

Disclaimer

This note has not been internally reviewed by the DØ Collaboration. Results or plots contained in this note were only intended for internal documentation by the authors of the note and they are not approved as scientific results by either the authors or the DØ Collaboration. All approved scientific results of the DØ Collaboration have been published as internally reviewed Conference Notes or in peer reviewed journals.

The DØ Detector

Scott Snyder
SUNY Stony Brook
Stony Brook, NY 11794

March 16, 1995

Contents

1	Coordinate Systems	3
2	Of Luminosities and Cross Sections	4
3	The Beam	4
4	DØ Overview	7
5	Central Detector	11
5.1	Drift Chamber Principles	12
5.2	Vertex Chamber	13
5.3	Central Drift Chamber	15
5.4	Forward Drift Chambers	17
5.5	Transition Radiation Detector	17
5.6	Central Detector Readout	20
6	Calorimetry	20
6.1	Calorimetry Principles	20
6.2	Calorimeter Geometry	22
6.3	Calorimeter Readout	27
6.4	Calorimeter performance	29

7	Muon System	29
7.1	WAMUS System	31
7.2	SAMUS System	33
8	Triggering and Readout	35
8.1	Level 1	35
8.1.1	Level 0	38
8.1.2	Main Ring Vetoes	38
8.1.3	Calorimeter Trigger	39
8.1.4	Muon Trigger	40
8.2	Level 2	41
9	Host Processes	43
9.1	Run Control	43
9.2	Data Logging	44
9.3	Downloading	48
9.4	Monitoring	49
10	Offline data processing	49
	References	50

This note contains a description of the DØ detector. The ‘official’ reference for the detector as a whole is [1]. The reader should consider a reference to [1] implied in almost everything in the remainder of this note.

1 Coordinate Systems

In what follows, a right-handed coordinate system will be used, with the positive z -axis aligned along the beam in the direction of the protons and the positive y -axis pointing up. Cylindrical (r, ϕ, z) coordinates are sometimes used, as are spherical (r, ϕ, θ) coordinates. The angular variables are defined so that $\phi = \pi/2$ is parallel to the positive y -axis, and $\theta = 0$ is coincident with the positive z -axis. Instead of θ , it is often convenient to use the pseudorapidity η defined as

$$\eta = -\ln \tan \frac{\theta}{2}. \quad (1)$$

The pseudorapidity approximates the true rapidity,

$$y = \frac{1}{2} \ln \frac{E + p_z}{E - p_z}, \quad (2)$$

in the limit that $m \ll E$.

It is also often convenient to use instead of momentum, the ‘transverse’ momentum, which is the momentum vector projected onto a plane perpendicular to the beam axis:

$$p_T = p \sin \theta. \quad (3)$$

This is particularly useful due to the fact that in a $p\bar{p}$ collision, the momenta along the beam of the colliding partons are not known (since many of the products of the collision escape down the beam pipe). However, their transverse momenta are very small compared to their momenta along the beam, so one can apply momentum conservation in the transverse plane.

One can also define a ‘transverse energy’ by

$$E_T = E \sin \theta. \quad (4)$$

If treated as a vector, the direction of E_T should be taken to be the same as the p_T vector.

2 Of Luminosities and Cross Sections

At a collider, the flux, or *luminosity*, of the beam is defined as the number of particles passing through a unit area at a certain point along the beam per unit time, and is typically measured in units of $\text{cm}^{-2} \text{s}^{-1}$. The luminosity at the Tevatron during run 1A was typically in the range $1 - 10 \times 10^{30} \text{ cm}^{-2} \text{s}^{-1}$. A reaction rate measurement is expressed as a *cross section*, the interaction probability per unit flux. Cross sections are usually expressed in *barns*, where $1 \text{ barn} = 10^{-24} \text{ cm}^2$. The reaction rate R , the luminosity \mathcal{L} , and the cross section σ are related by

$$R = \sigma \mathcal{L}. \quad (5)$$

It follows that the number of events of a specific type expected after running an experiment for a period of time is found by integrating the luminosity with respect to time:

$$N = \sigma \int \mathcal{L} dt. \quad (6)$$

The quantity $\int \mathcal{L} dt$ is called *integrated luminosity*; the total integrated luminosity for run 1A is $13.5 \text{ pb}^{-1} \pm 12\%$. The luminosity is determined by measuring the rate of a subset of minimum-bias events, which have a known observable cross section of 48.2 mb at the Tevatron energy [2, 3]. (This was derived by averaging measurements from E710 and CDF.)

3 The Beam

The DØ detector is located at the Fermilab Tevatron [4, 5, 6, 7], presently the world's highest-energy hadron collider, with a center-of-mass energy of 1800 GeV . A schematic of the accelerator complex is shown in Figure 1.

The Tevatron is a storage ring, composed of superconducting magnets. The ring is filled with bunches of protons and antiprotons, which circulate in opposite directions. At the B0 and D0 experimental areas, these beams are made to collide with each other. The process of filling the ring is quite complicated; a summary of the major steps is given below, but the reader should consult the cited references for more details.

The beams originate in the preaccelerator. There, H^- ions are formed and accelerated to 750 keV by an electrostatic Cockroft-Walton accelerator. The preaccelerator operates in a pulsed mode with a frequency of 15 Hz . The ions are bunched and transported to the start of the Linac. The Linac is a 150 m

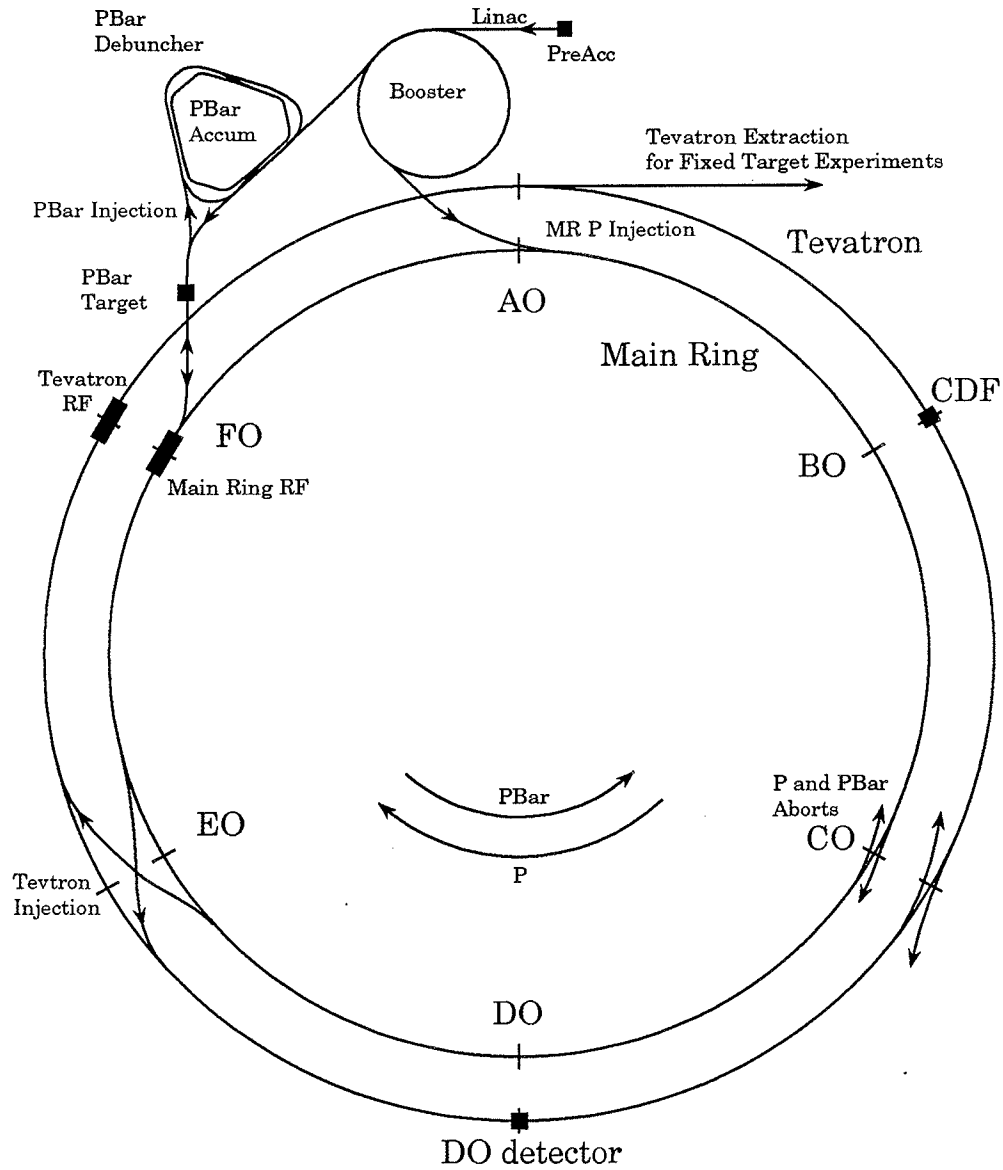


Figure 1: Schematic of the Fermilab accelerator complex (not to scale).
 [8, p. 112]

long linear accelerator, which boosts the energy of the ions to 200 MeV¹. After emerging from the Linac, the ions are passed through a carbon foil which strips off the electrons, leaving bare protons. The protons are then injected into the Booster, a 151 m diameter synchrotron. (A synchrotron is a device which confines charged particles in a closed orbit using bending magnets. RF cavities can be used to increase the energy of the stored particles; when this is done, the field of the bending magnets must also be increased in a synchronous manner in order to keep the particles in the same orbit.) One of the interesting features of the Booster is its rapid cycle rate of 15 Hz. To achieve this, the magnets are combined with capacitor banks to form LC circuits which resonate at 15 Hz. The Booster accelerates the protons to an energy of 8 GeV. The protons are then injected into the Main Ring, a large (1000 m radius) synchrotron composed of conventional magnets. The Main Ring lies mostly in a plane, except at the B0 and D0 experimental areas where it is bent into overpasses to allow room for the detectors (the separation between the Main Ring and the Tevatron is 19 feet at B0 and 89.2 inches at D0). Protons in the Main Ring can be used to make antiprotons (see below), or they can be accelerated to 150 GeV and injected into the Tevatron.

The Tevatron is a proton synchrotron made from superconducting magnets [5, 7]. It lies just below the Main Ring in the accelerator tunnel, and has a maximum beam energy of 900 GeV. (Upgrades to the cryogenic system are expected to raise this to 1000 GeV.) The Tevatron can be operated in one of two major modes. In fixed-target mode, the Tevatron is filled with protons which are accelerated and then extracted and directed towards numerous experimental areas. This cycle repeats with a frequency of about once per minute. In collider mode, the Tevatron is filled with six bunches of protons and six bunches of antiprotons, traveling in opposite directions. The beams are accelerated to the maximum energy of 900 GeV each and allowed to collide at the B0 and D0 experimental areas. (At other points where the beams would collide, they are kept apart by electrostatic separators). The beams are typically kept colliding for a time on the order of 20 hours, after which the machine is emptied and refilled with new batches of protons and antiprotons.

