Comparison of different solutions to the dynamic range reduction for the Fluorescence Detector front end electronics

Paolo Walter Cattaneo

INFN Sezione di Pavia and Dipartimento di Fisica Nucleare e Teorica, Via Bassi 6, Pavia, I-27100 Italy
Email: Paolo.Cattaneo@pv.infn.it

A number of different solutions to the problem of matching the large (16 bits) dynamic range of the current signal with the limited (12 bits) dynamic range of the ADC has been proposed. They are presented and reviewed critically, assessing their relative advantages and disadvantages.

1 Introduction

In the AUGER fluorescence detector the predicted dynamic range in current is about 16 bits using as LSB approximately the sky noise [2].
The signal waveform is sampled at 10 MHz so that technical and financial considerations on the available ADC model requires the range to be reduced to 12 bits before being digitised.
In order to achieve this reduction, several possibilities have been proposed, namely

- Separate paths from successive phototube dynodes
- Virtual channel summing ≈ 20 channels with low gain
- Switched amplification paths
- Three separate amplification paths based on 8-bit ADC with staggered sampling
- Dynamic compression

In the following we describe in some detail these approaches and evaluate their advantages and disadvantages.
It is important to note that the first three of these techniques rely on separating the signal in two or more paths with different amplifications at a certain points and applying some digital processing to choose the right path.

2 Quality parameters of the architectures

We have to define which are the quality parameters relevant for the choice of the architecture. That is which are the criteria that will be used in the following to classify the architectures.
A tentative list can be the following :

- Compression capability
- Simplicity
- Connectivity
- Analog-digital coupling
- Modularity
- Calibration
- Cost
- Compatibility

Compression capability is the amount of compression possible with a given architecture without degrading significantly the performance of the detector.
Simplicity is chiefly dictated by the extra circuitry involved.
Connectivity is the requirements on the cables between the very front end and the front end board, both in number and in dynamic range transmission.
Analog-digital coupling is the degree of coupling between the digital and the analog parts. That is the amount of digital signal, clock and switches used in the analog section before the ADC.
Modularity is the degree of coupling of any channel with the others.
Calibration is the degree of complexity both in hardware and algorithms required to keep the channel calibrated.
Cost is an estimation of the money required to equip one channel.
Compatibility is the degree of interchangeability of the different solutions in a common read-out set up: more specifically a set up on the front end board with channels each receiving a single analog line and outputting a 12 bit digital result.

2.1 Dynamic range transmitted on the cables

Particular attention will be paid to the dynamic range transmitted on the cable. That because transmitting a 16 bit dynamic range on the cables, assuming a 10V ADC amplitude means using e.g. (0.08mV,5V) or (0.16mV,10V) voltage range. That is a very hard task, both on the high and the low sides. For the smallest signal, controlling the pickup on the cable and preserving signal integrity, especially in presence of switching noise when the analog and digital parts are coupled, may be extremely challenging. Therefore solutions reducing the transmitted range to 12 bits will be considered with great favour.

3 Technique descriptions

3.1 Separate paths from successive phototube dynodes

Splitting the signal is performed at the earlier possible stage, taking the current signals both at the last dynode and at a previous one. The dynode gap depends on the amplification ratio desired. For example, in a 8 dynode tube the current at the anode and the 7th dynodes may be measured. The signal polarity at the dynodes and at the anode are opposite. The two signals are fed into separate preamplifiers.

The basic idea is that when the signal is large the current from the anode saturates the preamplifier, while the signal on the preamplifier connected to the dynode does not saturate.

An important problem that needs to be addressed with this configuration is that the saturated preamplifier does not have anymore a virtual ground, or equivalently its input impedance is not low anymore, due to the failure of the feedback. That means that the electron current will flow elsewhere. The exact pattern will depend on the biasing scheme, but it will probably flow through the biasing network. That might modify the bias of the tube in this signal region in a hard to predict fashion, unless particular care is taken in designing the bias network.

Furthermore saturation can be slow to recover and the saturated preamplifier might not be able to follow correctly the signal even after it falls below the nominal saturation threshold for a significant amount of time. This time can easily be longer than the clock period of 100ns. For example when the top of the trapezoidal signal saturates, samples of the leading and trailing edges with equal amplitudes can be processed differently, either with different amplifications or through a not fully linear amplifier. That would create major problem in the calibration. The same problem would emerge in any architecture based on saturating amplifiers.

