Hardware Trigger System for Fermilab E866

C. A. Gagliardi, E. A. Hawker, R. E. Tribble

Cyclotron Institute, Texas A&M University, College Station, TX 77843 USA

D. D. Koetke, P. M. Nord

Department of Physics, Valparaiso University, Valparaiso, IN 46383 USA

P. L. McGaughey

Physics Division, Los Alamos National Laboratory, Los Alamos, NM 87545 USA

C. N. Brown

Fermi National Accelerator Laboratory, Batavia, IL 60510 USA

Abstract

E866 is the fourth experiment to be performed with the Fermilab Meson East spectrometer. Originally constructed in 1980 for E605, this spectrometer has since been used for E772 and E789. During these three experiments, the hardware trigger has depended on components dating from the early 1970's. The trigger system was rebuilt for E866 to improve its reliability and increase its flexibility. The new system incorporates several new modules designed and built specifically for this experiment. Notably, a custom Track Correlator combines the features of a 16-bit memory lookup and a 12-bit prescaler into a single extremely fast CAMAC module. The new hardware trigger system and modules are described, and the system performance during E866 is discussed.

PACS Codes: 7.50.Ek, 29.90.+r, 13.85.Qk.

Keywords: Electronic Trigger System; Programmable CAMAC Trigger Module; Drell-Yan Muon-Pair Production.

Contact Author: Carl A. Gagliardi

Mailing Address: Cyclotron Institute
Texas A&M University
College Station, TX 77843
USA

Telephone: 1-409-845-1411

Fax Number: 1-409-845-1899

E-Mail: cggroup@comp.tamu.edu

1. Introduction

The NuSea Collaboration (Fermilab E866) has performed a high-precision investigation of the anti-quark composition of the nucleon through studies of Drell-Yan muon pair production in 800 GeV/*c* *pp* and *pd* collisions [1]. The experiment utilized the Fermilab Meson East spectrometer [2] that was constructed in 1980 for Fermilab experiment E605, and has since been used for experiments E772 and E789. A schematic view of the spectrometer is shown in Fig. 1. The spectrometer includes three dipole magnets, all of which bend in the vertical (y) direction, and four detector stations, each of which contains both wire chambers and scintillator hodoscopes. The scintillator hodoscopes provide the only inputs for the hardware trigger. There are seven hodoscope planes overall — Y1, X1, Y2, X3, Y3, Y4, X4 — that are divided into separate left and right sides for trigger purposes.

The hardware trigger for the Fermilab Meson East spectrometer was rebuilt for E866. In order to place the changes and improvements made in context, we first briefly outline the prior trigger. The hardware trigger decision was divided into two levels. The first level, constructed principally from commercial NIM electronics, identified simple patterns of hits typical of one or two muons passing through the spectrometer. When it observed a suitable pattern, it started the second-level hardware trigger, as well as critical elements of the read-out electronics. In parallel with the first level, a set of custom Trigger Matrix CAMAC modules [3] examined the hits in the Y1, Y2, and Y4 hodoscope planes to identify triple coincidences that were candidates for high p_T muons originating from the target. The outputs from the first-level trigger, together with the outputs from the Trigger Matrices, provided the inputs to the second-level hardware trigger. The second-level trigger system consisted entirely of custom-built modules which were designed in the early 1970's [4] and utilized MECL II circuitry. It was configured by setting a collection of jumpers and dip switches within the modules by hand. It could be used, for example, to select those events that contained muons on both the left and right sides of the spectrometer. In this case, it might be configured to accept all events including opposite-sign muon pairs, together with a prescaled sample of those events containing like-sign muon pairs to study the random coincidence rate. It could also be configured to accept a prescaled sample of other triggers for study purposes. If the event failed to satisfy any of the defined trigger criteria, the second-level trigger would issue a fast clear to those segments of the read-out system that had been started by the first-level trigger. For accepted events, the second-level trigger would start the rest of the read-out electronics.

