DOE/ER/03130--55

DE90 013321

**PROGRESS REPORT OF A RESEARCH PROGRAM
IN EXPERIMENTAL HIGH ENERGY PHYSICS**

Task C
Progress Report
for Contract Period January 1, 1990 to December 31, 1990

Co-Principal Investigators
Robert E. Lanou, Jr. and David Cutts
Department of Physics
Brown University
Providence, Rhode Island

July 1, 1990

Prepared for the U.S. Department of Energy under Contract
DE-AC02-76ER03130 (TASK C - Experimental)

MASTER

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

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.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

NOTICE

This report was prepared as an account of work sponsored by the United States Government. Neither the United States nor the United States Department of Energy, nor any of their employees, nor any of their contractors, subcontractors, or their employees, makes any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, product or process disclosed or represents that its use would not infringe privately owned rights.

ABSTRACT

An experimental program in strong and electro-weak interaction physics of elementary particles is being carried out using electronic detection techniques. Experiments have been performed at Brown, Brookhaven, and Fermilab.

The work described in this report by the Electronic Detector Group addresses the following: (1) neutrino interactions and intrinsic properties, (2) preparations for experiments ("D-ZERO") at the FNAL 2 TeV $\bar{p}p$ Collider, (3) new detection techniques for neutrino properties.

Progress Report of a Research Program
in Experimental High Energy Physics

The Department of Physics of Brown University presents herein a report of the progress made in the experimental program of the Electronic Detector Group during the present contract period (1 January 1990 – 31 December 1990). This work has been supported by the U.S. Department of Energy under Contract DE-AC02-76ER03130 (Task C – Experimental).

This report is divided as follows:

- I. Introduction to the Program
- II. Experiments in Progress and in Preparation During this Year
- III. Associated Matters
- IV. Scientific Personnel
- V. Papers Produced During this Year

I. INTRODUCTION.

During the past year work has progressed in three areas:

- A.) The extensive program to study the weak neutral current scattering of neutrinos and of neutrino properties (BNL-AGS-E734) was fully completed and the final paper published.
- B.) The Brown Task C projects for the DØ collider detector (FNAL-TeV-I-E740) — the Level-0 and Level-2 triggers — made much progress in prototyping, bench tests, beam tests, design and construction.
- C.) Ideas for other particle physics applications of new, cryogenic detection techniques were further developed and construction work begun for experimental tests of the ideas.

Each of these topics will be expanded on below.

II. EXPERIMENTS IN PROGRESS AND IN PREPARATION.

A.) The Brookhaven Neutrino Program.

As mentioned in the Introduction (Sec. I), we have now completed a very successful program of neutrino physics using the detector and beams we created at BNL for this purpose with our collaborators.

This year saw the publication of our final analysis of the data on neutrino (anti-neutrino) scattering from electrons. This publication has just appeared in the 1 June issue of Physical Review D (this is reference one in our Section V). To date, this is the definitive work on these reactions. Their importance lies in the fact that the interaction involves only leptons, has low Q^2 and is free from effects due to strong interactions or thresholds. It thus provides measurement of $\sin^2 \theta_w$ or alternatively the vector and axial vector couplings of the electron to the Z^0 intermediate vector boson (g_V, g_A). In the accompanying figure, we illustrate the present state of knowledge of g_V and g_A . The dotted curve labeled D is our result. The other curves A, B, C are, respectively, CHARM, PEP-PETRA, reactor electron-type anti-neutrinos. The solid

Figure A-1

ellipse is a new LEP result. We also, in our experiment on these reactions, set new, best limits on the magnetic moment (or alternatively, the electric charge radius) for the muon- type neutrino - 0.85×10^{-9} Bohr magnetons. This experiment is unique, and will probably remain so, in being able to complete the difficult measurement of the differential cross sections.

Over the duration of our program in neutrino physics at the AGS a great deal of physics was produced: 14 papers in Phys. Rev. Lett., Phys. Lett., N. I. M., Phys. Rev.; 28 papers in conference proceedings and 10 Ph.D. theses - 3 of them from Brown.

In addition to the physics mentioned above there were also results on: elastic neutrino-proton scattering (it's probe of heavy quark content in the proton plays an important role in the present controversy on the spin structure and content of the proton), neutrino mass and mixing, search for heavy neutrino decay, search for second class currents and determination of the axial vector mass in anti-neutrino quasi-elastic scattering.

It is with great satisfaction that we undertook and have completed this program. It has produced much valuable physics.

B.) The "D-ZERO" Detector

When the neutrino program at Brookhaven moved into its mature phase, with our efforts focussed on physics analyses following a succession of experimental runs, we entered a complementary phase of detector prototyping, design, and construction associated with the "D-Zero" ($D\bar{0}$) project at Fermilab. As noted in Section II-A, in this period our program of accelerator-based neutrino experiments has been completed. Related work in prototyping a novel detector for use in non-accelerator neutrino physics has been underway (see Section II-C), but is not at present a principal focus of this task. The $D\bar{0}$ project has moved into a very intensive phase in its run-up before first data taking (at the start of the next Fermilab Collider run in June 1990), and has received our primary attention. At the same time, as work related to mounting the $D\bar{0}$ experiment has made considerable progress, we have given some attention to potential projects in the future (see associated report for Task J, by Professor Partridge, for a description of SSC R&D activities).

The $D\bar{0}$ Detector is a major high energy physics facility, and is dedicated to the study of proton-antiproton collisions at the world's highest available energy, Fermilab's Tevatron-I collider. $D\bar{0}$ is the second detector at the Tevatron Collider, following CDF which was designed considerably earlier. The design of $D\bar{0}$ has thus been able to focus on providing a complementary set of strengths to CDF. Additionally, the $D\bar{0}$ design has evolved in response to the experience from the $S\bar{p}pS$ Collider at CERN, in both the performance of the detectors and the general physics themes which emerged. The key to maximizing the physics at the Tevatron's higher energy and luminosity is to build a detector which does an excellent job at measuring the relevant "particles" at 2 TeV-jets and leptons. In particular, leptons (electrons and muons and "neutrinos") can provide very important signatures to both old and new physics. For these particles the $D\bar{0}$ detector will measure accurately their identity, direction and energy, over the largest possible angle.