There are at least two possibilities to transmit information to the front end board: either the choice between high and low gain path is done just after the preamplifier or both signals are transmitted for processing at the front end board. That gives a non-trivial connectivity problem: in the first case a digital signal with information on the chosen path has to be transferred to the front end board, in the second two signal cables are needed per channels and also the connectors on the front end board are doubled.

The extra circuitry required consists at least of the doubling of the chain up to the point of comparison, of a comparator and a switch, plus extra digital circuits to handle the extra bit.

– Compression capability:
The compression capability of this system is limited by the quantisation error of the low gain path in the point of switching between the two paths. Following arguments similar to the ones in [1], we can calculate the maximum input dynamic range that can be compressed to 12 bits with limited loss of precision as almost 20 bits. The spread in amplification ratios for different tubes is known to be high, probably the design must allow for a factor of two, so that the actual range will be 19 bits.
– Simplicity:
The solution does not appear quite simple. It requires significant doubling of the analog circuitry, a dedicated tube bias network, a switching system somewhere along the path and dedicated digital circuitry to handle the extra bit arising from the choice.
– Connectivity:
The dynamic range transmitted on the cables is limited to 12 bits.
Every channel need two signals from the very front end to the digital board. Even the connector are doubled in size.
– Analog-digital coupling:
Some degree of coupling is unavoidable. At least a switch and a comparator are on the path of the analog signal and the digital bit on the chosen amplification must be transferred to the digital part. As a result the ADC samples need to be complemented by an extra bit before use.
– Modularity:
There is no coupling between channels.
– Calibration:
It has the problem of matching smoothly the two regions of low and high amplifications, that requires dedicated calibration. Problematic in particular in presence of saturation of the preamplifier.
– Cost:
It doubles the cost of the very front end part, plus a comparator and connected circuitry and of cables and connector. Some extra cost in the digital part.
– Compatibility:
This configuration requires at least an extra digital signal at the output to discriminate between high and low amplification. Depending on the configuration, also a double analog input might be required. That is not easily compatible with the standard implementation described above.

3.2 Virtual channel summing ≈ 20 *channels with low gain*

The basic idea of this architecture is that the signals are concentrated only in a restricted region whose projection moves on the tube surface. Therefore, at a given time, the relevant signal of large amplitude is available only on one

tube, while the others have only noise.
If a number of channels (say 20) are added together analogically with reduced gain in a virtual channel, when one tube delivers a signal large enough to saturate the 12 bit ADC, the sum supplies a signal basically dominated by a single large signal plus noise from the other 19. Under the assumption that two channels added together never saturate at the same time, the signal of the saturated channel can be recovered at the virtual channel after appropriate rescaling.
These conditions are too much ideal because due to the finite width of the light spot on the tubes and to the physical size of the shower, saturation can happen at the same time on neighbouring channels. Therefore signals should be summed from not adjacent pixels.
The modularity of the detector is 400 pixel per mirror. Their signals are fed to the front end boards, each of them might host 20 channels from non neighbouring pixels plus a virtual one adding the others.
The extra circuitry required consists of one extra channel including the ADC every 20. The front end board must include the adder for the 20 channels.

- Compression capability:
  The compression capability is similar to the previous one, that is the same of all the systems with low and high amplification paths which exploit the full 12 bits of the ADC for both paths. It is about 20 bits, with maybe one bit to be reserved for gain differences in the tubes.
- Simplicity:
  The solution requires a limited amount of extra analog circuitry. A detailed design would be required to asses the circuital requirements on the front end boards.
- Connectivity:
  The dynamic range transmitted on the cables is the full range 16 bits.
  The connections between the very front end and the front end boards need careful planning to avoid adding neighbouring pixels.
- Analog-digital coupling:
  The choice of the virtual channel in case of saturation is delayed to the digital section, so that no analog-digital mixing is required in that regard.
  On the other hand, on the front end board the input of the 20 pixel channels must be connected to the input of the adder. That might create some analog-digital coupling.
- Modularity:
  That is a possible weak point of the architecture. The channels are strongly coupled to each other. If a single channel becomes very noisy, it might render useless all the 20 channels connected to its adder. That means that a low number of failures might blind a significant fraction of the detector. In order to avoid it, each channel should be disconnectable singularly from the adder, through a digitally controlled switch.
  Furthermore the common input to the adder of 20 channels might create

cross-talk between them.

