The DØ Level 1 trigger system

Daniel Edmunds *, Philippe Laurens

Department of Physics and Astronomy, Michigan State University, East Lansing, MI 48824, USA

For the DØ Collaboration

A R T I C L E   I N F O

A B S T R A C T

This paper describes the principal components of the Level 1 trigger system for the DØ experiment during Run IIb of the Tevatron collider at Fermilab. This trigger system was designed to take advantage of the increased luminosity and shorter bunch spacing provided by the Tevatron in Run II. We place an emphasis on the key concepts that are potentially valuable for future experiments.

© 2008 Elsevier B.V. All rights reserved.

1. Introduction

The trigger system for the DØ experiment [1] includes three levels of decision making. Each trigger level rejects some events and passes on the remaining events to the following level. This layered system is illustrated in Fig. 1.

The first-level trigger (Level 1 or L1), which is described in detail in this article, examines every beam crossing with minimum dead time. It is based on custom hardware that is closely coupled to each detector subsystem. A component of the L1 trigger called the trigger framework (TFW) makes the final L1 trigger decision based on the information that it receives from the various L1 trigger subsystems. The thresholds and trigger requirements used in L1 reduce the input beam crossing rate of 1.7 MHz down to about 2 kHz of events that are sent to the second-level trigger (Level 2 or L2).

The L2 trigger uses a combination of custom hardware and commercial microprocessor boards. It consists of five L2 pre-processors, each associated with a detector subsystem, and a global L2 processor that forms the final L2 trigger decision based on reconstructed objects and combinations of these objects. The L2 trigger reduces the event rate by another factor of two down to about 1 kHz.

The third-level trigger (Level 3 or L3) is a farm of commercial computers. L3 makes the final event selection using sophisticated analysis on complete event data. L3 reduces to 50–100 Hz the rate of events that are recorded for offline analysis.

Fig. 2 describes the data flow and buffering in the DØ trigger DAQ system. The movement and storage of event data is controlled by the decisions from the three trigger levels. In the first step, the detector data from each beam crossing are simultaneously recorded in L1 pipeline buffers and processed by the detector-specific L1 trigger subsystems. The data from events that are given an L1 Accept by the TFW are preserved in L2 FIFO buffers and processed by detector-specific L2 pre-processors. Events accepted by L2 are read out by the DAQ system and processed by the L3 farm. All processing of a given event is done by one node in the L3 farm. While L1 and L2 decisions are always issued in chronological order, the L3 decisions order depends on the processing time for each event.

All L1, L2, and L3 components of the trigger and DAQ systems are programmable under the central control of COOR, the coordination program for DØ online operation. COOR is used to download the trigger list requirements for physics data collection runs and for commissioning and diagnostics runs.

2. L1 trigger architecture

The L1 trigger system is made of detector-specific L1 trigger subsystems and the trigger framework. The L1 trigger subsystems process their detector data to form a set of binary indicators that describe each beam crossing (e.g. whether a measured quantity is greater than a programmed threshold). The TFW combines these individual bits of information, called trigger terms, which characterize a given beam crossing to form the L1 trigger decision for that beam crossing.

2.1. L1 trigger subsystems

Fig. 3 shows the different detector components with their associated L1 trigger subsystem: the Calorimeter (CAL) with