The primary motivations for rebuilding the trigger system were to improve its reliability, while increasing its flexibility and simplifying its ability to be reconfigured. To achieve these goals, the rebuilt system eliminates all of the MECL II circuitry of the old system and is designed to be programmed and monitored through software. We also wanted to increase the acceptance for high p_T muon pairs, without compromising the acceptance for lower p_T muon pairs or increasing the overall trigger rate beyond the capabilities of the read-out system. To permit this, the new system introduces the ability to base trigger decisions in part on the correlations observed between hits in the Y4 and X4 hodoscope planes, in conjunction with the correlations that had always been available among hits in the Y1, Y2, and Y4 planes. This allows us to trigger on a large sample of events containing one high p_T muon in coincidence with a lower p_T muon. It also allows us to trigger on opposite-sign muon pairs when both muons are on the same side of the spectrometer. During the previous experiments, it was only practical to accept left-left or right-right pairs under special circumstances [5] because these same-side dimuons typically introduced a high trigger rate and corresponding dead time due to random coincidences. Finally, we took advantage of this opportunity to reconfigure the trigger as a single decision stage, thereby simplifying the dead time monitoring and read-out system gating significantly. This required the new trigger system to complete its entire decision making in no more time than used by the former first-level trigger alone, since the read-out electronics were not changed for E866.

This paper describes the new trigger system and discusses its performance. A forthcoming article [6] will describe the software that has been developed to configure and verify the trigger system.

2. Overall Structure

Figure 2 shows a block diagram of the new hardware trigger system. It consists of a collection of commercial and custom CAMAC and NIM modules, configured as a parallel pipeline, with all cable lengths adjusted to synchronize inputs to ±1 nsec at each stage in the pipeline. Signals from the various scintillator hodoscopes are sent to LeCroy 4416 discriminators. Their outputs are reshaped and synchronized to the accelerator RF clock by custom CAMAC pulse synchronizer/stretchers [7], providing single bucket time resolution at the Fermilab RF frequency of 53 MHz. The Y1, Y2, and Y4 stretcher outputs are sent to a set of 16 custom Trigger Matrix CAMAC modules [3]. The Trigger Matrices present the hits in eight Y2 counters and twelve Y4 counters to the address lines of a set of six 256 × 4-bit ECL SRAM circuits. The SRAM outputs are ORed together to form four bits, each of which is true to indicate that a specific Y1 counter hit constitutes a track candidate, when combined with the hits present in Y2 and Y4. The Trigger Matrix outputs are the ANDs of these internal signals with the four Y1 inputs. The Trigger Matrices search for hodoscope hit patterns characteristic of Drell-Yan muons produced at our target and discriminate against the intense flux of J/ψ muon pairs produced in the proton beam dump. These modules are configured into four groups according to the charge of the apparent muon track and whether it is on the left or right side of the spectrometer. The previous experiments with this spectrometer only utilized 12 Trigger Matrix modules, which was barely adequate to cover the Drell-Yan phase space in a limited p_T range. The additional modules for E866 permit complete coverage for Drell-Yan events produced in the target over a much broader p_T range, as well as Υ events produced in the beam dump.

The X4 and Y4 stretcher outputs are sent to a set of 4 Track Correlator CAMAC modules. Track Correlators were designed and built especially for E866 and are described in detail in the next section. Each Station 4 X-Y (S4XY) Track Correlator examines the eight X and seven Y hodoscope signals in a single quadrant of Station 4 — which is behind several meters of heavy shielding and is only sensitive to high energy muons. Identifying specific X-Y coincidence pairs provides a crude estimate in hardware of the p_T for each muon. Combining two such estimates with the outputs of the Trigger Matrices imposes a rough muon pair mass constraint at the hardware trigger level. This permits an increase in the high p_T Drell-Yan acceptance by factors of 2 — 5, without increasing the system dead time substantially. Each set of stretcher outputs ends at a Terminator OR; these provide separate logic ORs of all the scintillators on the left and right sides of each hodoscope plane.

The heart of the new trigger system is a pair of Track Correlators that identify our physics triggers. Each of these receives 14 input signals. Four signals from the Trigger Matrices indicate the presence of muon candidates, according to the charge and the side of the spectrometer in which the candidate appears. Eight signals come from the S4XY Track Correlators, tagging relatively low and high p_T individual muons in each quadrant. Two signals provide coincidences of the X1-X3-X4 Terminator OR outputs on each side of the spectrometer for diagnostic purposes. Each physics trigger Track Correlator also has two additional inputs. They were included to permit additional unanticipated logic to be incorporated into our trigger decisions as necessary during the run, but were not utilized during E866. In addition to the two physics Track Correlators, there is a third Track Correlator in the system that monitors all 14 individual Terminator OR signals. This provides a quick and simple means to construct multi-plane coincidences for diagnostics, testing, and cosmic ray studies.