To achieve these goals, and to maximize the physics potential of the detector, several basic design choices were made for $D\bar{0}$:

- (1) There is no central magnetic field. Since the relevant particles are jets and leptons (including non-interacting secondaries) calorimetric energy measurements should be better than energies deduced from track curvatures. Also, the non-magnetic central detector is smaller, which allows both calorimetry and muon detection to be better optimized.
- (2) Lepton identification, which has a key role in searches for high mass states, is a prime goal of the design. This requirement calls for the measurement of both electrons and muons over the largest possible solid angle. Seeing new effects in both channels is important, as the systematics in electron and muon measurements are quite different. Also, electron channels are better for precision mass determinations, while muons can provide a clean lepton signal within a crowded jet environment.
- (3) Missing transverse energy is an important signature, and achievement of the least possible E_T resolution is a fundamental goal. To achieve good E_T resolution, the detector must have calorimeter coverage over the full solid angle with minimal cracks. Good energy resolution is important, as is the near equality of response of the calorimetry to electrons and hadrons. Transverse segmentation should be sufficiently fine that angle (and P_T) errors are minimized.

The $D\bar{D}$ detector, designed to meet the above criteria, is shown in a cut-away perspective in Figure B-1. The 4760 ton $D\bar{D}$ detector has three principal components: closest to the interaction region is a tracking section with drift chambers and transition radiation detectors, surrounded by a uranium-liquid argon calorimeter and enclosed by a muon detector incorporating magnetized iron. Figure B-2 shows the layout of the $D\bar{D}$ calorimetry and the tracking section. Table B-1 provides the specifications of the $D\bar{D}$ detector.

The $D\bar{D}$ detector, following the criteria above, has moved from the initial stages of design and prototyping, to testing, production and installation. Members of this task at Brown have been very active during this critical period as design choices are finalized and implemented. As mentioned in the Introduction (I-B.), the Brown Group is responsible for two areas of the detector related to triggering and data

Figure B-1

FACILITY	
PROJECT	PHYSICS
DESCRIPTION	
DATE	
FERMI NATIONAL ACCELERATOR LABORATORY	
P.O. BOX 516, BATON ROUGE, LA 70803	

00 INTERACTION POINT

CENTRAL TRACKING
CALORIMETER

Figure B-2
Layout of DØ Calorimetry
and Tracking.

Table B-1

"DO Detector Specifications"

Proportional Drift Tube Chambers

Maximum length	228 in
Maximum width	100 in
Cathode pad pattern repeat length	24 in
Total number of wires	11,754
Total number cathode pads	23,503
Total number of readout channels	47,016
Total number of modules	176
Minimum angle covered	5 degrees

Central Region

Toroid z interval	-149 to 149 in
Toroid perpendicular distance	125 to 167 in
Toroid weight	2300 tons
Mean field	1.9 T

Central Region chambers -- before toroid

Number of wire layers	4
Number of modules	18
Number of wires	1728

Central Region chambers -- 2 stations outside toroid

Number of wire layers	3 + 3
Number of modules	76
Number of wires	5016

End Region (each)

Toroid z interval	176 to 236 in
Toroid perpendicular inner distance	36 to 164 in
Toroid weight (each)	900 tons
Mean field	1.9 T

End Region chambers -- before toroid

Number of wire layers	4
Number of modules	18
Number of wires	1032

End Region chambers -- 2 stations outside toroid

Number of wire layers	3 + 3
Number of modules	64
Number of wires	3978

Endcap Calorimeter

Number of Endcaps	2
Radial interval	1.7 to 88.2 in
Active length	85 in
Gross weight (each end)	500,000 lb
Total number of readout channels (each end)	15,968
Total absorption lengths	~ 9
Sampling gap	2 x 2.3 mm LAr

EW Section (per endcap)

Absorber	3 mm U
Longitudinal readout segmentation	4 times
Total radiation lengths	20.5
Transverse segmentation	$\Delta\eta=\Delta\phi=0.1$
layer 3: (approx. 2 x coarser for $\eta > 3.2$)	$\Delta\eta=\Delta\phi=0.05$
Total number of readout channels	7488

Inner Hadronic (per endcap)

Absorber	6 mm U/ 46.5 mm Cu
Longitudinal readout segmentation	4 times U/ 1 Cu
Transverse segmentation	$\Delta\eta=\Delta\phi=0.1$
(approx 2x coarser for $\eta > 3.2$)	
Total number of readout channels	5856

Middle Hadronic (per endcap)

Absorber	6 mm U/ 46.5 mm Cu
Number of modules in azimuth	16
Longitudinal readout segmentation	4 times U/ 1 Cu
Transverse segmentation	$\Delta\eta=\Delta\phi=0.1$
Total number of readout channels	1664

Outer Hadronic (per endcap)

Absorber	46.5 mm Cu
Number of modules in azimuth	16
Typical longitudinal readout segmentation	2 times
Transverse segmentation	$\Delta\eta=\Delta\phi=0.1$
Total number of readout channels	960

acquisition. One of these responsibilities is for the "Level-0" trigger, which uses scintillation counter hodoscopes near the interaction region to provide the initial signal of a $\bar{p}p$ interaction event. Because of excellent time resolution, it will also provide a measurement of the z -coordinate of the interaction, which is vital for subsequent E_T triggers. The other focus of this Task for D0 is the "Level-2" trigger and generally the full data acquisition system, from the digitization crates through the array of MicroVAX nodes which perform the Level-2 or software trigger and including the transmission of selected events to a host computer for reading. Subsequent sections in this report will describe the excellent progress we have made in both these areas.

i) The Level-0 Trigger System.

Brown University has assumed responsibility for design, construction, and testing of the Level-0 trigger system for D0. The initial design studies for Level-0 were conducted by Prof. Lanou. He was joined in this effort by Prof. Partridge, who has taken on leadership of this project. The remainder of this section will describe the requirements and design for the Level-0 trigger system; major discussion of progress on Level-0 is contained in the Report of Task J — it will be briefly summarized here.

Requirements for the Level-0 Trigger

The Level-0 trigger has four tasks to perform: identify interactions while rejecting beam-gas collisions, measure the location of the primary vertex for Levels 1 and 2 of the trigger, detect multiple interactions in a single beam crossing, and monitor the collider luminosity. The requirements placed on the Level-0 design by these goals are discussed below.

Identify Interactions

Inelastic $\bar{p}p$ interactions can be identified by detecting the two beam jets which result from the breakup of the incident beam particles. At 2 TeV, it is estimated that each beam jet will have ≈ 11 charged particles within a 10° cone about the beam axis. Single-diffractive interactions and beam-gas events originating in the central region will have a single beam jet and can be rejected by requiring detection of both beam jets. Beam-gas events that originate outside the central detector will be out of time and can be rejected by precise timing measurements. A key design goal for Level-0 is that it be at least 99% efficient for detecting inelastic collisions and reject 100% of upstream beam-gas interactions.

