February 2007

A Toroidal Charge Monitor for High-Energy Picosecond Electron Beams ⋆

Robert H. Simmons and Johnny S. T. Ng *

Stanford Linear Accelerator Center, Stanford University, Stanford, CA 94309, USA

Abstract

A monitor system suitable for the accurate measurement of the total charge of a 2-ps 28.5 GeV electron beam over a large dynamic range is described. Systematic uncertainties and results on absolute calibration, resolution, and long-term stability are presented.

Key words: beam charge, toroidal current transformer, absolute calibration
PACS: 41.85.Qg, 07.50.-e, 07.55.-w

1. Introduction

In a recent experiment a high-energy electron beam was used to measure the fluorescence yield in air from electromagnetic showers. Air fluorescence has been fundamental to the observations of cosmic rays at the energy scale of 10^{20} eV. Particles at these extreme energies may reveal yet-to-be discovered laws of nature, or they may be caused by mechanisms beyond our current understanding. Two recent observations with comparable statistics, however, showed an apparent discrepancy in the absolute flux and in the spectral shape. One experiment, the High Resolution Fly's Eye (HiRes), based its observation on air fluorescence[1]. The other experiment, the Akeno Giant Air Shower Array (AGASA), used large ground arrays of scintillation counters. The important implications of this disagreement motivated the FLASH experiment (FLuorescence in Air from SHowers), which aimed to make spectrally resolved measurements of air fluorescence excited by electromagnetic showers with an accuracy of 10% [3]. The fluorescence emission is directly proportional to the absolute beam charge, the number of particles interacting with air molecules. Therefore, accurate measurement of the beam charge was an important requirement.

FLASH was carried out at the Final Focus Test Beam (FFTB) facility at the Stanford Linear Accelerator Center (SLAC) over a two-year period. Its recently published results appear to validate the fluorescence measurement technique [4,5]. In the first (thin target) stage, the beam passed through a thin air gap. There was sufficient light for precision yield measurement at a beam intensity of approximately 2×10^9 electrons per pulse. Above this intensity, the yield became non-linear because of additional ionization by locally accelerated secondary electrons [5]. In the second (thick target) stage, the beam was directed at an alumina target with variable thicknesses up to 14 radiation lengths. The fluorescence was sampled in an air-gap behind the target. The showering process was expected to increase the number of particles per pulse by a factor of up to 100. Considerations for nonlinearity and radiation containment in the FFTB led to a nominal beam intensity limit of approximately 2×10^7 electrons per pulse. An absolute calibration uncertainty of 2% was specified, even at the lowest beam intensities. Detailed monitoring of long-term stability over the lifetime of the experiment was also required.

A ferrite-core toroidal current transformer design was chosen. This technique is well-known [6–9]. The specifications described above for the short pulse high-energy beam,

* Corresponding author. Tel.: +1-650-926-4490; fax; +1-650-926-5566.
Email address: jng@slac.stanford.edu (Johnny S. T. Ng).

Submitted to Nuclear Instrumentation and Methods (NIM) A
Work supported by the US Department of Energy contract DE-AC02-76SF00515

however, made it necessary to have a custom-designed toroid. In the rest of this paper, we describe the design (Section 2), absolute calibration (Section 3), systematic uncertainties (Section 4), and performance (Section 5) of the FLASH toroid.

2. The Toroidal Beam Charge Monitor

A block diagram of the toroid system is shown in Figure 1. For the diagnostics of picosecond beams, a broad radio frequency spectrum is involved. It was important to understand the high frequency properties of the ferrite for this toroid design. Careful considerations were given to the issues of mechanical setup and electrical ground isolation because of the complex electrical environment in an accelerator enclosure. Conventional pulse signal processing was used to achieve the required charge measurement accuracy.

2.1. *The Ferrite Core – Frequency Dispersion of Permeability*

It is well-known that the ferrite core permeability changes rapidly at high frequency. To mitigate potential problems, one approach has been to integrate the beam current pulse using capacitors implemented in the ceramic gap in parallel with the primary "winding" (beam current) [8]. This effectively cuts off the high frequency components before the signal was sensed by the transformer.

In the FLASH design, the beam current would be integrated initially via the capacitance across the terminals of the transformer secondary coil. The toroid's impulse response to a 2 ps beam might depend on the properties of the ferrite core in the RF regime. The core consisted of a stack of 8 ferrite rings mounted coaxially on a ceramic beam tube. A manganese zinc (MnZn) ferrite [1] made from sintered MnZn blended with iron oxide was used.

The ferrite's relative permeability is generally represented as a frequency-dependent complex parameter, $\mu(\omega) = \mu'(\omega) - j\mu''(\omega)$. The real part represents the reactive portion, the imaginary part the losses. The frequency spectra under excitation by an external sinusoidal magnetic field can be described by two magnetizing mechanisms: domain-wall oscillation and gyromagnetic spin rotation. This has long been a subject of investigation [10]. A recent measurement of the permeability of a similar MnZn ferrite was carried out in the frequency range from 100 kHz to 10 GHz [11].

