RF CALIBRATION OF CEBAF LINAC CAVITIES THROUGH PHASE SHIFTS*

A. Carpenter#, J. Benesch, C. Slominski, Jefferson Lab, Newport News, VA, USA

Abstract

This paper describes a new beam-based method of cavity energy gain calibration based on varying the cavity phase. This method can be fully automated and allows a larger range of momentum excursions during measurement than previous calibration approaches. Monte Carlo simulations suggest that a calibration precision of 2-3% could be realistically achieved using this method. During the commissioning of the Continuous Electron Beam Accelerator Facility's (CEBAF) energy upgrade to 12 GeV, 876 measurements were performed on 375 of the 400 linac cavities in Fall 2015 and applied December 2015. Linac optics appears to be closer to design as a result. The resulting ensemble proved to be 2% over the value needed to get the desired energy in the arcs. Continued offline analysis of the data has allowed for error analysis and better understanding of the process.

CEBAF 1995-2012

CEBAF is a recirculating electron accelerator using superconducting RF. It consists of an injector which provides fully relativistic electrons (130 MeV), two linacs of equal momentum gain, and ten recirculating arcs. The injector and linac energies are in the ratios: 0.1128:1:1 so electrons in arc 2 have ~90% more momentum compared to arc 1. With fully relativistic electrons, the extreme relativistic limit applies, so momentum measurements in the recirculating beam transport arcs are proportional to energy.

For 1995-2012 a momentum balance method was used to calibrate the cavity gradients (energy gain per unit length). The highest energy-gain cavity in each linac was calibrated against the magnets in the downstream arc using optics with dispersion ~6 m. The reference cavity in a given linac was then set to produce a 1.5 MeV/c change in momentum in the downstream arc. The gradient for the cavity under test was nominally set to produce a compensating 1.5 MeV/c change or turned off, and the reference cavity was used to restore the original orbit in the following arc. The change in the reference cavity's gradient was used to calibrate the gradient in the cavity under test. The associated momentum gain per cavity ranged from 1.5 to 6 MeV/c. Repeated arc 1 measurements on individual cavities in linac 1 had an RMS span of 5%. Measurements in arc 2 had a 9% RMS span as the momentum in the arc is higher but the change due to a single cavity is not. The RF control system is not stable for gains corresponding to 1.5 MeV/c or less. These errors were adequate for the forgiving optics of the original CEBAF layout and minimal emittance growth due to synchrotron radiation.

PHASE SHIFTS

In late 2014, it was realized that using the phase control on each cavity would allow a larger momentum shift than the gradient balance method as the cavity could be moved to a decelerating phase. This would increase the shift to almost twice the nominal or the maximum allowed by the following arc's acceptance. With assistance [1] Root was used for a Monte Carlo evaluation of the process. For momentum shifts limited to 0.4%, energy RMS errors were in the range 1.9-2.5% for all linac 1 cavities and for linac 2 cavities providing more than 2.5 MeV/c gain. For lower gain linac 2 cavities, amplitude RMS error was projected to reach 4.6% at 1.5 MeV/c. The cavity phase shift-based algorithm is summarized below. The analysis that follows considers cavities with nominally "low" (~1.5 MeV/c) to "medium" (~6 MeV/c) operating gradients and those with nominally "high" operating (~10.5 MeV/c) cavity gradients. This classification scheme distinguishes between the older, lower capability cavities, and the newer, higher capability cavities fabricated for the recent CEBAF upgrade.

Calibration and Shift Calculation Algorithm

Calibration of RF cavity gradient settings and proper phasing is achieved by shifting a cavity's phase and measuring the change in energy (momentum) in the nearest downstream dispersive region (beam transport arc). Since the only element changing is the phase of a single cavity, the change in energy can be calculated using the following model,

$$dE_i = E_i - E_{orig} = A\cos(\theta + \varphi_i) - A\cos\theta$$

where

dE_i ≡ change in energy measured in the nearest downstream dispersion region for shift φ_i
A ≡ amplitude (MeV) of the cavity energy cosine function to be estimated
θ ≡ cavity phase angle error to be estimated
φ_i ≡ the i^{th} cavity phase shift

This formula can be recast as $\delta E_i = X_{1i}\beta_1 + X_{2i}\beta_2$ where $\beta_1 = A\cos\theta$, $\beta_2 = A\sin\theta$, $X_1 = \cos\varphi_1 - 1$, and $X_2 = -\sin\varphi_1$. Using measurements made at multiple phase shifts, this model can be fit using ordinary least squares regression. Once estimates for β_1 and β_2 have been found, it is trivial to solve for A and θ where

