Towards the Development of an Integrated Testing Method for the Front End Board Electronics

Nevis Laboratories, Columbia University
Summer 2002 Research for Undergraduates Program

I. Stoica, J. Vigil

ABSTRACT

The Large Hadron Collider (LHC) will use new advancements in technology that will enable it to achieve unprecedented energies in order to probe the Standard Model and beyond. The front-end board (FEB) is the principal readout electronics for the liquid argon calorimeter of the ATLAS experiment that will be conducted at the LHC. In this paper the testing methods for the front-end board are discussed. An analysis of the performance of the FEB as well as possible improvements is presented.

TABLE OF CONTENTS

Background

Our current understanding of the physics governing the behavior of the universe at the tiniest scales is explained by the Standard Model of Particle Physics. Developed in the 1960's and 70's and formalized by Sheldon Glashow, Abdus Salam, and Steven Weinberg, the Standard Model attempts to explain the interactions of matter at very high energies that last existed only fractions of a second after the Big Bang. The Standard Model explains the interactions of twelve fundamental particles of matter called fermions (six quarks and six leptons) via their exchange of five force carrying particles called bosons. The Standard Model has undergone thorough testing at accelerators around the world and has thus far succeeded in explaining nearly all pheonomena observed in high energy physics experiments.

But despite its incredible success, the Standard Model is incomplete. It calls for 19 free parameters, all of which must be determined experimentally, and it fails to explain the gravitational interactions of matter. Moreover, the Standard Model requires the introduction of a sixth boson to explain the origin of mass and remain a closed a theory. This sixth boson, named for its postulator the Scottish physicist Peter Higgs who proposed its existence in the late 1960's, is necessary to complete the Standard Model. If it is determined that the Higgs boson does not exist, physicists will be forced to rethink their view of the universe. They will have to question the Standard Model and look for new theories. On the other hand, if it is determined that the Higgs Boson is indeed real the Standard Model will be more complete. Therefore, the search for the Higgs boson presents the ultimate triumph for the Standard model.

The Standard Model predicts that the Higgs boson is a massive particle, and by using measured values of some of the 19 free parameters in the model we are able to determine the phasespace in which we are most likely to find the Higgs boson. This phasespace is believed to extend from 90 GeV to 250 GeV with the most likely mass being about 113 GeV. A substantial portion of this phasespace has been covered with previous accelerator experiments. The four experiments at the Large Electron Positron (LEP) Collider at CERN have thoroughly explored the lower portion of this phasespace. Full analyses of the runs taken at CERN have shown that if the Higgs boson exists, its mass must be greater than 120 GeV. The collaborations at Fermilab plan future runs that will produce more data to attain 5-sigma certainty in the current results and extend the covered phasespace beyond 120 GeV. These runs will, however, take several years to complete and they will not eliminate all of the Higgs likely phasespace. Therefore, physicists around the world have agreed that a new higher-energy, higher-luminosity collider must be built. Thus, they agreed to construct the Large Hadron Collider (LHC). The LHC is currently under construction at CERN in the same tunnel that housed the LEP and thus, the decision to begin construction of the LHC was also the decision to end all experiments at the LEP.

Unlike the Tevatron, the LHC will be a proton-proton (p-p) collider and not a proton-antiproton (p-pbar) collider. Whereas the Tevatron can achieve a center of mass energy of 2 TeV, the LHC will achieve an amazing 14 TeV center of mass energy at a bunch crossing frequency of 40 MHz. The high CM energy coupled with the high bunch crossing frequency of the collider will permit experimenters to use it to search for a Higgs Boson having a rest mass as high as 1 TeV with 5-sigma certainty with relatively short runs when compared to those of the experiments at the Tevtron. As many physicists have said, if the Higgs boson exists and has a mass anywhere near what we expect, we will find it at the LHC.

Four experiments will be conducted at the LHC: A Large Ion Collider Experiment (ALICE), which will be used for heavy ion studies; LHC-b, which will be used for studies in B-physics; the Compact Muon Solenoid (CMS), a general purpose detector with a high-performance muon system; and A Toroidal LHC Apparatus (ATLAS), another general purpose detector designed to cover all physics that may lie on the horizon. All four of these experiments will be operated by some of the largest most international collaborations ever assembled. The ATLAS collaboration, for example, is already comprised of more than 1750 physicists and engineers from 144 institutions.