---

* Corresponding author. Tel.: +1 517 884 5521; fax: +1 517 355 6661.
  E-mail address: edmunds@pa.msu.edu (D. Edmunds).
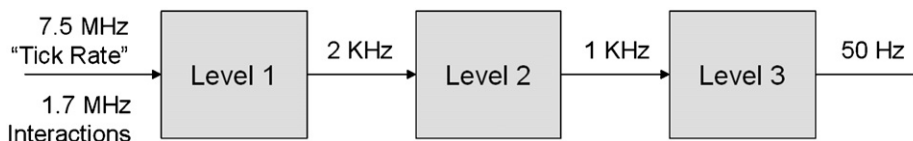
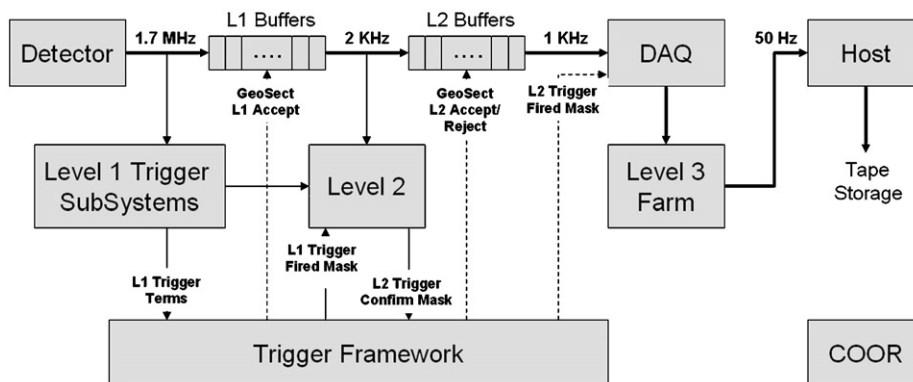**Fig. 1.** Three levels of triggering are used at DØ.

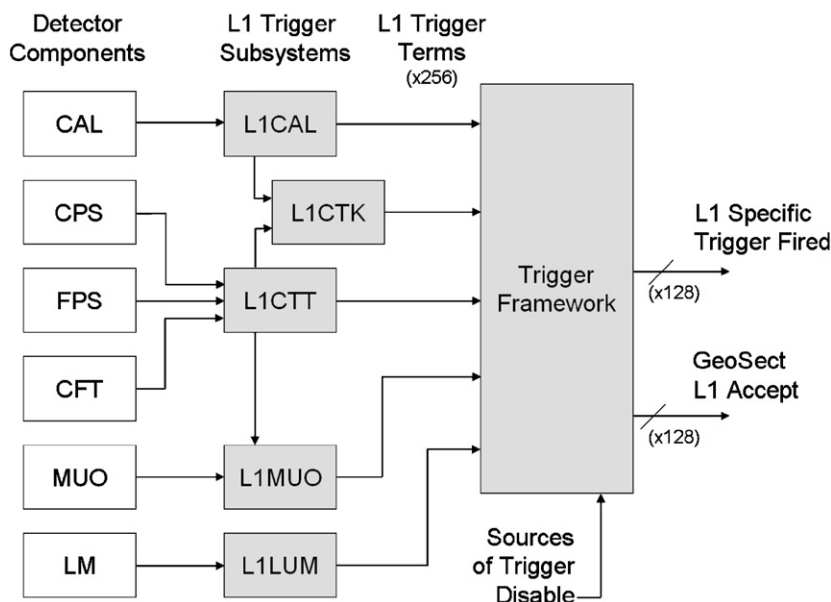**Fig. 2.** Event buffering and data flow in the data acquisition system.

Fig. 3. Level 1 trigger components.

the L1 Calorimeter Trigger (L1CAL), the Central and Forward Preshower Detectors (CPD and FPD) and the Central Fiber Tracker (CFT) with the L1 Central Track Trigger (L1CTT), the Muon Detectors (MUO) with the L1 Muon Trigger (L1MUO), and the Luminosity Detector (LM) with the L1 Luminosity Monitor (L1LUM). Information from L1CAL and L1CTT are combined in a L1 Calorimeter Tracker Match Trigger (L1CTK). Each of these L1 trigger subsystems sends a set of L1 trigger terms to the TFW.

2.1.1. L1 Calorimeter Trigger

L1CAL receives analog signals from the CAL front-end electronics cards. The signals from the electromagnetic and fine hadronic sections of 1280 trigger towers are received, digitized, and scaled to their transverse energy component. FPGAs use serial computation to run cluster algorithms on these transverse energies to build electron, jet, and tau objects, as well as global energy sums, and a missing energy vector. L1CAL applies thresholds on the total count and energy of such objects to form L1 trigger terms [2,3].

