Progress Report for 1995-96

Xiaochun He, William Lee, G. A. Petitt, Ziyang Zhang
Georgia State University

RECEIVED
MAY 31 1996
OSTI

Contents

1	Introduction	2
2	Participation in E866	2
2.1	Major Activities	2
2.2	Simulation Effort	3
2.3	Hardware Setup	4
2.4	System Design	6
2.5	Code Development	7
3	Muon Arm Simulations using PISA and PISORP	8
3.1	Track and Road Finding Schemes	11
3.2	Discriminant Analysis	11
3.2.1	Formulation	11
3.3	Discriminant Analysis Results	13
3.4	Trigger Rate Studies of the Muon Identifier Using Anode Wire Readouts	17
3.5	Test data preparation	18
3.6	Results	18
4	Direct VMEbus Interface with Linux	23

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.

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

OK Emd 5/30/96

1 Introduction

We have been involved in several projects during the present contract period. These include completion of our work on the RD94 test runs performed at the Alternating Gradient Synchrotron (AGS) at Brookhaven; the PhD. thesis project of Ziyang Zhang[1] which was completed during the year, was based on this work. We have continued our Monte Carlo simulation work. This includes studies of trigger rates in the muon identifier of the PHENIX experiment for RHIC. In addition to this we have continued our involvement in developing upgrades for the PISA and PISORP simulation codes for PHENIX.

We are most heavily involved in work on the E866 experiment at Fermilab[2]. GSU has taken on the task of modifying the trigger system for this experiment. A third level trigger based on digital signal processors (DSP's)[3] mounted on a VME bus is being developed for this. We feel that this project will be a valuable training ground for our work on PHENIX. We expect that the expertise that we acquire in development of the Level-3 trigger system for E866 will enable us to contribute significantly to development of the PHENIX trigger and online event processing system, especially for the PHENIX Muon Arms.

These projects are discussed in detail in the following pages.

2 Participation in E866

X.-C. He, G. Petitt and W. M. Lee

We have been active participants in the E866 experiment since the spring of 1995 and have been involved during the year in a variety of activities aimed at preparing for the upcoming run scheduled to begin in July. The experiment will be carried out using the magnetic spectrometer that has been used in several previous studies of dimuon emission from products of reactions produced by 800 GeV/c protons. (See "Progress Report for 1994-95" from GSU).

2.1 Major Activities

A listing of our main activities over the past year includes the following:

- Our group hosted the collaboration meeting held in Atlanta in April.
- GSU has purchased the Taxis Instruments DSP software which includes optimized C compiler, C source debugger, assembler, linker and simulator running on Sun workstation for the E866 level-3 code development.

- Members of our group participated in workshops on the Data Acquisition System (DAQ) for E866 at Los Alamos in July and at Fermilab in September.
- Bill Lee, the graduate student in our group spent the month of August at Fermilab setting up equipment for the upcoming run next June. He worked with Chuck Brown, and with the other graduate students involved in E866: Rusty Towell, and Derek Wise. They made several mechanical modifications to all the discriminator crates; the crates were then cleaned, tested, and all non-functioning power supplies were replaced. Several meetings were held at this time among the graduate students to discuss possible thesis topics.
- One member of our group, Xiaochun He, recently took a week-long training course in Santa Barbara covering the SPOX[4] operating system running on DSP's. Tom Carey (LANL) and Paul Stankus (ORNL) also participated. This was followed by a week in Los Alamos installing and exercising the "Blazer board"[5] which contains the 3 Digital Signal Processors (DSP's) that are the basis for the Level-3 trigger upgrade to the DAQ. (The Blazer board was purchased for the E866 collaboration by ORNL and was delivered about two months ago.)
- Based on discussions at the workshop in Los Alamos in June, we have worked out the design of the Level-3 processing system. The PDL (Program Design Language) for each of the tasks comprising the processing system has been written and technical notes describing the function of each task have been written and posted on the GSU Web pages[6] for review by other members of the collaboration.
- A program development system has been set up at GSU for testing the Level-3 software.