The permeability data of the ferrite used in this work as provided by the manufacturer are shown in Figure 2. The frequency spectra were fitted to curves described by the following model expression [12]:

$$\mu = 1 + \frac{\chi_d \omega_d^2}{\omega_d^2 - \omega^2 + j\beta\omega} + \frac{\chi_s \omega_s(\omega_s + j\alpha\omega)}{(\omega_s + j\alpha\omega)^2 - \omega^2}. \quad (1)$$

[1] Fair-Rite Products, Wallkill, NY, part number 5977011101

It represents a simple sum of the contributions from domain-wall oscillation and gyromagnetic spin. The fitted parameter values are $\beta = 6.8 \times 10^6$, $\omega_d/2\pi = 1.5$ MHz, $\chi_d = 427$, $\alpha = 1.1$, $\omega_s/2\pi = 4$ MHz, and $\chi_s = 1599$. This result is in agreement with the measurements of Ref. [11].

Although the model does not include eddy current effects, the description is satisfactory. Note that the magnetic resonances occur in the several MHz regime, and that the high frequency behavior (above 10 MHz) is relatively smooth. The high frequency cut off would be determined mainly by the transformer secondary's capacitance, limiting the signal spectrum to RF.

2.2. *Mechanical Setup and Electrical Ground Isolation*

The mechanical design was based on an existing monitor. The "donut shaped" ferrite has a thickness of 12.7 mm, and inner and outer diameters of 38.85 mm and 73.65 mm, respectively. An 8-turn "secondary" for signal pick-up was wound around the toroid. (An 1-turn winding for diagnostic purposes was also installed, but it was not used in this work.) The toroidal transformer was housed in a stainless steel casing. The assembly was mounted in a gap in the vacuum-tight beam line with standard $2\frac{3}{4}''$ flanges.

Electrically, the two most critical terminals were the input at the front-end amplifier, and the related "ground" or common point. The beam pipe, to which the monitor housing was connected electrically, was a potential source of ground noise current. The front-end electronics were mounted close to the beam line to minimize the electrical potential difference between these two points and noise pickup. The secondary winding was terminated in a balanced "Twinax" connector on the housing. It was connected to the front-end electronics box via two wires wound in opposite directions on a plastic support, forming a shielded twisted-pair of conductors. These were enclosed in a rigid 74 cm long aluminum tube, and terminated in two separate single-contact Lemo connectors on the box. The tube provided stable mechanical, as well as low-impedance electrical connection between the toroid and the electronics box. The box was located away from the beam line to minimize exposure to radiation associated with lost beam particles, and was further shielded in a lead-brick housing during beam operations.

Additional control input and signal output cables connections were made using single contact Lemo adapters. To reject common mode electrical noise in the frequency band of interest, each of these Lemo patch cables was wound 8 times through a ferrite toroid identical to those used in the transformer. This provided effectively hundreds of Ohms of reactance and essentially an open circuit for common mode noise. (It was important to use the same ferrites; those with the "wrong" magnetic properties proved ineffective.)

The AC to DC power supply assembly was housed in a separate, individually shielded and isolated box. The secondary voltages of the power supply were "floating" inside

the box, and were connected only at the front-end amplifier. The cable that provided this Ohmic connection was insulated and was also wound 8 times around a ferrite to reject common-mode pick-up.

The ground connections between the 117 VAC and the beam line was susceptible to low frequency noise. A resistor was installed at the appropriate place inside the electronics box to effectively reduce low frequency ground currents.

2.3. *Readout Electronics*

The current pulse induced in the secondary coil by the electron beam was first sent to a charge sensitive amplifier. A band-pass filter was used to improve the signal to noise ratio. The low cutoff was at several hundred kHz and the high cutoff at several MHz. The low cutoff removed noise at a resonant peak caused by the combination of transformer winding inductance and the capacitance at the input summing junction. The filtered signal was boosted by conventional op-amps and delivered to a remote data acquisition system approximately 100 ft away. The long-haul cable was the coaxial, low-attenuation type suitable for RF applications. The signal was digitized by an 11-bit ADC [2], and read out by a PC.

A dual-gain design was implemented to achieve the required accuracy for a range of beam intensity values spanning 4 orders of magnitude. A low-gain mode was used for high-current experiments, with beam intensities nominally at 10^9 electrons per pulse. For low beam intensities, nominally at 2×10^7 electrons per pulse, the readout was switched to the high-gain mode, with an approximately 60 times higher gain. The gains were adjusted such that, given the dynamic range of the ADC, one count represented less than 0.5% of the nominal charge value.

In the high-gain mode, the first-stage amplifier consisted of a low noise FET followed by a low noise, high sensitivity charge amplifier [3]. This configuration worked extremely well in the low charge environment. In the low-gain mode, the secondary winding was connected directly to the input summing junction of a low-cost, wide bandwidth amplifier [4]. This was wired in a charge sensitive configuration using a simple current feedback capacitor. The device gave satisfactory results, although the presence of parasitic inductance at the summing junction caused a 30 MHz "ringing" in the output. This can be mitigated in future implementations by the addition of a small series resistor.

3. Absolute Calibration

Accurate absolute calibrations of beam charge monitors in previous works have relied on the Faraday cup. Absolute charge accuracies of 0.2% have been achieved for low-intensity electron beams at 20 MeV [13]. At higher energies, the Faraday cups must be sufficiently massive to absorb the beam, and careful analysis of systematic effects associated with shower containment must be made. A one-meter long device was used to calibrate a wall-current charge monitor for beams up to 1.5 GeV[14]. The accuracy was determined to be 2% in that experiment; the Faraday cup contributed 1% to the systematic uncertainty.

The FLASH toroid was calibrated using charge injection methods. The charge injection capacitor and voltage source were measured against NIST standards. Detailed test measurements, analysis, and numerical modeling were carried out to obtain reliable estimates of the systematic uncertainties. Long term stability was monitored using data taken with a dedicated trigger timed in-between beam pulses.

