Precision Integrating HAPPEX ADCs

Robert Michaels

Thomas Jefferson National Accelerator Facility

Aug 11, 2003

I INTRODUCTION AND GENERAL DESCRIPTION

This is a specification for a precision integrating ADC to be used by the HAPPEX data acquisition.

Parity experiments require high precision measurements of signals integrated over a helicity pulse (33 msec), or over a fraction as small as 1/20 of a helicity pulse, this latter case is called "oversampling". The signals consist of PMTs from detectors (current sources) or voltage levels (voltage sources) from various devices such as beam position monitors (BPM) or beam current monitors (BCM). The reason we integrate is that in order to see the tiny parity violating signal in a reasonable amount of beam time the event rate must be far too high (e.g. gigaherz range) to count individual events, so we must integrate the signals.

The ADCs are designed to achieve high resolution (18 bits is desirable) with small nonlinearity, see specifications listed in table 1. Each ADC channel (fig 1) consists of an input amplifier, an integrating circuit, two sample-and-hold circuits, a difference amplifier, a summing circuit, and a 16 bit ADC chip. The input amplifier converts the input voltage to a scaled current which is integrated in the next stage; for current signals such as PMTs this amplifier stage is bypassed and the signal is integrated directly. The integrator output is sampled and held once 700 μsec after the beginning of the helicity pulse, and again 32 msec later near the end of the pulse, for non-oversampling mode. This timing is determined by an external ADC Timing Board built by the JLab Electronics Group. The difference between these two signals is the integrated result. To achieve the nonlinearity specification, a pseudorandom DAC voltage

("DAC noise") is added to this integrated result prior to digitization by the ADC, then subtracted later in analysis. The DAC noise smears the data over many ADC channels, thus reducing systematic errors from bit resolution.

II DESCRIPTION OF HARVARD/PRINCETON ADCS

Here we describe the the existing Harvard/Princeton ADCs which these new boards are intended to replace. The new board will be similar. In the subsequent section I will explain what is new. The reasons to replace the existing boards with a JLab design are:

- The existing boards use "officially" obsolete components and are no longer supported at Harvard or Princeton because key people have retired or moved on. A related issue is that the board manufacturing files and board layout were done on media which are now obsolete, so one would need to start from scratch anyway even if trying to make an exact copy.
- The existing boards are marginal in their performance characteristics for future parity experiments, and we believe we can improve on them. Table 1 reflects the desired new specifications.
- The boards are too customized and too few in number to be built by any outside manufacturer. Having the boards built and maintained by the JLab electronics group is helpful for long term maintenance.

Each board has 4 channels; the individual channels were briefly described in the introduction and in figure 1, namely it consists of an input stage, an integrator with reset switch, a sample-and-hold circuit, a difference circuit, a DAC summing circuit, and the ADC. There are two kinds of channels: voltage integrating and current integrating. For voltage integrating channels, the input stage is a balanced-input differential amplifier followed by a resistor. It converts the input voltage signal into a current which charges the integrating capacitors C1 and C2 in fig 1. For current-integrating channels, the differential amplifier is bypassed and the input current is integrated directly. Two different gain ranges are available by switching capacitor C2 in or out of the integrator circuit.

The integrator has less than 100 pA of leakage current and settling time less than 1 μsec. The integrator is sampled at the arrival of the "baseline" and "peak" signals from the Timing Board. Two sample/hold circuits hold these

levels and feed them into a difference amplifier. The output of the difference amplifier is sent into an analog summing circuit. Here, a DC level of a 16 bit DAC compressed by 5 bits is added to the signal. This is the "DAC noise" intended to reduce nonlinearities due to bit resolution. If the signal were extremely quiet so that it sat in only two or three ADC channels, one would see the effects of bit resolution when forming the difference between two helicity states; by adding the DAC noise in hardware and subtracting it in software will wash out this affect. In addition to the DAC noise, a -2.5 volt offset is added. The summed signal is fed to the ADC for digitization.

