16 Dec 1998

STAR Trigger Functional Requirements
Table of Contents

STAR Trigger Functional Requirements

1. Introduction

The heart of the STAR detector system is the TPC. This detector is designed to return data on dE/dx and particle trajectories in the pseudo-rapidity region of $-2<\eta<2$, from which we hope to glean an unambiguous signal of the formation of a quark-gluon plasma (QGP). The signature of the QGP is unknown, however, making it difficult to design a trigger to select collisions in which a plasma has formed. Therefore we have designed a trigger that is flexible, that can learn the normal distributions of particles within events, and that can be programmed to select events that are outside the normal distributions.

Event selection is based on fast detectors, which measure multiplicity and energy distributions in pseudo-rapidity η and azimuth angle ϕ. The multiple level structure of the trigger system is based on TPC timing considerations, and consists of a Level 0 selection followed by up to 4 additional levels of analysis and rejection. We expect to study the approach to the QGP through pp, pA and AA collisions.

The STAR baseline system consists of the TPC and the trigger detectors. Additional detectors include a silicon vertex tracker (SVT), an Forward TPC (FTPC), an electromagnetic calorimeter (EMC), a time-of-flight (TOF) system, Ring Imaging Cherenkov (RICH), and additional calorimetry at large pseudo-rapidity (η) regions. Our trigger must be able to treat each of the subsystems independently to allow us to extract information from detectors independent of each other as well as in coupled mode. This results from the fact that the TPC and SVT are slow devices and it should be possible to learn much about the distributions from the faster detectors to help in selecting events for the finer analysis afforded by the slow ones. Note that all of the detectors participate in the final trigger decisions concerning further analysis of a given event, but only fast detectors participate in the original event selection at Level 0.

The goals of the trigger system can be summarized as follows:

1. Select high multiplicity events since these typically have the largest number of interacting nucleons. The assumption is that new phenomena will result from the collective interaction of the largest number of participants.
2. Select events whose particle distributions suggest local inhomogeneity such as expected when a portion of the interaction volume encloses a plasma.
3. Select events having unusual concentrations of strangeness, since strangeness and anti-strangeness production levels may herald passage through the QGP phase.
4. Provide for a variety of different trigger types, such as for calibration events, cosmic ray events, high multiplicity, asymmetry, and others.
5. Be easily adapted to include new detectors
6. Be easily adapted to include new understanding of the physics
7. Reject beam-gas interactions
8. Select events in pp and pA for analysis, much as in AA

The constraints on the system include:

1. Provide a signal to open the TPC amplification grid within 1.5 micro-sec of the occurrence of the interaction

2. Select events based on particles produced in $-2<\eta<2$ interval
3. Allow understanding of any trigger bias introduced in event selection

1.1 Trigger levels

The timing structure of the trigger system is set to match the TPC time structure. This leads to 4 levels of trigger set by the conditions indicated.

Level 0: We need to open the TPC amplification grid within approximately 1.5μs of the occurrence of the collision. This is a soft number, set by the fact that the process of gating the grid requires ~0.5μs and by the fact that each μs of delay costs us ~2% (1.2% w/ HeC_2H_6) of the data on track length. All event selection must begin with a Level 0 trigger. Subsequent levels of trigger analysis must make a positive selection to allow an event to continue through the trigger selection chain.

Level l: The drift time of the TPC is ~50μs (80μs with HeC_2H_6), during which charge samples from each pad are strobed onto a switched capacitor array (SCA) which forms an analog pipeline for the TPC data. Bandwidth considerations for moving trigger information through the system suggest that summary data from sets of trigger detectors can be further analyzed during this storage period to decide whether a selected event should be rejected or passed to the next level of analysis. This time allocation is set somewhat arbitrarily at 100μs for Level 1 analysis. If this event is rejected, it takes a few μs for the TPC electronics to reset itself to be ready to accept another event.

Level 2: Once the SCA is filled, data begin moving through the TPC digitizer sequence, a process that requires 5 ms. regardless of occupancy in the TPC. During this time, network bandwidth allows all data from the trigger detectors to be moved to a central location for further analysis. Simultaneously, additional summary data from other fast detectors (EMC, TOF) can be added to this trigger data to determine whether the event should be kept or rejected. Time allocated for Level 2 analysis is again somewhat arbitrary. For events including TPC data, we want to make our decision to keep or reject an event in a time significantly less than the 5 ms. digitizing time so that we can increase the overall TPC event throughput. For events that include only fast detectors, we want to make the keep/reject decision in a time commensurate with the event build time or less. We thus allow duration for Level 2 processing to depend on event type - typically 4ms for a TPC type event, or 1ms for an EMC type event.

Level 3: At this point, TPC data is digitized and ready to read into the event builder. The hardware for this process is under DAQ control, with aid from the trigger group in defining the hardware architecture to optimize TPC and slow detector data analysis in further trigger keep/reject decisions. Level 3 analysis will be performed first on data from the subsystems separately such as track information from TPC and SVT systems. It will then be combined with summary information from fast detectors to give a quick look at the overall event. We expect it to be possible to move many events through the Level 3 processing to select the few events per second that are sent in raw form to hard media for archiving.

Level 4: We need a way to look at complete events before taping.

1.2 Trigger hardware

The STAR trigger system is divided into two parts. First are the trigger detectors; the Central Trigger Barrel (CTB), the Multi Wire Counter (MWC) and the Zero-degree Calorimeters (ZDC). The electromagnetic calorimeter (EMC) is also used throughout the trigger decision tree. Second are the electronics for data processing and triggering functions and the interfaces to the detector subsystems and DAQ.

