

Fermi National Accelerator Laboratory

FERMILAB-Conf-94/370

DØ Upgrade Muon Electronics Design

B. Baldin, D. Green, H. Haggerty and S. Hansen

*Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510*

November 1994

**Presented at the *IEEE 1994 Nuclear Science Symposium and Medical Imaging Conference*,
Norfolk, Virginia, October 30-November 5, 1994**

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

D0 Upgrade Muon Electronics Design

B. Baldin, D. Green, H. Haggerty, S. Hansen.

Fermilab, P.O. Box 500, Batavia, IL 60510

Abstract

The planned luminosity for the upgrade is ten times higher than at present ($\mathcal{L} \sim 10^{32} \text{cm}^{-2}\text{s}^{-1}$) and involves a time between collisions as small as 132 ns. To operate in this environment, completely new electronics is required for the 17,500 proportional drift tubes of the system. These electronics include a deadtimeless readout, a digital TDC with about 1 ns binning for the wire signals, fast charge integrators and pipelined ADCs for digitizing the pad electrode signals, a new wire signal triggering scheme and its associated trigger logic, and high level DSP processing. Some test results of measurements performed on prototype channels and a comparison with the existing electronics are presented.

I. INTRODUCTION

The D0 muon system [1] consists of 5500 channels of 3 cm diameter drift tubes installed in the small angle region and 12,000 channels of $4 \times 10 \text{ cm}^2$ rectangular drift tubes covering the wide angle region. The smaller tubes (SAMUS) are arranged in three layer planes of $312 \times 312 \text{ cm}^2$. These tubes provide only drift time information. The large drift tubes (WAMUS) are arranged in panels of 72 or 96 cells depending on the panel thickness (three or four layers). The drift times are used to measure the coordinate orthogonal to the wire. The wires of adjacent pairs of large tubes are tied together at the end opposite the preamplifier connections. By measuring the difference in arrival time of signals from adjacent wires, a coarse measurement of the coordinate along the wire can be made. The WAMUS tubes are outfitted with vernier pad electrodes that form a repeating pattern with a pitch of 61 cm. The transit time measurement (Δt) can resolve which section of the pads produced the signal. The pad charge ratio is used to refine the measurement to about 5% of the pad length, $\sim 3 \text{ mm}$.

The present time between collisions is $3.2 \mu\text{s}$ and the drift times are about 150 and 700 ns for the small and large tubes respectively. The planned luminosity for the upgrade is an order of magnitude higher than present with a minimum time between collisions of 132 ns. If this interval is less than the drift time, this gives rise to a series of problems in determining which collision is associated with a given drift chamber signal.

The present electronics [2,3], based on analog measurement and buffering of wire drift times and pad signal charges, can accept no more than two hits originating from one collision and the pad electronics only one. Time

measurements are based on time-to-voltage (TVC) converters and pad signals are integrated by charge sensitive preamplifiers. A two event deep analog buffering scheme is used. Analog data is multiplexed 96:1 into scanning ADCs. Sequentially scanning and converting so many channels takes too long to keep up with the trigger rates expected from the proposed upgrade of the Fermilab collider.

II. NEW FEATURES OF THE MUON ELECTRONICS

A. Wire Amplifier and Discriminator

A new wire amplifier based on the Analog Devices AD8001 and the Harris HF1105 monolithic operational amplifiers has been designed. The amplifier has a transformer coupled input which rejects common mode and low frequency noise. It has a rise time of 8 ns and an input noise level of 85 nA (RMS) for the bandwidth of interest. This allows us to operate with a discriminator threshold as low as $0.5 \mu\text{A}$, compared with the existing threshold level of $2.0 \mu\text{A}$. A discriminator using the Maxim MAX913 comparator with capacitive feedback to set the output pulse width at 40 ns is used. This comparator has excellent input overdrive versus delay characteristics. The 2x to 20x over threshold propagation delay difference is less than 2 ns compared to the more than 5 ns for our present comparator. The comparator has complementary outputs to match the differential inputs of the TDC chip we have selected. This arrangement has the additional benefit of less feedback to the amplifier input than single ended outputs. The power consumption has also been reduced significantly. It is less than 15 mA from $\pm 5 \text{ V}$ supplies as compared to 47 mA for our present amplifier and discriminator.