In addition to the need for a Higgs boson, there are many other problems with the Standard Model that have yet to be solved. Charge-Parity (CP) Violation, which may provide an explanation for the predominance of matter over antimatter in the universe, has yet to be understood fully in terms of the Standard Model. The quark coupling constants in the Cabibo-Kobayashi-Maskawa (CKM) matrix need to be measured more precisely in order to obtain a more coherent theory of quark mixing. Many physicists today believe each elementary particle in the Standard Model has a partner called its superpartner. This idea, known as Supersymmetry (SUSY), postulates that every fermion has a corresponding boson and every boson has a corresponding fermion. According to SUSY many of these superpartners should have masses within the observable range of the ATLAS detector, leaving open the possibility that if SUSY is real, we will discover it with ATLAS at the LHC.

Elementary particle physics is currently at a crossroads. The Standard Model may prove correct or we will learn that it is false and a whole new era in physics will begin. Whatever the outcome, this is an exciting period, for the physics that will be produced at the LHC and recorded by ATLAS will certainly bring us closer to understanding our universe.

Introduction

The LHC consists of four experiments, one of which is A Toroidal LHC Apparatus (ATLAS), a cylindrical colliding beam experiment. ATLAS is comprised of four core (fundamental?) parts: (1) the magnet system, that bends charged particles; (2) the inner track, that measures the momentum of charged particles that have been bent by the magnet system; (3) the calorimeters (electromagnetic (EM) and hadronic), that measure the energies of all the particles excluding muons (EM- detects and measures electromagnetic showers, hadronic- detects and measures hadronic showers (LAr)); (4) the muon detector, that identifies muons (Fig. 1).

Figure 1:
ATLAS Experiment

(1) Magnet System (gray)
(2) Inner Tracker (yellow)
(3) Calorimeter (orange/green)
(4) Muon Spectrometer (blue)

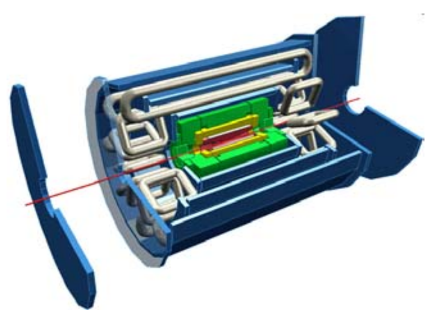

ATLAS has two calorimeters: an electronic calorimeter that detects and measures energies of electromagnetic showers as well as the initial stages of hadronic showers, and the hadronic calorimeter that detects and measures the energies of hadronic showers (hadrons are composite particles made up of a combination of 3 different colored quarks, baryons, or a quark and an anti-quark, mesons). ATLAS uses a method known as liquid argon calorimetry to identify and gauge hadrons. The collaboration at Columbia University is involved in the design, prototyping, and production of the principal readout electronics for the Liquid Argon (LaR) Calorimeter, the front-end boards (FEB) (Fig. 2).

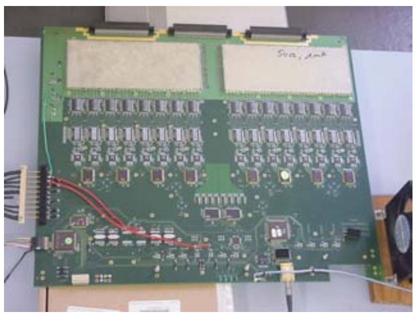

Figure 2:
Prototype of the FEB

There are 1524 front-end boards that process data from the liquid argon calorimeter; each board has 128 input channels for receiving and processing information. Data from the calorimeter first goes through a 4-channel preamplifier, then proceeds to a 3-gain shaper chip. This chip takes the incoming data, originally in the form of a triangular pulse, and uses a bipolar shaper function to shorten the pulse. The shaper chip further amplifies the resulting pulse using three different gain scales: high (x 100), medium (x 10), and low (x 1) gains. The shaper chip also sums the 4 channels originally received and sends this data to other electronics (not part of the FEB) that construct a first level trigger (L1) from this information. Next, the three amplified signals are sent to a Switched Capacitor Array (SCA) Chip that stores the information in analog form; the SCA samples the data received at 40MHz, thus storing 5 samples/event. The L1 trigger is received by the SCA, which proceeds to sort the desirable events from the rest. The events selected by the L1 trigger are digitized by a 12-bit Analog to Digital Converter (ADC) Chip, formatted, and then sent through optical fibers to the off-detector Read Out Driver (ROD).