The remaining major part of the accelerator complex is the antiproton source [6, 9], which is used to produce and store antiprotons for use in the collider. While collisions are occurring in the Tevatron, the Main Ring contin-

¹For run 1B, the Linac has been upgraded to an energy of 400 MeV.

ually runs antiproton production cycles at a rate of one every 2.4 s. Protons are accelerated to 120 GeV and extracted onto a nickel target. Each of these collisions produces a spray of nuclear debris, which includes some antiprotons. Immediately following the target is a lithium lens, a cylindrical piece of lithium through which a large (0.5 MA) current is passed. This generates an azimuthal magnetic field which acts to focus negatively-charged particles passing through it. Following the lens is a bending magnet which selects negatively-charged particles with energies of 8 GeV and transports them to the Debuncher. The Debuncher is a storage ring in which antiprotons are first ‘debunched’ (rotated in phase space from a configuration with a small time spread and large momentum spread to one with a large time spread but small momentum spread) and then stochastically ‘cooled’ to further reduce the momentum spread. Stochastic cooling [9, 10] operates by measuring the trajectory of collections of passing particles relative to the desired orbit; this generates a correction signal which is passed across the ring to kicker electrodes which apply a force on the particles to move them back towards the desired orbit. The effect on any single particle is very small due to the incoherent contribution of all the other particles near it in the beam, but when repeated over a large number of turns, the effect becomes significant. The antiprotons are kept in the Debuncher until just before the next pulse arrives, about 2.4 s later. They are then transferred to the Accumulator, another storage ring which lies inside the Debuncher. There, cooling continues for several hours, and eventually the antiprotons settle into a dense core near the inner radius of the Accumulator. When enough have accumulated to fill the Tevatron (typically on the order of $50 - 150 \times 10^{10}$), they are extracted from the Accumulator, accelerated to 150 GeV in the Main Ring, and injected in bunches into the Tevatron.

Some of the major parameters of the Tevatron for the 1A run are given in Table 1. A more detailed introduction to the accelerator may be found in [8, appendix A].

4 DØ Overview

DØ is a large, multipurpose detector for studying $p\bar{p}$ collisions which has been operating at the Fermilab Tevatron since 1992. The design was optimized for the study of high- p_T physics and high mass states, and stresses the identification and measurement of electrons and muons, the measurement of the

Accelerator radius	1000 m
Maximum beam energy	900 GeV
Injection energy	150 GeV
Peak luminosity	$\approx 10 \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$
Number of bunches	6 p , 6 \bar{p}
Intensity per bunch	$\approx 100 \times 10^9 p, \approx 50 \times 10^9 \bar{p}$
Crossing angle	0°
Bunch length	50 cm
Transverse beam radius	$43 \mu\text{m}$
Energy spread	0.15×10^{-3}
RF frequency	53 MHz
\bar{p} stacking rate	$\approx 3.5 \times 10^{10}/\text{hour}$
Beam crossing frequency	290 kHz
Period between crossings	$3.5 \mu\text{s}$

Table 1: Run 1A Tevatron Parameters. [11, p. 17][12, ch. 2][8, app. A]

direction and total energy of high- p_T jets, and the determination of missing transverse energy. Little emphasis is placed on identifying and tracking individual particles within jets, as the details of hadronization are not relevant to the underlying hard scattering.


Detectors for colliding beam experiments are composed of many different particle-detection devices, each with their own strengths and weaknesses. The general layout, however, is dictated by the physics of how particles interact with matter. Closest to the interaction point are the *tracking* detectors, which are devices designed to measure the three-dimensional trajectories of particles passing through them. Often, the tracking detectors are immersed in a magnetic field; this permits a determination of the momentum of the charged particles via a measurement of their bending radius. Surrounding the tracking detectors will be a *calorimeter*; this is a device which measures the energy of particles which hit it. A calorimeter should be ‘thick’ so that it will absorb all the energy of incident particles; conversely, tracking detectors should contain as little material as possible to minimize multiple scattering and losses prior to the calorimeter. A calorimeter is typically made thick enough to stop all known particles except for muons and neutrinos. Muons are identified by the use of tracking chambers outside the calorimeter; any charged tracks originating from the interaction point and penetrating the cal-

orimeter are likely muons. Neutrinos are not detected directly; their presence is inferred from an imbalance in the total detected momentum perpendicular to the beam.

The DØ detector, illustrated in Figure 2, follows the general plan outlined above. Surrounding the beam pipe are a set of tracking detectors. At DØ, however, there is no central magnetic field. Part of the reason for this is that the momentum resolution in a magnetic tracker is roughly proportional to the momentum, while the energy resolution of a calorimeter is roughly proportional to $1/\sqrt{E}$. Thus, for high- p_T objects, a calorimeter will yield a better momentum measurement than a tracking detector. A magnetic tracker would also tend to sweep low-momentum charged particles out of jets, degrading the calorimetric energy measurement. In addition, the material of the solenoid is absent. The absence of a tracking momentum measurement, however, implies the need for a very good calorimeter. DØ achieves this using a liquid argon sampling calorimeter made mostly from depleted uranium. To identify muons, an additional set of tracking chambers is installed surrounding the calorimeter. To provide a measurement of the muon momentum, magnetized iron toroids are placed between the first two muon tracking layers.

The detector is quite large; the entire assembly is about 13 m high \times 11 m wide \times 17 m long with a total weight of about 5500 tons [8, p. 26]. As indicated in Figure 2, the Tevatron beam pipe passes through the center of the detector. Also visible is the Main Ring beam pipe, which passes through the upper portion of the calorimeter, 89.2 inches above the Tevatron beam pipe. Not shown in the figure is the detector support platform, on which the entire assembly rests. This platform is mounted on rollers so that the entire detector may be rolled from the assembly area to the collision hall. In addition, the platform provides rack space for detector electronics and other support services.

The use of clocked electronics, such as digitizers, is kept to a minimum on the detector platform in order to reduce electronic noise. Analog signals are led from the platform, out of the collision hall, and into the Moving Counting House (MCH), a three-story structure containing the digitization electronics, the level-1 trigger, high-voltage supplies, and other services. In order to keep the cable runs short, the MCH is also mounted on rollers and follows the detector as it moves into and out of the collision hall. The detector data cables and other communications lines are led out of the MCH into the second floor of the assembly building (sometimes called the fixed counting house, or FCH).



DØ Detector

Figure 2: Cutaway view of the DØ detector. [1]




Figure 3: Side view of the DØ central detector. [1]

Some of the individual elements making up the DØ detector are described below.

5 Central Detector

The main goal of the central tracking chambers is to reconstruct the three-dimensional trajectory of charged particles which pass through them. This information can be used to decide if an electromagnetic shower in the calorimeter was caused by an electron or by a γ/π^0 . By also measuring the dE/dx of a track, one can determine if the track actually consists of several closely spaced charged particles, such as would result from a $\gamma \rightarrow e^+e^-$ conversion. The central trackers are also responsible for making a precise measurement of the location of the interaction vertex for each event. Precise position measurements are also useful for calibrating the calorimeter position measurements and for improving the accuracy of muon momentum measurements.

The central detector consists of four major subsystems (see Figure 3).

Immediately outside the beryllium beam pipe is the vertex drift chamber (VTX). Surrounding that is the transition radiation detector (TRD) for electron identification. Outside that is the cylindrical central drift chamber (CDC) capped on both sides by the two disk-shaped forward drift chambers (FDC). The entire assembly is contained within a nonmagnetic cylindrical volume of radius $r = 75$ cm and length $l = 270$ cm.

5.1 Drift Chamber Principles

When a charged particle passes through a gas, it will interact electromagnetically with nearby atomic electrons, resulting in the creation of electron/ion pairs along the path of the particle. The number of such pairs created depends on the energy of the particle and the type of gas, but for a typical gas at STP and a particle with unit charge, the mean number of electron/ion pairs formed will be on the order of 100/cm [13, p. 22].

If an electric field is applied, the electrons will start to drift through the gas towards the positive electrode, undergoing repeated collisions with the gas molecules. (The ions also drift in the opposite direction. However, their drift speed is much less than that of the electrons, so they can be ignored in this discussion.) If the electric field near the anode is strong enough, an electron can acquire enough energy between collisions to knock an additional electron free from a gas molecule. This additional electron can then go on to ionize more gas molecules; in this way, an avalanche is formed in which the number of electrons increases exponentially. When this avalanche reaches the positive electrode, it gives rise to a measurable current, the size of which is proportional to the original number of ions created. The ratio between the final number of electrons collected and the initial number deposited is called the *gas gain*, and for practical detectors is typically on the order of 10^4 – 10^6 .

The large electric field needed to obtain gas amplification is usually obtained by forming the anode from a very thin (20–100 μm) wire. Thus, an electron sitting in the gas far away from the anode will see a nearly constant electric field, and will drift towards the anode at a nearly constant speed. When it gets close to the anode, the electric field will start to rapidly increase, and the electron will initiate an avalanche.

The fact that an electron drifts with a nearly constant speed over most of the distance to the anode implies that one can turn a measurement of the time an electron took to drift to the anode into a measurement of the distance of the original source particle from the anode. A device designed for




Figure 4: End view of one quadrant of the VTX chamber. [1]

this type of measurement is called a *drift chamber*. In order to obtain a linear relationship between distance and time, it is necessary that the electric field be made as constant as possible over as large a volume as possible. Typically, additional field-shaping electrodes will be inserted into a drift chamber in order to make the field more uniform. One is also aided by the fact that the relationship between electron drift velocity and electric field tends to flatten out for sufficiently large electric fields. It is thus highly desirable to operate a drift chamber in this *saturation* region.

For further discussion of drift chambers and their application to high-energy physics, consult [14, 15, 16].

5.2 Vertex Chamber

The vertex chamber (VTX) [17, 18, 19] is the innermost of the tracking chambers. It can be used to accurately determine event vertex positions and complements the other tracking chambers by identifying conversions which occur in the TRD. It consists of four carbon fiber cylinders enclosing three concentric layers and occupies the region $3.7 \text{ cm} \leq r \leq 16.2 \text{ cm}$. The inner layer has a length of 97 cm, with each successive layer being about 10 cm longer.

An end view of the VTX is shown in Figure 4. The geometry of the VTX is called a *jet* geometry because the sense wires are arranged in planes which are parallel to the paths of particles emerging from the interaction

Active radius	3.7 cm – 16.2 cm
Number of layers	3
Active lengths of each layer	96.6 cm, 106.6 cm, 116.8 cm
Number of cells	16 inner layer; 32 in outer two
Number of sense wires	8 per cell, 640 total
Sense wire separation	4.57 mm radially with 100 μm stagger
Sense wire specifications	25 μm NiCoTin, 80 g tension
Sense wire voltage	+2.5 kV
Field wire specifications	152 μm Au-plated Al, 360 g tension
Gas composition	CO ₂ 95%, ethane 5%, H ₂ O 0.5%
Gas pressure	1 atm
Average drift field	1 kV/cm
Drift velocity	$\approx 7.3 \mu\text{m}/\text{ns}$
Maximum drift distance	1.6 cm
Gas gain	4×10^4
Position resolution	$r\phi \approx 60 \mu\text{m}$, $z \approx 1.5 \text{ cm}$

Table 2: Vertex Chamber Parameters. [12, p. 38][8, p. 34][18]

region. The inner layer of the VTX is divided into sixteen cells; the outer two layers contain 32 cells each. Between layers, the cells are staggered in ϕ to avoid dead regions and also to aid in resolving left-right ambiguities when linking together track segments in different layers. Within each cell, there are eight sense wires, which are staggered relative to each other by 100 μm in order to further resolve left-right ambiguities. Each cell also contains a set of aluminum traces mounted on the carbon fiber tubes plus an additional set of wires for field shaping. The $r\phi$ position of a hit is determined from the drift time and the wire hit. The z position is determined using a technique called charge division: the resistive sense wire is read out at both ends and treated as a voltage divider. Unfortunately, this method requires that the sense wire pulses be well separated and that the cell occupancy be low, a condition which is rare in a high-multiplicity environment such as DØ. The VTX is not operated at saturation, so extra care is required with high voltage to obtain a reliable time-distance relation.