3.1. *Inductive Charge Injection*

This method was employed to mimic the current induction process occurring in the beam environment. A calibration current pulse was injected onto an impedance-matched coaxial conductor inserted through the center of the toroid. Calibration was done after the monitor assembly was removed from the beam line; the same cabling plant as the experiment was used. A schematic layout is shown in Figure 3.

The stimulus generator consisted of a precision power supply providing a step voltage. It charged up the injection capacitor and produced the calibration pulse. The voltage source and the capacitor were calibrated to better than 0.1% against NIST standards. The rise time of the current pulse was approximately 2 ns.

This system provided accurate calibrations for the high-gain channel. A range of charge values could be generated by varying the values of the capacitor. For the low-gain channel, however, the 30 MHz "ringing" was made worse by the presence of large capacitors. Its accuracy was limited. A small resistor (5 to 20 Ω) was inserted in series to remove the "ringing" and investigate the associated systematic error.

3.2. *Direct Charge Injection*

Large dynamic range charge injection was achieved by directly injecting pulses at the summing junction of the first stage amplifier. A standard pulse generator [5] was used to trigger a fast CMOS step generator. It was designed to produce a stable 5 volt amplitude level with 2 ns rise time. A set of broadband (DC to 1 GHz) precision attenuators [6] was used to vary the magnitude of the test charge. This voltage was delivered to a 10 pF capacitor using a 50 Ω coaxial cable. Stable and clean output signals were observed for both gain settings. Since the amplifier input was at virtual ground, the current pulse discharged directly into

[2] LeCroy, Chestnut Ridge, NY, model 2249W
[3] Amptek, Bedford, MA, Preamplifier Model A250
[4] Analog Devices, Norwood, MA, model AD8056
[5] Systron-Donner, Model 101
[6] HP, Model 355 C (10-dB steps) and Model 355D (1-dB steps)

the amplifier; leakage onto the toroid branch was negligible. A schematic layout is shown in Figure 4.

The voltage and capacitor values were stable to better than 0.1%. This detailed data set allowed cross-checking and extension of the calibration obtained using the inductive method.

4. Systematic Uncertainties

4.1. *Calibration Charge Scale Accuracy*

The calibration charge accuracy was determined by the accuracies in the stimulus voltage and the injection capacitor. In addition to the DC voltage level [7], the waveform at the rising edge (within 50 ns) was also examined for distortions. It was found to be free of artifacts and stable at the 0.1% level using an "Infinium" 1.5 GHz oscilloscope. The performance of this instrument was verified against a Mercury relay source.

The dipped silver mica capacitors were compared against a three-terminal standard air capacitor [8] using a precision capacitance bridge [9] The capacitances, in the range of 50 pF to 4000 pF, were measured with 0.1% accuracy. They were constant in the frequency range of interest.

4.2. *Impulse Response*

For this work, the toroid's charge scale was calibrated using pulses 2 ns wide. It was to be applied to beam pulses 2 ps wide. A systematic error would arise if the impulse response of the charge monitor were sensitive to either the shape or the bandwidth of the signal. An estimate of this error was determined by analysis and modeling of the toroid and the front end electronics system.

An equivalent circuit of the toroid is shown in Figure 5. The source is the beam current derated by the turns ratio. The secondary is represented by an inductor L in parallel with a capacitor C_s. R_2 is a small secondary coil resistance. It is placed in series with L_2, representing the leakage inductance caused by imperfect coupling of the windings. Finally, the load, at the preamplifier input, is represented by the capacitance C_o across the twisted-pair connecting the transformer and the front-end electronics, and a load resistance R_o.

The trans-impedance, $v_o/(i_b/N)$, is given by

$$Z = \frac{Z_L Z_o}{Z_L + Z_o + Z_2}, \tag{2}$$

where $Z_L = j\omega L(\omega)/(1 - \omega^2 L(\omega) C_s)$, $Z_2 = R_2 + j\omega L_2$, and $Z_o = R_o/(1 + j\omega R_o C_o)$; $L(\omega) = L_0 \mu(\omega)$ is the toroid's inductance. Its frequency dependence is contained in the permeability of the ferrite described in Section 2.1.

At low frequency, the inductance is determined by the ferrite's permeability (approximately constant at 2000), $L_0 =$ 1.6 mH. The transformer capacitance value was determined by measuring the "self-resonance" frequency, $C_s = 51$ pF. The secondary coil resistance was small, $R_2 \sim 0.1\ \Omega$. The leakage inductance was given by a "self-inductance" coupling constant $k = 0.25$, $L_2 = (1 - k)^2 L(\omega)$. C_o was measured to be approximately 30 pF. Because of the current feedback configuration, R_o was small ($\sim 1\ \Omega$).

The toroid's trans-impedance is shown in Figure 6. At low frequency, it shows a roll-off typical of an LR circuit. The high frequency characteristics are determined by the ferrite core's frequency dispersion and the transformer's shunt capacitance. In the low-gain configuration, the appearance of a small parasitic inductance ($\sim 0.6\ \mu$H) at the summing junction caused an additional resonance at 30 MHz. In the mid-band, the trans-impedance is given by the load.

The impulse response of the long-haul coaxial cable can be described by $f_{cx}(\omega) = e^{-(1+j)\sqrt{\omega/\omega_e}}$, where the attenuation frequency $\omega_e/2\pi$ is in the 1-10 GHz range [15]. The coaxial cable turned out to have negligible effect in this analysis.