Finally, the outputs of the three trigger Track Correlators are routed to the Master Trigger OR, which was also built specifically for E866. This NIM module acts as the central switchyard for the entire trigger system. The Master Trigger OR accepts up to 16 independent trigger inputs, 8 ECL and 8 NIM. Each trigger input has an associated enable switch on the front panel of the module to specify whether or not it should be included in the higher-level logic for the run in progress. Four of the inputs have been used during E866, three for the prescaled-OR outputs of the physics and diagnostic Track Correlators and one for a pulser. The Master Trigger OR combines the input triggers that are enabled at any given time to form an All Triggers signal, generates a fast output for each input trigger, independent of the state of its enable switch, which is sent to scalers and latches in the data acquisition system for monitoring purposes, and outputs DC levels indicative of the actual switch positions, which are read by the data acquisition system to validate the current configuration. In parallel with the trigger inputs, the Master Trigger OR accepts up to 7 separate busy signals from the beam gating and read-out systems, 4 ECL and 3 NIM, each with an individual enable switch. It combines these external busy gates with internal busy signals that it generates at the beginning of each accepted trigger to produce the overall experiment busy. The Master Trigger OR clocks the overall busy gate by the accelerator RF to ensure that the experiment busy is always in a well-defined state at the time the beam strikes the target. The Master Trigger OR produces a Valid Trigger output for each enabled trigger that occurs when the system is not busy. For each valid trigger, the Master Trigger OR generates gates internally for those parts of the read-out system that must respond most rapidly. It passes the trigger to commercial NIM logic to generate other gates.

The overall propagation time of the new hardware trigger system from the RF clock inputs of the synchronizer/stretchers to the Valid Trigger output from the Master Trigger OR is 91 nsec. This is 6 nsec faster than the former first-level trigger system alone, while providing greater event selectivity than the former first- and second-level hardware trigger systems combined. It is about 15 nsec faster than required by the read-out electronics.

3. Track Correlator Design

The Track Correlator combines the features of a 16-bit memory lookup and a 12-bit prescaler into a single extremely fast CAMAC module. Figure 3 shows a simplified schematic of the data flow through the Track Correlator. Table 1 gives a complete list of the CAMAC functions that it recognizes. Whenever the front panel inputs to the Track Correlator are enabled, it presents the 16 input data bits to the address lines of an Integrated Devices Technology IDT10504S10C 64K × 4-bit, 10-nsec ECL SRAM.

This can be used to select four arbitrary combinations of the 16 inputs as candidate triggers. In parallel, the Track Correlator accepts an accelerator RF clock timing input, which it clips to a fixed width of 5 nsec. Each of the four SRAM outputs is ANDed with the clipped RF pulse to "deglitch" and synchronize it, forming the four raw trigger bits. Subsequently, each of the four bits is sent to the global clock input of an Altera EPM7032-6 erasable programmable logic device (EPLD) that acts as a programmable 12-bit prescale circuit. Each prescaler raises a flag whenever it concludes that the next event should be kept. In parallel with the prescalers, the Track Correlator ANDs each of the four raw trigger bits with the current state of its associated prescaler flag, as set after the previous event. (The 5 nsec widths of the internal raw trigger bits are chosen to be short compared to the propagation delay through the prescale circuits, plus their associated level translators.) The Track Correlator then outputs the four raw trigger bits, the four prescaled trigger bits, and an OR of the prescaled trigger bits, after reshaping all nine signals to fixed 12-nsec width pulses. These nine signals are in time with each other to better than ±1 nsec. By using the raw outputs alone, disregarding the various prescaled outputs, the Track Correlator can be used as a very high speed 64K × 4-bit memory lookup unit. The S4XY Track Correlators are used in this mode. A VETO input can suppress the generation of the clipped RF signal if it is NIM true, thus disabling the rest of the Track Correlator circuitry for that beam pulse. However, the VETO inputs have not been used during E866.

All of the glue logic in the Track Correlator data path is implemented in MECL 10HK to optimize its overall speed. It has a total propagation delay from the data inputs to the nine outputs of 20 nsec and can operate at frequencies over 60 MHz, thus providing single pulse trigger resolution at the 53 MHz Fermilab RF frequency. These speeds are determined by the 10-nsec ECL SRAM that we selected as a compromise between cost and speed. Had we utilized the 7-nsec version of the same SRAM chip, the Track Correlator would operate at frequencies of up to 75 MHz with a propagation delay of 17 nsec. Even at this speed, the worst-case manufacturer’s specification for the SRAM propagation, rather than the prescale circuitry, would dictate the overall rate capability of the module.

