The DØ Run II Trigger [1]

G. C. Blazey
Department of Physics, Northern Illinois University, DeKalb, Illinois 60115

Abstract

The general purpose DØ collider detector, located at Fermi National Accelerator Laboratory, requires significantly enhanced data acquisition and triggering to operate in the high luminosity ($L = 2 \times 10^{32} cm^{-2}s^{-1}$), high rate environment (7 MHz or 132 ns beam crossings) of the upgraded TeVatron proton anti–proton accelerator. This article describes the three major levels and frameworks of the new trigger. Information from the first trigger stage (L1) which includes scintillating, tracking and calorimeter detectors will provide a deadtimeless, 4.2 μs trigger decision with an accept rate of 10 kHz. The second stage (L2), comprised of hardware engines associated with specific detectors and a single global processor will test for correlations between L1 triggers. L2 will have an accept rate of 1 kHz at a maximum deadtime of 5% and require a 100 μs decision time. The third and final stage (L3) will reconstruct events in a farm of processors for a final instantaneous accept rate of 50 Hz.

I. Introduction

The Run I DØ triggering system (1992–1996) used two hardware triggers and a software trigger [1] to select the 3.5 Hz of events for further offline processing from the approximately 0.5 – 1.0 MHz collision rate. These levels were (a) L0 (a scintillator hodoscope trigger sensitive to all inelastic collisions); (b) L1 (hardware trigger decisions in 3.5 μs based on fast sums from calorimeter and muon detectors); and (c) L3 (software–based filters in a farm of approximately 50 processors). Between L1 and L3, a third hardware trigger, L2, refined the calorimeter–based trigger for electron candidates by examining the shape of the energy deposition and refined the muon trigger by using finer granularity hardware information. L2 imposed a bandwidth limitation as it interrogated only a subset of the L1 accepts and inhibited data taking while examining the event. The L1/L2 event accept rate was 150 Hz. (Event sizes were typically 450 kbytes.) The trigger system operated with overall data–collection efficiency of nearly 90%.

Run II will be characterized by a ten–fold increase of luminosity, to $2 \times 10^{32}\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and nearly a forty–fold decrease in the time between beam crossings from 4 μs to 132 ns. These two parameters require significant upgrade of the entire trigger system. For instance, at Run II luminosities a typical Run I trigger menu which includes high transverse momentum (p_T) jet, electron, muon and large missing transverse energy (E_T) triggers would require 2000 Hz for L1/L2 and 10 Hz for L3, well beyond Run I capabilities. In addition, the Run I Framework could not operate at a 7 MHz rate.

To meet the high rate demands of Run II, the DØ trigger will be improved in three basic ways. To increase rejection, the upgraded trigger will include (a) several new tracking detectors: the fiber tracker (CFT), the central preshower (CPS) and the forward preshower (FPS) and (b) a significantly modified muon detection system. Second, the L2 system, which will examine all events, will be strengthened by the addition of several new detector–specific preprocessing engines and a global stage to test for correlations between L1 triggers. Finally, the L3 system will undergo bandwidth and processor improvements to meet increased computational needs. (The L0 system will also be substantially upgraded and used primarily as a luminosity monitor.)

II. Level 1 or Hardware Triggers

A. Trigger Detectors and Operation

As shown in Figure 1, the upgrade L1 trigger detectors include the CFT, CPS, FPS, the calorimeter (CAL) and the muon scintillators (SC) and tracking chambers. Physically, the innermost L1 trigger detector is the CFT. This is encompassed by a superconducting magnet. Outside the magnet are the preshowers, CAL, and at the outermost radii the muon system. The muon system includes (in order of increasing radii) a layer of tracking and scintillator planes, a second magnetic spectrometer, and further layers of tracking and scintillator planes. The calorimeter, fiber tracker, and preshower detectors will provide electron triggering for $|\eta| < 2.5$ (η, the psuedorapidity of a particle, is equal to $-ln(tan(\theta/2))$ where θ is the polar angle of a particle relative to the beam line). The fiber tracker and muon systems will provide muon triggering in the region $|\eta| < 2.0$ [2].