The electronics must adhere to a stringent set of operational requirements; furthermore, due to the large number of total channels (over 190000), FEB electronics must minimize power consumption and cost, as well as being able to process data very quickly. The electronics must be highly reliable as the FEB boards will be situated in a limited access location; furthermore, they must be able to sustain up to 20Gy/yr in radiation. The FEB must be able to record and process all significant data; for this task, electrical noise must be minimized, especially coherent noise that cannot be extracted from the data stream. No time delays or dead time should be introduced by the electronics. Radiation tests have been done on various components of the FEB board; in July of 2002, the standard voltage regulators, the only remaining components on the FEB that were not radiation proof, were replaced with radiation-hard regulators. All the current FEB electronics can now withstand the specified radiation levels for an extended period of time.

This summer, we have designed and implemented algorithms that test the performance of the system. Some of the things we tested for include coherent and incoherent noise, and linearity of the data acquisition system.

Materials and Methods

At present, we are completing various hardware and software tests that will be used to completely characterize the performance of the FEB, as well as to understand how parameters such as temperature and electrical noise affect the quality of the data and how it is transmitted and processed. A calibration board is used to create and transmit pulse and noise runs to the FEB. Noise runs are used to determine the electrical noise of the system, while pulse runs are used to assess the filtering, shaping, sampling, and analysis capabilities of the FEB.

There are two basic types of data acquisition methods we have employed for the current research: dac scans, used to test the linearity of the system, and delay scans, used to measure the signal pulse shape.

A delay setting specifies the time between when the pulse is transmitted and when the first sample is recorded from the waveform. Only 5 samples, spaced 25 ns apart, from each received pulse are stored in the SCA chip; thus the measure of the pulse height using one fixed delay setting is not highly accurate. In order to achieve a better approximation of the pulse energy, it is necessary to sample the same waveform at different delay settings, thus obtaining as many samples as desired for a given pulse (specifically, 5×number of delay settings). This method of sampling the same pulse at different delay settings is what we refer to as the delay scan.

The delay scan is used to determine the precision of the pulse height measurement and the measurement error (the difference between the transmitted and received pulse), as well as the timing resolution. This information is vital for correct calibration of the FEB.

To measure the energy of a particle or of a shower, the data from several, sometimes hundreds or thousands of channels will be summed. For this reason we would like to ascertain that the channels act independently of each other, or with as little correlation as possible. We can measure the correlation between the channels by analyzing noise runs of a given FEB. White noise, or incoherent noise, is a random pattern of electrical signals that shows no correlation between the 128 channels. In practice, however, the 128 channels on an FEB are correlated due to fabrication and operational limitations such as the use of a single common power supply. Coherent noise reflects these correlations between the channels; we would like the coherent noise to be less than 5% of the total noise. Noise runs are useful in determining the percentages of each kind of noise; also, they are used to compute an auto correlation matrix (ACM). ACMs are calculated using an optimal filtering algorithm developed by W.E. Cleland and E.G. Stern[1]. The ACM shows the correlation between the five different samples in each event- the five samples are not completely independent due to the effects of the shaping function and possible presence of coherent noise.

The FEB is designed to detect, amplify, and analyze pulses of varying intensities, with a linear relationship between the initial waveform and the generated result. A dac scan allows us to test for the linearity of this system; a dac is a value that the calibration board uses to produce pulses of various heights (for example, a pulse with a given dac of

---

[1]Cleland, W., Stern, E. *Signal Processing Considerations for Liquid Ionization Calorimeters in a High Rate Environment*, Nuclear Instruments & Nuclear Methods in Physics Research A 338 (1994).

200 is lower in energy than a pulse with a dac value of 400). There is a practical threshold to the upper level of the dac value allowed, given by the point of saturation of the system. This value has been found experimentally to be about 250 dac for high gain data, 2500 dac for medium gain data, and 25000 dac for low gain data. This method of sending pulses of varying dac values (heights) to the FEB is what we refer to as the dac scan.