– Calibration:
  It shares with other architecture the problem of connecting smoothly the two regions of low and high gain, with the additional problem that the same high gain channel is used for 20 channels.
  In case of use of a pulser system for calibration, that requires the pulser to be addressable separately for each channels.
  Again the presence of cross-talk might complicate the procedure.
– Cost:
  It requires one extra channel every 20, starting from the point where the sum is taken. Depending on the actual implementation it might require substantial additional cabling. Also the pulser system can be quite complex.
– Compatibility:
  The requirements of this architecure are compatible with the standard implementation described above.

3.3 Switched amplification paths

This compression architecture consists of splitting the line in a low and high amplification paths selected according to the signal amplitude.
Two different schemes are considered. In the first the amplitudes are stored on capacitors, then a comparator converts the non-saturating value recording the chosen path as an extra bit. Depending on the implementation the capacitor may perform integration or track&hold.
In the second the signal is sent on a delay line and to a prevision circuit that switches the low gain if required with appropriate timing, so that after the line the signal is processed uniformly with uniform gain throughout the signal duration. The information on the chosen path is made available asincronously to the digital section.

– Compression capability:
  The compression capability is similar to the previous one, that is the same of all the systems with low and high amplification paths which exploit the full 12 bits of the ADC for both paths. It is about 20 bits, with maybe one bit to be reserved for gain differences in the tubes.
– Simplicity:
  The first solution does not appear particularly simple, it requires a certain amount of analog circuitry for the high amplification path, plus switches, capacitors, comparators and digital circuitry to handle the low-high amplification cases. Such a circuitry would be best implemented as an integrated circuit.
  The second requires a delay line of appropriate length and a delicate prevision circuit to handle the timing of the amplification switch. Furthremore

the switching signal is asyncronous with the ADC clock.

– Connectivity:
  For the first circuit the switch circuitry needs necessarily to be implemented on the front end board, therefore the full dynamic range needs to be transported on the cable. The cabling is simple: one cable per channel.
  For the second both options are possible, but in case of implementation on the very front end probably two cables per channel are required.
– Analog-digital coupling:
  For the first solution it might be a serious limitation. The digital and analog parts are tightly coupled and the switches are positioned on the signal paths. It might be difficult to prevent the low signals to be spoiled by the effects of the switches. Furthermore the comparator must treat analog and digital signals at the same time.
  For the second solution it is very limited, one switch is activated in presence of large signals. The prevision circuit must supply the bit with the chosen path.
– Modularity:
  There is no coupling between channels.
– Calibration:
  The first solution shares with other architectures the problem of connecting smoothly the two regions of low and high amplification.
  The second might, with an appropriate choice of the timing of the prevision circuit, allow to process a signal through a single path, simplifying its treatment.
  For both solutions the extra stage on the high amplification path has a stable gain but may add extra shaping, so that the sampled waveforms may differ.
– Cost:
  The requirements of the first solution are an extra fast amplifier, plus a set of switches, capacitor and driving digital circuitry. It is probably as costly as the other solutions.
  The second requires a delay line and some analog-digital circuitry for the prevision circuit. The cost of the delay line must be investigated.
– Compatibility:
  This configuration requires an extra digital signal at the output to discriminate between high and low gain. That requires specific arrangements of the read-out, different from the standard implementation described above.
  A possible compatible solution might use one of the ADC bits to store the high-low amplification bit and decoding it later digitally. That would require compression of the signal to 11 bits.

3.4 Three separate amplification paths based on 8-bit ADC with staggered sampling

This architecture makes use of three amplifiers with different amplifications and a multiplexer selecting alternatively the three outputs at three times the clock frequency. This output is sampled by a 8-bit ADC at 30 MS and in the digital section the 8-bit of the largest not saturated signal plus the 2-bit of the multiplexer are used to reconstruct the full dynamic range.
This approach is based on the assumption that the signal is smooth enough that three samples taken at 30MS may be effectively combined in a signle sample.

