



Design of the SNS BLM Analog Front End

BNL/SNS TECHNICAL NOTE

NO. 118

R. Witkover, D. Gassner, C. Mi

BNL, Upton, NY 11973, USA

October 31, 2002

COLLIDER-ACCELERATOR DEPARTMENT
BROOKHAVEN NATIONAL LABORATORY
UPTON, NEW YORK 11973

Design of the SNS BLM Analog Front End

Richard Witkover, TechSource Inc.,
David Gassner, Brookhaven National Laboratory
Chaofeng Mi, Brookhaven National Laboratory
October 31, 2002

The SNS BLM System is designed to measure beam losses from a maximum 1% local loss down to a 1 W/m operating loss tolerance. In fact, resolution of 1 % of the 1 W/m threshold has been requested. This amounts to a dynamic range of about 2×10^6 for the system, which would require at least 21 bits digitizing range. The approach taken is to exploit the difference in bandwidth between the 2 limits (the low end only requires low frequency monitoring) to reduce the noise and allow additional gain.

Estimate of the 1 W/m Dose Rate

The estimated dose rate corresponding to a 1 W/m loss ranges from 80 – 115 mR/hr at 1 foot¹. For simplicity a value of 100 mR/hr, which corresponds to the allowable dose rate for hands-on maintenance, will be used here. Note that this represents the beam-off activation dose rate. At BNL a multiplier of 1000 (based on “folk law”) has successfully been used for design purposes to estimate the “beam-on” dose rate from the activation.

Thus the Beam-On Dose Rate is:

$$1000 \times 100 \text{ mR/hr} = 100 \text{ R/hr}$$

or, since there are 2.16×10^5 pulses per hour, $= 0.463 \text{ mR/pulse}$.

Thus, during the 1 msec beam pulse the dose rate will be 0.463 R/sec peak at 1 ft.

Using the typical BLM sensitivity of 70 nA/R/sec, the 1 W/m loss corresponds to 32.4 nA peak BLM signal current during the pulse.

Estimate of the Upper End Loss Signal

The upper-end loss limit of 1% at a single location might be expected at such “Controlled Loss” points as the Collimators and Beam Dumps. For a 2 MW beam, a 1% loss would be 20 kW. Estimating the fast loss signal is complicated by the difference in transit times between the electrons and ion in the detector. Electrons will be collected within a few microseconds while the ions may travel for hundreds of microseconds. To be conservative only the electron signal will be considered, reducing the BLM sensitivity to 35 nA/R/sec. Then the upper end signal would be:

$$\begin{aligned} I_{\text{Max}} &= 16.2 \times 10^{-9} \times 2 \times 10^4 \\ &= 0.324 \text{ mA} \end{aligned}$$

Note that this assumes that the upper end loss occurs uniformly throughout the 1 msec Linac/HEBT beam pulse. Since the Linac pulse will be accumulated over 1000 turns, a 1% loss in the Ring could result in significantly higher signal current depending on when in the accumulation the loss occurred. Shafer² using relations based on work by Boag, estimated that the new SNS prototype design would tolerate an instantaneous 2 Rad local loss. Using the value for a 1 W/m loss (0.46 mRad/pulse) and scaling by 2×10^4 factor for a 20 kW (1%) loss, a 1% local loss would be equivalent to 9.2 Rad. Thus, in the Ring a loss of 1% of the beam locally on a single turn after accumulating a little over 200 turns, would show detector saturation.

The situation is more acute in RTBT, since the full accumulated beam is extracted in a single 645 nsec bunch. In this case the maximum loss tolerated without detector saturation would be only about 0.2% rather than 1%. Either we can accept the fact that this is a huge loss and will not be linearly detected above the 0.2% level, or, move the detectors away from the beam line so they do not saturate. To retain sensitivity at the low loss end of the range, and spatial localization, it is not desirable to locate the detectors further away from the beam line. In addition, the RTBT channels will require additional loading capacitance at the front end to soak up the huge charge generated by such a high loss so that the electronics will not be saturated.

