF system. By approaching the nominal phase lock loop RF system. By approaching the nominal frequency from above, the phase lock loop "jumps" the frequency across the multipacting barrier and the cavity can be switched on [5]. Thereafter, the gradient can be freely varied over the whole range. The effect limited our ability to process cavities in parallel, as a mobile RF system was needed rather than the installed RF equipment. After the helium conditioning was complete, the cavity was warmed up to 20 K, which outgassed the helium that had been adsorbed by the cavity as well as some hydrogen that had built up. There is some evidence that this thermal warm-up alone is beneficial to the cavity.

Operationally, helium processing of two cryomodules can be completed in 24 hours with the tunnel locked up for safety reasons. This is followed by a 16 hour warm-up and cool-down cycle, during which other routine maintenance can be performed. We have now processed 12 of the 41 cryomodules.

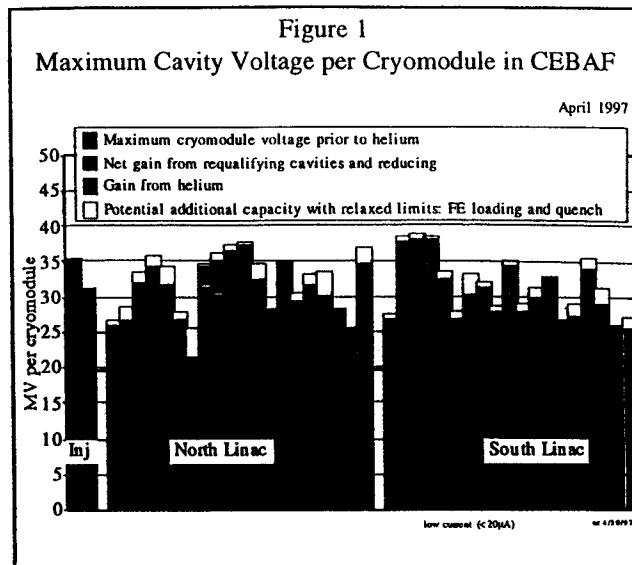
3.3 Waveguide vacuum processing

Waveguide vacuum processing is a form of pulsed RF power processing. We use it *in situ* to process the thermal transition region between the warm and cold waveguide windows. Gas is adsorbed in this area, and when the cavity reflected power increases due to cavity mis-tuning (in principle, beam loading can also produce this effect but this has not been observed), the RF wave pattern changes from a standing wave (most power reflected) to a traveling wave (most power transmitted to the beam). Shifting the thermal profile creates outgassing leading to arcing and waveguide vacuum trips.

Low duty factor, pulsed RF power is applied, maintaining the vacuum pressure below 10^{-7} torr. Four waveguides are treated at a time for 4 hours, using the installed RF power sources and specially adapted software. Following this, the waveguide vacuum is usually more stable. In particular, the cavities are more tolerant of mis-tunings, which makes operation easier. We have now processed the waveguide vacuum in 10 of the 41 cryomodules.

It is still unclear whether the vacuum improvement due to 20 K thermal cycling is equivalent to pulsed RF processing or whether there is a qualitative or quantitative difference between the two processes.

Using the two processing techniques, every cryomodule showed improvement (some cavities could not be processed as they were limited by other effects) and an overall gain of 63 MeV was obtained in the two linacs (i.e. 315 MeV for five passes). This includes some gain from relaxed limits. We estimate that continuing this treatment for all of the cryomodules would yield a further 135 MeV (675 MeV for five passes). Figure 1 shows the gains achieved.



4 RF/CAVITY SYSTEM PERFORMANCE

4.1 Optimizing RF Systems

Each superconducting cavity is individually fed by a 5 kW klystron with its own low level control module to feed back and stabilize the cavity phase and amplitude. There are 330 of these systems and maintaining each one at optimum performance requires tuning with RF power alone, as well as procedures using the beam. Careful file management is then required to maintain this information.

The cavities are tuned to the correct frequency using the AUTOTUNE program, which also provides an offset angle to normalize the phase angle between the input drive phase and the cavity gradient, compensating for differences in cable lengths and electronic delays. Over time, or under different conditions, the control modules may change and the AUTOTUNE data is no longer valid, so other checks must be made to ensure consistency.

An effective global check is to slew the master oscillator frequency and measure the forward power drawn from each klystron, as the cavity accelerating voltage is maintained constant by the feedback system [5]. An off-line script is then used to look for changes in forward power indicating those cavities whose center frequency is not correct.

A further check can be made using beam loading. The phase of the klystron forward power with respect to the cavity probe signal is monitored as the beam current is turned off and on [5]. Changes indicate that the cavity loading and the klystron drive power are not in phase.