2.2 Simulation Effort

For a better understanding of the E866 level-3 processing and its performance we have simulated the level-3 process in every step which includes decoding a "real" raw data records, finding the events within the data records, unpacking the events data, processing the events (track-finding) and analyzing the results. We started the simulation by executing the code on our Silicon Graphics (Indigo) workstation running IRIX 5.2, and then migrated the simulation code to

Our upgrade of the DAQ is focused on the Level-3 trigger system. The purpose of the new level-3 system is to tag events for speeding up the off-line analysis and to possibly filter out “bad” events. The computation time required to do this filtering will require that several processors, operating in parallel, be used. This can be accomplished by using multiple DSP’s mounted on VME boards. A block diagram of the E866 DAQ system, showing the transport bus, the VME bus and the DSP-based level-3 trigger module is shown in Fig. 1.

Several vendors make such VME boards on which are mounted as many

as six DSP's along with high speed memory. As indicated above, we have chosen the Blazer board from the Image & Signal Processing (ISP), Inc. for this purpose. The ISP Blazer sets a new price-performance standard for VME DSP boards based on the Texas Instruments TMS320C40. In its full hex C40 configuration three C40 processors, which are built into the motherboard, are accompanied by three additional C40 processors, which are implemented as TIM-40 plug-in modules. The full implementation delivers up to 300 MFLOPS of floating-point speed with a concurrent data transfer capability of more than 700 MB/sec. This system provides for concurrent operation of multiple processors on global as well as on local resources and for the transmission of data between a single processor or I/O device and any other processor. A Blazer layout of multiple processors and multi-master bus is shown in Fig. 2. From

Figure 2: The Blazer processors and multi-master bus for the E866 level-3 system

the simulations and the overall system performance study we have decided to use a Blazer board with three TMS320C40's and 8 Megabytes global DRAM. We have all components we need for the Level-3 study currently at LANL which will be shipped to GSU for software development and test within several weeks.

2.4 System Design

We have devoted a great deal of effort over the past year to designing the Level-3 processing system. The discussions at the workshop in Los Alamos last summer were very useful in planning the system. A block diagram showing the layout of the resulting system is shown in Fig. 3. In this diagram the data

Figure 3: Block diagram of Level-3 processing system

moves from left to right. The elements shown with solid bars across the top and bottom are data storage pools. The element farthest to the left labeled FPP is the "formatted packet pool" into which formatted data packets are moved from the high speed memory modules (labeled HSM in Fig. 1). The elements with oval shapes are processes residing on the single board computer (SBC) on the VME bus or on one of the DSP's in the Blazer board. The elements in the upper half of the diagram reside on the SBC (two SBC's can be used if necessary) while the elements in the lower half of the diagram reside on the Blazer board. The most important elements of the Level-3 system are the so-called "filter-slave" processes labeled S1, S2,...,Sn. These processes which will reside on the DSP's will process and tag events. Events can both be written to tape and sent via Ethernet to the workstations shown in the upper

part of Fig. 1 for analysis. As indicated in Fig. 3 the data stream will be divided into n separate streams and then recombined. This scheme removes the simple data-driven pipeline character of the DAQ system. There are, in fact, $n + 1$ parallel data streams. The data stream is comprised of non-data packets interspersed with data packets. The non-data packets will not pass through the DSP processors but will pass along the uppermost data path in the diagram while the data packets will be divided up into n streams for processing by the DSP's. When the streams are re-combined new packed header records will be created containing the tag information. The data packets are kept within spill boundaries for correct normalization.

One of our tasks in this development effort is to create procedures for communication between tasks running on the single board computers, which use the VxWorks[7] operating system, and the DSP's which use a system called SPOX[4]. Program development is done in a similar fashion for both of these systems. Members of the E866 collaboration with experience in previous muon experiments at Fermilab are very familiar with VxWorks but SPOX is new to all of us and we know of no other real-time applications using shared memory in which SPOX and VxWorks have been integrated. The main purpose of the SPOX training course was to glean as much information as possible from the developers of SPOX about potential stumbling blocks in this integration. Initial tests of the Blazer board in a combined VxWorks-SPOX environment were performed during the week that Dr. He spent at Los Alamos last month.