The L1 triggers associated with each of the trigger detectors examine each event and report their findings or trigger terms to the L1 Framework (LlFW). Each front–end digitizing crate will include sufficient memory to retain data from 32 crossings. This pipeline ensures deadtimeless operation for 7Mhz (132 ns crossings.) The L1WK will support 128 unique L1 trigger bits. Each bit will be pre–programmed to require a specific combination of trigger terms. A series of field programmable gate arrays (FPGAs) will examine the list of terms collected from the luminosity monitor, CFT/CPS, FPS, CAL and MUON L1 triggers to determine if a specific L1 bit has been satisfied. If so, the L1FW issues an accept and the event data is digitized and moved from the pipeline into a series of 16 event buffers to await a L2 trigger decision. The planned L1 accept rate is 10

[1] The author is pleased to acknowledge the support of the National Science Foundation.

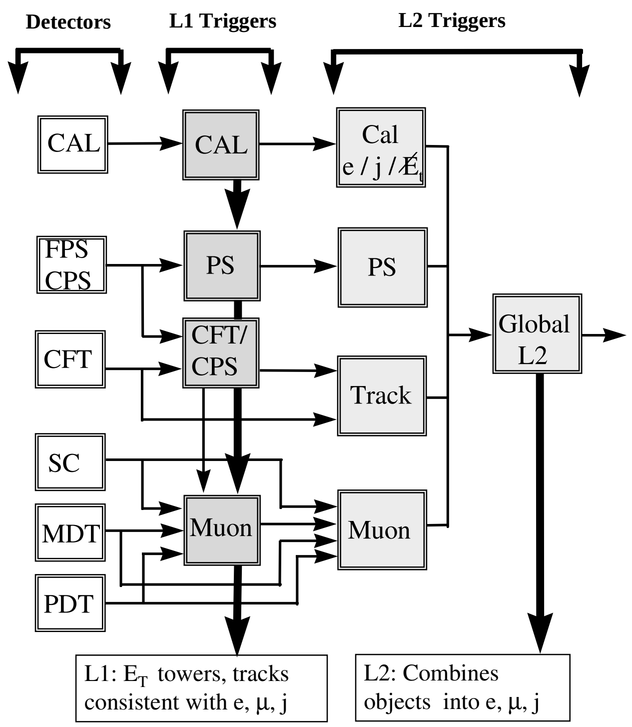

Fig. 1 L1 and L2 trigger elements. The horizontal arrows denote information flow.

kHz. Trigger decisions will be issued 4.2 μs after a beam–beam crossing.

B. Calorimeter

The L1 calorimeter trigger will be unchanged from Run II [1]. Both local energy deposition and global calorimetric characteristics can generate a trigger. The calorimeter is segmented into trigger tiles of $\Delta\eta \times \Delta\phi = 0.8 \times 1.6$ and trigger towers of $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ covering $|\eta|$ <4 (where ϕ is the azimuthal angle). A calorimeter trigger will require E_T above a preset threshold in one or more calorimeter trigger elements. A total of sixteen threshold sets are available. Additional trigger terms are provided for global quantities or trigger tower sums such as total E_T , total energy, and missing E_T.

C. Central Fiber Tracker and Preshowers

The central fiber tracker ($|\eta| < 1.7$) is composed of 77k scintillator fibers mounted on eight cylinders concentric with the beam line. Half of these fibers are mounted axially or parallel to the beam axis and the reminder are mounted with a stereo angle. The innermost cylinder is 20 cm from the beam and the outermost 50 cm. A two Tesla superconducting magnet is adjacent to the outer CFT cylinder. The central preshower detector is mounted just outside the magnet and is segmented into 6k channels with both axial and stereo orientations.

The trigger scheme for the fiber tracker is based upon the hit patterns in the axial fibers. Because of the central magnetic field a charged particle will bend in the azimuthal direction only. Thus, viewed along the beam, a charged particle will follow a curved trajectory. The magnitude of the curvature depends inversely on the particle momentum. Each of the CFT axial fibers (arranged in eighty 4.5 degree trigger sectors) are read out with visible light photon counters [3]. The counter outputs are discriminated for triggering and digitized for readout.

The trigger signals are fed into a series of large field programmable gate arrays resident on VME cards. The FPGAs are preloaded with the appropriate logic to form eight–fold coincidences from hits in the eight CFT axial layers. A trigger will be generated by any hit pattern consistent with track momentum above a software settable threshold. The lowest p_T threshold will be in the 1.5–3.0 GeV range and is limited, in part, by FPGA capacity. A decision can be obtained within 600 ns.