The DAC scan analysis consists of two primary steps. First, a data-taking run is completed in which the DAC value is progressively increased. For example, a DACscan run may involve 8000 events in which the first1000 events have a DAC value of 40, the second 1000 events have a DAC value of 70, the third 1000 events have a DAC value of 100 and so on. In such a run the DAC value would be incremented in equal steps of 30 resulting in a uniformly spaced set of data points that can be used to test for the linearity of the system. In addition to the DAC value, the relevant data in each event includes the pulse data for 128 channels. Since each pulse is sampled five times, each event consists of 5×128 numbers and the DAC value that are pertinent to the DAC scan analysis.

The energy of the pulse should be linear with the DAC value. However, since the pulse information obtained in the run comes in raw ADC counts, it is necessary to reconstruct the energy of each pulse using the raw ADC counts of the five pulse samples acquired in the run. This energy reconstruction is accomplished through the method of optimal filtering in which information acquired in a noise run is fed to a program implemented by Ioannis Katsanos and Jiamin Jin that generates five coefficients a_i such that the energy E of each pulse is given by the sum of the product of these five coefficients and the raw ADC count values S_i of the pulse, i.e. $E = \sum_{i=0}^{4} a_i S_i$.

The energies of the pulses for all 128 channels in each event are computed with the above formula. Since the energy behavior is likely to vary with channel number, the linearity of the energy as a function of DAC value must be checked for each channel separately.

The pulse energy of a specific channel with a given DAC value is taken to be the mean of the reconstructed pulse energies for all events for that channel with the given DAC value. For example, if 1000 events with a DAC value of 40 are recorded, then the pulse energy of a given channel, channel 34 say, is the mean of 1000 pulse energies of channel 34 reconstructed from 1000 events. The error in this mean can also be determined by computing the root-mean-square (RMS) of these 1000 pulse energies. A mean and RMS is computed for all DAC values taken in the run. The linearity of these data are then checked using a linear least squares analysis. For example, if eight different DAC values, $x_0,..,x_7$ are taken in a run, then eight mean pulse energies $y_0,..,y_7$ are computed for each channel. The best-fit line is determined by minimizing a vertically offset chi-squared, where chi-squared is defined as $\mathrm{X}^2 = \sum_i \frac{(y_i - (mx_i + b))^2}{\sigma_i^{\,2}}$, where m is the slope of the best fit line, b is the y-intercept of this line, and σ_i is the error in the mean pulse energy of a channel with a given DAC value, i.e. the quotient of the RMS and the square root of the number of events. The solution to this minimization problem is determined by solving the system $\frac{\partial \mathrm{X}^2}{\partial m} = 0$, and $\frac{\partial \mathrm{X}^2}{\partial b} = 0$ for the parameters m and b. An analytic solution to this system exists and it can be shown that the solution is

$$m = \frac{n\sum xy - \sum x \sum y}{n\sum x^2 - \left(\sum x\right)^2} \text{ and } b = \frac{\sum y \sum x^2 - \sum x \sum xy}{n\sum x^2 - \left(\sum x\right)^2}.$$

The above formulae permit the efficient computation of the best-fit line for the energy vs. DAC value for all 128 channels. If the distances of the data points (x_i, y_i) from the best-fit lines are well within the errors σ_i on the mean pulse energies y_i, then we can be confident that sufficient linearity in the pulse energy as a function of DAC value has been achieved.

The data taken in the DAC scan run can also be used to perform an energy resolution analysis of the FEB. A very high energy resolution is desired because the energies of thousands of channels must be summed in order to make the Level 1 trigger decision. A high energy resolution is required to make a reliable Level 1 trigger decision because otherwise the errors in the Level 1 trigger sums would propagate to significant errors in the energy measured in the Liquid Argon Calorimeter. The ATLAS collaboration at Nevis hopes to achieve an energy resolution better than 0.1%, although the current resolution stands at about 1%.

The first step in the energy resolution analysis is to compute the RMS of the pulse energies for every channel for each DAC value. These RMS values are then divided by the energy corresponding to the DAC value and plotted as a function of the energy. Since we expect the RMS of the pulse energies to be nearly constant with respect to the DAC value, we expect a plot of the quotient of the RMS of the energy and the energy versus energy to be a hyperbola and scale as 1/energy.

Discussion and Conclusions

Noise runs were used to calculate the total noise of the system, as well as to compute the fraction of coherent noise, a critical factor in determining FEB performance. The resolution in noise measurements is an indicator of the lower limit of the energy resolution of pulse runs.