2.1.2. L1 Central Track Trigger

L1CTT is based on information from the approximately 100,000 channels in the CFT, CPD and FPD. Tracks are searched for in the information from these detectors using combinatorial logic in FPGAs. The L1 trigger terms generated by this logic indicate the presence, number, location, and isolation of tracks above various energy thresholds [4,5].

2.1.3. L1 Muon Trigger

L1MUO looks for patterns consistent with muons using hits from the muon wire chambers, muon scintillation counters, and

tracks from L1CTT. Searching for these patterns is performed by combinatorial logic, implemented in FPGAs, on the information from approximately 60,000 channels of MUO and the L1CTT information. The L1 trigger terms generated by L1MUO indicate the presence, multiplicity, energy threshold, sign, and eta region of the various candidate muons [6–8].

2.1.4. Luminosity monitor

The luminosity monitor uses scintillation detectors placed around the beam pipe on both sides of the interaction point. Custom electronics measures the presence and relative timing of signals coming from both sets of detectors. Besides providing a measurement of the instantaneous luminosity delivered by the Tevatron, the luminosity monitor system also provides beam quality trigger terms to the trigger framework, and thus functions as a L1 trigger subsystem [9,10].

2.1.5. L1 Calorimeter Tracker Match Trigger

L1CTK is based on information that comes from the L1CAL and L1CTT trigger subsystems. This information is combined to produce L1 trigger terms that indicate a spatial match between various objects that were identified by L1CAL and tracks that were identified by L1CTT. The L1CTK trigger terms further indicate the type, multiplicity, and energy threshold of the object and track that were matched, as well as the eta region in which they were found, and an indication of their isolation [11].

2.2. Trigger framework

The TFW does not form a single L1 trigger decision; rather it forms 128 individual decisions called L1 specific triggers. A L1 trigger fired is issued if any of the 128 specific triggers has fired. Likewise the DAQ system is not controlled in a monolithic way; rather it is divided into 128 separate readout sections that are called geographic sections. Geographic sections include front-end crates, readout crates, as well as L1 and L2 trigger crates. The readout function of the TFW itself is also a geographic section within the DAQ system.

Each specific trigger is programmed to cause the event data readout from a list of geographic sections. Each L1 trigger fired thus causes L1 accepts to be sent to a subset of the 128 geographic sections. Each geographic section, when busy or unable to process an additional L1 accept, can disable all of the specific triggers that would cause an L1 accept to be sent to that geographic section. The trigger DAQ system can thus be partitioned into 128 triggers for different physics signatures, and 128 sections to manage the DAQ readout. This makes it possible for detector groups to work independently.

Of the 256 trigger terms [13] received by the TFW, most come from the L1 trigger subsystems while others come from the master clock (e.g. to identify beam gaps), from other systems (e.g. CAL pulser), and the TFW itself generates a few trigger terms (e.g. to skip crossings between successive L1 accepts). Different subsystems take different lengths of time to generate their trigger terms, but all terms must arrive at the TFW by 3.7 μs after the beam crossing that generated them. The TFW receives the trigger terms in FIFOs that automatically align themselves in depth so that all trigger terms from a given beam crossing are available to the TFW decision logic when it begins to make the L1 trigger decision for that crossing.

The trigger terms are used to generate separate "and–or fired" decisions for all 128 specific triggers (Fig. 4). Each specific trigger can be programmed by COOR to require a given trigger term to be asserted, negated (i.e. a veto term), or ignored. A positive and–or fired decision for a given specific trigger identifies the crossing as interesting for further analysis.

Additional conditions must be met before the and–or fired is allowed to become a specific trigger fired. COOR must have enabled the specific trigger, and its firing cannot be blocked by its prescaler. The TFW must verify that all the geographic sections that have been programmed to receive an L1 accept from this specific trigger are currently ready, and that the L3 system has not disabled the firing of this specific trigger to control its input rate. Finally, a second combination of trigger terms requirements, focusing on the beam quality trigger terms, is formed separately and used as an additional source of disable for luminosity accounting purposes (cf. Section 6). When none of these sources of disable is preventing a specific trigger from firing, the specific trigger is said to be exposed to the beam.