2.5 Code Development

We have a development system already in operation at GSU shown in Fig. 4. It is comprised of our Silicon Graphics workstation which communicates by Ethernet with the VxWorks system which allows us to write the codes at the VxWorks end for the interface between the VxWorks and the SPOX. This system, in turn runs on a Force33 single board computer on loan from Los Alamos. We also have a Motorola mvme167 single-board computer which we obtained as part of the SSC equipment distribution program. The single board computer can also be housed in a VME crate when we need multiple VxWorks system across the VME backplane.

This system is fully networked with the local and Internet networks shown in Fig. 5. With this system we have been able to exercise the VxWorks operating system using plain C programs downloaded from the Silicon Graphics workstation. We anticipate that, within the next few weeks the Blazer board will be shipped to GSU and installed in the VME crate for further develop-

Code Development and Simulation Setup

Figure 4: GSU development setup for the E866 experiment

ment work. We need a SUN workstation at GSU in order to have a completely self-contained development system. This item is required because the compilers and linkers for both SPOX and VxWorks are available only for the SUN operating system. We hope to get a SUN workstation from the second SSC equipment distribution. If not, our Los Alamos colleagues will loan us a SUN workstation.

Our plans for further work with the E866 DAQ are discussed in the accompanying proposal for 1996-97.

3 Muon Arm Simulations using PISA and PISORP

X.-C. He, G. Petitt

We have continued our work on RHIC as part of the group working on the muon arm for Phenix. Our work has consisted of Monte Carlo simulation work on the performance of the muon identifier. Our interest has been in the

Level-3 Simulation Network

Figure 5: Network connections of the GSU development setup

ability of the muon identifier to discriminate between muons and pions and on expected trigger rates. Over the past year interest μ/π discrimination has decreased as it has become clear that the pion backgrounds will not be as large as formerly believed. We have, therefore, completed this work and have posted a note summarizing all our results in the notes section of the Phenix Arm Web page[8].

Our main techniques for μ/π discrimination are Discriminant Analysis and a technique that we have named "Hit Pattern Recognition" which is a simpler, more transparent technique than Discriminant Analysis but which appears to be equally powerful. The simulation procedure used for our studies of μ/π discrimination is shown schematically in Fig. 6. Single muon or pion events are launched from the vertex. When pions are the primary particles, the tracks are ignored if the pion decays to a muon before reaching the muon identifier. Muons resulting from pion decay can only be identified as decay muons by putting cuts on the origin of the particle track and this can only be done off line. However, we are interested in studying the performance of the muon identifier as a first or second level trigger based on the hit pattern

Generation of Single Particle Event Mu/Pi Discrimination Data

Muon arm is "empty" except for the three tracking chambers and the six elements of the Muon Identifier. In particular, the central magnet, lead shield, etc. are not installed

Figure 6: Simulation setup for μ/π discrimination studies

in the muon identifier. Therefore, as indicated on the figure, pion events were rejected if the hits in the third tracking chamber were produced by decay muons. To minimize the number of decay muons, the lead and neutron shields and the “nosecone” were omitted. By launching the particles at the vertex the resulting tracks have trajectories very similar to those in an actual experiment.

3.1 Track and Road Finding Schemes

The algorithms developed for muon identification in PISORP locate hits in the first plane of the muon identifier and then search for hits in the second plane. If any hits are found in plane 2, the centroid of the hits is calculated and this point is used as the starting point for a search for hits in plane 3. This procedure is continued from plane to plane until plane 6 is reached or until no more hits are found in the search range.

In the following we show results for three different track and road-finding schemes. They are summarized in Fig. 7. At the top of the figure the track and road-finding angles are given for the three cases. Case 1, labeled “CDR” uses pad sizes identical to those specified in the PHENIX “Conceptual Design Report” (CDR)[9][10]; case 2 uses the same pad sizes but larger search angles than used for case 1; case 3 uses smaller pads but search angles very similar to case 2.