The trans-impedance function was convoluted with the CR-CR-RC bandpass filter and the coaxial cable's response, and the toroid's impulse response determined. The output signal in time-domain was a fast bipolar pulse.

4.2.1. *Pulse Shape Effect*

The average pulse shape measured in the high-gain mode is shown in Figure 7 for beam charges at 10^7 electrons, together with the charge-injection calibration pulse. Also shown are the simulated pulses, using response functions described in Section 4.2 and assuming Gaussian-shaped current pulses of two different widths. Similar results for the low-gain mode are shown in Figure 8.

While the simulated pulse shapes describe the data quite well, in particular, within the negative impulse signal portion of the bipolar pulse, it is difficult to account for detailed features. To determine the systematic uncertainty associated with pulse shape distortions, scale factors were calculated for a range of simulated pulses. Nominal as well as two extreme pulses were simulated to account for the variety of pulse shapes. The error was taken as the largest difference between the scales derived using the nominal 2 ps pulse compared with any 2 ns pulse, and the nominal 2 ns pulse compared with any 2 ps pulse. It was determined to be $\pm 0.1\%$ and $\pm 0.4\%$ for high-gain and low-gain channels, respectively. This uncertainty included errors in using 2-ns pulses to calibrate a 2-ps beam, as well as imperfect impulse response modeling.

[7] Measured with an Agilent 34401A multimeter, serial MY41046634, calibration traceable to NIST.
[8] General Radio Type 1403-A, serial 3562; calibration by IET Labs Inc. traceable to NIST.
[9] General Radio Type 716C, serial 4695, customized.

4.2.2. *Pulse Width Effect*

Although significant dependence of calibration on pulse width has been observed by other workers [14], the apparent pulse width insensitivity shown in the previous section can be understood. Pulse bandwidth sensitivity depends strongly on the design parameters of the transformer and the readout electronics. The beam charge can, in principle, be determined without any frequency information. In Fourier space, normalization of the signal pulse is determined by the zero-frequency component. In practice, however, a lower frequency cutoff is needed to reduce noise, and an upper cutoff is used to increase the signal to noise ratio.

Pulse width dependence could be caused by distortion of the pulse shape caused by the limited measurement bandwidth. For this monitor, the response of the ferrite core transformer rolled off at 10's of MHz. The response of the front-end electronics was narrower, cutting off at several MHz. This was well within the range where the signal spectrum changed little for both the calibration and the beam pulses. There was no significant error for calibration pulses narrower than 20 ns. For longer pulses, the signal spectrum would be changing significantly within the toroid's bandwidth, and strong pulse width dependence would be expected according to our calculations.

4.3. *Non-linear Response of the 2249W ADC*

The integration gate width of the 12-channel 2249W ADC was set at 1016 ns to accommodate a variety of signals. The digitization count was linearly proportional to the charge contained in the negative impulse portion of the bipolar signal. The wide ADC gate also included the positive "over-shoot" segment. The response of the 2249W ADC to positive signals turned out to be nonlinear, and dependent on the signal strength. The charge injection setup described in Section 3.2 was used to derive a correction function for this effect.

The ADC response as a function of gate width was mapped out for a range of charges. The value obtained using the nominal gate timing was compared to that obtained when the integration ended at zero-crossing time. The nonlinear response caused the ADC to output a value lower than the true signal. The amount of negative offset initially increased with beam charge; it flattened with a diode-like cutoff at 19-20 counts for the rest of the beam charge range of interest. This effect was also measured for a fixed charge, as the ADC gate width was increased by adjusting the gate closing timing. It was proportional to the gate width for a large portion of the "over-shoot" signal, as the closing gate was adjusted away from zero-crossing, indicating that it was insensitive to the pulse shape there.

The correction function was checked using a variety of beam data, correlating the toroid's response with photomultiplier counters monitoring the background rate and RF antennas picking up electromagnetic radiation emitted by the beam as it traversed an upstream air-gap. The residual offset after correction was found to be small, and included as an additional calibration factor.

4.4. *Long Term Stability*

The calibration was carried out in August 2004, after FLASH was completed. The first thin-target run took place in September 2003. This was followed by two more runs: a thick-target run in June 2004 and a combined thin and thick target run in July 2004. Constant monitoring of the toroid was needed to ensure the integrity of its data and to determine its stability over time.

A direct charge injection mechanism was built into the front-end electronics box. It applied a stable voltage to a 2 pF capacitor when triggered by a remotely generated control signal, timed in-between 10 Hz beam pulses. These data were available throughout the life time of the experiment. A variation in the measured reference charge was taken as a systematic error.

5. Results

5.1. *Resolution*

The toroid's resolution was measured by the spread in the ADC counts in the calibration charge injection data. For the high-gain channel, it was approximately 1 count, corresponding to 1% at 10^7 electrons. For the low-gain channel, it was approximately 0.5 counts, corresponding to 0.2% at 10^9 electrons. In both cases, the resolution scaled inversely with beam charge.

5.2. *Long Term Stability*

The toroid readouts of a reference charge are plotted in Figure 9 as a function of time. The RMS spread of all low-gain data was 1.8%. Excluding the initial test phase of the Sep. '03 run, which was used to tune the beam to reduce large beam-induced backgrounds, this reduced to 1.6%. Further restricting the analysis to an individual run period, within the Sep. '03 or June-July '04 data set, the spread was 1.3%. For the high-gain data, the spread was less than 0.5 ADC counts, responding to an error of less than 0.1%.