Further parameters of the VTX are given in Table 2.




Figure 5: End view of 3 of 32 CDC modules. [1]

5.3 Central Drift Chamber

The central drift chamber (CDC) [13, 20, 18, 21] is in the outer layer of tracking chambers and covers large angles (within a pseudorapidity range of about $|\eta| \leq 1.2$). Like the vertex chamber, it has a jet geometry. It consists of four layers of cells occupying the region $49.5 \text{ cm} \leq r \leq 74.5 \text{ cm}$, with a length of 184 cm. An end view of a portion of the CDC is shown in Figure 5.

The CDC is constructed from 32 separate identical modules, which are arranged in a cylinder. The cylinder is contained inside of an aluminum tube, with a carbon fiber tube sealing the inner radius. The ends of the chamber are capped with G10 walls followed by aluminum endplates which hold the wire tension. The module walls are constructed from Rohacell foam laminated with Kevlar and wrapped with Kapton. Within each cell, there are 7 sense wires, staggered by $200 \mu\text{m}$ relative to each other to help resolve left-right ambiguities. Each cell also has a set of field-shaping wires, plus a set of resistive strips silk-screened onto the Kapton cell surface. The $r\phi$ position of a hit is determined via the drift time and the wire hit. The z position of a hit is measured using inductive delay lines embedded in the module walls in the sense wire plane. When an avalanche occurs near an outer sense wire, a pulse is induced in the nearby delay line. By comparing the arrival times of the pulse at both ends, the z position can be determined.

Further parameters of the CDC are given in Table 3. See [13, 21] for a detailed description of the construction and testing of the chamber.

Active radius	51.8 cm – 71.9 cm
Number of layers	4
Active length	179.4 cm, 106.6 cm, 116.8 cm
Number of cells per layer	32
Number of sense wires	7 per cell, 896 total
Sense wire separation	6.0 mm radially with 200 μ m stagger
Sense wire specifications	30 μ m Au-plated W, 110 g tension
Sense wire voltage	+1.45 kV (inner SW) – +1.58 kV (outer SW)
Number of delay lines	2 per cell, 256 total
Delay line velocity	2.35 mm/ns
Field wire specifications	125 μ m Au-plated CuBe, 670 g tension
Gas composition	Ar 93%, CH ₄ 4%, CO ₂ 3%, H ₂ O 0.5%
Gas pressure	1 atm
Average drift field	620 V/cm
Drift velocity	\approx 34 μ m/ns
Maximum drift distance	7 cm
Gas gain	2×10^4 (inner SW) – 6×10^4 (outer SW)
Position resolution	$r\phi \approx 180 \mu\text{m}$, $z \approx 2.9 \text{ mm}$ (best case)

Table 3: Central Drift Chamber Parameters. [12, p. 41][8, p. 36][18]

5.4 Forward Drift Chambers

The forward drift chambers (FDCs) [11, 18, 20, 22] extend the outer tracking coverage from where the CDC leaves off down to an angle of $\theta \approx 5^\circ$ ($\eta \approx 3.1$). There are two sets of chambers, one located at each end of the CDC. Figure 6 shows an exploded view of one of the FDCs. Each FDC consists of three layers of chambers: two Θ layers sandwiching one Φ layer. The Φ layer is a single chamber divided into 36 azimuthal drift cells, each containing sixteen radial sense wires arranged in a plane containing the beam line. The two Θ cells consist of four separate quadrant modules, each of which is composed of six rectangular drift cells at increasing radii. Each cell contains eight sense wires oriented in a plane parallel to the z -axis and normal to the radial direction. The inner three cells are half cells in which the wires are placed at the edge of the cell; thus the electrons can drift in only one direction. Each Θ cell also contains a delay line of the same type as in the CDC to measure the position along the length of the cell. There are no delay lines in the Φ chamber. The two Θ chambers are rotated relative to each other by an angle of $\pi/4$. The construction of the FDCs is similar to that of the CDC, except that the FDCs use $25\mu\text{m}$ aluminum etchings instead of the CDC's epoxy strips.

Further parameters of the FDC are given in Table 4. See [11, 22, 23] for additional details on the construction and testing of these chambers.

5.5 Transition Radiation Detector

When a charged particle crosses between two materials with different dielectric constants, it radiates in the forward direction [15, p.136]. The intensity of this radiation is proportional to the relativistic gamma factor, $\gamma = E/(mc^2)$, and is concentrated on a cone with a half opening angle of $1/\gamma$. These characteristics suggest that one can use this effect to discriminate between particles which have similar energies but different masses. The radiation emitted from a particle crossing a single boundary is not readily detectable, so one must stack a large number of boundaries in order to give a reasonable signal. For highly relativistic particles, the radiation will be in the X-ray range.

The DØ transition radiation detector (TRD) [24, 25] is a device designed to distinguish electrons from heavier particles using transition radiation. It occupies the space between the VTX and the CDC and consists of three radial layers. Each layer has a radiator consisting of 393 layers of $18\mu\text{m}$




Figure 6: Exploded view of one of the FDCs. [1]

	Θ modules	Φ modules
Active radial extent	11 cm – 62 cm	11 cm – 61.3 cm
Active z extent	104.8 – 111.2 cm 128.8 – 135.2 cm	113.0 – 127.0 cm
Number of cells per layer	4 quadrants of 6 cells	32
Number of sense wires	8 per cell, 384/FDC	16 per cell, 576/FDC
Sense wire separation	8.0 mm radially with 200 μm stagger	
Sense wire specifications	30 μm Au-plated W, 50 – 100 g tension	
Sense wire voltage	+1.55 kV	+1.66 kV
Number of delay lines	1 per cell, 48/FDC	—
Delay line velocity	2.35 mm/ns	—
Field wire specifications	163 μm Au-plated Al, 100 – 150 g tension	
Gas composition	Ar 93%, CH ₄ 4%, CO ₂ 3%, H ₂ O 0.5%	
Gas pressure	1 atm	1 atm
Average drift field	1.0 kV/cm	1.0 kV/cm
Drift velocity	40 $\mu\text{m}/\text{ns}$	37 $\mu\text{m}/\text{ns}$
Maximum drift distance	5.3 cm	5.3 cm
Gas gain	2.3×10^4 (inner SW) 5.3×10^4 (outer SW)	3.6×10^4
Position resolution (drift)	$\approx 300 \mu\text{m}$	$\approx 200 \mu\text{m}$

Table 4: Forward Drift Chamber Parameters. [12, p. 45][8, p. 38][18]

polypropylene foil with a mean separation of $150\text{ }\mu\text{m}$. The gaps between the foils are filled with dry nitrogen. Surrounding each radiator is a cylindrical xenon-filled drift chamber to detect the emitted X-rays. At the Tevatron, electrons are the only charged particles likely to be produced with sufficient energy to produce detectable transition radiation. See [25, 26] for further information about the performance of the TRD.

5.6 Central Detector Readout

For the first stage of signal processing, the signals from the chamber wires are led into preamplifiers mounted on the ends of the chambers themselves. From there, the signals are fed into analog pulse shaping cards located on the platform underneath the detector. Finally the signals are sent to flash analog-to-digital converters (FADCs) located in the moving counting house. There the signals are sampled and digitized at a rate of 106 MHz, starting at the beam crossing. If the event is not accepted by the level-1 trigger (see section 8.1), the data are overwritten by the next crossing. Otherwise, the data are compressed by eliminating the flat portions of the signal between the pulses and sent on to the level-2 trigger. See [11, pp. 41–45] and [27] for more information on the CD electronics.

6 Calorimetry

Conceptually, a calorimeter is a device which stops particles in an absorber and measures the amount of deposited kinetic energy. Since $D\bar{O}$ has no central magnetic field, calorimetry is the only available method for measuring the energy of most types of particles. Thus, good calorimetry is a vitally important part of the detector. In-depth discussions of calorimetry in high-energy physics may be found in [15, 28, 29].

6.1 Calorimetry Principles

When a high-energy electron ($\gg 10\text{ MeV}$) passes through a high- Z material, the primary mechanism by which it loses energy is through Bremsstrahlung, in which a charged particle interacts with the Coulomb field around a nucleus and emits an energetic photon. A high-energy photon, on the other hand, will interact predominately by pair production, in which a photon converts into

an electron-positron pair in the vicinity of a nucleus. The particles emitted in these interactions can themselves undergo Bremsstrahlung or pair production. Thus, an energetic electron or photon passing through a dense material will result in a shower of secondary electrons, positrons, and photons. This process is called an *electromagnetic shower*. The shower will continue to develop until all the secondaries have sufficiently low energies that other energy loss mechanisms (mostly ionization) become important. The rate at which an incident particle loses energy is a constant of the material, and is usually specified as a *radiation length* X_0 :

$$\frac{dE}{E} = -\frac{dx}{X_0}. \quad (7)$$

As an example of a typical value, the radiation length for uranium is about 3.2 mm [30, page 1241].

Hadronic particles also cause showers, but they are qualitatively different from the electromagnetic showers caused by electrons and photons. Hadrons lose energy primarily through inelastic collisions with atomic nuclei. These collisions produce secondary hadrons, which can in turn undergo inelastic collisions. This process is called a *hadronic shower*, and it continues until all particles have either been stopped by ionization losses or absorbed by nuclear processes. The scale for the size of these showers is given by the nuclear interaction length λ for the material. For uranium, $\lambda \approx 10.5$ cm [30, page 1241]. Thus, hadronic showers are much more extended in space than electromagnetic showers of similar energy.

The showering process converts a single high-energy particle into many low-energy particles, all of which have about the same energy. The next step is to measure these low-energy particles. One approach is to make the absorber also function as an active material, such as a scintillator. Some typical materials for this are NaI, BGO ($\text{B}_4\text{Ge}_3\text{O}_{12}$), and lead glass. Such *homogeneous* calorimeters offer the best energy resolution; however, they are often not practical for large, high-energy detectors.

An alternate approach is to interleave layers of a dense, inert absorber with layers of a material which is sensitive to particles passing through it (such as a scintillator or some medium sensitive to ionization). This approach yields what is called a *sampling* calorimeter. Since most of the energy is absorbed in the inert material, only a portion of the incident energy can be detected. This fraction is called the *sampling fraction*. Since this is a statistical process, sampling degrades the achievable energy resolution.

Note that there is no *a priori* reason why the response of a calorimeter (i.e., the ratio of the measured signal to the energy of an incident particle) should be the same for electromagnetic and hadronic showers. In fact, the response will tend to be smaller for hadronic showers [31], since ν 's and μ 's produced by π and K decays will escape from the detector. The energy spent breaking up nuclei will also be invisible. This is quantified by the e/π ratio, the ratio of the calorimeter responses to electrons and pions. It is highly desirable that this number be close to 1. The reason is this: a hadronic shower will include not only hadronic components but also electromagnetic components deriving from γ 's from π^0 and η decays. The fraction of a hadron's energy which is deposited as electromagnetic showers can undergo large variations from shower to shower, but if the e/π ratio is 1, these fluctuations will not affect the energy resolution. A calorimeter with $e/\pi \approx 1$ is called a *compensating* calorimeter.

Even if a calorimeter is perfectly compensating, however, there are many additional effects which will tend to degrade the resolution. Since the showering and sampling processes are statistical in nature, one would expect the fractional error in the calorimeter signal to scale as $1/\sqrt{N_{ion}}$, where N_{ion} is the number of ionization electrons liberated. Since this quantity should be proportional to the total energy of the incident particle, one would expect the resolution to scale as $1/\sqrt{E}$. Some other sources of energy fluctuations in the DØ calorimeter include [12, p. 51]:

- Leakage of energy out of the calorimeter.
- Variations in high voltage, absorber thickness and spacing, electronics gain, LAr temperature, or LAr purity.
- Noise due to the natural radioactivity of the uranium plates.
- Electronic noise.