The ADC chip converts an analog input in the range -3 to +3 volts to a 16 bit digital value. The -2.5 volt offset puts the zero point of the integrated result 0.5 volts above the bottom of this range; this 0.5 volt offset avoids the danger of saturation at the lower end of the ADC's input range.

The "baseline" and "peak" signals can be interchanged by software so that positive as well as negative inputs can be accepted. This capability, the input to the DAC and the control signals to the ADC, are controlled by a programmable gate array on the board. The board accepts 4 control signals from the ADC Timing Board: "baseline", "peak", "reset", and "convert". We have already described that "baseline" and "peak" begin and end, respectively, the integration time. When "reset" is asserted the reset switch is closed and the integrating capacitor is discharged through a 1 kOhm resistor. Upon release of "reset", the input current begins charging the integrating capacitor. The integrator is allowed to settle about 15 μsec before "baseline". The width of "baseline" and "peak" determine the sampling time of respective sample-and-hold circuits, while the time delay between "baseline" and "peak" determines the integration time. After the "peak" sample has been taken, the "reset" is asserted by the ADC board to clear the integrator, and remains asserted until the next integration cycle, i.e. the next time "reset" is released, see fig 2.

Next the "convert" signal is strobed to start digitization. About 22 μsec is allowed between "peak" and "convert" to allow the difference and summing amplifiers to settle. Note that the components in this circuit have been selected for *low noise* rather than high speed.

Upon completion of the integration cycle the digitized values can be read over the VMEbus. In addition, the DAC value from "DAC noise" can also be read out. The new DAC value for the next integration cycle should also be set at this time in order to give the DAC and summing amplifier plenty of time to settle before the next conversion. When the conversion is complete it sets a "done" bit in the control/status (CSR) register of the ADC. Polling this bit allows CODA to tell when the ADC data are ready. After readout the bit is set to zero again.

The same control signals are fanned out to all 4 channels on a board. The various ADC boards in a VME crate are daisy chained, i.e. the control signals are sent to each along a ribbon cable. The source of these control signals is the ADC Timing Board built by the JLab Electronics Group. In most cases the signals operate at 30 Hz, but for "oversampling" runs we may run at a factor of up to 20 times higher rate, which is driven by the ADC Timing Board.

There is one analog output through which any of the 4 integrator outputs can be sampled on a scope for testing purposes. The choice of which output goes to the scope is done through VME control by setting the CSR register.

Setting bits in the CSR register controls the following features: 1) As mentioned above, one bit is the "done" bit for telling when conversion is ready; 2) Some bit(s) are used for setting the gain of the frontend; 3) As mentioned above, the CSR can determine which output goes to a scope output. This feature isn't necessary for the new board as long as we can somehow sample each output, e.g. they could have individual outputs. 4) A bit is used to enable or disable the DAC noise, i.e. to decide whether or not to add the noise to the input signal. The DAC itself is programmable from VME.

The ADC is an Analog Devices chip model AD7884 with a conversion time of 5.3μsec and a maximum throughput of 166 ksps. It has 16 bit accuracy with guaranteed no missing codes and integral nonlinearity of 0.003%. Intrinsic noise is 120 μV rms and signal to noise ratio for fast pulses is typically 82 dB. The power dissipation is typically 250 mW with a maximum of 325 mW.

The existing ADCs come in a 9U VME form factor with J1-J2 in the lower 2/3 of the crate. The specification for this crate is JLab Spec 96741-S-00144.

III WHAT IS NEW

The new ADC boards should have a lot in common with the old ones; the basic idea is the same except that the noise characteristics are a bit tighter; the specifications are listed in table 1. A desirable improvement would be to find ADC chips with 18 bit resolution. We believe suitable chips can probably be found. To reduce the pedestal noise, the E158 experimenters found that it helps to sample the signal several times before and after the integration time. Then the baseline is an average of several samples (say 10 samples) and the peak is also an average. Another improvement considered by E158 was to synchronize the sampling clock to the RF frequency of the accelerator, thus reducing noise due to the RF. I don't think this is necessary for the JLab design.