The change between runs could have been caused by drifts in the ADC system over the years, perhaps due to temperature sensitivity – an 0.3 counts drift would correspond to a 1.5% shift in the charge readout. The monitor system was dismounted from the beam line between runs, and that could have contributed to the variation. The two-band structure, seen in the Sep. '03 data set, was most likely caused by a timing shift in the two types of alternating triggers used. (They were shared with LED triggers for phototube monitoring.)

Parameter	High-Gain	Low-Gain
p_0	$(1.19 \pm 4.44) \times 10^{-2}$	0.08 ± 1.81
p_1	$(8.04 \pm 0.16) \times 10^{-3}$	0.345 ± 0.035
p_2	–	$(-2.28 \pm 8.21) \times 10^{-5}$

Table 1
Calibration curve fit parameters. Toroid charge in units of 10^7 electrons was fitted to the function: $N(1E7) = p_0 + p_1 \cdot ADC + p_2 \cdot ADC^2$.

5.3. *Calibration Curve Fitting*

The toroid's absolute charge scale was determined using two independent charge injection data samples. The accurately calibrated, but sparse, inductive charge injection data for the high-gain channel provided the reference scale for the more detailed calibration data set obtained by direct charge injection at the preamplifier summing junction. The inductive charge calibration data points were interpolated, and a difference was taken between the two calibrations after multiplying the direct-charge-injection charges by a scale factor. This relative scale factor was determined to be 0.974 ± 0.011 by minimizing the summed-squared differences. It applies to both the high-gain and low-gain data sets.

The calibration charge values were plotted against the average readout ADC counts, and the correlation fitted to polynomials of first and second orders. The MINUIT package [16] was used to maximize a binned-likelihood function. For the high-gain channel, the linear fit was sufficient, as shown in Figure 10. For the low-gain channel, the quadratic fit gave better-behaved residuals. In both cases, the fit residuals were within $\pm 1\%$ over the charge range of interest. The errors included correlations among the fit parameters, and the corresponding probability content within the one-standard-deviation likelihood contour was 68.3%. The results are summarized in Table 1. To check for systematics, the range of charges used in the fit was varied. The fitted parameter values were found to be stable. χ^2 fits gave the same values; the errors in this case would depend on normalization and were not used.

5.4. *Discussions*

A summary of the systematic uncertainties is given in Table 2. The largest contribution comes from the errors in the calibration curve fitting. This was mainly due to the limited calibration data set. The inductive charge injection for the low-gain channel has not been used in this analysis so far due to concern for the effect of "ringing" described in previous sections. The calibrations determined using this data sample, allowing the value of the series resistor to vary between 0 and 20 Ω, showed a spread of 3%. It agrees with the result shown in Table 1 when all uncertainties are taken into account.

	High-Gain	Low-Gain
Charge Scale:		-
- Inductive	0.1%	-
- Direct	1.1%	1.1%
Impulse Response	0.1%	0.4%
Stability:		
- Summer '04	0.1%	1.3%
- Fall '03	-	1.3%
- Fall '03 through Summer '04	-	1.6%
Fit Errors	2.1%	1.9%
	(at 2×10^7)	(at 10^9)
Total	2.5%	2.7%

Table 2
Summary of systematic errors quoted at nominal beam charges: 2×10^7 electrons for high-gain channel and 10^9 electrons for low-gain channel. Total is the quadrature sum of all uncertainties. (For low-gain, the 1.6% value for long-term stability was used in calculating the total.)

6. Summary

The toroidal beam charge monitor for the FLASH experiment showed good resolution over a large dynamic range. The long term stability was at the 1.6% level. Absolute calibration was obtained using economical charge injection methods, with accuracies at the 2.5% level at nominal beam charges. These performance figures may be improved in future implementations.

Acknowledgements

We appreciate the help from many colleagues at SLAC and in the FLASH collaboration. In particular, P. Chen, C. Field, K. Reil, S. Smith, P. Sokolsky, and R. Zdarko provided strong support for our work.

References

[1] R. Abbasi *et al.*, Phys. Rev. Lett. **92**, 151101 (2004).
[2] M. Takeda *et al.*, Astrophys. J. **522**, 225 (1999).
[3] J. Belz *et al.*, SLAC Exp. Proposal E-165 (2002).
[4] J. Belz *et al.*, Astropart. Phys. **25**, 57 (2006).
[5] J. Belz *et al.*, Astropart. Phys. **25**, 129 (2006).
[6] R. Talman, AIP Conf. Proc. 212, Accelerator Instr. Workshop, Upton, NY, 1989.
[7] R.S. Larsen, SLAC Report No. 63, May 1966.
[8] Bergoz Instrumentation (Saint Genis Pouilly, France); K.B. Unser, Proc. of IEEE Particle Acc. Conf., vol. 1, p.71, 1989.
[9] J. Hinkson, "ALS Beam Instrumentation", July 1998, (unpublished).
[10] G. T. Rado, Rev. Mod. Phys., **25**, 81 (1953).
[11] T. Tsutaoka, J. Appl. Phys., **93**, No. 5, 2789 (2003).

[12] Based on Eq. 1 of Ref. [11]; see also *Magnetism*, edited by É. du Trémolet de Lacheisserie *et al.*, (Kluwer Academic Publishers, 2003), Vol. 2, Chapter 16.
[13] J.O. Trier and K. Derikum, Proc. of Workshop on Diagnostics and Instr. for Part. Accelerators, Montrex, Switzland, 1993; CERN PS/93-35.
[14] T. Suwada *et al.*, Jpn. J. Appl. Phys., Vol. 39, 628(2000).
[15] S. R. Smith, in AIP Conf. Proc. No. 390, Beam Instrumentation Workshop, Argonne, Il, May 1996; also see R. L. Wigington and N. S. Nahman, in Proc. of the IRE, p. 166, Feb. 1957.
[16] F. James and M. Roos, CERN D506, Geneva 1994 (unpublished).