The trigger data is sent serially at 424 MHz to a CFT trigger manager (CFTTM) and to the MUON L1 trigger system. There is a limit of six CFT trigger candidates per sector. Each trigger candidate will be identified by momentum, charge, azimuth in the last axial layer, the presence of preshower energy deposition above threshold, and isolation. Note since only axial fibers contribute to the trigger no rapidity information is available for transmission. The CFTTM will provide sixteen L1 terms: the number of tracks above each of the four thresholds, the number of isolated tracks above the four thresholds (no other tracks within one sector), the number of tracks above the four thresholds with a coincident CPS energy deposition, and the number of isolated tracks above the four thresholds with CPS deposition.

The forward preshower detector will contribute to electron triggering in the region $1.4 < |\eta| < 2.5$. The L1 trigger will require a spatial match of preshower hits in each layer behind a layer of lead with corresponding track stubs in each layer before the lead. As with the CFT/CPS trigger hit decoding and triggering will be executed with FPGAs. At least two L1 terms (forward and backward FPS) will be available to the L1 Framework.

D. Muon

The central L1 muon trigger detectors include a scintillation counter layer segmented in azimuth and rapidity and a layer of proportional drift tubes (PDTs) before the muon spectrometer, two layers of PDTs after the magnet, and an outer layer of cosmic ray veto scintillation counters. The forward L1 muon detectors region include three layers of pixel scintillation counters segmented in azimuth and rapidity and three layers of mini–drift tubes (MDTs). The CFT is also an integral component of the L1 triggering scheme for both the central and forward regions. Central fiber trigger tracks, scintillation counter hits, and wire chamber centroids are sent to the trigger hardware via Gbit/s serial links from AMCC over coaxial cable [4].

The central L1 trigger combinatorical logic is performed locally in detector octants on MTC05 and MTC10 VME cards. As with the CFT trigger, trigger logic is implemented in FPGA's (in this case Altera Flex 10 K series FPGA's [5]). The MTC05 trigger algorithm matches CFT trigger tracks with scintillation counters for low p_T tracks. High p_T tracks also require cosmic ray veto scintillation counters since the tracks penetrate the muon spectrometer iron. The scintillation counter segmentation azimuthally matches the central fiber tracker trigger sectors and, in addition, has nine segments along the beam direction. A 25 ns timing gate used to define scintillation counter hits serves to reject background hits. A 50 ns timing gate defines cosmic ray veto scintillation counter hits. Up to four different muon p_T thresholds can be defined using the track information from the CFT.

The MTC10 trigger algorithm uses centroids found in the PDT's verified by an accompanying scintillation counter hit. The confirmation is necessary because the proportional tube drift time (600ns) is greater than the bunch crossing time (132 ns) and thus centroids can originate from several crossings. A low p_T trigger is defined using only the innermost tracking centroids. A high p_T trigger is defined using correlations between PDT centroids in any two of three layers.

Triggering in the end regions will also be done locally by octant in MTC05 and MTC10 VME cards. The forward MTC05 trigger algorithm matches CFT trigger tracks with the innermost layer pixel scintillation counters for low p_T tracks and with all three layer pixel scintillation counters for high p_T tracks. The pixel scintillation segmentation azimuthally matches the central fiber tracker and has radial segmentation of $\Delta\eta$ = 0.1. This segmentation ensures efficient triggering for muons of p_T greater than 3 GeV since a typical multiple scattering angle is about 3 degrees. A 25 ns timing gate defines pixel scintillation counter hits at the trigger level and serves to reject background. As with the central region, four different p_T thresholds can be defined using the track information from the CFT.

The forward MTC10 trigger algorithm will use centroids in the MDT's that have been verified by an accompanying pixel scintillation counter hit. Although the drift time of the MDT's (80ns) is less than the bunch crossing time, this confirmation is useful for rejecting centroids from background particles produced at points other than the interaction region. Similar to the central MTC10 trigger, a low p_T trigger is defined using only the innermost layer centroids while a high p_T trigger requires correlations between centroids in three layers.