Estimate of the Front End Impedance

Let us assume a standard cable length of about 100 m or 300 feet. For a capacitance of 16 pF/ft this gives 4.8 nF. If we use 2.2 “RC’s” as the rise time and require a 5 μsec rise, then the input load must be:

$$2.2 \times RC = 2.2 \times R_i \times 4.8 \times 10^{-9} = 5 \times 10^{-6}$$

and

$$R_i = 450 \text{ O}$$

We will use 470 O as the input resistance. Since the actual cable lengths vary it may be necessary to load the inputs with capacitors or to equalize the cable lengths to give equal rise times, however, the net effect on the risetime is small ($< 1 \mu\text{sec}$). If it is decided to do this balancing it should be done external to the circuit board to minimize changes when exchanging modules. A BLM-Cable Interface Panel containing these balancing capacitors and providing a transition between the coax from the detectors and mass termination connectors to the electronics modules is proposed. This same panel will house the input capacitors loading the RTBT line channels.

Estimate of the Feedback Resistance

Assuming a maximum input voltage for the ADC of 5 V, the feedback resistor will be chosen so that the signal will not exceed 2 V for a 1% (uniform) beam loss.

Then

$$R_f = 2 \text{ V} / 0.324 \text{ mA} = 6.17 \text{ K O}$$

Estimate of the Noise

Note that the environmental noise voltage gain is set by the resistors R_f and R_i , $6200/470 = 13.2$, but the signal is only set by R_f , since the ion chamber may be considered a true current source. In RHIC the total noise (environmental and circuit) was observed to be equivalent to 10 pA for a bandwidth of about 10 Hz. Since we will be using a 100 kSa/S ADC the bandwidth will be limited to 50 kHz. Scaling between these cases gives an equivalent noise current of 0.71 nA and an output noise voltage of 4.36 μ V. Calculation of the resistor Johnson and amplifier noise for the circuit gives 4.1 μ V. The first stage will be limited to 70 kHz with a roll-off capacitance of 390 pF for the 6.2kOhm resistor.

Estimate of the Output Voltage for 1 W/m Loss

The signal current for 1 W/m was found to be 32.4 nA. With a 6.2 kOhm feedback resistor this yields an output voltage of 200 μ V.

Estimate of the Voltage Equivalent to an LSB

For a 16-bit (15-bit plus sign) ADC with a maximum input of 5 V, an LSB is equivalent to 152.6 μ V. Thus, a 1 W/m loss is of the order of an LSB. Since the specification calls for resolution below 1 W/m, this is not sufficient.

Extending the Range below 1 W/m

Another 7-bits would be needed to achieve 1 % resolution of a 1 W/m loss, but at full 50 kHz bandwidth the noise would be excessive. The signal from the input stage is split 3 way. One path is to the MPS signal sensing circuit, one to the fast output for wideband monitoring and one for the low-level 1 W/m monitoring. By filtering the 1 W/m signal down to 1 kHz, the noise is reduced sufficiently to allow the measurement by sampling over a reasonable number of pulses. Equally important, the low pass filter will smooth out large, short duration losses, preventing saturation of the following gain-of-ten stage and ADC. The initial design used a filter with a cut off of a few Hz, but in this case the amplifier offset could not be distinguished from the BLM signal, which was of similar amplitude. While the first stage offset could be compensated to 1-10 μ V using a trim pot, the output following the gain-of-ten amplifier could still be as high as 100 μ V. In addition, a 1 degree C variation in temperature would change the offset by 50 μ V. This was un-acceptable. By setting the lowpass filter cut off at 1 kHz, the baseline, including temperature drift, could be measured and subtracted from the signal, while the noise would be reduced by about a factor of 7.