6.2 Calorimeter Geometry

The DØ calorimeter is a sampling calorimeter, with liquid argon (LAr) as an ionization medium. The primary absorber material is depleted uranium, with copper and stainless steel used in the outer regions. Since uranium is very dense, the physical size of the calorimeter can be kept small.




Figure 7: Schematic view of a calorimeter cell. [8, p. 44]

The calorimeter is divided into a large number of modules, each of which consists of a stack of interleaved absorber plates and signal boards. Figure 7 shows a cross-sectional schematic view of a section of this stack. The absorber plates are separated from the signal boards by a LAr-filled gap of 2.3 mm. The signal boards consist of a copper pad sandwiched between two 0.5 mm thick pieces of G10. The outer surfaces of these boards are coated with a resistive epoxy coating. During operation, the absorber plates are grounded, while a positive voltage of 2.0 – 2.5 kV is applied to the resistive coatings. As a shower develops in the calorimeter, charged particles crossing the LAr gap leave a trail of ionization. The liberated electrons are collected on the signal board after a drift time on the order of 450 ns, and induce a signal on the copper pad via capacitive coupling. In order to measure the transverse positions of showers, the readout pads are subdivided into smaller cells. The corresponding cells in adjacent signal boards are ganged together in depth to form *readout cells*; the details of this ganging vary from module to module.

An overview of the calorimeter is shown in Figure 8. Since the calorimeter uses LAr, it must sit inside of a cryostat in order to keep the argon cold. In order to allow access to the central detectors, the calorimeter is divided into


DØ LIQUID ARGON CALORIMETER

Figure 8: The DØ calorimeter. [1]

Module type	EM	FH	CH
Rapidity coverage	± 1.2	± 1.0	± 0.6
Number of modules	32	16	16
Absorber ^a	DU	DU-Nb	Cu
Absorber thickness (mm)	3	6	46.5
Argon gap (mm)	2.3	2.3	2.3
Number of signal boards	21	50	9
Number of readout layers	4	3	1
Cells per readout layer	2,2,7,10	20,16,14	9
Total radiation lengths (X_0) ^b	20.5	96.0	32.9
Total nuclear absorption lengths (λ) ^b	0.76	3.2	3.2
Sampling fraction (%)	11.79	6.79	1.45
Total readout cells	10,368	3000	1224

Table 5: Central Calorimeter Parameters. [12, p. 55] [8, p. 49] [11, p. 26] [32, p. 22]

^aOne of depleted uranium (DU), depleted uranium with 1.7% niobium (DU-Nb), or copper (Cu).

^bAt $\eta = 0$.

three major assemblies, each in its own cryostat: the central calorimeter (CC) and the two end calorimeters (EC).

The central calorimeter provides coverage out to a pseudorapidity of about 1.2. It is roughly toroidal, and consists of three concentric layers of modules. The inner layer consists of 32 electromagnetic (EM) modules, which are thick enough to contain most electromagnetic showers. The middle layer consists of 16 fine hadronic (FH) modules, which measure the showers due to hadronic particles. The final layer consists of 16 coarse hadronic (CH) modules, which measure any leakage of energy out of the FH layer, and which also serve to reduce punchthrough. The parameters of the CC modules are given in Table 5.

The two sections of the end calorimeter provide coverage on each side of the CC from a pseudorapidity of about 1.3 out to about 4. The EC is composed of three concentric layers of modules. Like the CC, the modules are divided into electromagnetic and fine and coarse hadronic types; however, the geometry is rather different. The center of the EC consists of a disc-shaped electromagnetic module, backed by the cylindrical fine and coarse

Module type	EM	IFH	ICH	MFH	MCH	OH
Rapidity range	1.3–3.7	1.6–4.5	2.0–4.5	1.0–1.7	1.3–1.9	0.7–1.4
No. of modules	1	1	1	16	16	16
Absorber ^a	DU	DU-Nb	SS	DU-Nb	SS	SS
Thickness (mm)	4	6	6	6	46.5	46.5
LAr gap (mm)	2.3	2.1	2.1	2.2	2.2	2.2
Signal boards	18	64	12	60	14	24
R.O. layers	4	4	1	4	1	3
Cells/r.o. layer	2,2,6,8	16	14	15	12	8
Total rad. len.	20.5	121.8	32.8	115.5	37.9	65.1
Total abs. len.	0.95	4.9	3.6	4.0	4.1	7.0
Samp. frac. (%)	11.9	5.7	1.5	6.7	1.6	1.6
Tot. r.o. cells	7488	4288	928	1472	384 + 64 + 896 ^b	

Table 6: End Calorimeter Parameters. [12, p. 57] [8, p. 49] [11, p. 29] [32, p. 25]

^aOne of depleted uranium (DU), depleted uranium with 1.7% niobium (DU-Nb), or stainless steel (SS).

^bMCH and OH cells are summed together at $|\eta| = 1.4$

inner hadronic modules. Arranged in a ring around this central core are the fine and coarse middle hadronic modules, and around them is a final ring of coarse outer hadronic modules. The parameters of the EC modules are given in Table 6.

In both the CC and EC, the area in η, ϕ space covered by a typical readout cell is 0.1×0.1 . However, in the third layer of the EM modules, where electromagnetic showers typically deposit the bulk of their energy, the readout cells have areas of 0.05×0.05 . In addition, cells with $|\eta| > 3.2$ have a ϕ size of 0.2 and are somewhat larger in η as well. See Figures 9 and 10 for an illustration of the calorimeter segmentation.

If one examines Figure 9, it is apparent that in the transition region between the CC and the EC ($0.8 < |\eta| < 1.4$) there is a relatively large amount of uninstrumented material. This is primarily due to the cryostat walls and the support structures for the calorimeter modules. Two additional devices are used in this region to recover some of the energy deposited in this dead material. The first of these is the massless gaps (MG). These are simply rings of two signal boards mounted on the end plates of the CCFH, ECMH,





Figure 9: Side view of the calorimeters. [1]

and ECOH modules. The second device is the intercryostat detector (ICD). This is a ring of scintillation counters mounted on the exterior of the EC cryostats. Both the MG and the ICD have the standard segmentation of 0.1×0.1 in η, ϕ space.

6.3 Calorimeter Readout

The signals induced on the readout pads are pulses with widths on the order of 450 ns [8, p. 48]. These signals are led out through four ports in the cryostats to charge sensitive preamplifiers mounted on top of the cryostats. From the preamplifiers, the signals are led to base line subtractor (BLS) modules located in the platform below the detector. The BLS modules perform analog shaping and split the signal into two paths.

The first path is used for triggering. The signals from all the fine hadronic cells within a 0.2×0.2 tower are summed; the signals from the electromagnetic cells are similarly summed. These signals then form the input to the level-1 calorimeter trigger (see section 8.1.3).

Figure 10: Layout of calorimeter channels in depth and η . [33]

The second path is used for the data readout. The incoming signal is sampled just before the beam crossing and again $2.2\ \mu\text{s}$ later, when the shaped signal peaks. The difference between these two samples is sent to the ADCs where, if the event is accepted by the level-1 trigger, the signals are digitized and sent on to the level-2 trigger.

6.4 Calorimeter Performance

The response of the calorimeter modules to single electrons and pions has been studied extensively in test beams [1, sect. 3.7], [34, 35, 36, 37]. The response for both is found to be linear to within 0.5%.

The resolution is parameterized as

$$\left(\frac{\sigma}{E}\right)^2 = C^2 + \frac{S^2}{E} + \frac{N^2}{E^2}, \quad (8)$$

where the constants C , S , and N represent calibration errors, sampling fluctuations, and noise contributions, respectively. For electrons, the measured resolutions are

$$C = 0.003 \pm 0.002, \quad S = 0.157 \pm 0.005(\text{GeV})^{\frac{1}{2}}, \quad N \approx 0.140\ \text{GeV} \quad (9)$$

and for pions,

$$C = 0.032 \pm 0.004, \quad S = 0.41 \pm 0.04(\text{GeV})^{\frac{1}{2}}, \quad N \approx 1.28\ \text{GeV}. \quad (10)$$

The position resolution for electrons is found to be about 0.8–1.2 mm, varying approximately as $1/\sqrt{E}$.

The e/π ratio of the calorimeter falls from about 1.11 at 10 GeV to about 1.04 at 150 GeV.

7 Muon System

Although muons decay into electrons, their lifetime of $2.2\ \mu\text{s}$ is so large compared with the scale of the detector that for all practical purposes a relativistic muon may be regarded as stable. Muons do not interact strongly, and moreover, their mass is too large ($\approx 200m_e$) to readily initiate an electromagnetic shower at Tevatron energies. Thus, a charged particle which makes it out of the calorimeter is likely to be a muon.




Figure 11: Side elevation of the muon system. [1]

The DØ muon system [38] consists of five magnetized iron toroids which are surrounded by three layers of proportional drift tubes (PDTs). See Figure 11. The PDTs measure the trajectory of muons before and after they traverse the magnetized iron; thus a measurement of the muon momentum can be made. The five magnets are the CF (Central Fe), covering the angular range $|\eta| < 1$, the two EFs (End Fe), covering the angular range from $|\eta| = 1$ out to about $|\eta| = 2.5$, and the two SAMUS (Small Angle MUon System) magnets, covering the range from about $|\eta| = 2.5$ out to about $|\eta| = 3.6$. The CF and the two EFs together are referred to as the Wide Angle Muon System, or WAMUS. Each section has one layer of drift tubes (the A layer) just inside the magnet, a second layer (B layer) just outside the magnet, and a third layer (C layer) 1 – 3 m further out. Due to the necessity of supporting the calorimeter, some regions underneath the detector are not covered by the full compliment of three drift tube layers.

The more material which must be traversed prior to entering the muon system, the smaller the background due to hadronic punchthrough will be. Figure 12 shows how the thickness of the detector in nuclear interaction




Figure 12: Number of nuclear interaction lengths as a function of polar angle. [1]

lengths varies with polar angle. At $\eta = 0$ the minimum muon momentum required to make it through both the calorimeter and iron is about $3.5 \text{ GeV}/c$ [12, p. 59]. At higher η , this rises to about $5 \text{ GeV}/c$.

7.1 WAMUS System

The PDTs making up the WAMUS system are rectangular in cross section, with one sense wire per drift cell. They are formed out of aluminum extrusions into rectangular modules containing either four (A layer) or three (B and C layers) planes of drift tubes (see Figure 13). There are 164 modules, each containing between 14 and 24 drift tubes, the lengths of which are between 191 and 579 cm. The tubes are oriented roughly parallel to the direction of the magnetic field in the iron toroid. In this manner, the deflection due to the magnet is measured by the drift time. The distance of a hit from the sense wire can be determined with a resolution of about 0.5 mm.

In order to facilitate access to the chamber electronics, all readout is done from one end of the drift tubes. To accomplish this, tubes are jumpered together in pairs at the far end. The front end electronics measures the arrival times of pulses at the end of each wire, as well as the time difference between pulses arriving at the ends of each jumpered pair of sense wires.