3.2 Discriminant Analysis

3.2.1 Formulation

The technique of Discriminant Analysis [11],[12] is one of the methods we have used for discriminating between muons and pions. To perform this analysis one creates in an ad hoc fashion a set of variables that are functions of the hit patterns in the detector pads. These variables are chosen to hopefully produce different values depending on whether they are produced by a pion track or a muon track. This method can be thought of as treating each variable as the axis of a P-dimensional space, where P is the number of variables used. A particle event is then described as a point or vector with P components in this space. A series of events of one type (μ) will form a cluster of points in this space while events of another type (π) will hopefully form a separate cluster. These clusters can be separated by a hyper-surface lying between them. It can be shown [12] that for variables with a Gaussian distribution the best such surface is described by the discriminant function DF with $DF = 0$

Pad Search Sizes (Degrees)

Case 1 "CDR"

Plane #	Delta-Theta			Delta-Phi		
	Pad	Track	Road	Pad	Track	Road
1	1.0	---	2.2	2.8	---	7.0
2	1.0	1.1	2.2	2.8	5.7	7.0
3	1.0	1.1	2.2	5.6	5.7	7.0
4	1.0	1.1	2.2	5.6	5.7	7.0
5	1.0	1.1	2.2	5.6	5.7	7.0
6	1.0	1.1	2.2	5.6	5.7	7.0

Case 2 "CDR plus"

Plane #	Delta-Theta			Delta-Phi		
	Pad	Track	Road	Pad	Track	Road
1	1.0	---	3.2	2.8	---	7.0
2	1.0	2.1	3.2	2.8	5.7	7.0
3	1.0	2.1	3.2	5.6	5.7	7.0
4	1.0	2.1	3.2	5.6	5.7	12.0
5	1.0	2.1	3.2	5.6	5.7	12.0
6	1.0	2.1	3.2	5.6	5.7	12.0

Case 3 "1.0,2.5"

Plane #	Delta-Theta			Delta-Phi		
	Pad	Track	Road	Pad	Track	Road
1	1.0	---	3.0	2.5	---	7.0
2	1.0	2.0	3.0	2.5	5.0	7.0
3	1.0	2.0	3.0	2.5	5.0	7.0
4	1.0	2.0	3.0	2.5	5.0	7.0
5	1.0	2.0	3.0	2.5	5.0	7.0
6	1.0	2.0	3.0	2.5	5.0	7.0

Figure 7: Pad search parameters.

where

$$DF = (\vec{x} - \vec{\bar{x}}_\mu)' C_\mu^{-1} (\vec{x} - \vec{\bar{x}}_\mu) - (\vec{x} - \vec{\bar{x}}_\pi)' C_\pi^{-1} (\vec{x} - \vec{\bar{x}}_\pi) \quad (1)$$

where the $\vec{x}$ are the values of the P-dimensional vectors and the $\vec{\bar{x}}$ are their average values. The quantities C^{-1} are elements of the inverse of the covariance matrix given by

$$C_{ik} = \frac{1}{N-1} \sum_{j=1}^N (x_{ij} - \bar{x}_i)(x_{kj} - \bar{x}_k) \quad (2)$$

where x_{ij} is parameter- i from event j .

Applied to the present data, if the clusters of parameter values for muon (μ) and pion (π) events are well separated we will have $DF \leq 0$ for muons and $DF \geq 0$ for pions. In practice, the clusters are not completely separated and some of each type of particle will be mis-identified. It is frequently desirable to choose some value $DF \neq 0$ to produce an acceptable combination of pion rejection and muon acceptance. We will refer to the value of DF used to achieve a particular combination of pion rejection and muon acceptance as the "bias".

3.3 Discriminant Analysis Results

In order to test the discrimination power of the DF one set of PISA output data at each energy was used for calculating the values of C^{-1} and $\vec{\bar{x}}$. The other set of output data was then used for calculating the DF.