Three ADCs were studied³ for this application: a 24-bit, 32 channel 100 kSa/s VME module, the ICS-110B, a 16-bit, 16 channel 100 kSa/s module, the VMIVME-3123 and a HyTec8401 16-bit, 16 channel 100 kSa/s module. Tests showed that the ICS-110B was able to deliver a true 18-19 bit output but suffered from thermal drift and noticeable thermal gain variation. The VMIVME-3123, even though delivering only 16 bits,

exhibited better gain stability. The HyTech module exhibited unacceptable performance and will not be considered. Using a pulse simulator to give the equivalent of the 1 W/m signal, the ICS and VMI units were able to resolve the signal to 1% when at least 600 samples were taken. One hundred samples (at 100 kSa/s) will be taken prior to the pulse to establish the baseline, which will then be subtracted from the 100-200 samples during the beam. In 1 second, 6000 baseline samples will have been taken. Data will be accumulated over a 10 second or longer interval and processed to compare against a 1 W/m reference. This data will be processed in the IOC. Warnings will be sent through the Control System as this level is approached. If the 1 W/m-level loss is sustained over a several minute period then a more emphatic warning will be generated. A signal may also be sent to the MPS through the IOC.

While both ADC's appear to meet the design requirements, these designs are now at least 5 years old and will be closer to 10 years old at the time the facility becomes operational. The question arises if they be around much beyond that? Certainly spares will be purchased, but eventually a newer replacement may have to be found. Many new ADC chips are now available offering 16-bit and better resolution at tens of MSa/s rates. It is recommended that a long term project be initiated to design a new board offering 16-32 channels of 100 kSa/s 16-24 bit ADCs.


Gain Changing for the AFE

The upper limit of 1% loss was requested to monitor the "Controlled Losses" expected at the Collimators, Injection and the Beam Dumps. This high limit will not be required or desirable for the majority of channels. By making the feedback resistor of the front end amplifier jumper-selectable, BLM channels can be set to different sensitivity. The circuit has been designed so that a double jumper will be used to change the gain and set a bit which will be read by the IOC to read back the channel gain setting.

A second stage of gain will be available for the fast output data following the point where the signal is split. Thus only the data going to the fast output ADC will be affected by this gain stage. The gain will be program controlled with settings of X1 and X10 and affect only the fast ADC readings. Computer readback of the gain state selected for each channel will be included. Three gain states will be provided.

Fast Loss Trip Sensing (MPS Signal)

Experience at the BNL Linac and at LANSCE has shown that it is more appropriate to shut off the beam based on accumulated loss during the macro-pulse rather than loss rate since hardware damage is normally due to amount of energy rather than rate of energy deposited. Excessive uncontrolled loss levels will be sensed by the Machine Protection System (MPS) based on the output of an integrator following the first stage. Several types of integrators (Figure 1) have been considered: The Gated Integrator, The Up/Down Integrator and the Leaky Integrator. The relative merits of each will be discussed.



The Gated Integrator is a common circuit around pulsed accelerators. The FET across the integrator capacitor is turned on except during the beam time, effectively shorting out the capacitor. Because the on-resistance is low (a few Ohms to several hundred Ohms) it may be necessary to add a current limiting resistor. The advantages of this circuit are that the offset of the first stage is not amplified as in the case of the Leaky Integrator, and an accurate integration of the input signal is obtained. The charge in the leaky integrator bleeds off during the pulse, becoming most significant for long pulses. On the negative side, (a) the FET switch may introduce a significant error due to charge injection, and (b) timing is required to trigger the switch. The charge injection can be compensated either with additional circuitry or in software. The requirement for timing, however has more serious implications. Since the output of this circuit will be used as a primary MPS input, the signal must be developed in hardware only. That is, it must be designed such that software processes cannot affect the generation of the signal. The timing system is dependent on software and so must not be involved with generating the MPS inputs.