Locate the Primary Vertex

The luminous region of the Tevatron collider will have a large spread along the beam direction ($\sigma \approx 30 - 40$ cm). If uncorrected, this beam spread would seriously degrade the detector resolution for quantities such as the transverse energy in the event, E_T . The Level-0 trigger can make a fast measurement of the location of the

primary vertex by measuring the difference in arrival time for particles in opposite beam jets. If t_+ and t_- are the arrival times for particles in the $+z$ and $-z$ beam jets, the location of the primary vertex is given by

$$z_v = \frac{c}{2}(t_- - t_+) \cos \theta$$

where θ is the angle of the counters with respect to the beam direction and all particles in the beam jet are assumed to have velocity c . Both Level 1 and Level 2 of the trigger make use of the vertex information provided by Level-0.

The Level 1 trigger will use calorimeter signals to make hardware E_T and p_T sums by weighting the energy in a tower with $\sin \theta$ for the tower. The $\sin \theta$ tables are calculated assuming a particular location for the primary vertex; the vertex location from Level-0 is used to select the appropriate $\sin \theta$ table. To insure selection of the proper $\sin \theta$ table, a vertex position resolution of $\sigma_v < 7$ cm is required and must be available 400 ns after beam crossing.

The Level 2 trigger algorithms require precise calculation of calorimeter quantities such as E_T and p_T . These quantities will be calculated by hardware processors using the digitized calorimeter signals and precision vertex information from Level-0. To minimize the error in E_T and p_T due to the uncertainty in vertex position requires $\sigma_v < 3$ cm.

Detect Multiple Interactions

The precise time information obtained by the Level-0 trigger is useful for quickly identifying events with multiple interactions. In the high luminosity environment anticipated for D0, many of the events will have more than one interaction depending upon the bucket structure. Such events will be difficult to analyze and increase rates for certain triggers. Multiple interactions also induce ambiguities in the vertex position determined by the Level-0 system, thus degrading the resolution on E_T and p_T at the trigger level. The desire to efficiently identify multiple interactions places a premium on obtaining the best possible time resolution and high segmentation of the Level-0 detectors.

Measure the Luminosity

The Level-0 detector will also be used to monitor the luminosity of the collider. The design goal for the luminosity monitor is a relative accuracy of 1%, independent of shifts in the beam position or direction. The absolute calibration of the luminosity monitor will depend on the ability of the accelerator group to determine the collider luminosity.

Design of the Level-0 System

Two Level-0 detectors will be built and installed on opposite ends of the central detector. Each Level-0 detector will consist of an array of plastic scintillation counters with photomultiplier readout for detecting charged particles in the beam jets. The counters will make precise measurements of charged particle arrival times for locating the position of the primary vertex and detecting multiple interactions. The design of the Level-0 counters and photomultiplier readout is described below.

Detector Design

The lack of space in the central region of DØ leaves few alternatives in the design of Level-0. The detectors must fit into the 4 cm wide gap between the forward drift chamber and the face of the endcap calorimeter cryostat while staying within the 66 cm radius of the forward drift chamber (see Fig. B-3). Furthermore, the detectors must maximize their rapidity coverage to insure high efficiency for detecting interactions while maintaining good time resolution for reconstructing the primary vertex. To meet these requirements, each Level-0 detector will be made of two layers, both having the design shown in Fig. B-4. The two layers will be oriented orthogonal to each other, to give essentially complete coverage of the region around the beam pipe (see Fig. B-5).

Two types of counters are used: short counters surrounding the beam pipe and long counters completing the coverage. All counters will be made from radiation resistant plastic scintillator such as Bicron BC-408. Both ends of the long counters will have photomultiplier readout, with the standard technique of averaging the times used to determine the arrival time of the charged particle. Long counters cannot be placed

Figure B-3

The DØ central detector and calorimeters. The Level 0 detectors fill the narrow gaps between the forward drift tracking chambers and the endcap cryostat.

Figure B-4 One of two layers in a Level 0 detector. The shaded regions show the position of scintillator material; unshaded regions show light guides and photomultipliers.

Figure B-5 Superposition of the two layers in a Level 0 detector.

close to the beams because of the large probability for multiple hits which would degrade the time resolution. Instead, short counters with photomultiplier readout on one end only are used next to the beam pipe. The short counters are insensitive to multiple hits since the time variation across the counter is small.

Photomultiplier Readout

The Level-0 counters will employ photomultiplier readout to detect the scintillation light. Long counters will have photomultipliers on both ends; short counters will have a light guide to bring the light out to a single photomultiplier tube. Testing of various photomultiplier tubes and base designs were made with final selection of an Amperex (see Task J discussion).

Calibration Systems

In order to determine the constants for each of the 72 channels of photomultiplier readout, a fast, reliable system must be used. We have chosen a laser based one in which the laser is triggered synchronously by the accelerator clock and the laser light is fanned out to each channel through fiber optics.

We have selected a Laser Sciences, Inc. nitrogen laser with 75KW power in a 3 nanosecond pulse. In order to be used simultaneously with the low noise analogue electronics of the calorimetry, the laser system itself must be very low noise. To achieve this the laser is encased in an r.f. shield and can be battery operated.

A portion of the fiber optics is in a high radiation area and must therefore be radiation resistant. For this reason, we have chosen fused silica as the optical medium. We have a number of options under consideration for the configuration of the final fiber optics harness. We have purchased some sample fibers of different types and sizes as well as samples of connectors. Thus a prototype calibration system was constructed to test concepts, components and overall performance. This work was carried out under the direction of Professor Partridge (see Task J Progress Report) and the collaboration of Professor Lanou.

ii) D0 Data Acquisition/Level-2.

The development, implementation, and installation of the D0 data acquisition system is an explicit responsibility of our group (Brown University, Task C). Considerable progress has been made in the past year, as will be described below. Small scale versions of the data acquisition system have been in operation at several sites at Fermilab continuously for several years, with major uses beginning in June 1987 in the NWA test beam and (since September 1988) at the D0 Detector Building. A number of features have been implemented as part of this deployment. Additionally, we have used a data acquisition development setup at Brown to work on different aspects of the system which were then successfully incorporated at Fermilab in the active data taking systems. Before discussing explicitly the progress in the last year, we first provide an overview of the basic system design.

Overview of D0 Data Acquisition

The D0 data acquisition system is based on an array or "farm" of MicroVAX computers which will handle the high data rates and trigger rate reduction required for the D0 experiment at the Tevatron-I collider. The D0 detector, which will become fully operational in the upcoming collider running period, uses excellent calorimetry and lepton identification to study the new physics at 2 TeV. Because of the high potential data rate (100,000 interactions/second each on average supplying 250,000 bytes of digitized information), the data collection system uses a hierarchical trigger structure. The "Level-0" trigger provides from scintillator arrays an accurate interaction time as well as a rough measure of the interaction position in the z-direction. Using this position measurement, the "Level-1" trigger can calculate transverse energy; and with this and other quantities obtained from fast outputs from the calorimeters, muon counters, and other elements of the detector, Level-1 forms a hardware-based selection in the 3.5 microsecond time between crossings. This selection results in a reduction of the trigger rate to 200-400 events/second; these events are then fully digitized subject to a final, software-based filter process (each event in a single MicroVAX) which reduces the events recorded by the host system to 1-2 per second. The D0 data acquisition, the collection of digitized data into an event, packaged in one of many parallel MicroVAX systems, and transmission of filtered events to the

host, is the responsibility of our group at Brown.