A muon trigger crate manager in each crate of MTC05 and MTC10 cards collects the octant trigger information from each card and produces trigger information for the central, north, or south regions. This crate trigger information is then sent to the muon trigger manager (MTM) which produces global muon trigger information. There are many terms contributing to this decision as four thresholds, three regions, three track qualities, and multiplicity are available to the MTM. The global muon trigger decision is subsequently sent to the L1FW where it is included in the global physics trigger decision.

III. Level 2

A. Operation and General Architecture

The L2 trigger will reduce the 10 kHz accept event rate by roughly a factor of ten within 100 μs using multi–detector correlations of objects found in the event [6]. Two distinct L2 stages are planned. The first stage, or preprocessor stage, prepares data from each L1 trigger for use in the second or global processor stage. The global processor represents the first opportunity for the combination of L1 trigger objects from different detectors. The relationships between L1 and L2 trigger elements are shown in Figure 1.

A L2 Framework, utilizing the same FPGA logic as the L1 Framework will coordinate the operation of L2 and report trigger decisions to L3. Upon receipt of a L2 accept from the global processor, L3 will initialize detector readout and move the event data into eight transfer buffers. There will be a one–to–one mapping between the L1 and L2 bits. For more detailed information on L2 operation see the contribution, "The DØ L2 Trigger" in these proceedings by J.Linnemann.

In the preprocessor phase, each detector system separately builds a list of trigger information. Individual preprocessors are planned for the calorimeter, CFT, muon tracker, and the preshower detectors. The calorimeter and muon preprocessors are nearly designed and discussed below. For each subsystem, the L1 trigger information is retained and transformed into physical objects such as energy clusters or tracks. The time required for the formation of preprocessor objects is about 50 μs. The condensation of information (a few hundred to a thousand bytes of information from each subsystem) allows rapid transmission of the information to the second part of the L2 system – a single global processor.

The global processor receives preprocessor information on data highways operating at 320 Mbytes/s to provide trigger decisions within 75 μs. The decisions are based upon correlations amongst multiple detector systems. For example, spatial correlations between track segments, preshower depositions, and calorimeter energy depositions may all be used to select electron candidates.

The operation of the overall trigger system, and L2 in particular, has been extensively modeled using the RESQ [7] simulation package. The deadtime and time budgets are evaluated with the constraint that a maximum of 16 events may reside in the L2 trigger input buffers, preprocessors, or global processor. The simulations show that the system deadtime for the highest data rates expected in the upgrade is 5% or less for preprocessor and global processor timing budgets of 50 and 75 μs, respectively [6]. In addition, the deadtime has a relatively weak dependence on the timing budgets.

Detailed GEANT detector simulations have been performed of the event rate reduction afforded by L2. The simulated event sample, generated with ISAJET, includes multiple hadron–hadron interactions within a single beam crossing. A 2500 Hz input rate can be reduced to a 500 Hz output rate for a full mix of high p_T triggers.

B. Calorimeter Preprocessor and Global Processor

The calorimeter preprocessor consists of several sub–units: an electromagnetic preprocessor, a jet preprocessor, and perhaps a missing E_T preprocessor. The preprocessor input is the full array of 1280 trigger tower E_Ts for both electromagnetic and the electromagnetic plus hadronic sums. Clustering algorithms build electron or jet candidates, calculate their position and energy and test them for shape and transverse energy requirements. When combined with information on the L1 triggers, the tower information results in a minimum of 4 kilobytes/event at 10 kHz, or 40 Megabytes/s.

Each calorimeter preprocessor will be implemented as a Digital α processor in VME [9]. This is a development project, shared between the CDF collider experiment [8] and DØ . The technology uses the processor board design and chipsets available from Digital Semiconductor. The chipset will include a memory cache large enough to accommodate both program code and event data. Timing studies from prototype C code indicate that currently available 500 MHz α processors should be adequate to meet the target mean execution time of 50 μs for the jet and electron processors. Simulations have also shown that rejection factors of five or more can be expected for electron candidates.

Data will be transported from the L1 calorimeter trigger to the preprocessors by approximately ten serial data highways onto a 128–bit wide bus. These serial highways will be a universal component of the trigger system. The 128–bit "Magic Bus" drives a 64–bit PCI through a DMA engine into the main memory of the Digital–α processor board.

One advantage of the serial processor architecture is that the same architecture is also natural for the L2 global processor. The "Magic Bus" will be used to interface the calorimeter, tracking, and muon preprocessors to the single global processor. Simulations indicate that sufficient bandwidth, memory, and CPU exist for testing correlations between the various detection elements.