- Compression capability:
  The compression capability is limited to approximately 16 bits if the precision has to be affected within the limit assumed in [1].
- Simplicity:
  It involves a not negligible amount of extra circuitry, both analog and digital and a three times faster clock.
- Connectivity:
  The circuitry needs necessarily to be implemented on the front end board, therefore the full dynamic range needs to be transported on the cable.
  The cabling is simple: one cable per channel.
- Analog-digital coupling:
  Likely to be a problem. The digital and analog parts are tightly coupled and a mixed digital-analog circuit like a multiplexer may render difficult the detection of low signals.
- Modularity:
  There is no coupling between channels.
- Calibration:
  It has the problem of connecting smoothly three regions of different amplifications, that maybe highly not trivial. Furthermore the calibration will be complicated by the fact that the three 8-bit samples refer to different part of the signal. The requirements on the bandwidth limitation of the signal may become very strict.
- Cost:
  It requires an three amplifiers, plus a multiplexer, but uses a cheaper ADC. Some digital circuitry is needed to combine the information and take care of the saturated paths.
  A cost estimation has already been done and the solution is comparable to (and maybe a little cheaper than) other solution.
- Compatibility:
  The output at 8 bits plus 2 bits for the amplification paths and the requirements on the clock frequency makes this architecture of difficult compatibility with the other solutions.

A compatible implementation would require on board generation of the clock, a combination of the 2 amplification bits with the other 8 and digital decoding later.

3.5 Dynamic compression

This architecture requires a circuit performing a dynamic non-linear compression of the signal just before sampling. It is demonstrated in [1] that a bilinear compression function suffices to guarantee adequate resolution. A pieceweise linear compression function seems preferrable to other functions because of ease of calibration and stability of the compression function.
A possible block diagram consists of the preamplifier followed by the anti aliasing filter, a variable gain amplifier, the compressor and the ADC. A driver and a receiver must be included at the extremes of the cables. Probably the functions of receiver and variable gain amplifier can be integrated in the same circuit.
The position of the compressor is an important point, if it is at the very front end side it will allow to reduce the dynamic range transmitted on the cable significantly. On the other hand, that requires that no significant shaping takes place between the compressor and the ADC.

- Compression capability:
  As demonstrated in [1], a bilinear compression function is adequate to compress up to 18 bits .
- Simplicity:
  It requires a specifically designed compressor circuit. Its design does not seem particularly difficult.
- Connectivity:
  If the compressor is positioned on the very front end board, only 12 bits need to be transferred.
  Cabling is simple: one cable per channel.
- Analog-digital coupling:
  There is no coupling between analog and digital part.
- Modularity:
  There is no coupling between channels.
- Calibration:
  The question of stability and calibrability of a compression function is a delicate point. It is discussed at length in [3] and reference therein.
  The conclusion is that using an adequate pulser the compression function can be calibrated within 1%. That is probably better than the needs of AUGER.
  The requirement on the pulser and the effect of having pulses of variable forms may need some specific studies.

– Cost:
  It requires the compressor circuit in the analog section and a lookup table to invert compression in the digital part.
  If the compressor is on the very front end board, several requirements on cabling, driver and receiver can be relaxed so that cheaper solution could be implemented.
– Compatibility:
  The required set up is fully compatible with the standard implementation proposed above.

4 Conclusions

Several architectures suggested for compressing the dynamic range from 16 to 12 bits have been shortly revised. Some of them are presented with different implementations. Their advantages and disadvantages are assessed.
Some of the proposed solutions seem to be particularly critical from the point of view of connectivity or of dynamic range transmitted on the cable. Also the strong coupling of analog and digital circuitry may be a delicate problem in presence of very low level signals.

References

[1] P.W. Cattaneo. The compression function for the fluorescence detector of the auger experiment. GAP-98-??

[2] B. Dawson. Amplitude dynamic range in auger fluorescence electronics. GAP-97-064.

[3] Z. Ajaltouni et al. Evaluation of fermi read-out of the atlas tilecal prototype. *NIM*, A403:98–114, 1998.