* Notice: Authored by Jefferson Science Associates, LLC under U.S. DOE Contract No. DE-AC05-06OR23177 and DE-AC02-06CH11357. The U.S. Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce this manuscript for U.S. Government purposes.
# adamc@jlab.org

$$\theta = \tan^{-1}(\beta_2 \div \beta_1)$$

$$A = \beta_1 \div \cos\theta.$$

In general, larger phase shifts produce more accurate estimates of phase error and amplitude. However, a sufficiently large phase shift could exceed the dispersive region's energy error tolerance, producing a normalized momentum error that steers the beam into the beam pipe of the beam transport arc. The strategy is to use the largest possible phase shift without exceeding the arc's energy error tolerance. During normal operations, each recirculation beam transport arc at CEBAF with 7 m dispersion optics can tolerate normalized momentum errors (dp/p) of up to 1e-3 without tripping the machine. During beam studies, 2 m dispersion optics were utilized in order to increase the momentum error tolerance. In addition, a positive energy offset was applied to push the beam to the positive boundary of the arc's momentum error tolerance. Since applying a phase shift to a reasonably well-tuned cavity generally lowers the cavity's energy and momentum gain, this effectively doubled the range of momentum errors that beam studies were able to utilize. Measurements were made for phase shifts corresponding to both positive and negative normalized momentum errors (dp/p) between ±2e-3.

Calculation of the maximum phase shift is predicated on the arc's energy tolerance, the energy being produced by the cavity of interest, and the current phase error of that cavity. The first two of these are typically known within tolerance, but the phase error is an unknown quantity being measured. This calculation requires a conservative estimate of largest likely phase error of a cavity. Historical experience at CEBAF suggests that 10° is a reasonable phase error estimate after manual phasing. Once the energy tolerance, the cavity energy, and the current phase error values are determined, the maximum shift, φ_{max}, is calculated such that the resulting energy change, δE, is less than the arc's energy error tolerance. Individual phase shifts are chosen evenly across the interval, [-φ_{max}, φ_{max}], excluding the zero phase shift.

Each energy reading used in this algorithm is actually the average of five sequential measurements made at a given phase shift and after a brief quiescence period in the software implementation. This approach allows for online removal of outlier readings, for lowered variance in energy readings, and time for energy variations due to the phase shift to settle out. Testing revealed that 12 phase shifts with an energy reading averaged over five measurements provided a sufficiently good fit.

OFFLINE DATA ANALYSIS

The transformation used in the algorithm from (A,θ,φ) to (β_1,β_2,X_1,X_2) invalidates the classical OLS approach for statistical inference for the parameters of interest, namely amplitude and phase error, so instead, the calibration measurements were fit using a Nonlinear Least Squares (NLS) method as implemented in R's nls routine [2]. The amplitude and phase error estimates from the code above were used as initial starting points in this routine. The raw data was refit first using the phase and amplitude variables only and then with an additional constant (error) term. The addition of the last provides a means for identifying invalid calibration data sets. The results shown below were fitted with phase, amplitude and constant terms:

$$dE_i = X_0 + A\cos(\theta + \varphi_i) - A\cos(\theta) + \varepsilon_i$$

where the terms are defined as above with

X_0 ≡ Error indicator

ε_i ≡ Random error associated with the ith measurement. Distributed i.i.d. $N(0, \sigma^2)$

Of the 876 measurements, twenty-one values of X_0 outside [-0.4, 0.2] were found to correspond with invalid data and will be omitted in the subsequent analysis, leaving 855 measurements to consider. This range of X_0 matched a manual review of the data and identified additional invalid calibration measures initially deemed acceptable. Statistical significance of X_0 was not a useful indicator of invalid data as relatively small X_0 values could be statistically significant while having very limited impact on the parameter estimates of interest. Figure 1 shows the full range of error indicator terms. This approach for identifying invalid calibration data relies on qualitative expert review, but future work aims to incorporate the error indicator term into the automated calibration software tool.

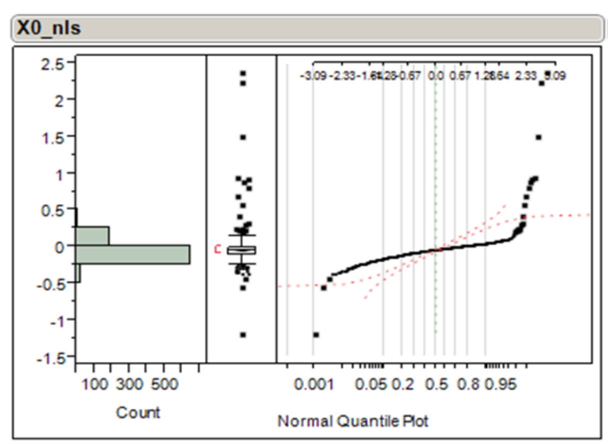

Figure 1: The values of error indicator terms displayed in histogram, box plot, and Normal quantile plot.

Figure 2 shows the amplitude estimates (MeV/c) derived from the NLS fit against the estimates' t values, with two distinct patterns emerging. The high gradient group highlighted (blue) are the modules that were fabricated for the recent 12 GeV energy upgrade. The other block contains cavities purchased in 1991-1993. Most notably, errors on high momentum gain cavities were larger (lower t values) than the older, lower capability units. While reasons for this are currently unknown, all estimates are highly significant.