The D0 data acquisition system comprises an event's path from the 100 VME crates, where the digitization of the detector's signals takes place, to the data management structure (ZEBRA) in the memory of a specific MicroVAX Level-2 node, and out again, if the event is accepted, to a corresponding structure in the host computer. The key feature of this system is that the data, after being digitized in parallel, flows at high speed over parallel lines to a single, selected MicroVAX which then performs the complete Level-2 filter analysis. As shown in Figure B-6, the digitization crates are divided by detector type into 7 sections; the crates in each section share the use of a data cable which operates at 40 MBytes/sec with a simple data dump protocol. Each of these 7 cables, plus an eighth which carries data from the Level-1 hardware trigger, is connected to an external port of a dual-ported memory in every MicroVAX Level-2 node. Thus the data flows from the VME crates directly into MicroVAX memory at 320 Mbytes/sec. As indicated in Figure B-6, this data flow is controlled by a separate MicroVAX system, the acquisition "Supervisor", which has interface lines both to the VME data sources and to each of the nodes, and which directly enables the dual-port channels in an available node before initiating the event transfer.

The data acquisition system is centered on this simple yet high speed transfer of an event from the digitization crates to a specific MicroVAX node. There are a number of specialized hardware components, principally ZRL dual-port memories, which make this scheme possible. All of the components are in active use by various D0 groups for data acquisition. We describe the characteristics of each of the principal components below. A new generation (advanced versions of these components), appropriate for 1991 rather than 1987, is in the design and development stage, and is discussed in a later section.

Data Cable — The data cable is the unidirectional bus connecting the VME crates with the Level-2 MicroVAX nodes. It is specified to run at up to 40 Mbyte/sec., and has 32 data lines, 2 control lines (SYNCH and BUSY), 2 parity lines, and 9 expansion lines. The data cable protocol calls for BUSY to

Figure B-6: Diagram of the DØ Data Acquisition System.

be asserted at the start of transfer and held throughout, and data is latched by the receiving memory on the trailing edge of each SYNCH pulse. Signals are transmitted using Futurebus open collector transceivers.

Dual-port Memory — The dual-port memory provides event buffering into the Level-2 MicroVAX nodes. Each quad-width Q-Bus dual-port board consists of two channels of 256 KBytes memory. The external ports are 32 bits wide, use the Futurebus transceivers, and can be operated as either input or output, with 100 nsec read/write cycle time (for 40 MByte/sec. transfer speed). On the Q-Bus side, the dual-port boards are transparent Q-Bus memory, and with an associated controller, achieve a 3 MByte/sec transfer into the MicroVAX's private memory. This transfer does not degrade the performance of the cpu, so that event analysis processing of one aspect of the event (say the Level-1 trigger data block) proceeds while other data is transferred. An advanced version of this module, a multi-port memory, is being developed. This board will have a number of additional ports - one of which, for example, will provide direct processor access to the data. See a later section for a discussion of this development.

VME Buffer/driver — This module buffers event data in a VME front end electronics crate and outputs it on a data cable. To the crate electronics the buffer/driver (VBD) appears as normal VME memory. The VBD supports 8/16/32 bit transfers, and can act as a DMA master — incorporating list processing capability and supporting up to 30 MByte/sec transfer rate within VME. The board incorporates dual 256 KByte buffers, toggled on alternate events. The output port is 32 bits wide and compatible with the data cable protocols, with parity generation and a 40 Mbyte/sec transfer rate. A non-DMA version has been operational for several years, as part of the initial systems deployed. This spring the new VBD became available and three boards are now in use, one associated with each of the three variants of VME digitizing crates (calorimeter, tracking, and muon). This new element of the data acquisition is discussed more fully in a later section.

Sequencer — For each data cable a Sequencer module allows the use of the cable by multiple VME crates (each with a buffer/driver). This module will be interfaced with the Level-1 trigger system and the Level-2 Supervisor as well as each of the buffer/drivers in its readout section. It contains a downloaded lookup table, so that given the 32-bit pattern of “specific triggers fired” from Level-1, it obtains a list of the crates to be readout. After handshaking with the Supervisor (to have the dual-ports of a selected Level-2 node enabled), the Sequencer sequentially enables the use of the data cable by each VME crate. This specific operation is described in a later section. Presently the systems at Fermilab use a first-generation version of this module which simply manages sequential use of a data cable by several VME crates; more complicated functions (such as event synchronization) are being handled by software in the Supervisor MicroVAX. The full-function Sequencer will be deployed later this year.

Progress in Data Acquisition Deployment

Much of the progress made in the last year has been related to work done in deploying acquisition systems and in supporting systems actually in use. These efforts have included both improved hardware and related software. We have made considerable progress in the development of advanced components, needed for data acquisition with the full D0 detector. This work will be described in a later section.

Actually, data acquisition systems based on the above design have been in operational use since June 1986, as part of 6 different D0 systems. In addition, at Brown a number of MicroVAX computers of various flavors have been used for data acquisition development work and as a dedicated Monte Carlo farm. Indeed, large-scale production running of GEANT-based simulations of the D0 detector, so vital for design decisions, for several years was run entirely on a MicroVAX analysis farm at Brown. The computing facilities at Fermilab and at our collaborating institutions, including supercomputer and other non-conventional systems, were unable to provide an implementation of the D0 GEANT package that is competitive with a MicroVAX farm. In the last year or so, our Fermilab D0 collaborators and others have begun

operating a MicroVAX farm for Monte Carlo (GEANT) event generation. The Fermilab DØ group is using software for control of the farm operation that was originally developed at Brown. Additionally, towards the end of the last collider run we encouraged our sister experiment, CDF, to explore MicroVAX farms for the production analysis of their raw data tapes when the existing processor farms at Fermilab proved inadequate. CDF's success with their MicroVAX farm shows the importance of this system architecture for offline analysis as well as online data acquisition.

Although Level-2 systems have been operating in a number of places, a real milestone was passed in September 1988 when a data acquisition system was installed in the 2nd floor counting room of the DØ assembly building. This acquisition facility initially consists of three Level-2 nodes plus a Supervisor MicroVAX, as are also installed at the NWA test beam. These four MicroVAX occupy part of the first rack of the full DØ system; additional nodes will be added as needed. Figure B-7 shows the nucleus of the Level-2 system which was installed at DØ.