Each prescale circuit contains a 13-bit counter register and a 14-bit preset register. With each event, the counter increments. The highest order bit of the counter is the flag that instructs the Track Correlator to accept the next event. When that event arrives, the prescale circuit reloads the counter with 13 bits stored in the preset register. Thus, the prescale circuit can be programmed to accept 1 in *N* events, where *N* can take on any value from 1 to 4093. The fourteenth bit in the preset register instructs the prescale circuit to freeze the counter register at its current value. It is used to program the prescaler to pass no events and provides an alternative way to program the prescaler to pass all events.

During testing with noise and cosmic ray triggers prior to beam turn-on, we found that rare, short, out-of-time inputs are presented to the Track Correlators, in spite of the use of the pulse synchronizer/stretchers. These can lead to very short (<2 nsec wide) clock inputs to the prescale circuits that fire some, but not all, of the internal counter flip-flops. In extreme cases, this led to a very long count until the next event was accepted when the circuit had been programmed with a small prescale factor *N* greater than one. Logic was added to the prescale circuits to inspect the counter value routinely. If a value is found that is inconsistent with the programmed prescale factor, the prescale circuit enters a special mode in which it accepts the next event, regardless of the number of events since the last one was accepted. With this modification, the prescale circuitry has been very reliable.

Three CAMAC commands provide the ability to read and write the preset values and to read the counter value for each prescale circuit. There is no direct command to load the counter registers with specific values, but there is a special command that instructs the Track Correlator to load all four prescale circuit counter registers with their respective preset values simultaneously. Thus, the standard programming sequence for the Track Correlator prescalers is to load all four preset values, then transfer these values to their counters with a single F8 command.

Six CAMAC commands involve the SRAM circuit. An EPM7032 EPLD acts as a 16-bit register to hold the address that is used for CAMAC commands that access the SRAM. The 4-bit SRAM contents may be read or written in either “single” or “block” mode. In single mode, the address register contents are unchanged during the CAMAC instruction. In block mode, the address register automatically increments after the SRAM read or write takes place to prepare for the next command. When the address register increments past 64K, it sets a flip-flop that rejects all further SRAM commands until a new value is written into the address register explicitly. This facilitates block mode SRAM reads and writes using the standard CAMAC Q-repeat protocol.

One additional EPM7032 EPLD acts as the master control circuit for the Track Correlator. It decodes and responds to all CAMAC commands, serves as the master for the internal 16-bit bidirectional bus that is used for all CAMAC reads and writes, and enables and disables the other circuitry, as appropriate. The Track Correlator has two primary modes to insure that the module is in a well-defined state whenever it is being programmed. It must be issued a CAMAC F24 instruction before it will respond to any other CAMAC commands. In this state, the front panel inputs are disabled and the contents of the address register are presented to the SRAM address inputs. The RF clock is also disabled to insure that the prescaler counters will remain constant, rather than responding to signals arising from SRAM reads or writes. In this configuration, it is possible to issue any other desired CAMAC commands. Once the Track Correlator is issued an F26 instruction, the front panel inputs are presented to the SRAM circuit, the address register outputs are disabled, the RF clock begins cycling, and the module is in normal data-taking mode. A front panel NIM input can also be used to disable the front panel inputs of the Track Correlator, for example between beam spills, but it doesn’t activate the CAMAC programming capabilities. The Track Correlator has two status LEDs, one that indicates whether the front panel inputs are enabled or disabled and one that indicates that it is responding to a CAMAC command.

Each Track Correlator is packaged as a 2-wide CAMAC module to provide sufficient front-panel space, even though it only consists of a single printed circuit board. It consumes 0.8 Amp of +6 V and 1.9 Amp of −6 V during normal operation.

4. Trigger System Performance

The new trigger system has now been used to recommission the Meson East spectrometer and to take data for nearly a year while Fermilab E866 has run, approximately six months to study *pp* and *pd* collisions investigating the anti-quark composition of the nucleon and the balance studying other issues involving muon pairs produced in *pA* collisions. The *pp* and *pd* studies [1] illustrate some of the typical ways we have programmed the system. The *pp* and *pd* data were taken with three different spectrometer magnet configurations — a low mass setting focusing on the J/ψ, the ψ', and the Drell-Yan continuum below 7 GeV, an intermediate mass configuration providing good acceptance from a virtual photon mass of approximately 5 GeV to the Υ, and a high mass configuration providing good acceptance from 7 GeV to 15 GeV.