We find that the DF produces very good discrimination between muons and pions of known equal (or sufficiently close) momenta. However, the results depend on the "bias" that is chosen. For a large negative bias, all pions are correctly identified but all muons are incorrectly identified as pions. For large positive bias the reverse is true. As the bias is varied from negative to positive values around zero, the fraction of correctly identified muons increases and the fraction of correctly identified pions decreases. We present our results in a format that shows this variation with changing bias. The results of the Discriminant Analysis for the three schemes shown in Fig. 7 are summarized in Figs. 8, 9, and 10.

Each of these shows plots of the fraction of muon tracks correctly identified as muons ("Muon ID efficiency") versus the fraction of pion tracks incorrectly identified as muons ("Fake Pion ID efficiency"). In all cases five discriminant variables were used:

1. Last plane hit

Mu/Pi Discrimination efficiency

seg(CDR) (10 - 30 theta range)

5 Disc. Variables 1,2,4,9,18

Figure 8: Discrimination efficiency for pad layout described in the PHENIX CDR (Case 1 in Fig. 7).

Mu/Pi Discrimination efficiency

seg(CDR plus) (10 - 30 theta range)

5 Disc. Variables 1,2,4,9,18

Figure 9: Discrimination efficiency for pad layout shown for Case 2 in Fig. 7.

Figure 10: Discrimination efficiency for pad layout shown for Case 3 in Fig. 7.

2. Total hits
3. Total hits divided by the number (1-6) of the last plane to register any hits
4. Ratio of total hits in first plane to register any hits to total hits in last plane to register any hits.
5. Angle θ between the beam axis and the straight-line fit to the road.

The last variable listed above has very little effect on the discrimination power except to smooth out the curves shown in Figs. 8,9, and 10.

This set of discriminant variables is not exactly the same as used by Saskia Mioduszewski [13] but the results are very insensitive to the choice of variables. Roughly speaking, any three or four variables are sufficient and the use of more than four or five does not improve the discrimination power.

As indicated in the figures, the discrimination power is nearly perfect for all the incident particle momenta above 2.25 GeV/c. The abrupt degradation in performance below that momentum can be explained on the basis of the fact that, at about this momentum, some of the muons are not able to completely penetrate all six layers of the muon identifier. At 1.5 GeV/c none of the muons reaches beyond plane 5. At momenta for which the muons are able to penetrate all six layers, muons and pions are easily distinguished from each other. At these momenta there is no need to closely match the momentum of the particles used to produce the training data with the momentum of the test particles. For example, the results for 5 GeV/c are equally good if either 3 GeV/c or 8 GeV/c training data are used for calculating C^{-1} and $\bar{x}$.

3.4 Trigger Rate Studies of the Muon Identifier Using Anode Wire Readouts

The routines in PISORP which simulate tracking through the muon identifier in the muon arm of PHENIX have been revised by Ken Read, Saskia Mioduszewski, and Ajit Mohanty to reflect the switch from pad readout to direct readout from the anode wires of the Larocci tubes. We reported on our use of the newly revised PISORP. The results have been posted on the Muon Arm Web page[14]

A very brief report on the differences in trigger rates brought about by the omission of the neutron and lead shields is given. In addition we have also studied the effect on the trigger rate of omitting both the neutron and

lead shields and the nosecone. Normally simulations should have both shields removed and the nosecone installed.

Finally, since the tracking algorithm permits the user to simulate tracking through strip detectors without hodoscopic ambiguities we have compared the trigger rates with and without these ambiguities.

3.5 Test data preparation

The results presented in the following sections were obtained using both the UA1 and RV_JPSI event generators. Three sets of output were obtained for each event generator. In one case the lead shield and neutron shield were installed; in the second case they were omitted. In the third case both the shields and the nosecone were omitted. In addition, a set of runs was made for 500 events. Due to the long running time for ua1 events this was done only for the "shields out, nosecone in" case. For the RV_JPSI generator angle and momentum filters were used to guarantee that all particles were emitted into the muon arm with sufficient momentum to penetrate beyond plane 3 of the muon identifier for the "shields in" cases. The angle range was 12° to 25° and the minimum momentum was 2 GeV/c.