It should be mentioned that the real-time software framework used in all the nodes, VAXELN, which provides the structure for exercising the data acquisition hardware, has been very successful. Code that runs on the VMS host, in various high-level languages (including FORTRAN and C), is easily packaged and downloaded to MicroVAX nodes, with concurrent programming, DECNET support, and full symbolic, remote debugging included. Real-time performance is excellent, and creation of device drivers is straight-forward. Additional features include remote displays, a remote command facility, and a remote performance analyzer. The rapid development and deployment of the data acquisition systems were possible only because of the ease-of-use and strengths of the VAXELN environment.

To provide a framework for discussing our work, we describe the readout control of the acquisition system at DØ. To direct the data flow, the Supervisor has communication lines with the Level-1 trigger, with the VME readout crates, and with the Level-2 nodes. In addition to these control lines, it is interfaced to each of the data cables and thus able to itself send data to a Level-2 node's dual-port memories, for diagnostic testing as well as for error recovery. These connections, via

FIGURE B-7. The nucleus of the Level-2/data acquisition system at DØ. The rack of Level-2 nodes is on the right.

DRV11-J parallel interfaces, allow the Supervisor to maintain a list of available nodes (by trigger type) and to control the data flow by enabling dual-port memories and VME readouts. Communication with the host VAX is supported for run control and status/diagnostic functions, via a separate job in the Supervisor MicroVAX.

The typical event cycle is initiated by a readout request from the VME crate (TRGR) associated with the Level-1 trigger. The Supervisor reads the trigger type and the "SYNCH" bits (the 3 low order bits of the event number), accompanying the TRGR readouts request. Using the trigger type, the Supervisor assigns the Level-2 node (enables its dual-ports), initiates the TRGR readout, and proceeds to process other readouts (determined by trigger type) to complete the event. At DØ there are three additional data cables: for muon chamber data, for central tracking detector data, and for calorimeter data. (We will expand to the full 8 data cables later this year.) For readout requests from each of these sources, the Supervisor reads the accompanying SYNCH bits and compares them with the SYNCH bits associated with the TRGR readout. If the bits agree (the two data blocks are from the same event) the Supervisor enables the appropriate dual-port and VME readout. If the bits indicate that the VME data is for the next event, the Supervisor enables the dual-port but itself sends a short dummy block (with error flag) to complete the present event in the Level-2 node. Or, if the SYNCH bits indicate that the VME has old data, the Supervisor enables the VME readout but not the dual-port — thus dumping that buffer "on the floor" — and waits for a further readout request. Various additional timeouts prevent data flow from being hung in any abnormal condition associated with the Supervisor, and all the data flow actions are logged and visible via a status display.

Since the deployment at DØ, the four data cables that have been installed at present are in nearly continual use. Figure B-8 is a closeup of two Level-2 nodes, showing the 4 data cables (two per bundle) now in place. These data paths connect the Level-1 trigger framework, muon PDT electronics, central detector FADC's, and calorimeter ADC's to the Level-2 system. Figure B-9 shows two VME crates holding FADC's for digitizing forward drift chamber data, and the data cable which connects

FIGURE B-8 Closeup of two Level-2 nodes installed at D0.

FIGURE B-9 VME crates with FADC's for forward drift chamber at DØ, with data cable connection to the Level-2 system.

to the Level-2 system. As is discussed later, the data cable Sequencer is associated with the VME buffer/driver (VBD) boards as a presently active development project. The "Sequencer" allows the use of a single data cable by multiple VME crates.

Over the last year there has been extensive use of the DØ system by many groups. The calorimeter electronics has undergone an extensive checkout (the "5000 channel test") and is now gearing up towards final connection with the completed central calorimeter. [The central calorimeter cryostat has this month been sealed!] Various parts of central tracking have begun to take data at DØ : the Forward Drift chamber, the Central Drift chamber, and now the Transition Radiation detector have all been connected to the system. The main muon chambers are rapidly completing their installation and have been heavy users for some time; the small angle muon system (SAMUS) is just now coming online. Most of these groups have used multiple digitizing crates and thus employed a first generation Sequencer.

The Supervisor MicroVAX serves as the data-flow control for the data acquisition: enabling the dumping of event data from the VME crates directly into the memories of a selected Level-2 node. Being at the focal point it is sensitive to a number of acquisition problems which we traced to aspects of the triggering and the VME readout. To aid in these studies, we have developed a "software logic analyzer", which was successfully deployed at both the D0 Hall and the NWA testbeam systems. This package monitors the interface registers, and on every change it records them together with the time - keeping a history over some period of the dataflow. With this tool we are able to study the realtime structure of the data acquisition, which leads to corrections to the triggering and readout hardware and suggests improvements in the data acquisition software. At Brown a simple realtime clock was built, consisting of a 32 bit counter being driven at 1 MHz, and coupled via a buffer/latch to a DRV11-J parallel interface. With this device a time can be latched and read in a few microseconds, much more quickly and simply than the standard clock interface, (which is encumbered by many features). We use this clock in the Supervisor diagnostic described above, and will employ it in Level-2 filter analysis timing measurements.

Progress in Data Acquisition Software

In parallel with the developments in hardware has been considerable progress in the software associated with the data acquisition system, both in the control software associated with the Supervisor MicroVAX as well as in the event handling routines which resides in each Level-2 node. We describe these packages below, after describing the environment, VAXELN, in which they operate.

The MicroVAX computers in the data acquisition farm are running programs under the VAXELN dedicated real-time run system. This system includes EPASCAL, a "dialect" of PASCAL with major additions to support real-time activity. ELN is supplied with VMS utilities to compile EPASCAL, link (with files produced by EPASCAL or other compilers, such as FORTRAN or C) and build dedicated, real-time applications to run on stand-alone VAX processors. DECNET/Ethernet communication can be automatically included in the system. ELN has a small run-time KERNEL which controls scheduling and access to resources which allows much flexibility with multiple concurrent "jobs" and subprocesses within jobs. There are special data types for signalling and message passing between jobs, processes and other computers (also ones running under VAX/VMS). An example of our use of ELN's multitasking capability is the operation of independent drivers for the dual-port memories, which are synchronized with the analysis routines.

The program sequence works as follows. The Data Acquisition Supervisor receives signals from the readout Sequencers when the digitization of an event is complete. The Supervisor finds an available Level-2 processor, enables its dual-port memories, and tells the front-end digitizing electronics to start transferring data. Although most of the Level-2 processors will be running identical versions of the filtering programs, the Supervisor may direct special events to particular processors that are assigned to specific tasks, such as calibration. One program, SMAIN, handles these control tasks, as well as an event simulation process which writes data to the data cables via a DEC Q-Bus register interface board (DRV11J). In order to control the Supervisor functions from the host VAX computer, a separate job, STALK, is included. It communicates with SMAIN via the VAXELN concept of an AREA which is shared virtual memory

space between the two jobs. The communication to and from the host VAX uses a separately developed package of routines (ELNCON), running under the VAX/VMS operating system. It makes transparent to the 'average' FORTRAN programmer all the system service calls needed to perform these tasks. This same communication package is also used in other cases when a program on the host VAX/VMS system needs to talk directly to a computer running under VAXELN, as for diagnostics and monitoring. ELNCON has proved to be a convenient packaging of the communication tools which are part of the VAXELN environment. See Figure B-10 for a diagram of the software in the Supervisor.