TABLE 1. ADC Specification

Quantity	Specification	Comments
Bit Resolution	18 bits	
Rate Capability	10 kHz	
Input Signals	PMT signals (currents) or Voltages	2 flavors: current or voltage integrator
Input Polarity	One sign, but reversible	
Gain Adjustment	Factor of 100	
Max Input (Low Gain)	10 V	
Differential Nonlinearity	$\leq 2 \times 10^{-5}$	Definition - Sec III
Integral Nonlinearity	$\leq 5 \times 10^{-5}$	Definition - Sec III
Pedestal Noise	$\leq 100\mu$Volt or $\leq$2.5 ADC chan FWHM	
VME Form Factor	9U with J1-J2 in lower 2/3	Section IV

In the original design the choice of components sacrificed speed for accuracy, i.e. low-noise. While the same compromise is probably needed in the future, we would like to understand how the circuit will respond to fast photomultiplier signals and what the compromises are. This is more of a question than a specification, and this paragraph will be rewritten after we find out. The system should provide a true integral of the input signal, for signals with rise time $\tau \geq 5$ nsec and widths between typically 30 and 150 nsec (FWHM). Ideally the circuit would respond well to signals at rates as low as 1 kHz, but this might be impossible given the noise specs. The circuit should respond well to rates between 10 MHz and 10 GigaHz.

Here I explain how *we define* differential and integral nonlinearity. Ideally the response to an input signal should be a perfect line; however, the ADC has a bit resolution which produce a "staircase" response. The input signal drives the lowest order bit from a zero to a one repeatedly at various points in the magnitude of the input. Let these points be called the "bit toggle points" of the input. Consider the ADC the response at each bit toggle point. Call the response signal at the bit toggle point S and the input voltage V. Ideally we have $S = aV + b$ where a and b are constants. The differential nonlinearity spec says that we want the deviation of this response dS at any bit toggle point to be $dS \leq 2 \times 10^{-5} S$.

Now let I be an integral of the signal S over any interval in the dynamic range (i.e. add up all the ADC signals from many events where the input is fluctuating). We compare this to the integral $I_{\rm lin}$ of the ideal linear response over the same interval. The integral nonlinearity spec says that the deviation must be $dI = I - I_{\rm lin} \leq 5 \times 10^{-5} I$.

Note, we are not concerned with the bit resolution of the ADC since we use DAC noise to reduce its affect to an acceptable level.

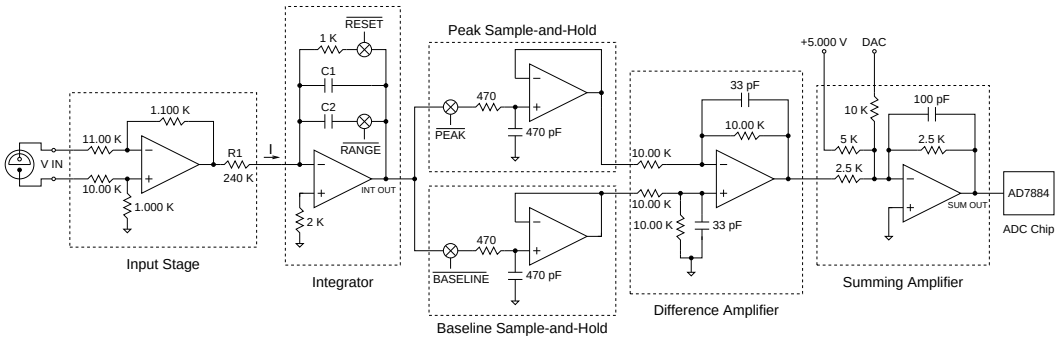

FIGURE 1. Circuit diagram of one channel of the existing design of the Harvard/Princeton 16 bit integrating ADC.