The pedestal for a given channel is given by $\frac{1}{events}\sum_{n=1}^{events} ADC_n$ (summing only the first sample in each event). Pedestals are computed and are then used in resolution measurements as demonstrated below. Figure 3 shows a plot of the pedestals for a high gain noise run.

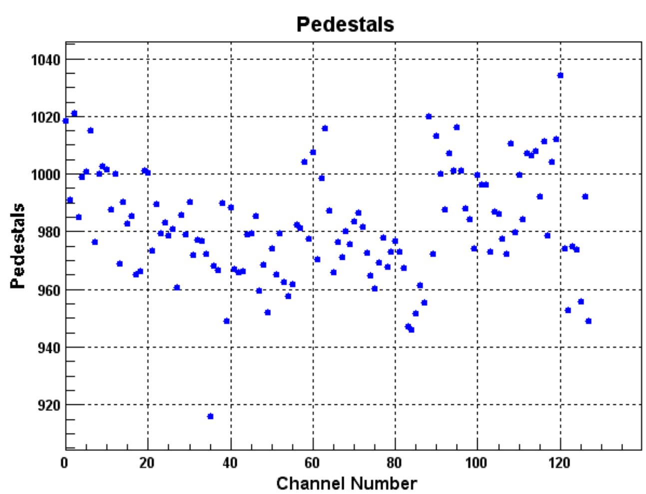

Figure 3: Pedestals for a High Gain Noise Run

The total noise of the system is defined as $\sigma_{total}{}^2 = \sigma_{white}{}^2 + \sigma_{incoherent}{}^2$. σ_{white} is obtained by assuming zero correlation between all the channels. We measure σ_{white} for each individual channel and add all 128 channel errors in quadrature to obtain the total incoherent noise: $\sigma_{white} = \sqrt{\sum_{n=0}^{127}\left(\sigma_{white(n)}{}^2\right)}$, where σ_{white} is

$$\sqrt{\frac{1}{(5\times events)}\sum_{i=0}^{(events-1)}\sum_{j=0}^{4}\left(ADC_{ij} - Pedestal\right)^2 - \left(\frac{1}{(5\times events)}\sum_{i=0}^{(events-1)}\sum_{j=0}^{4}\left(ADC_{ij} - Pedestal\right)\right)^2}.$$

σ_{total} is given by

$$\sqrt{\sum_{i=0}^{(events-1)}\left(\sum_{j=0}^{127}\sum_{k=0}^{4}\left(ADC_{ijk} - Pedestal_j\right)\right)^2 - \left(\sum_{i=0}^{(events-1)}\sum_{j=0}^{127}\sum_{k=0}^{4}\left(ADC_{ijk} - Pedestals_j\right)\right)^2}$$

Assuming that the coherent noise is spread evenly over all 128 channels, we can obtain a measure of the noise per channel: $\sigma_{channel}^{coherent} = \frac{1}{128}(\sigma_{coherent})$. We then use this to calculate the fraction of total noise that is coherent. Our results show a coherent noise fraction of .03586 which meets the FEB requirements that the fraction be at most 0.05. A summary of our noise results for high gain data is given in Table 1. A plot of the noise is shown in Figure 4.

σ_{white}	99.3086
σ_{total}	107.167
$\sigma_{coherent}$	40.2804
$\sigma_{channel}^{coherent}$	0.314691
$\sigma_{coherent}/\sigma_{total}$	0.0358593