Once the TFW has determined that any of the 128 specific triggers has met its triggering and–or requirements, and verified that no source of disable is currently preventing it from firing, a global L1 trigger fired is issued with a 128-bit mask listing all L1 specific triggers that fired for this beam crossing. This 128-bit mask is translated by a mapping table (programmed to associate each specific trigger with a subset of the geographic sections) into a 128-bit mask of geographic sections that will be sent a L1 accept. This mapping table is shown on the right side of Fig. 4. A similar but reverse mapping table is used to translate the 128 L1 busy signals from the geographic sections into the 128 L1 busy disable signals for the individual L1 specific triggers.

The TFW includes many scalers to measure in detail the operation of the DØ trigger DAQ system. Scalers are provided for all trigger resources including all trigger terms, all specific triggers, all sources of disable, all geographic sections, as well as those required to monitor the overall performance of the system (cf. Section 5).

The TFW reads out and is handled by the DAQ system in the same way as any other geographic section. The information readout includes the state of all the TFW inputs and outputs for the beam crossing that caused the L1 trigger and for one
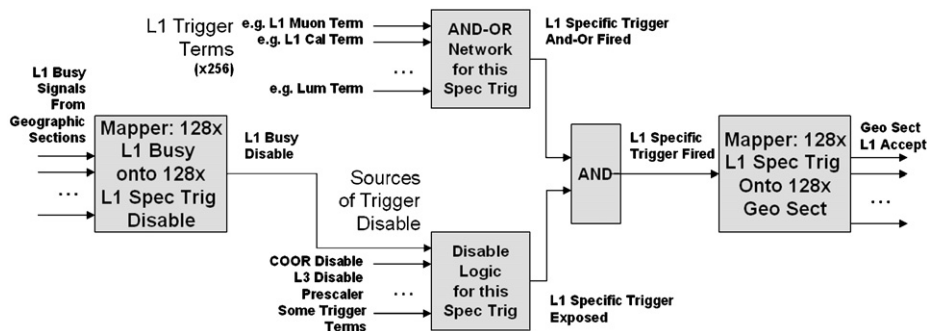

Fig. 4. Trigger framework decision logic and lookup tables.

preceding and one subsequent beam crossing. This information makes it possible to verify the proper operation of the TFW. Data from the adjacent beam crossings provide a record, at the spatial and energy resolution of the L1 trigger subsystems, of the activity in the DØ detector during the time surrounding the recorded event. The TFW data are also used during L3 filtering to tag each event record with its L1 event number, its Luminosity Block Number (LBN) (cf. Section 6), and to retrieve the mask of L1 specific triggers fired that directs the L3 filter scripts.

The TFW is implemented in 9U 400 mm cards and customized VME crates. One core design with 128 input and 64 output signals, 16 main FPGAs, one FPGA for the VME interface, and one FPGA for global board support functions is used for all cards in the TFW. Only five card sub-types, based on this core design, are needed to implement all of the TFW's functions. These sub-types differ mainly in the routing of inputs and output signals to, from, and between the 16 main FPGAs on each card [12].

3. Serial command link

The serial command link (SCL) [14,15] is the common hardware path between the TFW and all of the geographic sections. It carries all of the control and timing information needed to operate a crate in the DØ Run II trigger DAQ system. Each link consists of one serial cable carrying information at 1 Gbit/s from the TFW to the geographic section, and one optional 8-bit-wide parallel cable that carries status information from the geographic section back to the TFW.

Each turn of the beam in the Tevatron is divided into 159 ticks of 132 ns (cf. the 7.5 MHz tick rate in Fig. 1). In Run II the Tevatron provides 36 beam crossings per turn, arranged in 3 groups of 12 crossings. Within a group the crossings occur for every 3rd tick. For each tick the SCL delivers 75 bits of information to each geographic section and returns 8 bits of status information from each geographic section.