Another circuit, the “Up/Down” integrator, uses a fixed bias to down integrate the charge accumulated on the integration capacitor. A diode, which can be connected in several configurations, is used to keep the signal from going negative. The bias must be large enough to down-integrate the largest possible input in the inter-pulse period, or about 10 V over 15 msec. The loss signal would be of the order of a few volts over 1 msec, making the bias a significant fraction of the signal. Worse, since the losses can occur over any portion of the pulse, the bias cannot be compensated. Thus losses occurring over a long duration would be cancelled by the bias whereas the same charge occurring in a short time would trip the MPS.

The simplest circuit is the “Leaky Integrator” in which the integration capacitor is bridged by a resistor to drain the charge during the inter-pulse period. No timing is required, charge injection compensation is not necessary and down-integrating bias current is not a problem. Analysis of the circuit indicated that a 3 time constant leakage would leave a residual of 5-6 % while 4 time constants would leave 2 %. The peak signal would reach only 96.2% (due to the leakage resistor) for the 3 time constant case, but only 91.7% for the 4 time constant case. Note also that the residual (and the peak) will depend on the average prior losses. Figure 2 shows the results of the simulations.




Figure 2. Simulation of Leaky Integrator for 3 and 4 time constant decay.

The large drain resistor may cause a significant offset due to the high dc gain for any first stage offset voltage. An offset trim resistor will reduce the contribution from the integrator stage but the input stage offset will still be amplified. Tests indicate the first stage can be trimmed to $1 \mu\text{V}$ but $10 \mu\text{V}$ might be more realistic in practice.

In the RTBT line the beam is only 645 nsec long and, while ion transit time, cable capacitance and circuit bandwidth will extend this signal, it will still be too short to properly acquire. In addition, the loss signals will be high current pulses and could require a separate circuit design. Rather than do that it was decided to pre-integrate the signal by loading the input with lumped capacitance ($0.1 - 0.47 \mu\text{F}$) to knock down the signal to acceptable levels. Since only the total charge is of interest, the signal from the input stage will be fed to the leaky integrator using an on-board jumper, before going to the viewing gain stage. This jumper setting will be read back through the control system to assure configuration control.

Detailed AFE Circuit Design

The AFE circuit is shown in Figure 3. The choice of the amplifiers is critical since the circuit is dc coupled and offsets will be amplified or integrated by the later stages. Several amplifiers were considered and the Burr-Brown OPA627BP was selected for the front end and Leaky Integrator. It has a low voltage offset, $V_s = 40 \mu\text{V}$ typical, $100 \mu\text{V}$

maximum, a bias current, $I_b = 1 \text{ pA}$ and a $\text{GBW} = 16 \text{ MHz}$. The voltage offset temperature coefficient is $0.8 \text{ } \mu\text{V}/^\circ\text{C}$ max, $0.4 \text{ } \mu\text{V}/^\circ\text{C}$ typical.