The phase error estimates vary in statistical significance much more widely than do the amplitude estimates. This is because the hypothesis test uses a null hypothesis that the phase error equals zero. The phase error standard errors are of similar magnitudes across parameter estimates, while the estimates themselves can range greatly in magnitude. As the phase error estimate approaches zero it is dwarfed by its standard error, leading to a lack of statistical significance. Figure 3 displays this pattern for the t values of phase error estimates for all cavities. As phase error estimates approach five degrees, they achieve a statistical significance at the p=0.01 (t=3.25) level. Thus the program calculates phase errors with more than adequate precision to minimize longitudinal emittance growth.

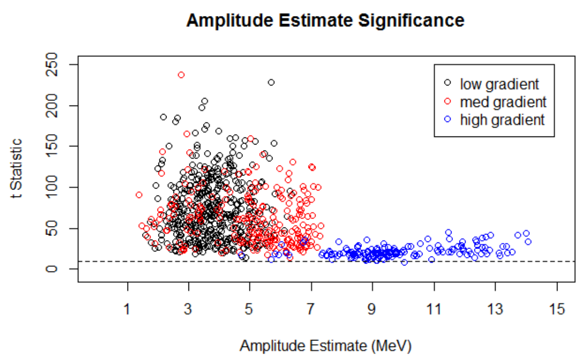

Figure 2: The statistical significance (t values) of amplitude estimates are displayed by cavity gradient classification. The t=10 level is noted with a dashed line.

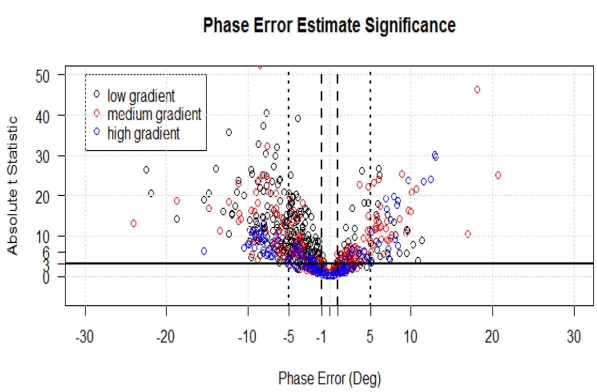

Figure 3: Shows the t value of phase error estimates for all cavities.

The fractional error of the parameter estimates (Fig. 4) was expected to be negatively correlated with the magnitude of the amplitude estimate. This is because higher gradient cavities produce larger momentum shifts for a given phase shift which allows the arc momentum acceptance to be fully utilized.

The positive slope of the low/medium linac 1 cavities has low significance. This trend should have been more noticeable in linac 1 given the lower momentum in the dispersive region following it. However, as the trends in Fig. 4 show, this effect was only seen in the high gradient cavities.

For the higher gradient cavities, fractional errors are higher than expected based on simulation results. For high gradient cavities in the linac 2, mean error is 5.5%; the simulation suggested 1.5%. For high gradient cavities in linac 1, the mean error is 4.6%, not the 1.5% expected from the Monte Carlo. While the reason for this higher than expected error is unknown, this error is still comparable to the historical value of 5%.

For the low and medium gradient cavities, fractional error was similar to simulation predictions or better. Linac 1 and 2 cavities had a fractional error of 1.9% and 2.0%, respectively. This compares favorably with both the simulation predictions, 1.1%-2.0% and 1.9%-4.6%, respectively and the historical fractional errors of 5% and 9% respectively.

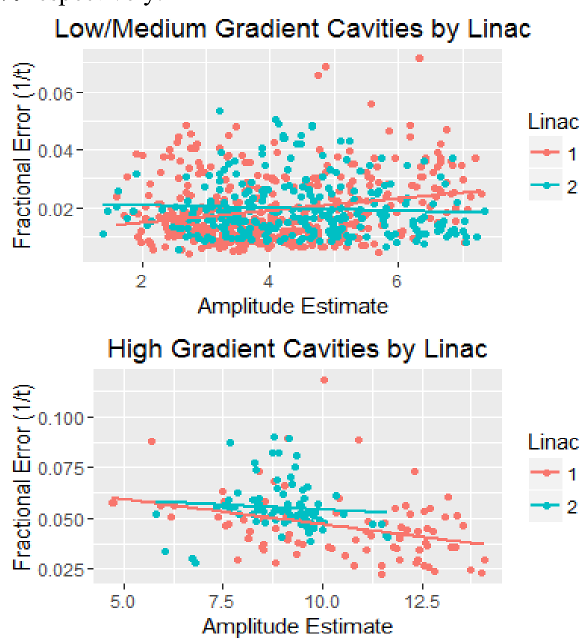

Figure 4: The fractional error of amplitude estimates separated by cavity type and linac. The trend lines are simple linear fits.

CONCLUSIONS

The new method provides a factor of 2.5 improvement in error over the old momentum balance method for cavities in old lower gradient cryomodules. It matches the old method in linac 1 and improves on it by a third (from 9% to 6%) in linac 2 for the new higher gradient modules. In addition, the RfPhaser software tool used to automatically phase cavities has been upgraded to additionally provide a fully automated means of gradient calibration. The improvement in gradient calibration has made it easier to match the optics in the machine to design.

REFERENCES

[1] ROOT assistance provided by Luke Myers, now at Bluffton University.

[2] `https://www.r-project.org`