B. Wire Signal Triggering

The WAMUS pads have 700 pF of capacitance which limits the noise performance and bandwidth of their preamplifiers. Because the pad signals are unique to one tube while the wire signals are not, due to its jumper wire, the present triggering scheme is based on the discrimination of a differentiated pad signal integrator output. The 400 ns width of this signal is not a problem at a crossing time of $3.2 \mu\text{s}$, but is too long for one of 132 ns. The wire signal is much faster than the pad signal, but the tube pair ambiguity problem must first be resolved in order to use it as a trigger. The signals coming from two adjacent wires need to be separated in order to determine which of the drift cells is hit.

A wire signal will always arrive first at the near end of the tube in which it occurred before crossing the jumper and traveling the full length of the adjacent tube and appearing at its near end. To determine first arrival, a simple two D-type flip-flop separator circuit clocked by the wire signals with cross-connected Q-bar outputs to D inputs is used. Both flip-flops are reset upon generation of the output signal. When the difference in arrival time is less than the propagation delay and setup time of the flip-flop this scheme breaks down. Replacing the present far end jumper with a delay line of 20 ns solves this problem. The custom delay line, a DL3855 from Dataonics, is matched to the 330 ohm tube impedance and is AC ground referenced using HV capacitors. Fig. 1 shows the response of the wire signal separator with this delay installed on a pair of test tubes measured using the present Δ TVCs on cosmic rays. Scintillator counters were located at the end of the drift tubes which guarantees the minimum time difference between two wire signals. Fig. 1a shows the Δt distribution without regard to the separator output. Figs. 1b and 1c show the same data qualified by the separator outputs. There is clean separation between the two tubes.

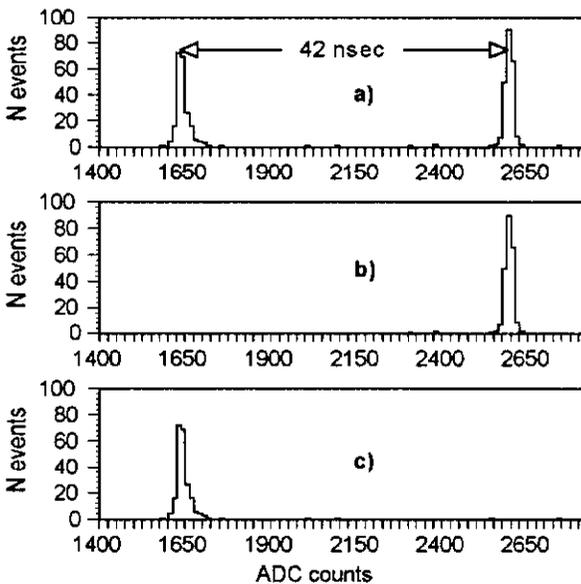


Fig. 1 Δt distribution with a $10 \times 10 \text{ cm}^2$ trigger scintillator at one end. a - all signals, b - wire 1 tag, c - wire 2 tag.