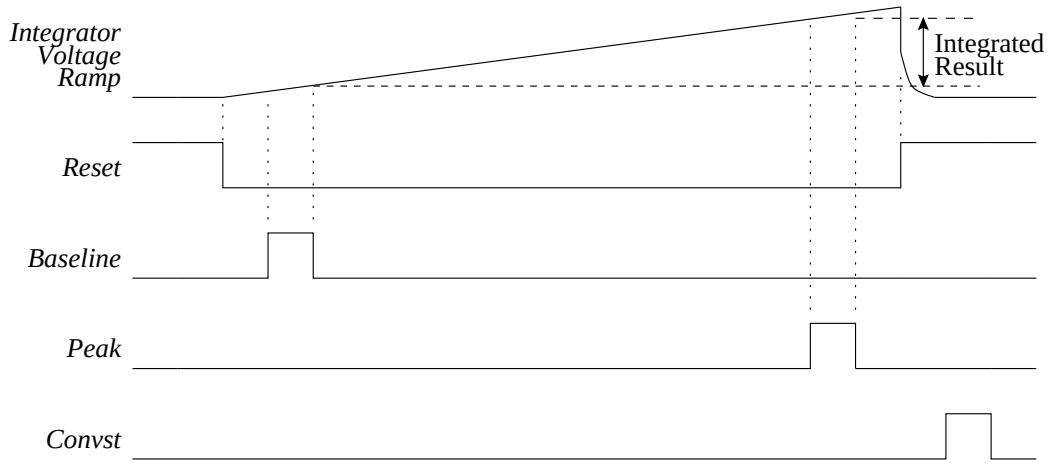

FIGURE 2. Timing diagram for one ADC integration cycle. Usually this occurs at 30 Hz, but for "oversampling" it may be as much as 20 times faster.

IV MISCELLANEOUS SPECIFICATIONS

The VME form factor should remain 9U with J1-J2 in the lower 2/3 of the crate, as per JLab Spec 96741-S-00144. This allow us to make use of existing crates, also 9U provides more space for channels than 6U. The number of channels in one unit should be at least four but not more than ten. The number of boards we'd like is at least 6. We need 6 units to spread them about the lab. If it is not too difficult, we'd like to have two 6U versions of the board as well.

The input connections for the "baseline", "peak", "reset", and "convert" should be compatible with the output of the ADC Timing Board, i.e. a ribbon cable with a particular ordering of inputs. The signal inputs are twinax connectors.

It would be desirable to have scope pickoffs to check the timing signals from the ADC Timing Board, and to check the integrated result. The old board had a provision for selecting via programming which channel was sent to the scope pickoff for checking the integrated result; we either need this feature, or should be able to check each channel separately.

The gain should be adjustable by a factor of 100. By this I mean that the "standard" inputs would be no larger than 10 Volts, but we should be sensible to signals of 0.1 Volts through gain adjustment. Part or all of the adjustment (at least a factor of 5) should be achievable through VME programming. It is desirable but not required to have the entire dynamic range adjustable via VME. Another part could, if necessary, be achievable through replacement of components on the frontend, to be carried out infrequently by a qualified technician. An adjustment by technician would typically be done only once at the start of an experiment or test period. The input polarity will also be fixed for long periods of time, but it should be possible to reverse this polarity; it is not necessary that this be a programmable feature, it may be done through a hardware intervention since it is infrequent.

The boards come in two flavors: voltage integrating and current integrating, and it is desirable to be able to switch between these modes by removing a stage at the input. Normally all the channels on one board can be of the same flavor, though if it is easy to select the flavor on a per-channel basis, this would be desirable. It is desirable but not required to make the switch between flavors a VME-programmable feature. It is sufficient to make the switch a hardware jumper. In the old design it was necessary to unsolder and resolder components; we prefer to avoid this.

Regarding the time schedule for producing these boards, we anticipate

needing them for experiments in the latter half of 2005. We would like to have one or two boards – they could be prototypes for testing the specs – by next summer (2004), but this is not crucial. The prototypes would be very helpful to commission a highly accurate new detector, the luminosity monitor, which we are installing. Therefore we desire to have one or two boards by next summer, but would require to have all the boards by summer of 2005.