The timing signals carried by the SCL include: a 53 MHz clock, a 7.59 MHz (132 ns) tick clock, a 16-bit turn number and an 8-bit tick number to identify the current crossing in the Tevatron. These timing signals, and all other signals on the SCL, are locked to the Tevatron RF system by the DØ master clock. The SCL is the only system of timing signal distribution that is used within DØ.

The SCL carries all of the trigger information to the geographic sections. For each tick, and separately for each geographic section, the SCL specifies if an L1 accept has been issued. L1 accepts are issued with a fixed delay from the beam crossing that caused them. When an L1 accept is issued, the beam crossing that it belongs to is specified by its tick and turn number.

Every L1 accept issued to a geographic section is followed by a matching L2 accept or L2 reject command. This L2 trigger decision information is specific to each geographic section. Tick and turn numbers are used to specify the event that a given L2 decision belongs to. The L2 latency varies from one event to the next but the L2 decisions are always issued in the same order as the L1 accepts.

In addition to timing and trigger information, the SCL also carries an initialization signal that is used to synchronize the data buffers throughout all of the DØ trigger DAQ system.

The parallel status cable carries information from the geographic section back to the TFW. This information includes: the L1 busy signal which indicates the geographic section's readiness to respond to a L1 accept, the L2 busy signal which indicates its readiness to respond to a L2 accept, L1 and L2 error signals sent upon detection of an error in the processing of SCL requests, and a signal that indicates that the geographic section has responded to an initialization request.

The SCL is controlled from a central crate called the SCL hub end that is located in the TFW. All of the geographic sections use a common SCL receiver mezzanine card which presents all of the signals carried on the SCL in a parallel format. The SCL serves geographic sections that are located in the detector platform, in the main DAQ equipment areas, and in remote test stations.

4. Control of the TFW

COOR [16] is the centralized coordination program used in DØ to download and control all of the components in the online system. COOR is controlled by a set of configuration files describing the details of all resources in the online system. The DØ physics trigger list is managed via a trigger database which produces the programming requests given to COOR. COOR then communicates with all the online systems using the necessary paths and protocols to carry out the user requests.

The TFW is a complex system that was designed for flexibility. It is programmed via a large number of control registers that are accessed through a VME interface on each card. To implement a clean interface to the TFW hardware the L1 trigger control computer (L1 TCC) was placed between COOR and the complexity of this hardware. L1 TCC presents a high-level interface to COOR for control of the TFW resources. COOR sends human readable text commands describing how the logical TFW resources need to be programmed to match the trigger list requirements. For example, COOR specifies a prescaler ratio for a given specific trigger and L1 TCC computes the required data values to implement this request and writes these data into the appropriate registers on the card that controls this TFW function. L1 TCC always replies to each request from COOR with an acknowledgement message, either "done" for the normal completion of a request, or "error" in the case of a failure.

The trigger control software (TRICS) [17] that runs on L1 TCC is the interface to all of the functions and resources in the TFW. As shown in Fig. 5, TRICS connects with a number of client programs, including COOR. In addition to its control functions TRICS also collects and serves monitoring information (cf. Section 5), collects and serves luminosity information (cf. Section 6), and presents a graphical user interface (GUI) for special control or tests.

Programming the TFW to the point where DØ is ready to record physics data takes place in three separate phases. The first phase is called configuration and only needs to be done once after power-up. During the configuration phase, TRICS loads firmware into the 1500 FPGAs in the TFW. After configuration, the TFW hardware possesses all its digital logic, but in an undefined state. The second phase is the initialization of the system, which needs to happen at least once after configuration. During initialization, all registers in the system are programmed and the system is brought to a fully defined, default, idle state. COOR may request an initialization any time that it wants to bring the TFW back to this known state, free of any history dependency. This is typically done before each store in the Tevatron. Because the initialization phase programs and then reads back every register in the TFW, it also provides a quick sanity check of the system. The third phase is the actual run programming, when COOR defines the trigger requirements, and ultimately enables a subset of specific triggers to start a new run.