The road finding scheme used to obtain our results starts in plane 1 of the muon identifier. The beginning of a road is signified by hits in both a vertically and a horizontally mounted proportional tube. The intersection of these tubes defines the x, y coordinates of the beginning of the road. The tubes in plane two whose intersection is closest to the same θ, ϕ coordinates, measured with respect to the vertex, as the hits in plane 1 are then checked for hits. If no pair of hits is found, then adjacent tubes on either side are checked for hits. When a pair of hits is found, the algorithm continues to project forward to check plane three for hits in exactly the same fashion as for plane 2. This search area in plane 2 thus includes 9 square regions, each 8 cm X 8 cm in area. This scheme is indicated in Fig. 11.

3.6 Results

Figure 12 shows the effects of removing the lead and neutron shields and of removing the shields and the nosecone. The portions of the histograms that correspond to valid trigger-roads are shown cross-hatched. Note that 69 of the 100 events produce no tracks to plane 3 and 22 events produce only one track to plane 3.

Figure 11: Schematic of the trigger road finding algorithm used in this study. The third plane of the muon identifier is searched for hits over a 24 cm X 24 cm area exactly as indicated here for the second plane.

As indicated by the arrows in the figure, removing the shields has only a slight effect, increasing the trigger rate from 9 to 10 %. However, removing the nosecone has a significant effect, increasing the trigger rate to 35 % or, assuming a Au + Au central collision rate of 1.4 kHz, it gives a trigger rate of 490 Hz. Figure 13 shows the effects of turning "on" the hodoscopic ambiguities. This is shown for 500 UA1 events run with the nosecone installed but with both shields omitted. The results for ambiguities "off" is shown by the solid lines and the results for ambiguities "on" is shown by the dashed lines. The set of histograms on the left and right hand sides of the figure are the same except the right hand side shows the distribution of roads for which the number of roads is ≥ 2 . For the 500 event case the number of trigger roads with ambiguities "off" is 9 % and with ambiguities "on" the rate increases only slightly to approximately 11 %.

The results obtained for 100 J/Psi events is shown in Fig. 14. The

Figure 12: Distribution of trigger roads for 100 UA1 Events. The histograms at the top show superimposed the number of trigger roads for the cases where the lead and neutron shields are installed and the number of roads produced when these shields are omitted. The histograms at the bottom show superimposed the number of roads produced when both shields and the nosecone are installed and the case where all three items are omitted. For all cases the hodoscopic ambiguities in the strip detectors are turned "off". The shaded areas represent the numbers of roads (≥ 2) that produce a trigger.

Figure 13: Distribution of trigger roads for 500 UA1 Events. The histograms on the left and right are the same, except that the one on the right shows only the cases where the number of trigger roads equals at least 2. Both the lead and neutron shields were removed for the PISA run used to generate this data.

Figure 14: Distribution of trigger roads for 100 J/Psi events. The two upper histograms show the same results except that the one on the right shows only the distribution where there are at least two roads. A similar remark applies to the two lower histograms. In all cases the nosecone is installed and hodoscopic ambiguities are “on”.

momentum and angle cuts specified for the event generator should guarantee 100% trigger efficiency. However, for the case where the shields are installed, two of the 100 events produce only one track reaching to plane 3, presumably because multiple scattering has caused two of the tracks to be deflected to the piston or the lampshade magnet before reaching the third plane of the muon identifier. Therefore, with the shields installed the trigger rate is 98% as indicated in the upper right-hand section of Fig. 14.

When the shields are removed, however, all particles produce tracks reaching to plane 3 so that the trigger rate is 100% as indicated in the lower right hand section of the figure.

In Fig. 15 we show the result of superimposing the UA1 and J/Psi events using the MERGE facility of PISORP. Now, the superposition of the UA1 and J/Psi tracks produces a trigger rate of 100% in all cases.

Hodoscopic ambiguities should of course always normally be "on". We have looked at the "off" case to show how performance would vary if outside information (such as timing across a tube to provide crude horizontal resolution) were used to resolve hodoscopic ambiguities. Of course, such information would not be available to a LVL-1 trigger. Moreover, hodoscopic ambiguities will be "broken" or resolved by means of extrapolation from the muon tracker both offline and in the LVL-2 trigger.