A block diagram of the logic used to form hit maps for the first level trigger system is shown in Fig. 2. The amplifiers (A) and discriminators (D) are used to generate logic pulses from the wire signals. The complete hit map for a particular collision cannot be determined until the maximum drift time has elapsed because the tube signals are delayed by an unknown amount up to the maximum drift time. The strobe for the trigger logic must then be a copy of the crossing signal delayed by this amount. The separator outputs fire retriggerable digital one-shots clocked by the 53 MHz

accelerator RF, whose output widths are set equal to the maximum drift time of the detectors. The purpose of the one-shots is to lengthen the discriminator signals to such an extent that the earliest arrivals from a particular collision persist until the sampling time. The centroid logic is an AND/OR network that takes advantage of the staggered

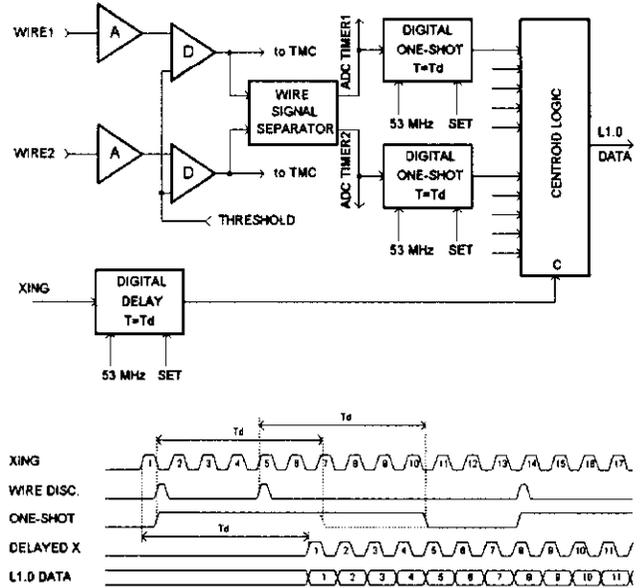


Fig. 2 Wire signal trigger logic block diagram and timing.

layering of the muon tubes to convert the physical multi-layer hit map that has a pitch of 1 tube width to a logical map that is one layer deep with a pitch of $1/2$ tube. Since the storage time of the tubes is greater than one crossing interval, the trigger logic cannot correlate hits and crossings. It can only produce a list of trigger candidates. Additional high time resolution detectors have to be included in the Level 1.0 trigger system in order to determine the crossing from which the hits originated.

C. Pad Signal Gated Integrator

The existing pad electronics is the commonly used charge sensitive preamplifier followed by a dual base line subtractor. The integration time constant of the preamplifier is very long ($40 \mu\text{s}$) and limits the counting rate of the channel to a few kHz. This scheme only works when the beam crossing intervals are greater than the drift time, and the likelihood of multiple hits in this interval is small, which is clearly not the case for the upgrade collider mode.

In the new design, each pad channel is triggered by its corresponding separator output. This results in a sampling interval that is fixed with respect to the pad signal. Studies done on the present pad data which is digitized to 12 bits of precision show that 10 bits of ADC resolution is sufficient. A

plot of coordinate resolution versus ADC resolution is shown in Fig. 3. This allows us to use a newly available low cost pipelined ADC followed by a more reliable digital buffer instead of the analog buffers presently used. The delay

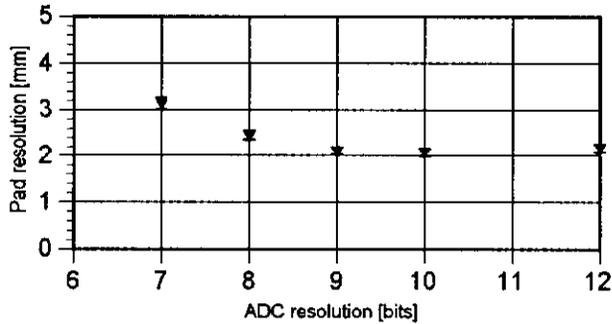


Fig.3 Pad coordinate resolution versus ADC resolution.

between the sampling point and valid data out of the Analog Devices AD875 is three clock cycles, which allows baseline sampling without any external delays.