Each setting required a different trigger optimization to maximize the data-taking efficiency while minimizing the dead time, so we can achieve less than 1% systematic normalization uncertainty between the data with our liquid hydrogen and deuterium targets. (The data rates with the deuterium target were typically a factor of ~2 higher than with the hydrogen target.) For the low mass setting, previous experience from E772 indicated that we would have large corrections for random coincidences and rate-dependent reconstruction efficiencies if we increased the proton beam intensity much past 5×10^{11} per 20 sec spill. At this low beam intensity, we were able to operate with a relatively simple physics trigger. It required either a pair of Trigger Matrix hits, one on each side of the spectrometer, or an opposite-sign pair of Trigger Matrix hits on the same side of the spectrometer, either left-left or right-right, together with X-Y coincidences observed by the S4XY Track Correlators in both the upper and lower quadrants of that side of Station 4. The latter guarantees that there are two separate muons present in Station 4. This significantly reduces the probability that the Trigger Matrices fired on a single muon track in Station 4 in random coincidence with track fragments in Stations 1 and 2 from the intense flux of low energy muons originating in the proton beam dump. This trigger configuration led to an average of 2200 events per spill with the liquid deuterium target, including events that passed one of the physics trigger conditions and additional events that satisfied any one of a number of study triggers, as described below. The overall trigger rate was approximately quadratic with the beam intensity. The average dead time was only 1%. Removing the up-down S4XY coincidence requirement for left-left and right-right pairs would have increased the overall trigger rate by a factor of 2, while making a negligible increase in the acceptance, given the spectrometer configuration that was used for these data.

For the intermediate and high mass spectrometer configurations, the proton beam intensity was restricted to 2×10^{12} per spill due to beam-induced heating of the cryogenic liquid targets and radiation safety concerns. For the intermediate mass data, the S4XY Track Correlators were programmed separately to identify X-Y coincidences anywhere in a given quadrant and in a restricted region of that quadrant containing most of the lowest p_T muons reaching Station 4. The latter region includes 14% of the area of each quadrant, but contained 93% of the muon tracks. The trigger conditions of the low mass configuration were augmented with the additional requirement that the event include no more than one low p_T muon, effectively vetoing most muon pairs from J/ψ production. A prescaled sample of events containing two or more low p_T muons was also taken for study purposes. The average event rate was 3200 events per spill with the deuterium target, with an average dead time of 2.1%. The average event rate would have been 3 times higher at the same beam intensity without the constraint on the number of low p_T muons in the event. Notably, early in the run when the proton beam duty factor was poor, the event rate with the same trigger was over twice as high at an average beam intensity of only 10^{12} per spill.

The detector rates throughout the spectrometer were sufficiently low in the high mass configuration that a very simple trigger accepting any event containing at least two Trigger Matrix hits sufficed. This led to an average of 2100 events per spill with the deuterium target at an average dead time of 1.5%.

In contrast to the measurements with the liquid hydrogen and deuterium targets, the thicknesses of the solid targets used during the *pA* collision studies were chosen so the average detector rates throughout the spectrometer were approximately equal, independent of target type. This allowed us to accept data at average trigger rates approaching 40,000 events per spill without encountering significant rate-dependent systematic uncertainties in our measured cross section ratios. Such high rates led to alternative trigger optimizations. For example, during one study of J/ψ production, the trigger system was programmed to accept all events that contained Trigger Matrix hits on both the left and right sides of the spectrometer, regardless of the charge states of the candidate muon tracks. Events were also accepted if they contained a pair of opposite-sign Trigger Matrix hits on the same side of the spectrometer whenever the S4XY Track Correlators reported observing either X-Y coincidences in both the upper and lower quadrants of the appropriate side of the spectrometer or an X-X-Y triple coincidence in either of the two quadrants. Figure 4 shows a contour plot of the reconstructed horizontal angles for opposite-sign muon pair coincidences when utilizing this trigger configuration. The small gap between the two sides of the spectrometer causes the reduced intensity near $\theta_x(\mu^+) = 0$ or $\theta_x(\mu^-) = 0$. The requirement that two separate X hodoscopes must fire to accept pairs with both muons in the same quadrant of Station 4 causes the shallow minimum near $\theta_x(\mu^+) = \theta_x(\mu^-)$. Elsewhere, the spectrometer acceptance for left-left, left-right, and right-right pairs is smooth.