The sequence of the different jobs and processes running in each Level-2 node is as follows. The MicroVAX node receiving a new event starts the filtering process as soon as the data from the first data-cable is available in memory. A separate subprocesses configures the data into the standard DØ memory management format, utilizing the ZEBRA program package developed at CERN. ZEBRA is a FORTRAN utility which manages data through COMMON BLOCKS with an embedded pointer structure. The special subprocess adds a header with the pointers around the data (which was already placed in the correct memory locations by the dual-port driver). Thus the data is immediately available to all the FORTRAN routines which have been included in the filter program. Since this is a uniform way of handling the data, all routines for filtering and other manipulation of the raw data may be tested out in an offline mode and can be transferred directly to the online filter without any rewriting. The same structure of subprocesses which is set up for the Level-2 filtering program in VAXELN will also be used to make special processing programs for calibration and other purposes. A complete structure for implementing such special programs, downloading them to a select set of Level-2 processors and managing this system is being implemented. See the "Filter Framework" discussion below.

Also included in the main Level-2 program are: 1) a subprocess to handle slow communication with the Supervisor, obtaining run and trigger information and returning status information, 2) filling routines and a subprocess to interact with a host program for a simple, real-time histogramming system, 3) dump routines and

SOFTWARE in SUPERVISOR

FIGURE B-10

a subprocess to send event-dumps to a separate host program as well as 4) an error trapping and reporting mechanism. Figure B-11 is a schematic of these procedures.

Other stand-alone programs (jobs) in the each Level-2 MicroVAX include a job to set the system clock synchronized with the clock on the host VAX (TIME_GET) and a job which passes status messages to a sever (MSGJOB).

There are several control tasks which run on the host under VAX/VMS. To support this system of programs (which run in many separate computers, all under the VAXELN real-time system) various utilities running under VAX/VMS were provided with the VAXELN package (EBUILD for assembling separate .EXE files into systems and EDEBUG for full symbolic, remote debugging of these programs). Other utilities were developed especially to meet the needs of the DØ implementation of this system. Many of the host programs incorporate a menu type user interface, COMPACK, which has been selected as the general DØ command interpreter. Examples of these utilities, based on COMPACK, include a direct control interface to the Supervisor via a program called SUPCON, and a monitor facility, L2_MONITOR, which displays an updating screen of information from each processor in the system. The vital parameters which define the state of the Level-2 system are maintained in files which are managed via the utility L2STATE. Through L2STATE one monitors (or modifies) the Level-2 node definition file and hardware setup file, as well as several files having to do with the state of the Filter Framework, described below. Procedures for limiting the ability to change these vital parameters have been implemented. An overall control program, DØRUN, has many of these facilities such as menu items as well as special commands for DECNET operations needed to control the Level-2 MicroVAX array. The host software routines related to the Level-2 system are shown schematically in Figure B-12.

As part of our continued effort to make monitoring/diagnostics an integral part of the Level-2 system, we have developed several new facilities which perform such tasks.

The method of "trapping" specific run-time errors in the Level-2 filter code which has been present in the system for a long time, was upgraded to make it easier to

SOFTWARE in Level-2 Node

Figure B-11

SOFTWARE on Host.

Figure B-12
35

use. The system now has a settable flag which makes it behave in one of two ways when a run-time error is encountered. The first way is to trap the error with the VAXELN debugger (EDEBUG), which does NOT have to be connected to the node at the time of the error. When the debugger connects, the actual error and traceback information will be available to study the error. In the running system such errors should be extremely rare, and it may also be important to keep the particular Level-2 node running even after such an error occurs. The second way of catching such errors simply logs an entry in a data-base file giving specific information about the error, skips the rest of the event and sets up the program to receive the next event. A facility to make searches in this data-base of errors has been developed. It is able to search using either the time of the error, the routine which caused the error, the type of the node, the name of the node or a specific sequence of run-numbers as keys.

Another important way to monitor and study the Level-2 system is to look at the time-sequence of the events passing through a node. A hook was put into the Level-2 program such that it will record the time when particular pre-determined points are reached. This recording only takes place when the program which will receive this information is running. That program runs on any VAX/VMS node; after being told which node and how many events to use for the display, it will request the information from the node. When the information is received, it will be set up to be filled into pseudo-histograms of non-equidistant bins such that the display looks like a time-sequence of on/off edges for the various states as shown in Figure B-13.

The filter framework:

Significant progress has been made recently in implementing a software environment for managing and controlling the filter routines which will perform the event selection in the Level-2 MicroVAX nodes. At DØ, the Level-2 filter will be responsible for removing over 99 percent of all the events passed to it. This means that precise control and monitoring of the filter's operation is very important. Additionally, it is important to have flexibility with respect to both the order in which the algorithms are applied and the threshold values used in each of the various algorithms.

We base the filter framework on the concept that each algorithm forms a 'TOOL',

L2 Logic Analyzer Display

Figure B-13

which is as much as possible independent of all other TOOLS. A single filter (many will be running at the same time) may then be built up of a sequence of these TOOLS. Both the sequence of the TOOLS defining one filter and the parameter values used for each call to a TOOL may be changed in the running system when a new run is started. Each TOOL may take part in several filters with different sets of thresholds and other parameters for each call. Each Level-1 trigger (up to 32) may activate more than one filter sequence (for a total of 128 possible filters). Each of used filter-bit will have a 'script' associated with it. The scripts are lists of TOOL-numbers and parameter set numbers in the order they should be employed to form the complete filter. A definition file, L2TOOL.DAT, was set up to include information for each possible TOOL in the system. The program L2STATE is used to manipulate these definitions. The TOOL algorithms used will have to be tested thoroughly before being certified for use in the actual system. The Level-2 nodes will also be split into groups of various 'TYPES' each with a different executable program for the filter. A file, L2TYPE.DAT, was set up to include the definitions for the TYPES. The routines needed to make sure the executable for each TYPE has the correct TOOLS will be automatically made by the L2STATE program.

The filter framework has been worked on and upgraded in response to needs and requests by the various groups working on filter algorithms for this system. In particular an important change was to allow sorting of the filters according to importance and/or like impact of the efficiency if one filter was executed before another. This sorting takes place when all the filter-scripts are assembled on VAX/VMS before starting a run in the Level-2 system. Each Level-2 node will then look for any filters to be activated in that order, actually improving the efficiency of this search since only bit numbers actually in the ordered list will have to be checked.