Fig.4 shows the block diagram of the pad signal processing. The preamplifier uses RF bipolar transistors and is optimized for operation with the high capacitance of the pad. The ADC and memory buffer are continuously clocked at 53/4 MHz. Signals synchronous with the ADC clock are generated by the ADC timer which is triggered by the asynchronous separator output. The timer generates PEDESTAL and SIGNAL bits which are flag bits appended to the data stream to facilitate later zero suppression and pedestal subtraction. The integrator is reset after its signal

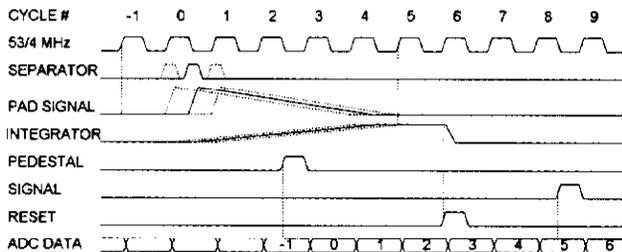
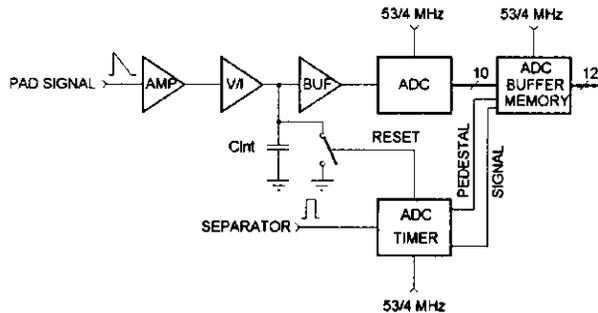


Fig.4 Pad signal processing block diagram and timing.

has been sampled by the ADC. The memory buffer is a FIFO memory with external logic that allows the difference between the read and write pointers to be programmable. Level 1.0 triggers are formed while the data is clocking through the buffer. The length of the buffer is adjusted so that the data emerges approximately in time with the arrival of the L1.0 trigger accept.

D. Digital Time Measurements

The requirement on time resolution of the muon system is driven by the Δt measurement. The wire signal traverses one 60 cm vernier pad in about 3 ns. We would like to have a resolution of at least 30 cm or about 1.5 ns. Time measurements are made by a four channel TDC chip (TMCTEG3) developed for the SSC by a group at KEK [4]. The TMC chip has a bin width of 0.78 ns at its nominal clock frequency of 40 MHz. The chip provides five bits of interpolator data and a hit flag bit for each channel. The data is stored in an internal 128 deep memory with independently settable read and write pointers. We are going to operate the chip at a frequency of 53/2 MHz which will give us ample resolution (1.2 ns binning) and sufficient storage time for the Level 1.0 decision time. At this frequency, the data can be stored for as long as 4.8 μ s before it is overwritten, which is less than the Level 1.0 decision time of 4.1 μ s. This makes possible a deadtimeless synchronous data readout which is very attractive.

The Δt measurement is done indirectly by subtracting the two drift times for adjacent tubes as opposed to the present system, which performs the Δt measurement in hardware. Our studies using the existing analog electronics show that there is no significant difference between the two methods for determining Δt .

E. Deadtimeless Readout

The principle advantage of the new system is its storage of digitized data in circular buffers pending a trigger decision, with no data loss. The TMC chip has its own 128 deep buffer controlled by its read and write pointers. The difference in the pointer settings allows the data to be delayed for the necessary trigger decision time. For ADC data an external memory buffer is used which runs in a manner similar to the TMC internal buffer. For the sake of simplicity, all the components of the front end system are synchronized to subharmonics of 53 MHz. Because of the lower bandwidth of the pad channel, a 53/4 MHz frequency is used for its synchronization. Fig.5 shows the block diagram of the data readout and data processing scheme. Six bits of data from an external coarse time counter are appended to the TMC output in order to extend its range to span the full drift time interval. The bunch spacing occurs at an uneven subharmonic (1/7) of 53 MHz which makes it necessary to run the coarse time counter at 53 MHz in order to resolve the ambiguities

between bunches occurring at the beginning or the middle of a TMC clock cycle.