TRICS is written in C# and uses object-oriented programming to manage both the low-level view of the TFW hardware and the high-level view of its logical resources. Before modifying the contents of any control register in the framework TRICS first reads the register and compares the data that it finds with its knowledge of what had previously been written in that register. After writing a new value in a register TRICS immediately reads
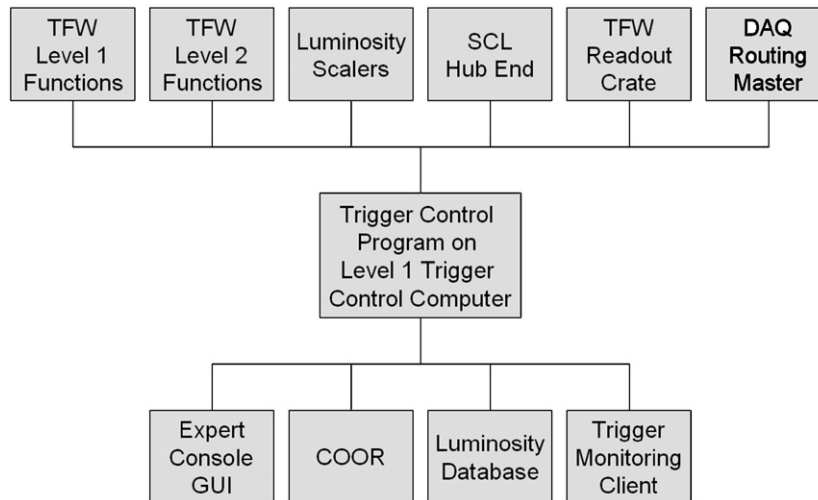
Fig. 5. L1 TCC trigger control software.

back the register to verify that the appropriate data have been retained. In this way TRICS continuously verifies the health of the TFW hardware. All communications with, and actions taken by, TRICS are displayed on a screen and recorded in a detailed log file with time stamps.

5. Monitoring of the TFW

By design the TFW is the central point for monitoring the operation and performance of the whole DØ trigger DAQ system. The scalers mentioned in Section 2.2 provide the bulk of the raw monitoring information. This information is acquired as a series of snapshots and a new snapshot is typically collected every 5 s. These monitoring data snapshots allow computation and display of quantities such as the L1 trigger term rates, the rates of L1 specific triggers firing, and the percentages of beam crossings for which a given geographic section is L1 or L2 busy. This information allows both the monitoring of the physics triggers and the monitoring and diagnosis of problems in the trigger DAQ system. Monitoring clients provide text and graphics displays of these data both in the control room and on the web [18,19].

To be most useful the monitoring data snapshots must capture all values simultaneously throughout the whole system and this information should preferably be acquired from beam crossings that result in a specific trigger firing. A combination of hardware and software provides this functionality.

In this system L1 TCC arms a central monitoring control function so that it will cause all of the distributed monitoring resources to simultaneously capture and locally store their monitoring data the next time that a specific trigger fires. L1 TCC reads out all of the data in this monitoring snapshot, and then re-arms the central monitoring control function when it is time to collect the next monitoring snapshot. A provision is made so that L1 TCC can immediately force the capture and readout of monitoring data when there is a problem and none of the L1 specific triggers have fired within a 5 s period. This provision is necessary to provide the monitoring information needed to diagnose a problem that has blocked the normal firing of specific triggers and flow of event data.

L1 TCC serves blocks of monitoring data to monitoring clients that use successive snapshots to calculate quantities such as current counts and rates and percentages. The time base for these calculations comes from information within each monitoring data snapshot that shows the exact number of ticks since the previous snapshot. Monitoring clients at DØ are thus able to make the distinction between measurements rounding up to 100% and conditions present during every tick of a time interval. This can be very useful, e.g. when diagnosing certain symptoms of crates remaining L1 busy.

6. Online luminosity measurement

The instantaneous and integrated delivered luminosity is computed through the number of inelastic collisions. Poisson statistics on the count of beam crossings with no collision are used to calculate the luminosity. Because this calculation is highly non-linear it must be carried out separately for each of the 36 bunch crossings at DØ. For the current beam structure this requires 36 per-bunch scalers for each luminosity indicator signal. To allow for changes in the Tevatron beam structure the TFW design had to provide 159 per-bunch scalers for each luminosity indicator signal. Up to 10 luminosity indicator signals are handled by the TFW.