At this point a fair amount of time is being spent in making sure that physicists implementing algorithms which will be used in Level-2, find it easy to have their code included in the system and verified. The most important part of this is the test facility running under VAX/VMS. The actual filter code which will run in the VAXELN system may be tested completely in VAX/VMS this way. The VMS-FILTER facility

builds on libraries of routines supplied by other members of the DØ collaboration. It has the ability to define specific filter-scripts for testing particular algorithms, read in events from a file (the events could be data taken by the real DØ-detector or generated Monte Carlo data), and use the filter framework to activate the filter-scripts.

Data Acquisition Development Projects

A number of DØ data acquisition facilities have been in operation for some time, at Fermilab and elsewhere. These systems have been functioning very successfully, using the specialized hardware and software which we have described above. As we prepare for the operation of the DØ detector in the upcoming collider run (expected in less than a year!), we are very active in several data acquisition development projects. Some of this activity is related to the assumption of the responsibility for the VME readout modules whose initial version, currently in use, was designed and built at the University of Pennsylvania. Brown University (Task C) now has complete responsibility for the data collection path, from the output of the ADC/FADC digitizers in VME crates, through the Level-2 system, and on to the host. A new version of the VME buffer/driver module (VBD) has been built, tested in our acquisition test system at Brown, and as mentioned above now in active use by each of the three detector groups (calorimeter, tracking, and muon). This board is a VME master, driven by a list processor, as opposed to the previous boards which were essentially dumb memory and required a 68K processor module to be used in the VME crate. The new VBD builds the data blocks itself, via DMA transfers into its memory banks, and constructs the checksum word, length word, and other header/trailer information. It has various sophisticated capabilities, such as the ability to handle multi-level interrupts (so that "Level-1.5" information can flow to the digitizers to veto one event while at the same time the VBD is transferring across the VME bus data from the previous event). At this point all the different digitizing electronics groups have "signed-off" on its performance, and we are initiating full production.

The new VBD (VME buffer/driver) will also incorporate hardware to perform automatically the readout "synchronizing" function now carried out by software in the Supervisor MicroVAX, as described previously. "Synchronization" refers to the

important aspect of data collection which insures that every data block from each VME crate, when readout, originated from precisely the same event. The DØ acquisition is particularly efficient (and a model for future SSC systems) as the event is never "built" and then moved – but rather, on many parallel paths, is fed into memories associated with a specific, selected Level-2 filter processor. The new VME readout hardware will replace software in the Supervisor which ensures that these parallel paths all carry data for the same event. A readout Sequencer will control the use of a given data cable; it will circulate a token to the VBD boards in its readout group. This token will include bits which describe the VME crates required for the specific trigger (as received from the Level-1 framework). The token will also will carry the low order bits of the event number; each VBD in turn will check the event number and load the appropriate buffer onto the data cable (or an error pattern).

Other development work in progress involves the hardware in the MicroVAX systems. The replacement next year of the dual-port memories with newly designed multi-port memories will be a major improvement to the functionality of each Level-2 node. These new modules will have several additional ports beyond the 40 MByte/sec input port and a standard backplane port: a port that will allow the processor direct access to the data, a high speed output channel, and a special function port. With these boards we need not move the data in memory, but simply build the ZEBRA data structure in the multi-ports themselves. Details of these boards and their impact on the Level-2 system are given in the accompanying Renewal Program.

One particular aspect of this development is the use of one channel of the multi-port memories as a high speed data channel for accepted events, surplanting the use of ethernet for the event path from the Level-2 system to the host computer. This development requires both software in the event-handling portion of the Level-2 system plus hardware and software in the BI-host. From our experiences at NWA and DØ it has been clear that this much-enhanced output bandwidth will be needed in the future.

Another important feature of the multi-port memories will be the ability of attached coprocessors to access the event data. We anticipate that a number of

compute-intensive filter algorithms can be greatly speeded up by using specialized processors. In particular, we have made significant progress in studying neural networks and their ability to perform typical Level-2 filter algorithms. Neural networks have a crucial property in that as they are trained on a particular problem they force a parallel solution. Thus algorithms based on neural networks could be implemented in the Level-2 nodes, using array processor chips attached to the multi-port memories' special function ports. Other use of attached coprocessors, such as for reducing calorimeter raw data to energy sums, looks very promising. We discuss these exciting possibilities further in the Task C Renewal Proposal.

C.) Neutrino Calorimetry in Superfluid ^4He .

There are several problems in particle physics, mostly related to neutrino properties which would benefit from being able to look at the interactions at low energies and in a low intrinsic background medium. Among these are solution of the "solar neutrino problem"; neutrino mass and mixing at neutrino sources ranging from either accelerators, reactors or the sun; neutrino magnetic moments; detection of dark matter and even precision gamma ray spectroscopy where large detectors are required.

As we have discussed previously, we have considered ways to try to accomplish this. We came up with a new idea which may very well provide a suitable solution. This idea grew out of discussions and calculations we carried out with two of our condensed matter/low temperature colleagues at Brown. The idea involves the use of superfluid helium as the detector material. A typical temperature of operation for the helium would be 100 mK. At these temperatures there are two very valuable properties of the helium. One is that no other element can remain in solution in it, thereby making it uniquely free of radioactive material. The second is that much of the energy deposited by an interaction rapidly appears in the form of rotons, a bulk excitation whose properties are stability against decay and high propagation velocity. These rotons would be detected by bolometry on thin silicon wafers.

In 1987, we published a paper in Physical Review Letters describing its possible application to solar neutrinos. Other applications, such as the detection of astro-

physical dark matter, have been published in various conference proceedings. An overview of the whole field of dark matter detectors was carried out by this group for the "Snowmass 88: High Energy Physics for the 1990's". These ideas have attracted much interest and favorable review; we continue to receive many invitations for colloquia and seminars on this subject at universities and national laboratories in the U. S., Canada and Europe.

Following guidance from the Department of Energy, we applied for and received (1 May 1988) a grant from the Medium and Low Energy Nuclear Physics Division to carry out an R & D program on the technique. The grant number is DE-FG02-88ER540452 and Professor Lanou is designated as Principal Investigator. The funds are largely for equipment and construction; no individual currently on Task C of this contract receives any funds from the grant.

During the last year several crucial milestones were reached. The construction of the dilution refrigerator with capability of 500μ watts at ≤ 200 mK was completed and its operation tested. The refrigerator has performed stably for long periods at its design temperature of 35 mK. Following these tests the cell in which the roton tests are to be conducted was instrumented and installed. The instrumentation consists of three sources, two thin constantan phonon sources and a 10 nanocurie Americium-241 Alpha source followed by two 8 mil silicon wafers equipped with germanium and composite thermometers as well as calibration pulsers.