I1, 2 = OPA627BP
 I3, 5 = OPA2277U
 I4 = OPA277P

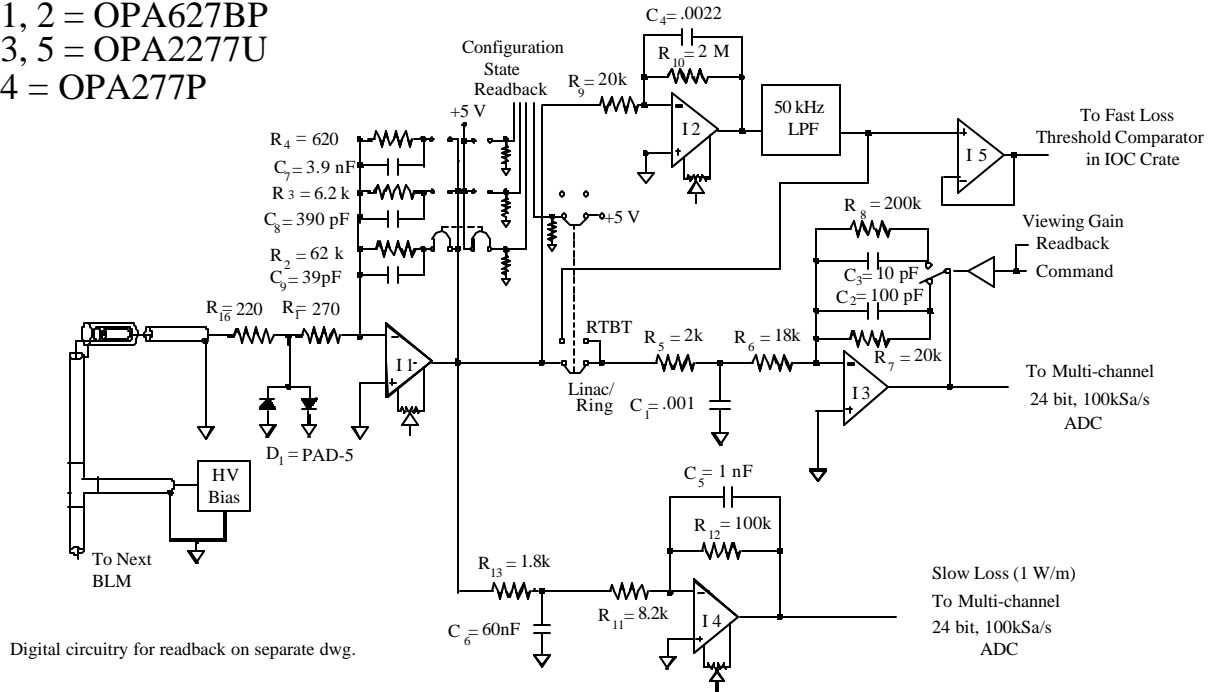


Figure 3. Prototype BLM Analog Front End Design

Resistor choices for the first stage have been made in an earlier section: $R_1 = 470 \text{ } \Omega$ and $R_2 = 6.2 \text{ k}\Omega$. The Linac may require a resistor 10 times this or, $R_3 = 62 \text{ k}\Omega$, and high loss areas a smaller resistor, $R_4 = 620 \text{ } \Omega$.

Consider the Fast Loss Integrator. For a first stage output voltage of 2 V , corresponding to a 1% $20 \text{ } \mu\text{s}$ sustained local loss, the combined electron and ion current for the new SNS ion chamber would have risen to 84.5% and the “Leaky Integrator” output would be 0.909 V for a single pulse. Twenty μs was chosen since the MPS should shut off the beam by this time. For the Leaky Integrator the residual causes the signal to rise, reaching an equilibrium value of 1.05 V after 3 pulses.

The output of the integrator will go through the unity gain buffer to a comparator on the MPS Interface Module in the IOC crate. The leakage resistor was chosen as $2 \text{ M}\Omega$ with a 2.2 nF capacitor to obtain an RC between 3 and 4 time constants. The 50 kHz filter was added to reduce high frequency noise observed in field tests of the prototype. The OPA2277 amplifier has a very low offset voltage ($10 \text{ } \mu\text{V}$) and a 500 pA bias current.

The 1 W/m gain stage is the final portion of the circuit to be designed. Here again a low pass filter is required. As stated earlier, a 1 kHz filter will be used. It will allow the BLM

signal to reach its peak value in about 350 μ sec. The in-situ noise observed in the RHIC BLM system, scaled for a 6.2 kOhm first stage feedback resistor and a 50 kHz bandwidth was estimated to be 0.71 nA or 4.1 μ V. Including the amplifier and resistor noise and multiplying by the gain, the total noise is 6 μ V output, sufficient to allow a 1% measurement of the peak 1 W/m signal.

Estimates of Amplifier Offsets

We will now calculate the offsets (at 25 deg C) using “typical” values.

Output of the First Stage:

$$V_o(I1) = V_{os}(I1) + I_b(I1) * R_f$$

For the Burr-Brown OPA627BP : $V_{os} = 40 \mu$ V and $I_b = 1$ pA. So for the highest gain case of $R_f = 62$ kO and infinite input impedance to ground,

$$V_o(I1) = 40 \times 10^{-6} + 1 \times 10^{-12} * 6.2 \times 10^4 = 40 \mu$$
V

Because the input bias current is so low there is no measurable difference in offset as a result of selecting a different feedback resistor.