The cavity gradients were calibrated using the beam and compared with the expected values from the measured Q_0 values of the cavity. This consistency check was used to correct calibration and other errors. The gradients are now correct to within 5%. This also provided cross-checks on cavity dissipation, required to compensate the cryogenic loading as a function of gradient.

4.2 Cavity Phase

The phase of the cavity fields with respect to the beam are calibrated using the KREST program. The concept is simple: move the klystron phase (or more exactly the phase of the cavity probe signal) by $\pm 10^\circ$; use the energy feedback system to measure the energy change; and set the phase to equalize the energy change on either side of the nominal phase. The program is used regularly (at least daily) as the phase distribution system is still not sufficiently stable. However, KREST works so well that we can tolerate this lack of stability, at least for now.

The net result of this preparatory work is that we can set the cavity phase and amplitude correctly, independent of gradient and beam current.

4.3 Operating Cavities

The klystron power is limited operationally by lowering the high voltage and by using the modulating anode to minimize the electric power. The klystron filament voltage is used to limit the cathode current to extend the filament lifetime. The parameters are set at the beginning of an experimental run, based on the beam energy and on the total current traversing the cavities (beam current multiplied by the number of passes through the linacs). An overhead of 15% is added to the klystron power to provide headroom for phase and amplitude control.

An off-line computer program is used to calculate the maximum gradient limit of each cavity, given the cavity limitations, the required current in the linacs, and the klystron power limits. The gradient limits are used by the Linac Energy Management (LEM) program in the control room to set up and maintain the linac energy. These programs are being integrated into one package, the Momentum Management System (MMS), which will set the arc magnets as well as the RF systems. The gradient of each active cavity is set proportional to the maximum allowed gradient for the cavity to maintain the accelerating voltage in the linac constant. With a 5% overhead at the start of the run, problem cavities can be taken off-line to extend the time between maintenance periods.

5 LINAC PERFORMANCE

5.1 Preparatory Tests

During breaks in routine beam operation for physics, the gradients of individual cavities were slowly increased to determine the maximum acceptable operating gradient. This required operating sections of the accelerator at full klystron power and was carried out over a period of months, using successive maintenance periods to change the klystron voltage tap settings until each zone had been tested. The maximum gradient was usually limited by the arc rate, and one per hundred hours was considered acceptable initially. These numbers were then used to define the maximum operating gradient.

5.2 High Gradient Operation

The first attempt in May 96 was to set the machine for 1 GeV per turn (corresponding to 5 GeV total, 25% above nominal). All the cavities were simultaneously brought up to their nominal setting. After reducing the gradient in troublesome cavities and allowing them to condition, most (but not all) could be returned to their original setting. After about a shift, the gradients in the linac were adjusted to obtain exactly 500 MeV per linac and pulsed beam was established. The tune was rechecked, the feedback systems were switched on and the gradient calibration was checked with beam.

The 1 GeV run was repeated in December 96 when CW beam was established. Initial set-up of the gradients took many hours due to waveguide outgassing and fixing RF problems. Pulsed beam was set up and apertures were swept to ensure that the optics was properly loaded. 100 μ A CW beam was then established and maintained for several hours while working on the RF systems to improve the trip rate.

Following this success, further cavity processing was scheduled, leading to a repeat test. This time, the linacs were set up for 1.12 GeV (40% above nominal). Following approximately the same scenario, 90 μ A CW beam was established (limited by the capacity of the beam switchyard dump) and after a 2 hour run, the trip rate was 45/hour – sufficiently good for an accelerator test but not yet good enough for physics. The average gradient of the active cavities was 7.76 MeV/m, compared to the nominal 5 MeV/m.

An interesting effect was observed that had been predicted for some time. The ponderomotive force of the RF fields tends to expand the cavities and the cavity tuner changes to maintain the correct operating frequency. If the cavity trips off at high gradient, the resulting frequency shift can be sufficiently large that the cavity goes out of resonance. This was observed in only a few cavities. The present work-around is to use manual tuning or slow ramping of the gradient to allow the tuner to follow the frequency changes. This works, but is both time-consuming and suffers from mechanical hysteresis of the tuner. A better solution would be to slew the drive frequency in the control module, but this will take some time to develop.

The result of the tests is that first experiments at 5 GeV have been scheduled for November 1997, and we expect to be scheduling experiments at 5.5 GeV next year.

5.3 High Power Operation

The nominal beam current to be divided amongst the three experiments is 200 μ A, i.e. 1 μ A total current in the linacs with five passes. The beam current is regulated as part of our NEPA (National Environmental Protection Agency) permit and initially we were only permitted to operate up to 180 μ A, due to DOE imposed safety limits.