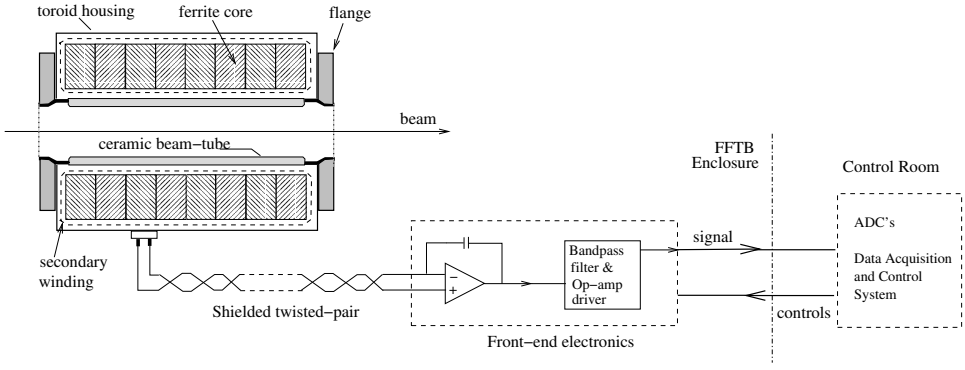

Fig. 1. Block diagram of the toroid system. A schematic center-cut view of the toroid is shown. The key components are indicated.

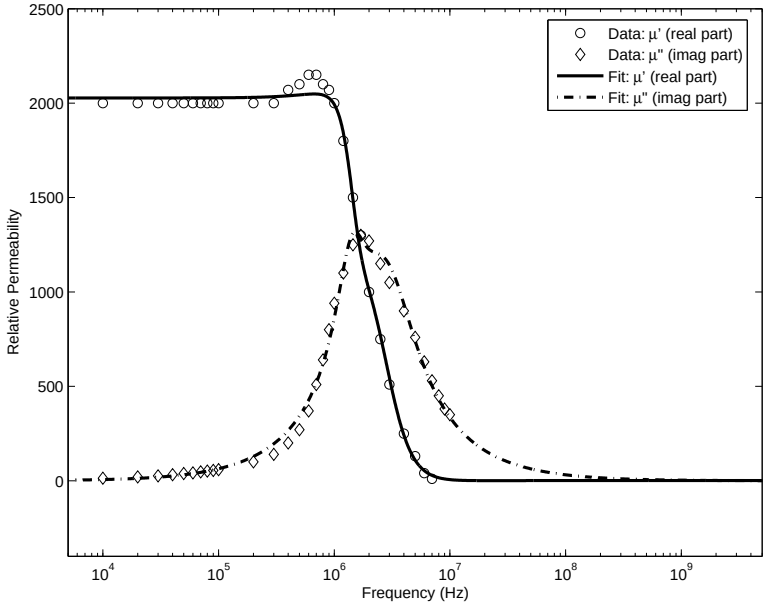

Fig. 2. Frequency dispersion of complex relative permeability for the ferrite. Points were taken from the manufacturer's data sheet. Curves are fits to the model described in the text.

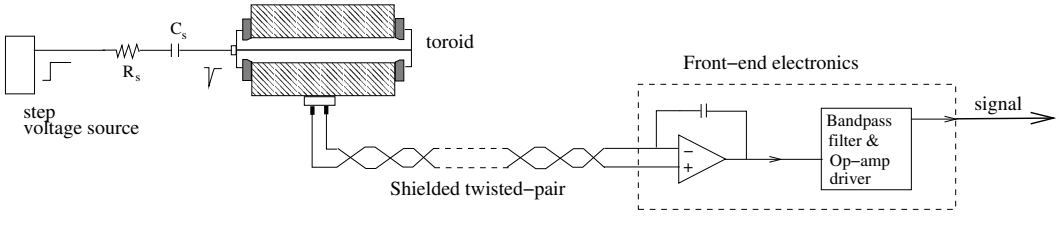

Fig. 3. Schematic layout of the inductive charge injection calibration setup. See text for more discussion.

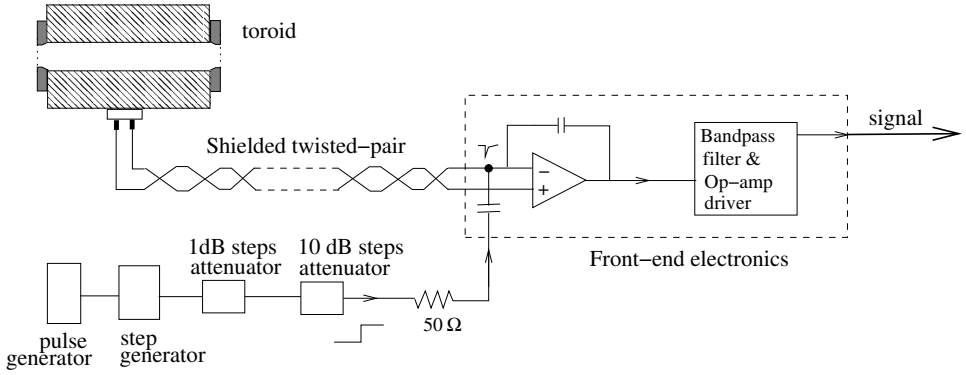

Fig. 4. Schematic layout of the direct charge injection calibration setup. See text for more discussion.