Note that this is the untrimmed offset voltage. Tests on the prototype circuit have




Figure 4. Thermal tests results for the AFE Circuit.

shown this can be trimmed to 1 μ V and maintained at level at room temperature over many days and repeated on/off cycles. Tests in a temperature controlled oven have shown

that the offset variation is only a few μV from 27 to 47 degrees C. (Figure 4.) However, for design purposes an effective offset of $10 \mu\text{V}$ will be used.

Viewing Gain Stage Output:

The input to the ADC used for data logging is taken through the Viewing Gain Stage, which has gains of X1 and X10. For this analysis we will use the worst case of the X10 gain. The amplifier is the OPA2277U, which has a typical offset voltage of $10 \mu\text{V}$ and a bias current of 500 pA . There is no provision for a trim pot.

$$V_o(\text{VG}) = V_o(\text{I3}) = (V_o(\text{I1}) + V_{os}(\text{I3})) * A_{\text{VG}} + I_b(\text{I3}) * R_f$$

Where $A_{\text{VG}} = R_f/R_i = 10$ and $R_f = 100 \text{ k}\Omega$.

The results for various configurations are:

Gain	Without 1 st Stage Trim Pot	With 1 st Stage Trim Pot
	W/O 2 nd Stage Trim Pot	W/O 2 nd Stage Trim Pot
10	550 μV	250 μV
1	55 μV	25 μV

Clearly, a trim pot on both the first stage results in the smaller offset. With a trim pot on the first stage the offset for a gain of 10 will be between 1 and 2 LSBs for a 15-bit plus sign ADC (5 V max input).

Integrator Stage Output:

To estimate the static offsets of the “Leaky Integrator” consider it as a heavily rolled-off amplifier, in which case its output would be:

$$V_o(\text{LI}) = (V_o(\text{I1}) + V_{os}(\text{I2})) * R_{10}/R_9 + I_b * R_{10} + V_{os}(\text{I5})$$

The leakage current of the buffer does not contribute since the feedback resistance is zero. Using $R_9 = 20 \text{ k}\Omega$ and $R_{10} = 2 \text{ M}\Omega$, the offsets are:

Without 1 st Stage Trim Pot		With 1 st Stage Trim Pot	
W/ 2 nd Stage Trim Pot	W/O 2 nd Stage Trim Pot	W/ 2 nd Stage Trim Pot	W/O 2 nd Stage Trim Pot
10 mV	20 mV	202 μV	10.1 mV

The results of the thermal tests in Figure 4 show an offset of nearly 50 mV for the case of a 1st stage trim pot, no 2nd stage pot. This was due to a damaged PC board which had

excessive leakage at the input to the integrator stage. However, the variation with temperature is still only a few mV. A trim pot will be used on the Leaky Integrator stage.

1 W/m Output Stage:

The 1 W/m loss is only considered significant if it occurs over an extended time period. The fast beam loss signal is limited by noise, but the slow response needed for the 1 W/m measurement can be used to advantage by limiting the bandwidth and using a high resolution ADC. Choosing a 1 kHz bandwidth results in a signal rise and fall of about 350 μ sec. This will allow a measurement of the baseline prior to the beam pulse which can be subtracted from the signal. It would be desirable to have the offset as small as possible to minimize the errors. An X10 gain inverting amplifier will be used to increase the signal to the ADC. An OPA627BP or an OPA 277P will be used here. As before, it is assumed that the first stage will use a trim pot, resulting in an offset of 10 μ V. The offset, with no trim pot for I4 is then:

$$V_o(1W/m) = V_o(I4) = (V_o(I1) + V_{os}(I4)) * R_{12}/(R_{11} + R_{13}) + I_b(I4)*R_{12}$$