Figure 13: End view of muon drift tubes. [1]




Figure 14: Muon cathode pad. [1]

Using this time difference, a crude measure of the position of a hit along the wire can be found (within 10 – 20 cm). A more precise determination of the hit position is made by inserting at the top and bottom of each tube a set of vernier cathode pads (see Figure 14). These pads are insulating Glasteel coated with copper cladding and forming a diamond pattern which repeats with a period of ≈ 60 cm. The cladding is separated into an inner region and an outer region. Electron avalanches on the sense wire will induce pulses on the cathode pads. The ratio of the charge deposited on the inner and outer pads can be used to localize the hit to within about 3 mm modulo the half-repeat period of about 30 cm. This ambiguity is resolved using the coarse time difference measurement. The cathode pads in adjacent layers of PDTs are staggered by about 1/6 of a repeat period relative to each other to further resolve ambiguities which occur near the corners of the diamond pattern. Each tube also has a single bit output which is set if there were any hits on the cathode pad in that tube (the *pad latch*). These bits form the input to the muon trigger (see section 8.1.4).

To measure the momentum of a muon, its trajectory before and after passing through the magnet must be known. Tracks through the B and C layers give the trajectory after the magnet, while tracks through the A layer give the trajectory before the magnet. If possible, the A layer tracks are matched to tracks in the central detector and to minimum ionizing traces in the calorimeter in order to increase the lever arm. The primary interaction vertex found by the central detector may also be used to define the incoming track. The momentum (P) resolution is most easily formulated in terms of the inverse momentum $k = 1/P$ as

$$\left(\frac{\delta k}{k}\right)^2 = (0.18)^2 + \left(\frac{0.01}{k \text{ GeV}}\right)^2. \quad (11)$$

Further parameters of the muon system are given in Table 7. For further details, see [1, 38, 39].

7.2 SAMUS System

Due to the high occupancies in the forward region, the SAMUS system uses smaller drift tubes. It is composed of three stations each consisting of three drift tube planes. Each plane is composed of two subplanes, offset by half a tube diameter. The three layers are rotated with respect to each other.

	WAMUS	SAMUS
Rapidity coverage	$ \eta < 1.7$	$1.7 \leq \eta \leq 3.6$
Magnetic field	2 T	2 T
Nuclear interaction lens.	≈ 13.4	≈ 18.7
Number of modules	164	6
Number of drift cells	11,386	5308
Sense wire specifications	50 μm Au-plated W, 300 g tension	50 μm Au-plated W, 208 g tension
Maximum sagitta	0.6 mm	2.4 mm
Sense wire voltage	+4.56 kV	+4.0 kV
Cathode pad voltage	+2.3 kV	—
Gas composition	Ar 93%, CF ₄ 5%, CO ₂ 5%	CF ₄ 90%, CH ₄ 10%
Bend view resolution	± 0.53 mm	± 0.35 mm
Non-bend view resolution	± 0.3 mm	± 0.35 mm
Average drift velocity	6.5 cm/ μs	9.7 cm/ μs
Maximum drift distance	5 cm	1.45 cm

Table 7: Muon System Parameters. [1, 38][12, p. 61]

The drift tubes themselves are constructed from stainless steel tubes with a 3 cm diameter, each containing a single sense wire with a $50\text{ }\mu\text{m}$ diameter. Some additional parameters of the SAMUS system are given in Table 7. For further details, see [1, 38, 40].

8 Triggering and Readout

At the Tevatron, beam crossings occur at the interaction region at a rate of about 290 kHz. At a luminosity of $5 \times 10^{30}\text{ cm}^{-2}\text{ s}^{-1}$, an inelastic collision will occur in about 3/4 of these crossings [1, sect. 5.1]. However, the processes which are of the greatest interest are much rarer. Because it is infeasible to record and process data from every crossing, there must be some mechanism to select out the small fraction of interesting events for permanent storage. This process is called *triggering*.

The overall layout of the DØ trigger system is shown in Figure 15. It can be conceptually divided into two hierarchical pieces: level-1 and level-2. The level-1 trigger is a collection of dedicated hardware processors which operate on a coarse subset of the event data. Most level-1 trigger decisions can be made within the $3.5\text{ }\mu\text{s}$ interval between beam crossings, permitting operation without deadtime. However, some triggers, called level-1.5 triggers, may require longer. The goal of the level-1 trigger is to reduce the event rate from the beam crossing frequency of 290 kHz to a rate of 200 – 300 Hz.

Once an event has been accepted by level-1, the complete event is digitized and the data transferred to one of 48 level-2 nodes. These are general-purpose computers which process events in parallel. They perform a fast reconstruction of the event, and can use arbitrary software filters to decide whether or not an event should be kept. If level-2 passes an event, it is transferred to the host system, where it is permanently recorded on magnetic tape.

The components of the trigger system are further described below.

8.1 Level 1

At the heart of the trigger system is the level-1 trigger framework [41, 42, 43]. This is a special-purpose hardware processor which is responsible for combining the results of individual level-1 components into a set of global decisions, commanding the readout of the digitization crates, and interfacing

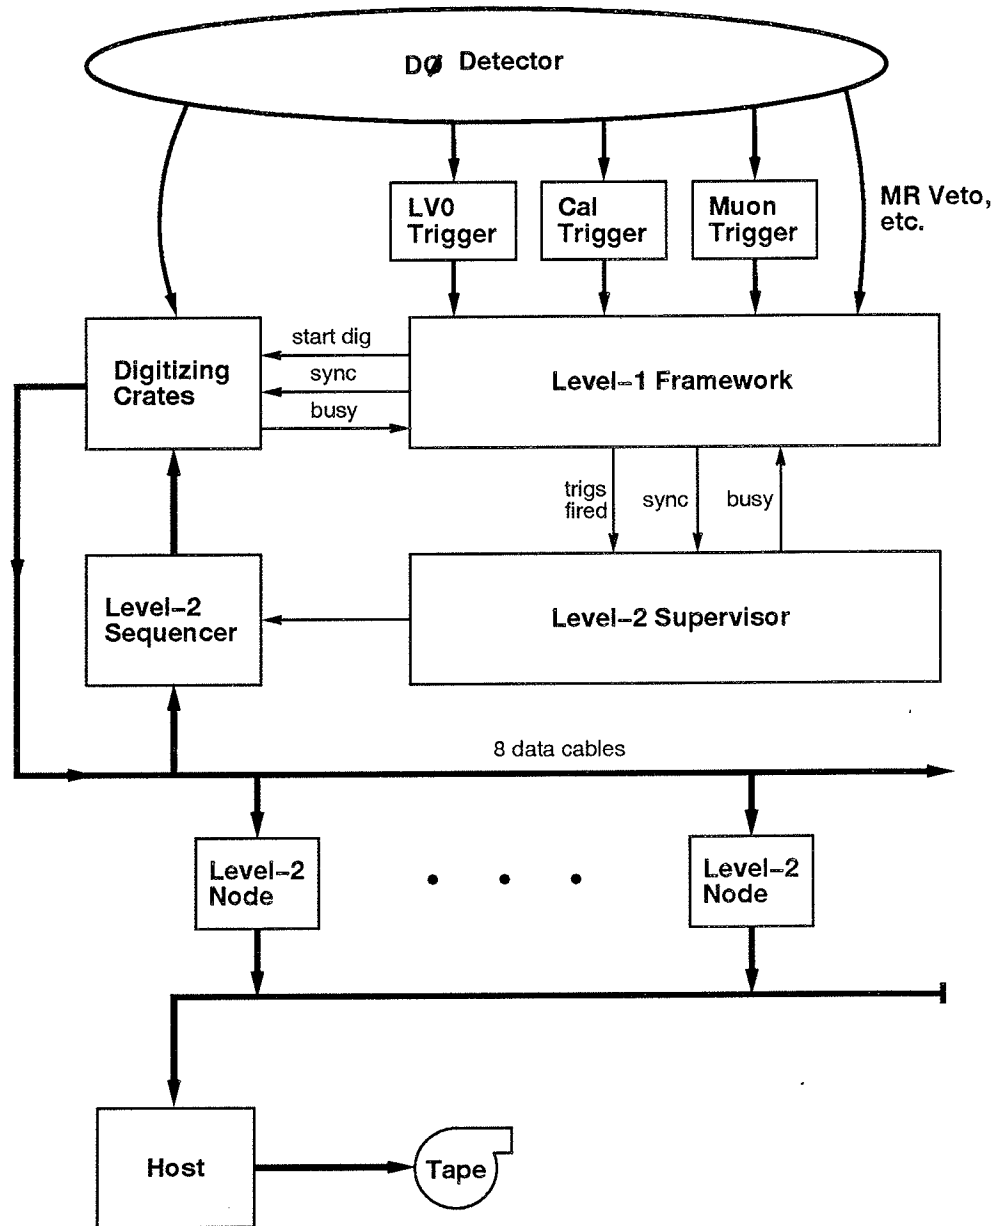


Figure 15: Overall layout of the DØ trigger system.

with level-2. The primary input to the framework consists of 256 *trigger terms*. Each of these is a single bit, indicating that some specific requirement is met for the present event (such as ‘at least two muons’, or ‘ $E_T > 20 \text{ GeV}$ ’). Most of these inputs come from detector-specific level-1 processors, but some come directly from sources such as scintillators or from accelerator timing signals.

The 256 trigger terms are reduced to a set of 32 level-1 *trigger bits*, also known as *specific triggers*, by an and/or network. Each trigger bit can be programmed to require that certain trigger terms be set and that certain other terms not be set. Each trigger bit also has a programmable prescale; setting the prescale to some value N means that the trigger will only actually fire once in every N times that its trigger term conditions are satisfied.

The digitizing hardware is located in 86 front-end VME crates in the moving counting house. For purposes of readout, these crates are grouped into 32 *geographic sectors*. For each specific trigger bit, the framework has a list of geographic sections to be read out. When a trigger fires, the framework commands the appropriate set of sectors to begin digitizing. The crates are double buffered, so they can start digitizing an event while a previous one is still being transferred to level-2. Each geographic sector also sends a busy signal back to the framework when any crate in it is unable to accept another trigger. This will inhibit any further triggers involving that geographic sector.

At the same time that the framework commands the front ends to begin digitizing, it sends a signal to the level-2 system containing the mask of specific triggers which fired. The level-2 supervisor then manages the transfer of the event data from the digitizing crates to a level-2 node, as described below. The level-2 supervisor also sends a busy signal to the framework when there are no free nodes to accept an event.

If a level-1.5 decision is required, the framework starts the digitization cycle as usual, but delays notifying level-2. When level-1.5 processing is complete, the framework either sends the event on to level-2 or aborts the digitization cycle, depending on the result.

The trigger framework also maintains numerous counters to monitor system dead time, the luminosity seen by each trigger, and other performance measures.

The trigger framework is controlled by a dedicated Vaxstation 4000/60, called the trigger control computer (TCC). The TCC does not directly participate in trigger processing, but instead is responsible for programming the level-1 hardware at the beginnings and ends of runs, and for monitoring the

performance of the system. The TCC is the interface through which the rest of the data acquisition system talks to the trigger framework.

The sections below describe some of the major inputs to the trigger framework.

8.1.1 Level 0

The level-0 system [44, 45] is designed to detect events containing an inelastic scattering and to provide a fast measurement of the location of the event vertex. It consists of two separate detectors located at each end of the central detector between the FDC and the EC. Each detector consists of two layers of rectangular scintillator bars, read out with photomultipliers. Due to the rectangular geometry of the scintillators, the coverage is not uniform in ϕ . There is nearly complete coverage in the range $2.2 < |\eta| < 3.9$ and partial coverage extending out to $1.9 < |\eta| < 4.3$. An inelastic $p\bar{p}$ collision will typically include a large amount of activity in the far forward regions (from the spectator quarks); thus, one looks for a coincidence between the signals from the two scintillator arrays. The level-0 trigger is $> 99\%$ efficient for nondiffractive inelastic collisions. By comparing the arrival times of the signals from the two arrays, the approximate position of the interaction vertex may be found. A fast vertex determination with a resolution of ± 15 cm is available within 800 ns after the collision. A more accurate determination with a resolution of ± 3.5 cm is available within $2.1 \mu\text{s}$ [8, p. 59]. The vertex position is available as several level-1 trigger terms; it can also be used in level-2 processing. The level-0 system can also identify events which are likely to contain multiple interactions.