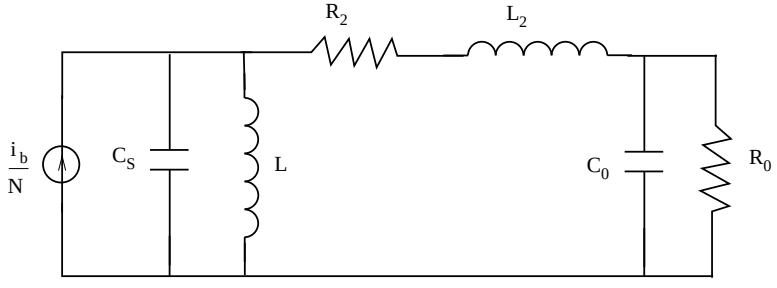

Fig. 5. Equivalent circuit model of the toroid.

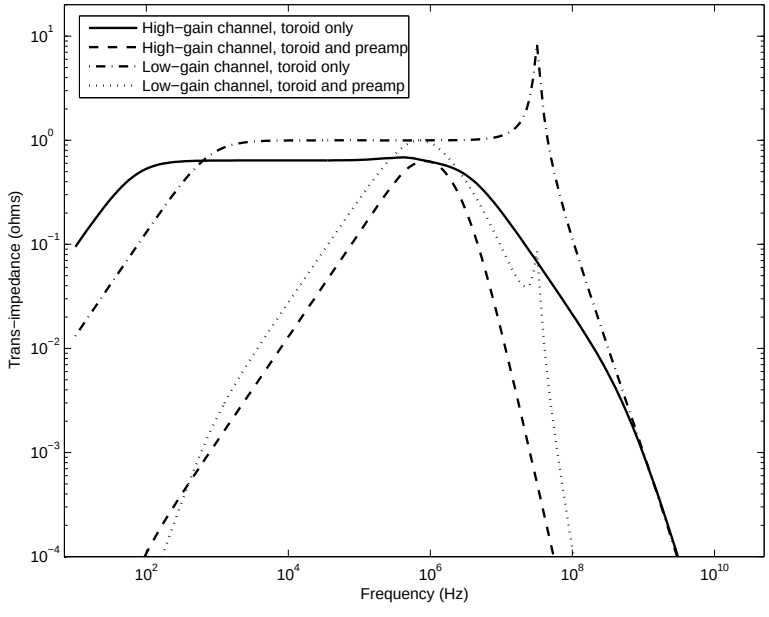

Fig. 6. Toroid trans-impedance for both gain settings calculated from the equivalent circuit.

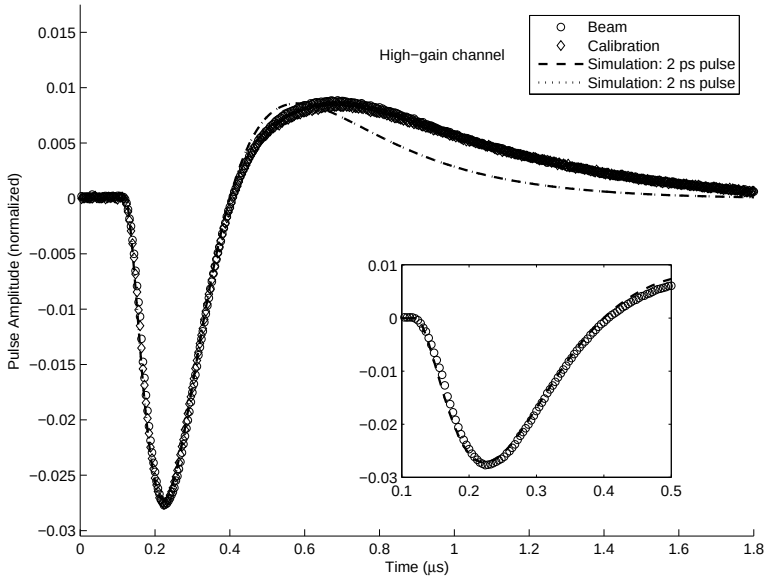

Fig. 7. Toroid impulse sampled by a Tektronix TDS380 for the high-gain channel. Beam and calibration data points are shown. The curves represent simulated responses assuming Gaussian pulses of 2 ps and 2 ns widths.

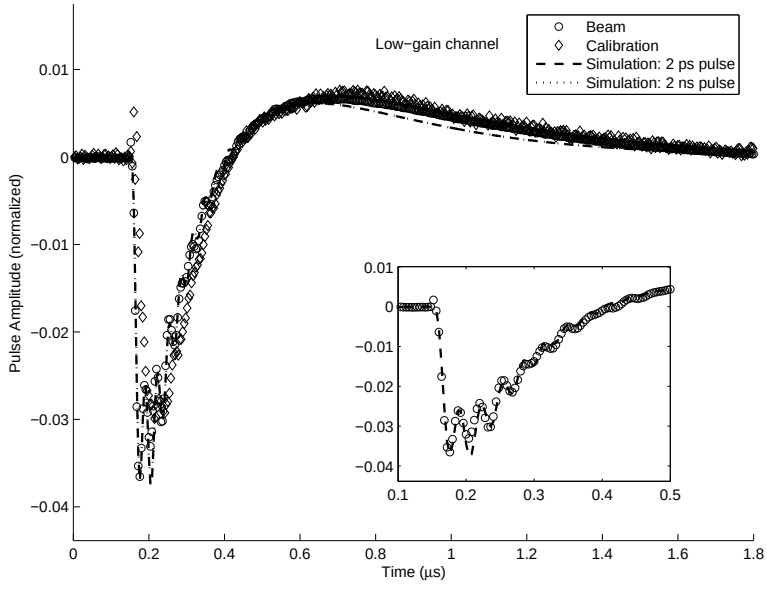

Fig. 8. Same as Figure 7 for the low-gain channel.