C. Muon Preprocessor

The muon preprocessor will improve muon identification by repeating the L1 muon trigger calculation with greater resolution and more information. The L2 trigger will collect information from each muon chamber and scintillator detector and repeat the L1 calculation but incorporating calibration information and more precise timing information from the scintillators. The preprocessor input arrives at a rate of 10 kevents/s. The data, averaging 5kB/event, is spread across roughly 140 sources on fast 160 Mbit/s serial links. The preprocessor will calculate transverse momentum, rapidity, azimuth, and quality for the muon candidates.

The muon preprocessor will use a massively parallel architecture to provide a deterministic algorithm with execution time independent of the number of detector hits. The geographic segmentation of the detector allows reduction of track finding into small well–defined regions, one per processing element. The translational symmetry of the detector elements allows each of the many parallel processing units to execute the same algorithm, concurrently [10]. Timing and algorithm simulations based on the multi–processor Adaptive Solutions, Inc., CNAPS–128 [11] card indicate that 40–45 μs will be adequate for a full reconstruction of muon candidates.

The preprocessor will be implemented in two parts. The interface crate and the processor crate. The processor crate will house approximately 15 CNAPS processor cards which consist of 128 parallel processors with attendant memory that communicate through either a PCI bus or an external 20 Mbyte/s link. A Pentium processor and ethernet card on the PCI bus will provide the means for software transfer and monitoring. The VME interface crate will hold boards each designed to receive up to 16 serial link inputs and format the data into bytes for transfer into the CNAPS card via the external link.

D. Tracking Preprocessors

The L2 trigger system will also include at least two central tracking preprocessors. A CFT preprocessor will assemble and perhaps p_T or azimuthally order a list of trigger tracks before transmission to the global processor. To improve forward electron triggering the FPS preprocessor will compute the azimuth and rapidity of forward electron candidates. The global processor will then correlate the CFT and FPS candidates with calorimeter electron candidates. No technology has yet been selected for these preprocessors but variants of the serial and parallel architectures discussed above are under consideration.

Future expansion of the L2 trigger system might include a silicon impact parameter trigger (SVT). This preprocessor would utilize hits from the million channel silicon vertex tracker, located inside the CFT, to search for track vertices some distance from the interaction vertex. These secondary or displaced vertices are characteristic of long–lived particles and as such can be used as heavy quark tags. The physics benefits from the inclusion of such a trigger are currently the subject of intense study. There has not been a technical study of an SVT.

IV. Level Three

The Run II L3 system is a straightforward upgrade of the existing data acquisition and trigger. The enhanced system will have an input event rate of 1 kHz and a 50 Hz accept rate with further increases in bandwidth possible. (Run II event sizes will be 250 kbytes.) The increased rates will be achieved with a highly parallelized data pathway and fast processors.

Both Run I and Run II systems are characterized by parallel datapaths which transfer data from the front–end crates to a farm of processors. Each event is examined by a processor with a suite of filters. Because of these fundamental similarities the upgrade can utilize much of the Run I cable plant and infrastructure. The L3 upgrade is briefly described below; however a more complete description can be found in the contributions "The DØ Data Acquisition System" by G.Watts and "DØ Data Transport Architecture" by R.Zeller.

Figure 2 shows the upgraded Run II data acquisition path. Each front end crate generates one block of data per event. These data blocks move independently through the data system and are recombined into single events at their assigned L3 processor node. The system makes its routing decisions based on information contained in the data block header, L2 trigger information, and pre–loaded routing tables. Data blocks from multiple events flow on the datapaths at the same time.