For the case of the OPA 2277U,

$$V_o(1W/m) = V_o(I4) = (10 \mu V + 10 \mu V) * 10 + 5 \times 10^{-10} \times 10^5$$

$$V_o(1W/m) = V_o(I4) = 250 \mu V$$

And for the case of the OPA 627BP (with no trim pot on the second stage),

$$V_o(1W/m) = V_o(I4) = (10 \mu V + 40 \mu V) * 10 + 1 \times 10^{-12} \times 10^5$$

$$V_o(1W/m) = V_o(I4) = 500 \mu V$$

And for the case of the OPA 627BP (with a trim pot on the second stage),

$$V_o(1W/m) = V_o(I4) = (10 \mu V + 10 \mu V) * 10 + 1 \times 10^{-12} \times 10^5$$

$$V_o(1W/m) = V_o(I4) = 200 \mu V$$

It would seem that using an OPA277P for the second stage without a trim pot is nearly as good as using the OPA627BP with a trim pot, and will be used in the design.

The data will be analyzed in the IOC but should not present a computing burden since the base line data would only be acquired for 1 msec before the pulse, and the signal data for 1-2 msec during and after the pulse, a total of 300 samples per channel. Even for an installation of 32 channels (one-half Ring) in one crate, this would amount to 36 kbytes per pulse. The processing would consist of averaging the baseline data, summing the

signal data and taking a single difference between the two. The result would then be stored and used to compute a many-pulse average against which the 1 W/m threshold would be applied. As stated earlier, tests with a signal simulator and the various ADCs showed that data averaged over 10 seconds should meet the requirement to measure the 1 W/m loss to 1 percent.

Control and Readback

The first stage amplifier gain is set by means of ganged jumpers to select from among 3 feedback resistors and indicate which has been chosen. Jumpers will be used because the selection is intended to be “permanently” associated with the specific BLM location and not to be changed with different operating conditions. Jumpers do not allow casual changes to the signals which are used for the MPS loss protection. Read back of the selection is required for the Control System to know the gain setting. Since channels on the same board may require different gains due to local shielding or high controlled loss conditions, individual channel readback is required.

The signals presented to the ADC for data logging will have computer controlled gain change capability to allow for different operating conditions. This is done in the Viewing Gain stage, with gains of 1 and 10 available. This will also be done on a channel by channel basis.

Since the beam in the RTBT will be only 645 nsec duration, putting the direct amplifier output into the 100 kSa/s ADCs would not provide useful data acquisition. A jumper allows selection of either the front-end amplifier or leaky integrator output to be passed through the Viewing Gain stage to the ADC. A ganged jumper is used to identify the mode configuration for each channel.

The most straight forward approach would be to use a separate bit for each gain setting and each state readback. However, a single 8-channel board would require 40 read back bits and 8 command bits. Each AFE module will have on-board logic to allow the 5 states for each channel to be read back under control of the IOC. Bits will be provided to select which channel on which module is being read. Similarly, the 8 Viewing gain bits will be encoded to 3 lines and applied when the module code matches that board. The module identifying codes will actually be wired on the back plane connector and simply passed through the module. That way each board does not have to be jumper configured to identify what channels it supports.

Conclusions

An acceptable design for the BLM AFE can be made without resort to exotic amplifiers. Acceptable offsets can be achieved with the use of an offset adjusting pot for each front end amplifier and for each leaky integrator. However, the offset calculations are based on the manufacturers data sheets and deviations from these figures are expected. A prototype of a single channel has been built and tested for thermal drift over a reasonable range of

temperatures with good results. The circuit should meet the requirements of the SNS BLM system.

¹ NSNS Conceptual Design Review, June 1997, Chapter 3, Page 97

² Shafer, R., Personal Communication (email) dated 5/13/02

³ Yongbin Leng, Lawrence Hoff, "ADC Module Evaluation Report for BLM System", July 18, 2002, BNL
Non-published report