Luminosity measurements must also be carried out with good time resolution to follow the time dependency of the Tevatron luminosity. To do this the recorded beam time is divided into luminosity blocks that are controlled by L1 TCC. A new luminosity block is always started at every run boundary (start, stop, pause, resume), and has a maximum length of 60 s. All recorded events are tagged with the never-repeating 32-bit LBN.

In addition to the delivered beam luminosity, the experiment also needs to determine the actual recorded luminosity that each L1 specific trigger has been exposed to. Maintaining a set of 159 per-bunch scalers for each of the 128 L1 specific triggers was deemed prohibitive. Such a brute force approach was replaced with a compromise that requires a much smaller number of sets of per-bunch scalers while still providing flexibility in the programming of the specific triggers independently of each other.

The sources of specific trigger disable come in two classes; some are potentially correlated with the bunch luminosity profile (e.g. L1 busy and the beam quality trigger terms), while others are intrinsically non-correlated (e.g. COOR or L3 disable). In this approach prescalers had to be implemented as non-correlated to the bunch luminosity. This was accomplished using a duty ratio design based on a pseudo-random number generator. COOR assigns each of the 128 specific triggers to one of 8 exposure groups [20] that are defined by their common sources of correlated disable. Each exposure group is fully instrumented

with a set of per-bunch scalers, and each specific trigger includes scalers for measuring its individual non-correlated disable fraction. An accurate measurement of recorded luminosity is thus provided for each of the 128 L1 specific triggers.

7. Conclusion

The DØ Run II detector and its trigger system have been in operation since 2001. It is expected that the Tevatron will continue delivering collisions through 2009 or 2010. The L1 trigger system has been serving DØ well at peak Tevatron luminosities exceeding $3 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and over $3\,\mathrm{fb}^{-1}$ of integrated luminosity have been recorded so far in Run II.

Acknowledgments

References

[1] V.M. Abazov, et al., Nucl. Instr. and Meth. A 565 (2006) 463.
[2] M. Abolins, et al., Nucl. Instr. and Meth. A 584/1 (2007) 75.
[3] http://www.nevis.columbia.edu/~evans/l1cal/.
[4] J. Olsen, IEEE Trans. Nucl. Sci. NS-51 (2004) 345.
[5] http://www-d0online.fnal.gov/www/groups/cft/CTT/online/ctt_main.html.
[6] V.M. Abazov, et al., Nucl. Instr. and Meth. A 552 (3) (2005) 372.
[7] http://www-d0online.fnal.gov/www/groups/calmuo/.
[8] http://atlas.physics.arizona.edu/~johns/l1muo/l1mu/l1mu.htm.
[9] John Anderson, et al., FERMILAB-CONF, 06-470-E, 2006.
[10] Chyi-Chang Miao, Nucl. Phys. Proc. Suppl. 78 (1999) 342.
[11] http://atlas.physics.arizona.edu/~johns/l1muo/l1caltrack/l1caltrack-top.htm.
[12] http://www.pa.msu.edu/hep/d0/l1/framework.
[13] http://www.pa.msu.edu/hep/d0/l1/framework/andor_terms.html.
[14] http://www.pa.msu.edu/hep/d0/l1/scl.html.
[15] http://www-ese.fnal.gov/d0trig.
[16] http://www-d0.fnal.gov/d0dist/dist/packages/coor/devel/doc/coorover.ps.
[17] http://www.pa.msu.edu/hep/d0/l1/tcc.
[18] http://www.pa.msu.edu/hep/d0/l1/tcc/trigmon_snapshot_frame.html.
[19] http://www-d0online.fnal.gov/www/groups/tm/daqrates/current_store.html.
[20] http://www.pa.msu.edu/hep/d0/ftp/l1/framework/exposure_groups.