Additional prescaled study triggers selecting events with a coincidence between the left and right X1-X3-X4 triple coincidences and events that contain at least one Trigger Matrix hit were taken in parallel with all of the various physics trigger configurations. Together, they provide a continuous monitor of the scintillator hodoscope efficiencies. Special dedicated runs to measure the hodoscope efficiencies accurately were taken from time to time, utilizing the diagnostic Track Correlator to select events containing at least five hodoscope Terminator OR signals on the same side of the spectrometer — including at least two of Y1, X1, Y2, and at least three of X3, Y3, Y4, X4. The special runs demonstrate that the Y1, Y2, and Y4 hodoscope planes that provide the inputs to the Trigger Matrices typically averaged ~99% efficient.

The study triggers also provide independent triggers that can be used to verify the performance of the trigger system as a whole on an event-by-event basis. This has been done by comparing the hodoscope hits observed in an event with the results of the trigger decision at several stages in the pipeline that were also recorded with each event. We find that the average Trigger Matrix efficiency was 98.5%. In contrast, the S4XY Track Correlators were better than 99.99% efficient, and the physics trigger Track Correlators, which needed to respond to narrower input pulses than the S4XY Track Correlators and experienced more of the out-of-time noise hits mentioned above, were 99.75% efficient. Thus, the individual hodoscope channels provided the largest contribution to the inefficiency of the overall trigger decision.

In conclusion, we have built and installed a new hardware trigger system, consisting of a mixture of commercial and custom CAMAC and NIM modules, for use with the Fermilab Meson East spectrometer during E866. The new system satisfies all of the original design criteria. It has now been used for data-taking for nearly a year, and there have been very few problems with the hardware during that period.

Acknowledgments

We would like to acknowledge the full NuSea Collaboration for the many useful suggestions during the design of the trigger system. This work was supported in part by the U.S. Department of Energy under Grant Numbers DE-FG03-93ER40765 and DE-FG02-88ER4016.

References

[1] E. A. Hawker *et al*., Phys. Rev. Lett., submitted (1997).

[2] G. Moreno *et al*., Phys. Rev. D **43** (1991) 2815.

[3] H. D. Glass, Ph.D. thesis, State University of New York at Stony Brook (1985).

[4] H. Cunitz and W. Sippach, Nevis Laboratory, Columbia University (unpublished).

[5] M. S. Kowitt *et al*., Phys. Rev. Lett. **72** (1994) 1318.

[6] P. M. Nord *et al*., Nucl. Instrum. Methods, to be submitted.

[7] R. Gray and J. P. Rutherfoord, Nucl. Instrum. Methods **A244** (1986) 440.

Table

Table 1: CAMAC commands recognized by the Track Correlator. All function codes except F24 are disregarded unless the front panel inputs have been disabled in software with a previous F24 command.

Function	Address	
0	0 — 3	Read Prescaler Preset Value
1	0 — 3	Read Prescaler Counter Value
2		Read SRAM Contents (single)
3		Read SRAM Contents (block)
4		Read SRAM Address Register
8		Load Prescaler Counters with Presets
16	0 — 3	Write Prescaler Preset Value
17		Write SRAM Address Register
18		Write SRAM Contents (single)
19		Write SRAM Contents (block)
24		Disable Front Panel Inputs
26		Enable Front Panel Inputs

Figure Captions

Fig. 1: Schematic view of the Fermilab Meson East spectrometer. The ring-imaging Cherenkov detector and the electromagnetic and hadronic calorimeters were not used during E866.

Fig. 2: Block diagram of the new hardware trigger system. Parallel ECL signal buses are indicated according to the number of channels they contain. All others are either NIM or single ECL signals. Only the logic for the left side of the spectrometer is shown. The logic for the right side of the spectrometer is identical. Additional logic for pulser and RF signal distribution is not shown.

Fig. 3: Data flow through the Track Correlator. The solid lines represent ECL signals, the long dashed lines represent TTL signals, and the short dashed lines represent NIM signals. Those buffers and level translators that are labeled with a C are enabled or disabled by the EPM7032 main control integrated circuit. To simplify the view, the control lines connecting the main control circuit to the prescalers, the SRAM, the address register, and the various buffers and level translators are not shown.

Fig. 4: Contour plot of the correlation between the reconstructed horizontal angles for opposite-sign muon pairs observed during a study of J/ψ production in *pA* collisions.