Level 1.0 and Level 2.0 decisions come from the D0 trigger framework down to the front ends mounted on the chambers. In order to reduce the likelihood of synchronization problems and identify the crossing at which the trigger has occurred correctly, the trigger framework provides an eight bit crossing number to the front end

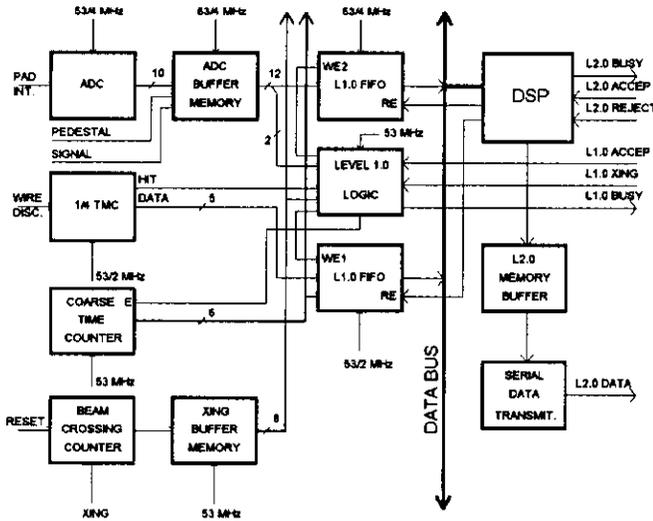


Fig.5 System level block diagram.

electronics. This number is compared to a local crossing counter value, which is stored in a buffer memory of the same type as those used within the TMCs and attached to the ADCs. The data appearing at the output of the circular buffers must be directed to a second level of FIFO memory buffers used to store Level 1.0 accepts at the point exactly correspondent to the crossing generating the accept. This could be done by performing timing adjustments on all the trigger signals; however achieving a timing uniformity of 18 ns over the entire muon system is a daunting prospect.

By comparing the values of the local counter and the Level 1.0 crossing number, the time of arrival of the trigger is no longer critical. However, the buffer delay setting must be greater than the Level 1.0 decision time. The duration of the data transfer to the Level 1.0 FIFOs is programmed within the Level 1.0 logic and is equal to the maximum drift time for TMC data. For the ADC data, because of the rise time of the integrator signal, and a one clock cycle uncertainty in the arrival time of the asynchronous pad signal, the transfer takes about 0.5 μ s longer. Based on the flag bits previously appended to the data stream, only valid hits are stored. There are event boundary marker words embedded in the data written to the Level 1.0 FIFOs, so several events can be stored in a buffer.

F. DSP Processing

The outputs of the FIFOs are connected to a data bus, controlled by a DSP processor. We are at present evaluating various processor candidates. On the receipt of a Level 1.0 accept, the processor starts its readout sequence, performing the necessary checks of the FIFO flags and event boundary markers. Level 2.0 buffering is performed by the DSP. For each Level 1.0 accept, one event worth of data is transferred to the DSP memory to await the outcome of the Level 2.0 decision. The DSP keeps track of the pointers to the event boundaries. On a Level 2.0 accept, the DSP adds the event to its output queue. High level processing is performed on the events in this queue before being sent on to the data acquisition system. On a Level 2.0 reject, the buffer pointer is advanced to the next event in the Level 2.0 queue.

The time to transfer an event into the Level 2.0 queue should not be too large a fraction of the total processing bandwidth available from the processor; otherwise no significant preprocessing can take place. The stated maximum rate of Level 1.0 accepts is 10 kHz and for Level 2.0 accepts 1 kHz. This sets the scale for the time available to do the Level 1.0 and 2.0 operations. Our initial estimates show that the transfer to the Level 2.0 queue can be done in about 30 μ s, or 30% usage of the processor bandwidth. There is a small additional overhead involved in managing the queues. We believe this still leaves a reasonable fraction available for the as yet undefined operations.