8.1.2 Main Ring Vetoes

During normal operation of the Tevatron, the Main Ring is used to produce antiprotons, with a cycle period of 2.4 s [8, p. 129]. Since the Main Ring passes through the DØ detector, losses from the Main Ring will show up in the detector and must be rejected. The largest losses occur when beam is injected to the Main Ring, and again 0.3 s later when the beam passes through transition². These losses are dealt with by vetoing on the MRBS_LOSS

²Consider a bunch of non-relativistic particles travelling in a circular orbit. The particles with a larger than average momentum will also have a larger than average velocity and will pull ahead of the rest of the bunch. In order to keep the bunch from blowing

trigger term. This term is asserted as a possible veto during a 0.4 s window starting at injection, continuing through transition, and allowing time for the calorimeter and muon high voltage to recover from the large losses. This results in a dead time of about $0.4/2.4 \approx 17\%$ [46].

It is also possible to have small losses whenever a Main Ring bunch passes through the detector. These can be mostly eliminated by vetoing on the microblank trigger term, which is asserted whenever a Main Ring bunch is present in the detector during the livetime of the muon system (± 800 ns around the beam crossing). This adds an additional $\approx 8\%$ deadtime [46].

8.1.3 Calorimeter Trigger

The level-1 calorimeter trigger [42, 43] is responsible for making trigger decisions based on calorimeter information. It is a pure level-1 trigger, and thus must be able to return a decision within $2.4 \mu\text{s}$ [47]. The inputs derive from the trigger pickoffs in the calorimeter BLS cards, which sum cells into towers of size 0.2×0.2 in $\eta - \phi$ out to a pseudorapidity of 4 (see section 6.3). Separate inputs are provided for cells originating from EM modules and FH modules, 1280 of each (the CH calorimeter modules are not used for the level-1 trigger). At the input of the calorimeter trigger, all the inputs are simultaneously flash digitized; all subsequent calculations are entirely digital. The trigger calculates a number of global sums of calorimeter inputs. These sums are:

- The total electromagnetic energy, $E(\text{em}) = \sum_i E_i(\text{em})$.
- The total hadronic energy, $E(\text{had}) = \sum_i E_i(\text{had})$.

up longitudinally, one must therefore arrange for the head of the bunch to be decelerated relative to the rest of the bunch, and for the tail to be accelerated (again, relative to the rest of the bunch). For highly relativistic particles, however, the situation is different. In this regime, the velocity of a particle is nearly constant (at c) regardless of its momentum. However, the path length is not constant: a particle with larger than average momentum will have a larger than average bending radius and will thus fall behind the rest of the bunch. So in this situation, one must accelerate the head of the bunch more than the tail. The point in the acceleration cycle at which the switch between these two descriptions occurs is called *transition*; the energy at which it occurs depends both on the mass of the particles being accelerated and the size of the accelerator ring. Properly rearranging the accelerating fields when passing through transition is tricky, and accelerators often experience extra losses at that point.

- The total scalar sum of electromagnetic transverse energy,
 $E_T(\text{em}) = \sum_i E_i(\text{em}) \sin \theta_i$.
- The total scalar sum of hadronic transverse energy,
 $E_T(\text{had}) = \sum_i E_i(\text{had}) \sin \theta_i$.
- The total transverse energy, $E_T(\text{tot}) = E_T(\text{em}) + E_T(\text{had})$.
- The missing transverse energy, $\cancel{E}_T = \sqrt{E_x^2 + E_y^2}$, where

$$E_x = \sum_i (E_i(\text{em}) + E_i(\text{had})) \sin \theta_i \cos \phi_i \quad (12)$$

and

$$E_y = \sum_i (E_i(\text{em}) + E_i(\text{had})) \sin \theta_i \sin \phi_i. \quad (13)$$

These quantities are then compared with a set of programmable thresholds. Each of these comparisons yields a trigger term which is input to the trigger framework. This allows one to specify trigger requirements such as ‘total scalar E_T above 140 GeV’, or ‘missing E_T above 20 GeV’.

Each individual trigger tower also has a set of associated thresholds. The EM energy in each tower is compared with four programmable thresholds; the hadronic energy is also compared with a separate set of four thresholds. For each tower and pair of thresholds, a bit is set if the EM energy is greater than, and the hadronic energy less than, their respective thresholds. In addition, the sum of the EM and hadronic energy is compared with another set of four thresholds and a separate bit set for each tower for which this threshold is exceeded. A count is then made of the number of towers for which each bit is set and the result compared to another set of thresholds. For each threshold that is exceeded, a trigger term is asserted. This allows one to specify trigger requirements such as ‘one EM tower above 5 GeV in the central region’ or ‘two towers with total energy above 10 GeV anywhere in the calorimeter’.

8.1.4 Muon Trigger

The inputs to the level-1 muon trigger [48] are the pad latch outputs from the muon system (see section 7.1). There is one bit from each tube, for a total of about 16,700 bits. The level-1 trigger divides the muon detector into five regions: CF, EF-north, EF-south, SAMUS-north, and SAMUS-south.

Within each region, the trigger looks for patterns of hits which are consistent with a muon emitted from the nominal interaction vertex. In most regions, hits are required in all three layers of muon drift tubes; however, in some CF areas without three layer coverage, tracks are allowed to pass with hits in only two layers. The level-1 trigger counts the number of track candidates in each region and compares the result with a set of preset thresholds. If any thresholds are exceeded, the appropriate trigger term is asserted to the level-1 framework.

The pure level-1 muon trigger is not capable of distinguishing between muons of different momenta because it matches tracks between layers very coarsely. There is also a level-1.5 muon trigger which is capable of imposing a p_T cut on a muon by requiring a finer match between layers. However, due to the combinatoric problems of doing this matching, especially in the busy SAMUS regions, this computation often takes longer than the $2.4\ \mu\text{s}$ allotted for level-1 trigger term decisions. (Decision times typically range from 1 to $5\ \mu\text{s}$ in the WAMUS regions, but can take up to $100\ \mu\text{s}$ in the busy SAMUS region.) Thus, enabling the level-1.5 trigger may cause extra dead time.

8.2 Level 2

The level-2 trigger [49, 50, 51] is primarily a large farm of general-purpose processors which run software filters using the complete data for an event. Before this can happen, however, the data must be collected together in a single place.

The digitizing hardware is located in a large number of VME crates. Each crate also contains a card called a VME buffer/driver (VBD). These cards are daisy-chained along one of eight high-speed (40 MB/s) data cables, each of which is a loop originating and terminating at a special sequencer card. When the crate has finished digitizing, the VBD copies the data into an internal buffer. It then waits to receive a readout token from the sequencer before sending the data out over the data cable.

Each level-2 node is a Vaxstation 4000/60 (some of which have been upgraded to 4000/90's for run 1B), running the VaxELN realtime operating system. The bus of each node is extended out to a VME crate. Each crate contains four dual multiport memory (MPM) boards which are connected to the eight data cables. Each level-2 node also contains a VBD which is used to transmit events up to the host system.

The process of transferring an event from the digitizing crates to a level-2

node is controlled by the level-2 supervisor. When the supervisor receives a trigger from level-1, it picks an idle level-2 node and enables its MPMs to receive data from the cables. The supervisor then causes the sequencers to start circulating tokens on the data cables. These tokens are used to arbitrate access to the data cable: a VBD may only transmit data while it is in possession of the readout token. When a VBD receives the token and has data to send, it grabs the token and dumps its data out on the cable. Any MPMs on that cable which have been enabled will receive and store this data. The VBD follows its data with another token, which has been modified to indicate that that crate has been read out. Once the complete event has been received by the level-2 node, it is converted to Zebra format [52] by adding appropriate headers to the data received from each cable. The node can then start doing filter processing.

Since the data for an event is initially distributed among many front-end crates, and since each crate may have several events buffered, it is conceivable that a readout error may cause pieces of different events to get mixed together during the readout process. To prevent this, the framework stamps each event with a unique number. This number is sent along with the trigger mask to the level-2 system. The lower four bits, called *sync bits*, are also distributed to each crate along with the start-digitize commands. Level-2 puts the sync bits into the tokens it circulates, and a VBD will only read out if the sync bits in the token match the sync bits associated with the data that that crate has buffered.

The level-2 node software [53] is modular and is composed of a framework, which is responsible for getting events into and out of the node and for communicating with the rest of the system, and a set of *tools*, which do the actual filtering. Each tool can take a set of parameters describing the cuts to be made. For example, the parameters used by the jet tool include, among others, the number of jets to require, the minimum E_T of the jets, the cone size of the jets, etc. Each tool then returns a boolean value telling whether or not the current event passed the cuts. Lists of tools with specific values for the parameters are collected together into *filter scripts*. Each script is associated with one level-1 trigger bit. During normal operation, the framework looks at the mask of level-1 bits for each event. It calls in sequence every script associated with the level-1 bits which were set. For every script, every tool in the script is called; if all the tools in the script pass, then the script as a whole passes. For each script which passes, a bit is set in a 128-bit mask of *filter bits*. If any filter bits are set, the event gets sent to the host, as

described below.

9 Host Processes

The DØ host cluster is a cluster of Vax computers running the VMS operating system. For most of run 1A, it consisted of three CI-connected nodes (an 8700, a 6410, and a 6620), a collection of HSC-served RA-series disks, and a group of satellite Vaxstations. Dedicated X terminals were also used to increase the number of available displays.

The online software is generally structured as a set of detached server processes which respond to requests from users, to requests from other servers, or to events taking place in the experiment. Most communication between processes is done using a locally-developed package called ITC (Inter-Task Communication) [54], which provides a convenient layer over the VMS DECnet and mailbox communication services.

The online processes can be broadly classed into four groups: run control, data logging, downloading, and monitoring.

9.1 Run Control

At the heart of the system is the run coordination program, COOR [55]. This program is responsible for three principal tasks:

- Allocating hardware and software resources to individual users of the system.
- Arranging for hardware and trigger systems to be programmed according to the requests of users of the system.
- Stepping the other parts of the system through the steps needed to begin or end a run.

Users typically interact with COOR through a program called Taker [56]. Through Taker, users choose a *trigger configuration* [57] which they want to use. A trigger configuration is a set of text files which describe to COOR how to configure the online system for a particular purpose. A configuration file will contain descriptions for one or more level-1 trigger bits (see section 8.1). For each bit, it specifies:

- The level-1 programming: what trigger terms are required and what thresholds to set for the individual trigger processors.
- The list of geographical sections to be read out, and any special mode settings for the digitization hardware.
- A list of level-2 filter scripts to be run for events which pass the level-1 trigger.
- For each level-2 script, a list of recording streams to which events passing the filter should be written.

When it receives a trigger configuration request, COOR examines its internal model of the online system to determine if the new request would conflict with any previous requests. If not, COOR issues the appropriate commands to configure the system for the new request and adjusts its internal model accordingly. Since a configuration can be set up to read out only a portion of the detector, it is possible to have multiple users using the system simultaneously, each reading out a different subdetector. This is useful for calibration and testing tasks.