The DOE has now granted permission to test the accelerator up to the full current.

We have had a series high power tests, scheduled together with the high gradient runs, as the RF and cavity preparation was the same. The additional problem is to ensure that the beam-loading power vector is aligned with the klystron power vector over the range of beam currents. The strategy adopted was to raise the current to 100 μA and re-check the cavity tunes by sweeping the RF frequency, as described above. After fixing any problem cavities, the current was increased until the trip rate became excessive again, monitoring the klystron drive for evidence of saturation. The RF frequency was again swept and problem cavities addressed. In a series of tests, the current was successfully increased to 180 μA . The final step to reach the maximum current is expected soon.

All the problems encountered were due to system interactions between control module, RF power source, and the cavity, which could all be (somewhat laboriously) addressed. There are no indications that there are any fundamental limitations. The primary difficulty was access to a suitable dump, as the only dumps capable of taking the power (up to 1 MW) are in the experimental Halls and the experimenters are usually either taking data or are taking access for fixing their equipment.

We are now able to schedule multiple Halls with a total current of 200 μA and expect that the availability, after initial conditioning, will be independent of beam power.

6 EXPERIMENTAL CAPABILITIES

6.1 Hall Status

Hall C has been taking data since November 95 at ever increasing current. The maximum power that has been delivered to an experiment is 4 GeV, 120 μA (480 kW) which was achieved with similar availability to any other run. Four experiments completed data-taking and data for 16 theses was obtained. Initially, solid targets were used, but the cryo-target was commissioned in August 96. In the first three months of 97, the detector was re-configured for the t20 experiment and data-taking has commenced.

Hall A started taking beam in April 96 using the electron spectrometer only. A problem with the superconducting coil of the hadron spectrometer delayed two-arm testing until February 97. Commissioning officially ended in April 97 and the first experiment is starting now.

Hall B took first beam in December 96 to evaluate backgrounds and determine whether the low currents required (see below) could be obtained with sufficient stability. A vacuum chamber misalignment was found in the experimental line and instabilities in the RF systems created beam loss. The misalignment has now been corrected and the improvements in the RF tuning that were driven by the high gradient and high power programs

are expected to lead to a significant reduction in background for the detector.

As far as the accelerator is concerned, every beamline in the machine has now been successfully commissioned.

6.2 Low Beam Current for Hall B

The range of beam currents required by Hall B is from 1 nA to 1 μA . A version of the Switched Electrode Electronics (developed for the linacs with a current range of 1 μA to 1 mA) was developed to improve the low current capability. The range achieved was from 100 nA to 200 μA and this electronics is installed on beam position monitors in the injector, and is now being installed for the feedback systems and all the experimental beamlines. They were used successfully for the initial phase of Hall B commissioning.

A new type of beam position monitor is being developed for use at the target [6]. It uses resonant cavities, operating on the lowest deflecting mode at 1497 MHz. The expected sensitivity of these monitors is 1 nA.

One important capability has been demonstrated. Currents as low as 200 pA have been produced in the Injector, observed on fluorescent screens in the Hall B beam-line, and used by the detector for stability and background studies. The ability to produce the range of beam currents required has been demonstrated. When the new 1 nA monitors become available, we will be ready to make operation at these low currents a standard operating condition.

It had been expected that the low current beams for Hall B would always need a "witness" beam to allow the feedback systems to lock on. In these studies, it was demonstrated that if a beam with more than a few μA is present on the first turn to enable the energy feedback systems to stabilize the linac energies, then the Hall B beam at any energy would be sufficiently stable. This will be a great help in scheduling multiple Halls.

7 ACKNOWLEDGMENTS

This paper presents an overview of the work of a large group of dedicated people in Jefferson Lab. I would like to thank each and every one for their tenacity and skill, which have resulted in the excellent performance obtained. This work was supported by the U.S. Department of Energy under contract number DE-AC05-84-ER40150.

REFERENCES

- [1] M. Poelker: *Appl. Phys. Lett.* **67**, 2762 (1995).
- [2] J.S. Price et al.: '5 MeV Mott for rapid, precise electron beam polarization measurements'; *Spin 96*, Vrije Universiteit Amsterdam, 10-14 Sept. 96.
- [3] C.K. Sinclair et al.: 'Atomic Hydrogen Cleaning of Semiconductor Photocathodes', this conference.
- [4] C. Reese et al.: 'Improvement of the Operational Performance of SRF Cavities via in situ Helium Processing and Waveguide Vacuum Processing', this conference.
- [5] L. Doolittle: private communication.
- [6] R. Ursic and C. Pillar: '1 nA Beam Position Monitoring System', this conference.