There are several possibilities for this preprocessing. The immediate goals are packing the data into an efficient format and optimal parameterization for the next level of processing. Pedestal subtraction, calculation of the longitudinal coordinate using Δt and digitized pad charges, the matching of hits through multiple layers of tubes, since each processor has access to three or four layers' worth of information, are examples of possible tasks that could be performed. There are several 32 bit DSPs available now with 40 Mips performance. In the WAMUS case, the aggregate capacity of 160 of these devices is 6400 Mips. The opportunities for SAMUS processing are not as varied, due to the detector structure. System uniformity considerations dictate that the same processors will be used for both.

G. High Speed Serial Links

The trigger logic generates a new candidate list at a 53/7 MHz rate. The need to send this list from the front end trigger logic to the Level 1.0 trigger processors within one bunch crossing interval sets the link bandwidth requirement. One WAMUS chamber can generate a maximum of 96 hit bits. The centroid logic reduces the maximum number of bits to 48. Six bytes must be sent every 132 ns which is a rate of 45 Mbytes/s as compared to our present Level 1.0 bandwidth of 0.6 Mbytes/s. Until recently the only transmission option for this bandwidth would be a ribbon cable with 50 or more conductors. Even with 0.5 mm pitch cable, routing 160 inputs

of this sort into a few crates is a very difficult mechanical problem. At most eight such cables could be routed into a standard electronics module. For each Level 1.0 trigger card, 24 such inputs are typically required.

Reasonably priced serial links with rates as high as 1 Gbits/s developed for telecommunications have recently become available. These links meet the Level 1.0 bandwidth requirements and carry data on a single coaxial cable or optical fiber. The plan is to locate the Level 1.0 system near the detector, with relatively short interconnects which means that conversion to optical is not necessary. The most attractive device to date is the Cypress CY7B923/933 Hot Link transmitter/receiver set. Its present data transfer rate is 330 Mbits/s, inadequate for the Level 1.0 rate, however a 600 Mbits/s version is due to be released shortly. We would run the link synchronously with the accelerator RF at 530 Mbits/s.

The Level 2.0 data transfers require much less bandwidth. As previously stated, the specified maximum event rate is 1 kHz. The bandwidth requirements are defined by the SAMUS tubes because they have the highest occupancy. For a fraction of the events, there are 96 hits from a SAMUS station. If no processing were performed other than formatting, the event would have a size of about 120 bytes. The bandwidth at the maximum Level 2.0 rate would be 0.12 Mbytes/s. The lowest frequency at which the Cypress chip will operate is 16 Mbytes/s. To send 120 bytes at that rate will take 7.5 μ s, an insignificant fraction of the interval between Level 2.0 accepts.

III. CONCLUSION

The operating parameters for the upgraded luminosity the next colliding beam run at Fermilab present many design challenges. Electronics technology has developed at such a rapid pace that it will be possible to run much faster with lower deadtime than in the previous run without a significant cost increase compared to the existing system. High speed serial links and low cost, high speed digitizers and memories now make practical a completely different design strategy that results in smaller, lower power, and higher reliability systems.

IV. REFERENCES

- [1] S.Abachi et al. "The D0 Detector." *Nuclear Instruments and Methods*. A338, pp.185-253, 1994.
- [2] D.Green et al. "Accurate Two-Dimensional Drift Tube Readout Using Time Division and Vernier Pads." *Nuclear Instruments and Methods*. A256, p.305, 1987.
- [3] M.Fortner et al. "The VME based D0 muon trigger electronics" *IEEE Transaction on Nuclear Science*. vol. 38, pp.480-485, 1991.

- [4] Y.Arai, F.Sudo and T.Emura. "Developments of Time Memory Cell VLSI's." KEK Preprint 93-49, June 1993 H.