The configuration files which COOR reads are rather cumbersome to work with for complex trigger configurations. To ease the construction and maintenance of these files, a utility called 'trigparse' has been written which takes as input a compact, physics-oriented specification of the triggers and produces as output the full set of configuration files required by COOR [58]. In order to allow reconstruction of past configurations and to track changes, all changes to configuration files are logged using DEC CMS.

9.2 Data Logging

Once a level-2 node has passed an event, it transmits it to the host cluster, where it is received by the data logger [59], a program running on one of the host computers. This program and others associated with it are responsible for receiving raw data from the level-2 system and copying it to magnetic tape, while performing all necessary bookkeeping tasks. Data may also be sent to the express line for rapid reconstruction, and to the online DAQ pool for online monitoring. The overall flow of data in the host system is illustrated in Figure 16.

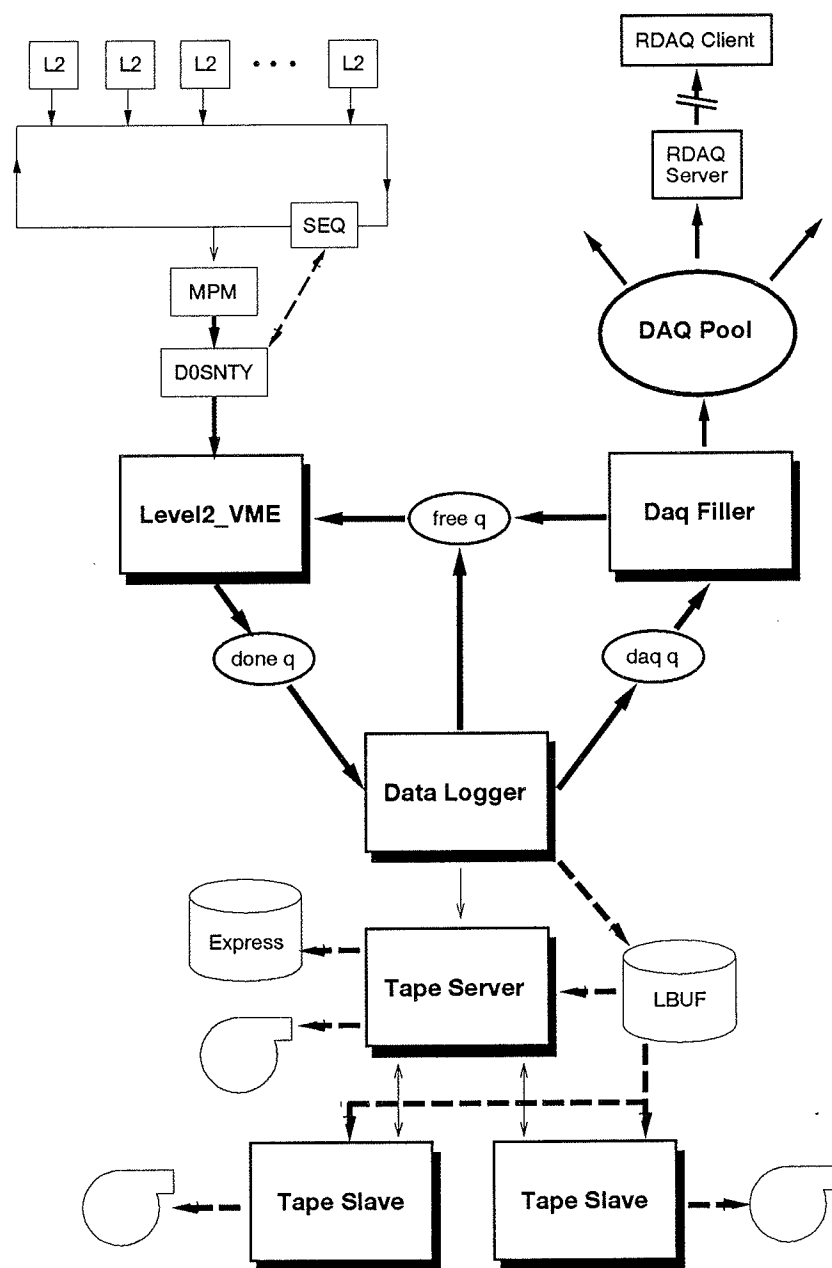


Figure 16: The DØ data logger. [59]

The process of reading out an event through the logger starts when a level-2 node has passed an event and wants to send it to the host. The node copies the event to the VBD (VME Buffer/Driver) card in its VME crate and tells the VBD to read out.

The VBDs are connected along a readout cable, similar to the cables used to convey data into the level-2 nodes. The sequencer for this cable is controlled by D0SNTY, a Vaxstation 4000/60 running VaxELN. While D0SNTY waits for an event to arrive, it circulates tokens around the readout cable. If the token reaches a VBD which has data to send, that VBD grabs the token and sends its event data along the readout cable. D0SNTY then ceases to circulate tokens until it is ready for another event. (Note that this implies that the level-2 nodes at the start of the cable are read out preferentially to those further down the cable.)

As the event is read out over the cable, it is stored in a multiport memory (MPM) board in the D0SNTY VME crate. When the readout is complete, D0SNTY copies the event from the MPM to a prearranged location in VME memory which is mapped to memory in the host VAX through a DWMVA interface. Once this copy is complete, D0SNTY notifies the host by sending an interrupt. When the host is ready for another event, it notifies D0SNTY and tells it where in VME memory the next event should be written.

On the host computer, there are three processes which comprise the core of the logger. These are called `Level2_VME`, `Data Logger`, and `DAQ Filler`. Between them, they maintain a small pool of event buffers.

The `Level2_VME` program is responsible for interfacing with the level-2 readout system. It finds a free buffer in the logger's internal pool, and arranges for that buffer to be mapped into VME address space. It tells D0SNTY to read out into this buffer, and waits for the interrupt signalling the completion of the transfer. It then tells `Data Logger` that a new event is available, and goes to find another free buffer. Once an event is in the event buffer, it does not need to be copied again within the logger.

Note that `Level2_VME` can be replaced with another program for diagnostic or testing purposes. For example, there is a version which reads raw data from a file instead of receiving it from level-2. There is also a version which implements the old DECnet protocol for receiving data from level-2.

When `Data Logger` gets an event from `Level2_VME`, it verifies that the event is intact by computing its checksum and by validating the Zebra structure. If the event appears damaged, it is thrown away. (Damaged events can optionally be dumped to disk for later analysis.) If the event is good,

the logger proceeds to process the event by storing the symbolic names of all the triggers and filters which fired into the event's TSUM bank. The logger then examines the mask of filters which the event passed, and on the basis of this decides to which output streams this event should be written. For each stream, there is a file present on one of several buffer disks to which events in that stream are written. Before writing an event, the logger stamps it with the proper run number and output event number, and performs other bookkeeping tasks. The disk files which are written have a fixed maximum length, which is usually set to 250000 MB. When a file reaches this size, it is closed and a new partition is opened. A command is then sent to the tape server to copy the old file to tape, and information about the file is recorded in a database. There is a fixed limit for any stream of 99 parts per run. When this limit is reached, the logger sends a command to the run coordination program COOR instructing it to stop the run.

The event is then sent to DAQ Filler. (Actually, this is only done if DAQ Filler is idle, in order to prevent hanging the entire system if a problem should develop with the DAQ pool.) DAQ Filler copies the event into the DAQ pool (a.k.a. the *Global Shared Common*) [60]. This is a pool of memory on the host which can hold a collection of events. DAQ Filler inserts events at the start of this list (it is a *producer*), and other processes (*consumers*) can copy events starting from the end of the list. When the pool gets full, old events are dropped from the end of the list to make room for new events. The events in the pool are tagged with the level-2 filter mask, so that events can be selectively read from the pool on the basis of what triggers and filters they passed.

Strictly speaking, the DAQ pool is local to a single node (since it is implemented using shared memory). However, it is possible to read events from the pool from another node by using the RDAQ (Remote DAQ) services. A program using RDAQ communicates over DECnet with a server process running on the node on which the DAQ pool resides. This server copies events out of the pool and ships them over the network to its client.

To return to the data written to the buffer disk: After completing a data file, the logger sends a request to the tape server to copy the file to tape. The tape server assigns the file to a spot on a tape, records this assignment in the tape database, and queues the file to be copied to that tape. If the file is in an express stream (signified by the stream name starting with EXP_) the tape server copies the file to one of the expressline disks. Files in a global monitor stream (starting with GMS_) are also copied to the expressline, but

are not written to tape.

The tape server is capable of writing data to any tape drive which is served to the cluster. However, using a tape drive which is not local to the node on which the server is running incurs a penalty of extra network I/O. In addition, the current VMS implementation of tape serving is not very robust, and has been linked to a number of system crashes. To solve these problems, it is advisable that tape copies be performed by a process running on the node to which the tape drive is attached. This is accomplished by running a tape slave process on such nodes. Before the tape server attempts to mount a tape on a remote node, it attempts to connect to a tape slave running on that node. If it is successful, the tape server instructs the slave to mount the tape and handles all further copies to that tape by sending commands to the slave. If the server is unable to connect to a slave, it does the mount and copy itself.

After a file has been successfully copied to tape, it is removed from the buffer disk. In addition, a small text file (the *drool* file) is written to a particular directory in order to communicate information about the file to the offline production group. If there is an error while copying a file to tape, the tape is ejected, and that file and all other files which were queued for that tape are requeued to another tape.

9.3 Downloading

Most of the programmable hardware at DØ resides in one of a large number of VME crates. Through the use of a device called a *vertical interconnect*, the address spaces of several VME crates are mapped into another VME bus. This master crate will typically contain a dedicated 68000-based microcontroller, which is linked with other such controllers by an experiment-wide token ring. By sending it the appropriate messages, the microcontroller can be made to read from or write to any addresses in the VME crates which it controls. Communication between devices on the token ring and the host cluster is handled by three gateways, dedicated microvax computers which are connected both to the token ring and to the local online cluster network. Programs running on the host cluster talk to the gateways (and thus to any devices on the token ring) via a protocol called CDAQ (Control and Data Acquisition).

One of the primary users of CDAQ is a process called `COMM_TKR`. This process receives hardware configuration requests from `COOR` and is respon-

sible for turning high-level, functional requests (such as ‘turn on the pulser in such-and-such crate’) into low-level, hardware-oriented requests (such as ‘write 0x00080000 into VME address 0x00801240 in the front end at token-ring address 747’). COMM_TKR is also responsible for downloading calibration constants to the hardware. This data is kept in a number of DBL3 databases. These databases are populated by periodically taking special calibration runs using a program called Calib.

9.4 Monitoring

One final set of processes is responsible for monitoring the health of the experiment. Chief among these is the alarm server, which provides a centralized collection point for alarm messages and other significant occurrences. Besides handling downloads, the microcontrollers on the token ring continually scan sets of environment sensors, such as thermometers and readbacks from power supplies. If any of these readings fall outside of a preset tolerance, the microcontroller sends an alarm message to the alarm server. Any program running on the host cluster can also send error messages through the alarm system. Some alarms are deemed fatal; these indicate conditions which could compromise the quality of the data being taken. When a fatal alarm occurs, any physics run in progress is automatically paused until the condition is corrected.

Also used for monitoring incoming data are the examine processes. These are typically run on the satellite workstations. They receive events from the data stream via RDAQ, run some subset of the full $D\bar{O}$ event reconstruction on the data, and produce a set of histograms for each run which can help to identify potential hardware problems.