In a recent preliminary test we have been able to detect 5.5 MeV alpha's impinging directly on to the cold silicon wafers and producing thermal pulses detected in the germanium thermometer. This is an essential milestone. At the present time this "commissioning run" is still in progress. We expect to fill the cell with superfluid helium soon and test for roton production and subsequent detection via quantum evaporation. So far the R&D program is on track and at least one necessary physics condition has been met. We are hopeful that neutrino detection by this method will be demonstrable.

The new results on solar neutrinos are very tantalizing. At the recent "Neutrino '90" conference held at CERN the SAGE (Soviet-American Gallium Experiment)

presented first results on their search for p-p branch neutrinos from the sun. The preliminary result: the observed rate is < 75 S.N.U. at 68% CL and < 135 S.N.U. at 95% CL. In fact their best fit is for 0 events observed with 14 events expected. Additionally, the Homestake (Cl-37 experiment) presented evidence for a significant time dependence of flux over a 20 year period. Whether or not it is correlated with the sunspot cycle is less clear. A Kamioka result on a one year time span showed little or no confirmation; however, the statistical significance is marginal.

Both these results argue for the increased importance of the technique we are developing for both solar neutrino experiments and for accelerator experiments with superfluid helium to test for neutrino magnetic moments. It may turn out that the neutrino sector (as well as the Higgs sector) is our best chance to go beyond the "standard model".

Further testing of the cryogenic devices is continuing. This work is being carried out by ourselves and our collaborators in the low temperature and condensed matter groups here at Brown under the direction of Professors G. M. Seidel and H. J. Maris; none of those persons receive any compensation from Task C of this contract.

Further plans are discussed in the accompanying Renewal Proposal.

III. ASSOCIATED MATTERS

In this section we discuss several items which cut across the boundaries of several sections. Although we believe Section II describes significant progress in a program of important physics, in terms of operations and equipment funds, this has been another difficult year for us. The overall budget remains below earlier levels in spite of necessary travel to three laboratories (FNAL, BNL, and SLAC), strong R & D activity, and increasing time which must be devoted to construction and installation. It is unrealistic that we will be able to maintain a quality program indefinitely. We could not have managed without the significant, real financial support we received directly from the University and the Physics Department again this year. It is also unrealistic to expect this kind of University support to continue on a regular basis. Consequently, the budget presented in the accompanying Renewal Proposal contains

figures which represent more realistic figures than we have had to operate under recently.

The major contribution which we are making to the DØ project in the area of data acquisition has been possible only because of the receipt over the past several years of a significant portion of our requested R&D funds. With this support we have been able to stay in the forefront of this developing technology and maintain excellent relations with manufacturer's technical personnel. This forthcoming year we must develop the final version of the Level-2 system, including new multi-port memories and hardware filter techniques. R & D funds are important for us to continue our momentum in this area, so appropriate for use with the "next generation" detector at the end of the year.

Some other activities of the individuals in the group should be mentioned.

Professor Cutts is on the scientific advisory committee for the International Conference on Computing in High Energy Physics. This year's edition of the conference was held in Santa Fe where he and Professor Hoftun presented several papers related to D-Zero and applications of neural networks to high energy physics. He also served as a technical advisor on the DOE's review of the ZEUS U.S. collaboration. During the past year he returned from sabbatical leave at Fermilab where he was a Visiting Senior Scientist working on D-Zero; he has now resumed his position as Executive Officer in the Department of Physics at Brown.

Professor Lanou is on the scientific program committee for P.A.N.I.C. (the International Conference on Particle and Nuclear Physics) being held this year in Boston in honor of Viki Weiskopf's 80th birthday. He has also served this year on a number of panels in the field, among them: DOE technical reviews of the Argonne National Laboratory and of the U.S. Collaboration of the MACRO Project in the Gran Sasso, Italy. He also served on an advisory panel to the Director of Los Alamos National Laboratory concerning possible future physics directions for LAMPF. He also continues as Chairman of the Department of Physics here.

We have again been very fortunate to be able to have several Brown undergraduates work with us. Among them this year have been Andrew Sornborger, Chris Barter,

David Asher, Kevin Epstein, Steven Ramsey, Mario Triantafillou, Gaza Gyuk, and John Yun. This gives them an early introduction to research; our experience in the past is that many of these students go on to graduate school and enter the field of high energy physics professionally. Of this year's group, Chris Barter will go to graduate school at Harvard, Andrew Sornborger to Brown and David Asher will take a year off before attending graduate school.

IV. SCIENTIFIC PERSONNEL WITH THE RESEARCH PROGRAM.

R. E. Lanou, Jr.	Professor
D. Cutts	Professor
R. A. Partridge*	Assistant Professor
J. S. Hoftun	Assistant Professor (Research)
D. Nesic	Research Associate
L. Wang*	Electronic Engineer
D. Cullen-Vidal	Graduate Research Assistant
H. Xu	Graduate Research Assistant
F. Nang	Graduate Research Assistant

(* also on Task J)

V. PUBLICATIONS, REPORTS AND THESES IN THIS REPORTING PERIOD.

1. J. L. Callas, D. Cutts, M. V. Diwan, J. S. Hoftun, B. W. Huglock, R. E. Lanou, T. Shinkawa and collaborators, "Determination of Electroweak Parameters of Muon Neutrinos and Anti-neutrinos on Electrons", *Physical Review D* **41**, 3297 (1990).
2. J. S. Hoftun, "Level-2 System Programmer Users Manual", Brown University D-Zero Note # 909 (December 1989) and DOE/ER/03130-51C.
3. D. Cutts, J. S. Hoftun, D. Nesic, A. Sornborger, C. R. Johnson and R. T. Zeller, "Applications of Neural Networks in High Energy Physics", presented at the conference of "Computing in High Energy Physics - 1990", (Santa Fe, April 1990), D-Zero Note # 964 and DOE/ER/03130-52C.
4. D. Cutts, J. S. Hoftun, D. Nesic, C. R. Johnson and R. T. Zeller, "DØ Data Acquisition - High Speed Data Paths", presented at the conference of "Computing in High Energy Physics - 1990", (Santa Fe, April 1990), D-Zero Report # 972 and DOE/ER/03130-53C.
5. D. Cutts, J. S. Hoftun, D. Nesic, C. R. Johnson and R. T. Zeller, "Special Function Co-Processors for Level-2", D-Zero Note # 971 (April 1990) and DOE/ER/03130-54C.
6. R. E. Lanou and D. Cutts, "Progress Report of a Program in Experimental High Energy Physics - Task C", (July 1990) and DOE/ER/03130-55C.