Table 1: Summary of Noise Results for High Gain Noise Run

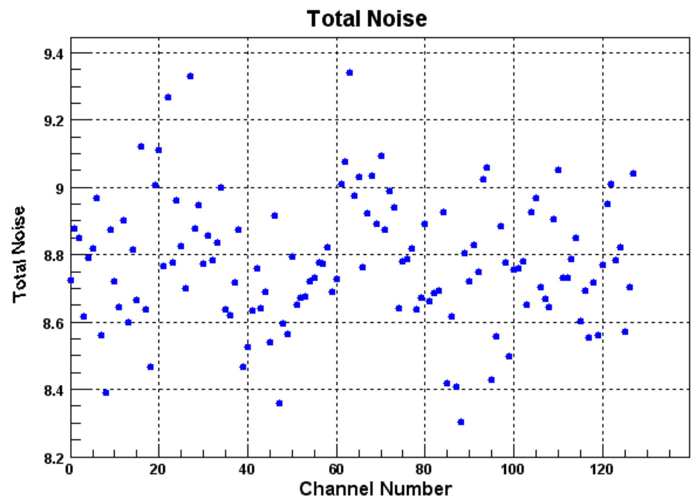

Figure 4: Total Noise for High Gain Noise Run

The results of the DAC scan analysis proved more difficult to obtain than expected. A preliminary analysis revealed that the linearity of the energy as a function of DAC value was better behaved at higher energies than at lower energies. This result was procured by making a plot of the differential nonlinearity, defined as the vertical distance of the data points (x_i, y_i) from the best-fit line through those points. We expect the differential nonlinearity to be greater at lower energies because the clock feedthrough, which is independent of the DAC value, is greatest in proportion to the energy at lower energies. This expectation was confirmed by the differential nonlinearity plot. The differential nonlinearity at low energies was so great that the decision to construct the best-fit line using only the last four data points instead of all eight was made. An increase in the differential nonlinearity at low energies was anticipated as a result of this decision. However, the use of only the last four data points gives a better picture of the linearity at this time because the clock feedthrough has not yet been subtracted in the analysis.

The energy versus DAC value plot for our favorite channel, channel 34, shown in figure 5 demonstrates good linearity within the resolution of the plot. The last four data points used to construct the best-fit line all appear to fall on the line as well as the third and fourth points, which were not used in the line's construction. The first and second data points, however, can be discerned to fall a small distance from the line of best fit, albeit it is difficult to tell exactly how far they are from the line.

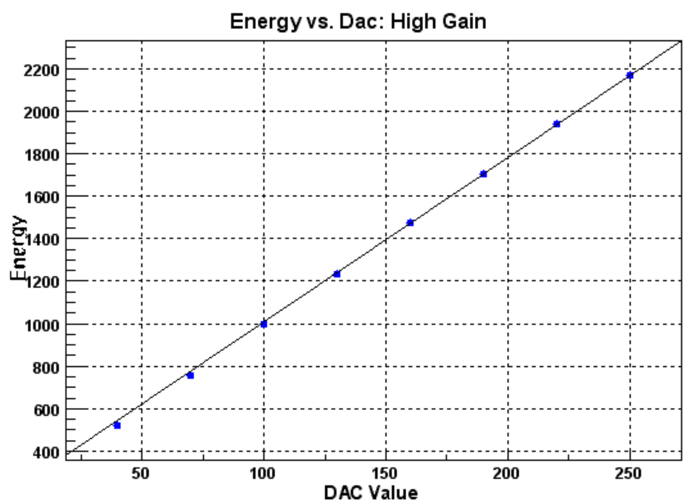

Figure 5: Energy vs. DAC plot for Channel 34 for High Gain DAC Scan

The differential nonlinearity (figure 6) plot reveals a substantial nonlinearity at low energies. The first data point is a distance of more than twenty units in energy from the best-fit line, while the second and third data points are still more than 5 units in energy from the line. Only the last 4 points appear sufficiently close to the line. The last four data points establish linearity of about 1% at high energies.