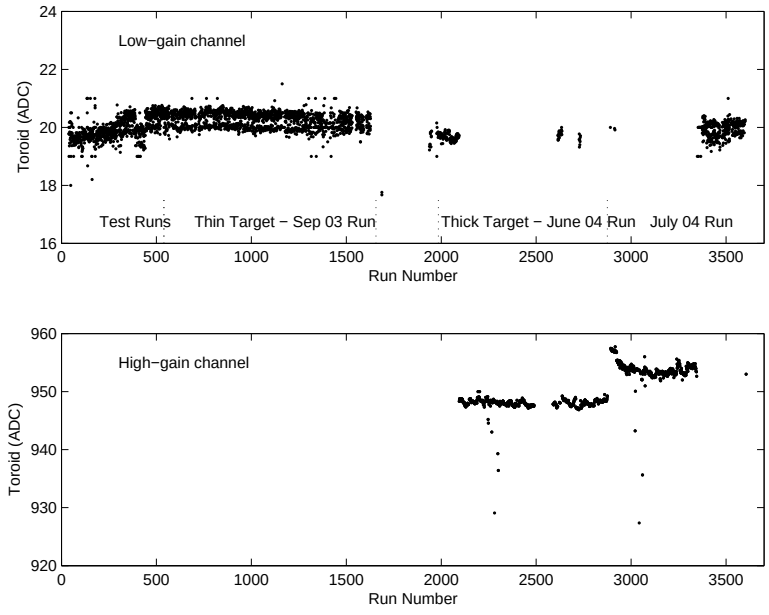

Fig. 9. Toroid stability over the lifetime of the FLASH experiment. (Note the zero-suppressed scale.)

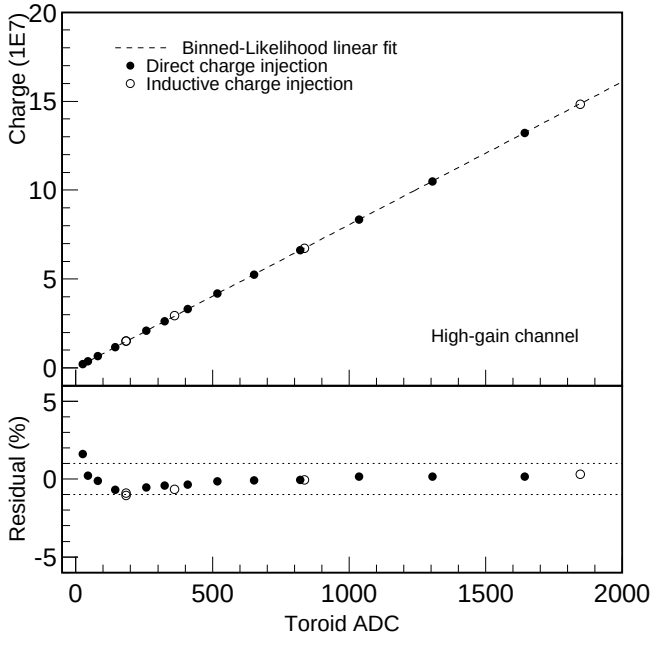

Fig. 10. High-gain channel calibration. A linear fit was made to direct charge injection data; inductive-charge-injection charges are shown for comparison. A ±1% range is indicated by the dotted lines in the residual plot.

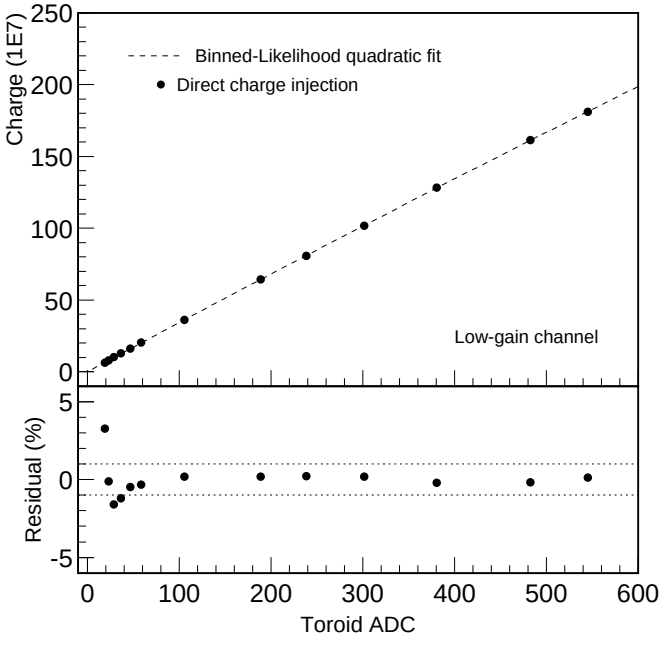

Fig. 11. Low-gain channel calibration. A quadratic fit was made to direct charge injection data. A ±1% range is indicated by the dotted lines in the residual plot.