4 Direct VMEbus Interface with Linux

X.-C. He and G. Petitt

We have been developing an affordable, portable and direct VMEbus access system. The system consists of an EISA bus based PC using the Linux operating system (a UNIX clone for the PC) and a VMEbus system. An EISA bus adapter mounted on the PC provides a direct hardware link to a VMEbus adapter installed in a VME crate. The VMEbus adapter can either be configured as a VME slot 1 module (functioning as a VME system controller) or as a slave module. This setup provides high-speed data transfers between the two systems and allows them to share memory; the memory appears to each system as if it were its own. Also, a card on one bus may be accessed and directly controlled by the other system.

We have been developing the system using a model 487-1 adapter pair (from the Bit 3 Computer Corporation) which supports bi-directional A32/D32 random access bus mastering from either system and also supports 32-bit data

Figure 15: Same as Fig. 14 but with each event being a superposition of one UA1 and one J/Psi event. The dashed histograms are the UA1 event distributions. They are identical to the superimposed distributions but shifted to the left by two roads.

transfers using the EISA system DMA controller. This controller enables the adapters to transfer large blocks of data from one system memory to the other at sustained data transfer rates up to 26 Megabytes per second.

Either adapter card can be configured to operate as I/O or memory slave, or as a bus master with control capabilities. Memory mapping controls random access to remote bus RAM, dual-port memory, and remote bus I/O; it also permits high-speed random access 32-bit reads and writes to either the VMEbus or EISA bus at speeds comparable to reads and writes to local memory.

Our setup is shown schematically in Fig 16 which can be easily expanded as needed. All of the software packages running under Linux including Linux itself are freeware, which includes a variety of UNIX-type tools such as C/C++ and "F77" compilers, XFree86 (free X Windows software), Tcl/Tk tools, TCP/IP network capabilities and CERN software.

We have developed some low level device drivers based on these software packages which are the interfaces to the EISA bus and VMEbus adapters. These drivers can easily be used for direct VMEbus access or diagnostics which is the main purpose of this project. Furthermore, a small portable data acquisition system can be easily built upon this system using these low level driver functions.

Figure 16: Direct link between VMEbus and PC Linux system

Since our group has been heavily involved in the online DAQ upgrade

work for the E866 experiment at Fermilab, this system will have a direct application on the E866 project by providing us direct VMEbus access without having to rely on a network connection or a parallel system diagnostic scheme.

References

- [1] "A Muon Identifier for Relativistic Heavy Ion Measurements", Ziyang Zhang, PhD. dissertation, Georgia State University, Atlanta, GA, June, 1995.
- [2] FANL Proposal P# 866, "Measurement of $\bar{d}(x)/\bar{u}(x)$ in the Proton".
- [3] Texas Instruments, Inc. Houston Texas 12345.
- [4] Spectron Microsystems, Santa Barbara CA 93117
- [5] Image and Signal Processing, Inc. Anaheim CA 92807.
- [6] <http://petitt.phy-astr.gsu.edu/L3.html>
- [7] Wind River Systems, Alameda CA 94501
- [8] <http://phnxmu.lanl.gov/Files/muon/notes/notes.html>, "Mu/Pi Discrimination with the Muon Identifier", G.A.Petitt and X. He, Note 95-14, August 10, 1995.
- [9] PHENIX Conceptual Design Report (CDR), January 29, 1993.
- [10] PHENIX CDR Update, November, 1993.
- [11] G. Capon, Nucl. Instr. and Methods A **276**, 58 (1988).
- [12] M. Kendall, A. Stuart, and J.K. Ord, The Advanced Theory of Statistics (Macmillan, New York, 1983).
- [13] <http://phnxmu.lanl.gov/Files/muon/notes/notes.html>, "Preliminary Results for Muon Discrimination", S. Mioduszewski, Note 95-9, 1995.
- [14] <http://phnxmu.lanl.gov/Files/muon/notes/notes.html>, "Trigger Rate Studies of the Muon Identifier using Anode Wire Readouts", G.A.Petitt and X. He, Note 95-19, November 2, 1995.