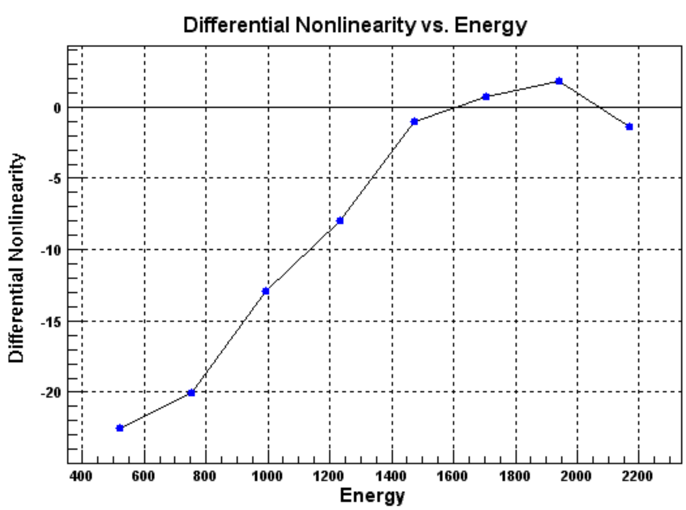

Figure 6: Differential Nonlinearity vs. Energy

The energy resolution plot should appear to scale as 1/energy. While the final energy resolution plot has this desired form, a special cut on the energy was applied to the original plot to obtain the correct form. Prior to the application of this special cut, the energy resolution plot, which consists of eight points of the RMS(E)/E versus E, had four points that appeared to obey the anticipated behavior, while the remaining four points had RMS(E)/E values significantly greater than expected. Histograms of the energies used to compute the RMS values in those points revealed that many of the energies used in these RMS calculations had values of zero. Upon discussing this problem with Professor Parsons, it was learned that the zero energies are a consequence of the calibration board's occassional failure to fire. After the application of the zero energies cut, the RMS(E)/E versus E plot acquired the desired form, as shown in figure 6. The first six points all appear to lie on a curve that scales as 1/energy, while the last two points appear to deviate from this curve by a small factor. The plot also demonstrates that the current energy resolution is just above 1%, still an order of magnitude greater than the 0.1% necessary to obtain a reliable Level 1 trigger decision.

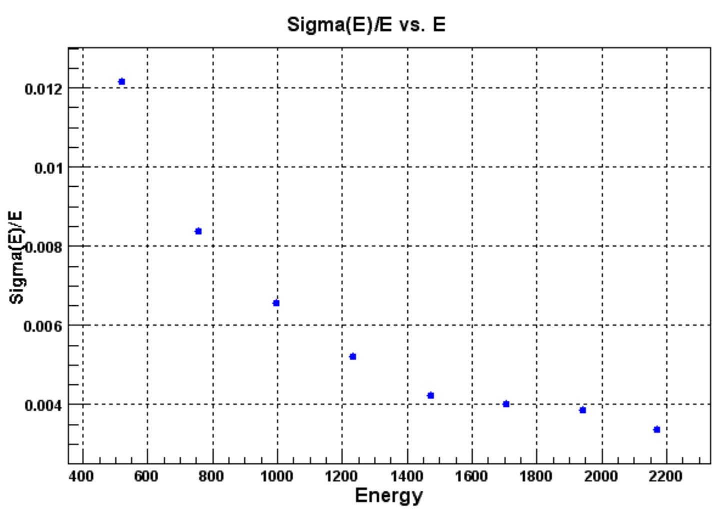

Figure 6: RMS(E)/E vs. E for High Gain DAC Scan

Future Developments

As previously mentioned, the FEB output is not a linear function of dac value due to the error introduced by the clock feedthrough. The next step in the software development is to account for this feedthrough by subtracting it from the output signal, and further investigate the linearity after this operation has been performed. This will be done within the upcoming weeks.

Further testing of the FEB will examine the temperature, currents, and voltages measured while the FEB is running and study the correlations between these parameters and the performance (e.g.: noise) and linearity of the board. The plan is to launch upon this task by the end of the month and complete it during the fall of 2002.

These tests will not be sufficient in determining the ultimate performance of the FEB. However, the board will need to be tested on real test beams. New parameters will also be introduced, such as pileup noise. Pileup noise occurs due to interference of several sequential events that are too closely spaced together in time, an effect that is not readily observable in the regular laboratory setting. Within the next 3 years, all the ATLAS FEBs will be produced, tested and installed in the ATLAS experiment.

Acknowledgements

We first want to thank the National Science Foundation for providing us with the financial support to participate in the Research Experience for Undergraduates program. Next, we would like to thank the members of the physics department at Columbia University for inviting us to Nevis Laboratories to learn about their research firsthand. We also want to thank the members of the ATLAS collaboration at Nevis, including Jiamin Jin for all her help with Root and understanding the Front End Board, Ioannis Katsanos for his production of the optimal coefficients and further help with Root and our understanding of the FEB and Liquid Argon Calorimeter, Sylvain Negroni for teaching us how to create the Root-tuples, and Stefan Simion for his assistance in interpreting our results. Last but not least, we want to extend our most sincere thanks to our mentor, Professor John Parsons, who as director of the REU program brought us to Columbia, and from whom we have learned a great deal about the Liquid Argon Electronics, the ATLAS experiment, and physics in general. Without his careful guidance, our successful completion of this project would not have been possible.