10 Offline Data Processing

During run 1A, events were primarily written to two streams. The ALL stream contained all events sent to the host and was directly written to tape. The EXPRESS stream contained only events from the small subset of filters which was of the most interest for rapid analysis (mostly top and W/Z triggers). Besides being written to tape, EXPRESS stream events were copied to the express line, where they were reconstructed and immediately made available for analysis. For run 1A, data was recorded at a rate of 2 –

3 Hz, with about 10% of the events going to the express line.

Periodically, the raw data tapes which were written by the online system were collected and transported to the Feynman computing center. There, the raw data was reconstructed on a cluster of Silicon Graphics machines [61]. The product of the reconstruction is two sets of files. STAs contain the full raw data of the event augmented with the results of the reconstruction, and are about 600-1000 kbytes/event. DSTs contain only the reconstruction results for high-level objects, such as electrons, muons, etc. They are about 15 kbytes/event. Most analysis starts with the DSTs. STAs are usually required only when one wants to re-reconstruct an event, or to examine an event in detail with the event display.

The reconstructed events were further split into a set of filter streams. This was done on the basis of filter procedures supplied by the various physics groups, which were designed to select out the events of most interest. The filtered event streams were then made available to users on the DØ file server.

References

- [1] S. Abachi et al. The DØ detector. *Nuclear Instrumentation and Methods in Physics Research A*, 338:185 – 253, 1994.
(URL: <http://d0wop.fnal.gov/d0pubs/journals/NIMA338.185/NIMA338.185.html>)
- [2] N. Amos, J. Linnemann, R. Partridge, and L. Paterno. Luminosity calculations for DØ. DØ note 2031, 1994.
- [3] N. Amos et al. Change to the DØ luminosity monitor constant. DØ note 2186, 1994.
- [4] Leon M. Lederman. The Tevatron. *Scientific American*, 264(3):48 – 55, March 1991.
- [5] Helen T. Edwards. The Tevatron energy doubler: A superconducting accelerator. *Annual Reviews of Nuclear and Particle Science*, 35:605 – 660, 1985.
- [6] FNAL. Design report Tevatron 1 project. FNAL internal note, September 1984.
- [7] FNAL. A report of the design of the Fermi National Accelerator Laboratory superconducting accelerator. FNAL internal note, May 1979.

- [8] William Joseph Thompson. *Search for the Top Quark in the Muon + Jets channel at $D\bar{O}$* . PhD thesis, State University of New York at Stony Brook, Stony Brook, New York, February 1994.
- [9] M. P. Church and J. P. Marriner. The antiproton sources: Design and operation. *Annual Reviews of Nuclear and Particle Science*, 43:253 – 295, 1993.
- [10] D. Möhl, G. Petrucci, L. Thorndahl, and S. van der Meer. Physics and technique of stochastic cooling. *Physics Reports*, C58:73 – 119, 1980.
- [11] Srinivasan Rajagopalan. *The dE/dx Capabilities of the $D\bar{O}$ Tracking System*. PhD thesis, Northwestern University, Evanston, Illinois, June 1992.
- [12] James Herbert Cochran, Jr. *Search for the Truth in the $e\mu$ Channel at $D\bar{O}$* . PhD thesis, State University of New York at Stony Brook, Stony Brook, New York, December 1993.
- [13] Ties Behnke. *The Central Drift Chamber for the $D\bar{O}$ Detector: Design, Construction, and Test*. PhD thesis, State University of New York at Stony Brook, Stony Brook, New York, 1989.
- [14] Richard C. Fernow. *Introduction to Experimental Particle Physics*, chapter 10, pages 234 – 258. Cambridge University Press, 1986.
- [15] Konrad Kleinknecht. *Detectors for Particle Radiation*. Cambridge University Press, 1987.
- [16] F. Sauli. Principles of operation of multiwire proportional and drift chambers. In Thomas Ferbel, editor, *Experimental Techniques in High Energy Physics*, pages 79 – 188. Addison-Wesley, 1987.
- [17] A. R. Clark et al. Test results from a precision drift chamber vertex detector prototype using dimethyl ether. *Nuclear Instrumentation and Methods in Physics Research A*, 261:420 – 426, 1987.
- [18] A. R. Clark et al. The central tracking chambers for $D\bar{O}$. *Nuclear Instrumentation and Methods in Physics Research A*, 279:243 – 248, 1989.

- [19] A. R. Clark et al. D0 vertex drift chamber construction and test results. *Nuclear Instrumentation and Methods in Physics Research A*, 315:193 – 196, 1992.
- [20] D. Buchholz et al. The use of waveform digitizers with drift chambers for D0. *Nuclear Instrumentation and Methods in Physics Research A*, 257:556 – 566, 1987.
- [21] Domenico Pizzuto. *D0 Central Tracking Chamber Performance Studies*. PhD thesis, State University of New York at Stony Brook, Stony Brook, New York, December 1991.
- [22] Jeffrey W. Bantly. *The D0 Detector Forward Drift Chamber Performance and Physics Capability in the 1990 FNAL Testbeam Run*. PhD thesis, Northwestern University, Evanston, Illinois, June 1992.
- [23] Robert E. Avery et al. Performance of the forward drift chambers for the DØ detector. *IEEE Transactions on Nuclear Science*, 40(4):573 – 577, 1993.
- [24] J. F. Detœuf et al. The D0 transition radiation detector. *Nuclear Instrumentation and Methods in Physics Research A*, 265:157 – 166, 1988.
- [25] J. F. Detœuf et al. Status of the transition radiation detector for the DØ experiment. *Nuclear Instrumentation and Methods in Physics Research A*, 279:310 – 316, 1989.
- [26] Y. Ducros et al. Monte carlo simulation of the performance of the D0 transition radiation detector. *Nuclear Instrumentation and Methods in Physics Research A*, 277:401 – 406, 1989.
- [27] R. G. Angstadt et al. A working, VME-based, 106 MHz FADC data acquisition system for the tracking detectors at D0. *IEEE Transactions on Nuclear Science*, 39(5):1297 – 1301, 1992.
- [28] C. Fabjan. Calorimetry in high-energy physics. In Thomas Ferbel, editor, *Experimental Techniques in High Energy Physics*, pages 257 – 324. Addison-Wesley, 1987.
- [29] Richard C. Fernow. *Introduction to Experimental Particle Physics*, chapter 11, pages 259 – 284. Cambridge University Press, 1986.

- [30] Particle Data Group. Review of particle properties. *Physical Review D*, 50(3):1173 – 1876, August 1994.
- [31] Richard Wigmans. On the energy resolution of uranium and other hadron calorimeters. *Nuclear Instrumentation and Methods in Physics Research A*, 259:389 – 429, 1987.
- [32] Jaehoon Yu. *Determination of the Strong Coupling Constant (α_s) and a Test of Perturbative QCD Using $W + \text{jets}$ Processes in the $D\bar{O}$ detector*. PhD thesis, State University of New York at Stony Brook, Stony Brook, New York, August 1993.
- [33] Anthony Spadafora, July 1992.
- [34] Stephen J. Wimpenny et al. The hadron and electron response of uranium/liquid argon calorimeter modules for the d-zero detector. *Nuclear Instrumentation and Methods in Physics Research A*, 279:107 – 113, 1989.
- [35] Anthony L. Spadafora et al. Test beam results from the D0 liquid argon end calorimeter electromagnetic module. *Nuclear Instrumentation and Methods in Physics Research A*, 315:279 – 284, 1992.
- [36] Paolo Franzini et al. Performance of the $D\bar{O}$ uranium-liquid argon calorimeter modules. *Nuclear Instrumentation and Methods in Physics Research A*, 289:438 – 445, 1990.
- [37] Hiroaki Aihara et al. Performance of the D0 end calorimeter electromagnetic module. *IEEE Transactions on Nuclear Science*, 38(2):398 – 402, 1991.
- [38] C. Brown et al. $D\bar{O}$ muon system with proportional drift tube chambers. *Nuclear Instrumentation and Methods in Physics Research A*, 279:331 – 338, 1989.
- [39] J. M. Butler et al. Study of fast gases, resolutions and contaminants in the $D\bar{O}$ muon system. *Nuclear Instrumentation and Methods in Physics Research A*, 290:122 – 130, 1990.
- [40] Yu. M. Antipov et al. Drift tubes for the SAMUS muon spectrometer of the D0 detector at FNAL. *Nuclear Instrumentation and Methods in Physics Research A*, 297:121 – 125, 1990.

- [41] Maris Abolins et al. The level one framework: D0 note 328 revised. DØ note 705, 1988.
- [42] Maris Abolins et al. A high luminosity trigger design for the Tevatron collider experiment in D0. *IEEE Transactions on Nuclear Science*, 36(1):384 – 389, 1989.
- [43] Maris Abolins, Daniel Edmunds, Philippe Laurens, and Bo Pi. The fast trigger for the DØ experiment. *Nuclear Instrumentation and Methods in Physics Research A*, 289:543 – 560, 1990.
- [44] G. S. Gao and R. Partridge. High speed digital TDC for DØ vertex reconstruction. *IEEE Transactions on Nuclear Science*, 38(2):286 – 289, 1991.
- [45] J. Bantly et al. The level 0 trigger for the DØ detector. DØ note 1996, 1993. Submitted to *IEEE Transactions on Nuclear Science*.
- [46] John M. Butler. Main ring deadtime. DØ note 1682, 1993.
- [47] Daniel Edmunds. Changes to the first level trigger timing. DØ note 827, 1989.
- [48] M. Fortner et al. The VME-based D0 muon trigger electronics. *IEEE Transactions on Nuclear Science*, 38(2):480 – 485, 1991.
- [49] D. Cutts et al. A microprocessor farm architecture for high speed data acquisition and analysis. *IEEE Transactions on Nuclear Science*, 36(1):738 – 739, 1989.
- [50] David Cullen-Vidal et al. DØ level-2/data acquisition; the new generation. In *Proceedings of the Conference on Computing in High-Energy Physics*, Tsukuba, Japan, March 1991.
- [51] Dave Cutts et al. Operation of the D0 data acquisition system. In *Proceedings of the Conference on Computing in High-Energy Physics*, pages 262 – 264, Annecy, France, 1992. CERN 92-01.
- [52] M. Goossens et al. *Zebra*, 1992. CERN Program Library Number Q100.

- [53] James T. Linnemann et al. The DØ software trigger. In *Proceedings of the Conference on Computing in High-Energy Physics*, pages 199 – 201, Annecy, France, 1992. CERN 92-01.
- [54] John Featherly. *ITC User's Guide*, 1991. (unpublished).
(URL: <http://d0hs19.fnal.gov/rndocs/itc.html>)
- [55] Bruce Gibbard et al. Run control and resource management in the D0 run time system. In *Proceedings of the Conference on Computing in High-Energy Physics*, pages 265 – 268, Annecy, France, 1992. CERN 92-01.
- [56] Bruce Gibbard and Scott Snyder. “*Taker*” *A Generic Human Interface for D-Zero Data Taking*, 1993. (unpublished).
(URL: <http://d0hs19.fnal.gov/rndocs/taker.html>)
- [57] Bruce Gibbard. *Working with COOR's Configuration Files*, 1993. (unpublished).
(URL: http://d0hs19.fnal.gov/rndocs/coor_configurations.html)
- [58] Scott Snyder. *Trigparse*, 1994. (unpublished; available at `d0$configs$coor_sim:trigparse.doc`).
- [59] Scott Snyder. *The DØ Data Logger*, 1993. (unpublished).
(URL: http://d0hs19.fnal.gov/daqdocs/html/logger_toc.html)
- [60] David Quarrie. Data acquisition system programmers reference manual. Computing Division manual PN251.04, FNAL, 1990.
- [61] Adam Para et al. The D0 offline production and analysis. In *Proceedings of the Conference on Computing in High-Energy Physics*, pages 593 – 596, Annecy, France, 1992. CERN 92-01.