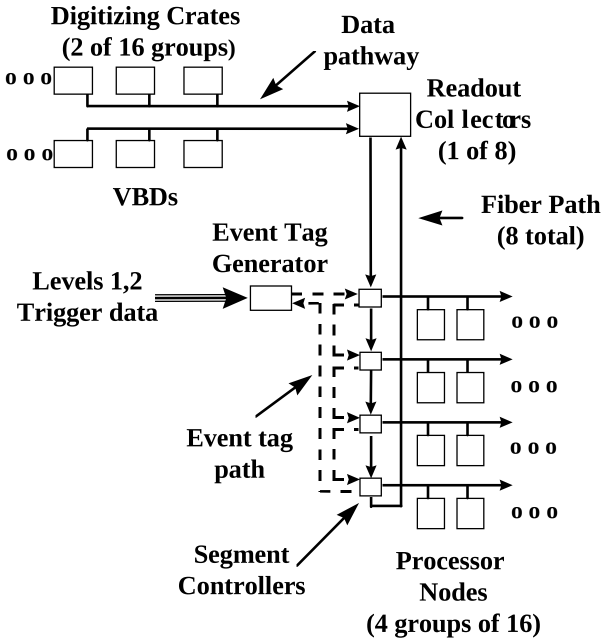

Fig. 2 L3 architecture. See the text for details.

Once a L2 accept has occurred, digitized data is loaded from the L3 transfer buffers onto one of 16 high speed data pathways (48 mb/sec) by driver modules which sit in VME front–end crates. The various cables are arranged to distribute the data load evenly. The datacable carries the data to a VME Receiver Collector, the VRC. The VRC sends the data from either of two datacables on fiber using low level Fibre Channel hardware, to the L3 Farm Segment Controllers (SC's) for the long–haul transmission between the D0 counting rooms. Each fiber is connected to each SC in a daisy chain fashion.

The Event Tag Generator (ETG), also notified at L2 accept, uses trigger bits to assign an event to a specific event class. A SC accepts the tag and assigns the event to a specific L3 node in the attached segment. Each SC examines each passing data block and extracts those data blocks belonging to an event assigned to a node on that SC. The remaining data blocks are passed onto the next SC. Data blocks not extracted by the last SC are returned to the VRCs for recirculation. This acts as a self limiting rate mechanism. Data blocks move from the SC's into processor nodes through multi–port memories for final event assembly.

Queuing simulations with the program MODSIM–III [12] show that the planned readout/control system easily accommodates 1 kHz. The simulation includes the front–end crates, data pathways, controllers, and processor segments. With additional controllers and processors the system bandwidth can be increased to 10 kHz [6].

The existing 48 processors will be replaced with the best price/performance multi–processor systems available in 1998–1999. The operating system will be Windows NT. High level programming or event filtering will be executed in C or C^{++}. Rate rejection will be obtained by filtering each event with "physics tools". These tools will have access to all event data to search for electron, muon, and jet candidates as well as interesting event topologies. Any event meeting filter requirements will be transferred to tape storage for later off–line reconstruction.

V. Conclusion

The DØ Trigger system requires significant upgrade to operate in the high rate environment of Run II. With the addition of new trigger detectors and greatly strengthened L2 and L3 stages the trigger will provide a rejection factor of one million. The widespread use of FPGAs at L1 and in the Framework; fast serial and massively parallel processors at L2; and parallelization and increased computation at L3 will provide the flexibility needed for efficient operation of the detector into the next millennium.

VI. References

[1] S. Abachi *et. al*, "The DØ Detector", *NIM, vol. A338, 1994* pp. 185.

[2] Upgrade documentation can be found on the DØ WWW page *http://d0sgi0.fnal.gov/.*

[3] M. Adams *et. al*, "A Detailed Study of Plastic Scintillating Strips with Axial Wavelength Shifting Fiber and VLPC Readout", *NIM, vol. A360, 1995* pp. 263.

[4] Manufactured by AMCC. 6195 Lusk Boulevard, San Diego, CA.

[5] Manufactured by Altera Corp. 2510 Orchard Parkway, San Jose, CA.

[6] Trigger documentation can be found on the DØ WWW page *http://d0sgi0.fnal.gov/* under "DØ Trigger".

[7] C.Sauer *et. al*, "The Research Queuing Package Version 2", *RA138, RA139* IBM Research Div., San Jose CA.

[8] F. Abe *et. al*, "The CDF Detector", *NIM, vol. A272, 1988* pp. 376.

[9] Manufactured by Digital Equipment Corporation. Information can be found on the WWW page *www.digitial.com/semiconductor/index.html*.

[10] M. Fortner, "A Massively Parallel Processor for Level–2 Muon Triggers at D0", Proceedings of the AIHENP '96 conference, Lausanne, Switzerland, September, 1996.

[11] Manufactured by Adaptive Solutions, Eugene, Oregon, USA.

[12] Manufactured by CACI International, La Jolla, CA.