1.2.1 Detectors:

CTB: There are 120 trays of the central trigger barrel, each containing two scintillator slats with each slat viewed by a single photomultiplier. The trays are arrayed around the TPC in a barrel of radius ~2m and length ~4m. Each slat delivers light according to dE/dx in the scintillator. The overall track multiplicity in each slat is proportional to the total light generated. Consequently, an ADC on each PMT gives an indication of total charged particle multiplicity in the region from $-1<\eta<1$ in the CTB.
MWC: Each TPC endcap has ~4000 wires which respond like a normal multiwire chamber to passage of a charged particle. We take the prompt signal from each wire through a preamp and discriminator to determine whether each wire was hit. The wires are then taken in groups of 20 and a count of hit wires in each group is made and sent to the Level 0 electronics to give an indication of multiplicity passing through the TPC endcaps, thus in the region from $-2<\eta<-1$ and $1<\eta<2$.
ZDC: at the cave exit on each side, in the crotch of the first beamline dipoles, there are two simple Pb-fiber calorimeters whose purpose is to record passage of fragmentation region neutrons. These provide verification of centrality for AA collisions as well as a minimum bias signal.
EMC: This is not technically termed a trigger detector although the information provided is used in all levels of the STAR trigger.

1.2.2 Electronics:

Electronics: Data from every trigger detector is digitized each RHIC crossing and sent to a Data Storage and Manipulation (DSM) board, which forms a portion of a digital pipeline. A trigger control unit (TCU) controls fast event selection at Level 0, where information is available from a tree of field programmable gate arrays (FPGA) which perform the manipulations on the DSM boards. Summary information from fast detectors (Trigger, Time-of-flight (TOF), Electromagnetic Calorimeter (EMC)) is passed to Level 1 for initial topological analysis and possible event rejection. For events passing Level 1, full event information for fast detectors is passed to Level 2 for detailed topological analysis and possible rejection. DAQ is notified of any event passing Level 2 analysis. DAQ gathers slow detector information and begins track analysis for Level 3 trigger decisions based on separate detector segments. Finally, at Level 4 (called L3 global), all data for a given interaction is built into an event and analyzed for suitability for archiving.

1.3 Glossary of acronyms

CTB	Central Trigger Barrel
DAQ	Data Acquisition
DSM	Data Storage and Manipulation boards
EMC	Electromagnetic Calorimeter
FIFO	First-in-First-out storage memory
FPGA	Field programmable gate array
L0	Level 0 trigger
L1	Level 1 trigger (1-100 μs)
L1CPU	L1 processor
L1CTL	L1 control processor
L2	Level 2 trigger (0.1 - 5 ms)
L2CPU	L2 processor
L2CTL	L2 control processor
L3	Level 3 trigger
L4	level 4 trigger
LUT	Lookup Table
M	Particle multiplicity
mip	minimum ionizing particle
MWC	Multi-wire Chamber
PID	Particle Identification
PMT	Photomultiplier tube
RHIC	Relativistic Heavy Ion Collider
SVT	Silicon Vertex Tracker
TCD	Trigger - Clock Distribution
TCU	trigger control unit (L0)
TDI	Trigger-DAQ Interface
TOF	Time-of-Flight detector
TPC	Time Projection Chamber
TRIG	Trigger system
TSL	Trigger System Logic
ZDC	Zero-degree calorimeter
Z	Position align beamline

This document is arranged to have three sub-headings for each requirement: Requirement, Justification, and Status.

2. Executive Summary

The tables below provide an overview of the parameters for the trigger system as it is being implemented. It is our intention to produce 10% spares for all home-grown modules.

Tables of Parameters:

Central Trigger Barrel

parameter	**CTB**
Purpose	Measure charged particle multiplicity in $-1<\eta<1$
Cell coverage	$\pi/30$ in ϕ; 0.5 in η
Occupancy	10/cell for central AuAu
Accuracy	<3% at M>1000 ; single hits at low M
Channels (cells)	240
Duty cycle	every RHIC bunch crossing
Radiator	BC404 : 1x21x112.5,130 cm
Light Detectors	Hamamatsu R5946 PMTs
ADC (PMT amplitude)	8 bits
High voltage	LeCroy 1440
HV control	LeCroy 1445A controller
Data transfer	PMT base to platform via RG58 cable
Calibration system	Led on each slat
DSM boards	16

Multi Wire Counter

parameter	**MWC**
purpose	measure charged particle multiplicity in $-2<\eta<1$ and $-1<\eta<-2$
Cell coverage	variable in η and ϕ
Occupancy	0.1 - 0.2 per cell
Accuracy	<3% at M>1000 ; single hits at low M
Channels	7680 wires
Pre amps	1000 @ 8 channels each
Wire grouping	20 wires per group
Sectors	24
Digitizers	400 - 5 bits each
Duty cycle	every RHIC bunch crossing
Data transfer	via 24 fibers 1 GHz
Optical Transceivers	24
DSM boards	16

Zero-degree calorimeter

parameter	**ZDC**
purpose	verify centrality in AuAu collisions
cell coverage	tiny, near 0 degrees
occupancy	25 n for central AuAu collision
Channels	6
ADC resolution	8 bits
Duty cycle	Every RHIC bunch crossing
Radiator	Ck fiber
Converter/absorber	Pb plates
Detector	Hamamatsu 2490-05 PMT
Data transfer	analog via BNC to platform
ADC boards	1
DSM boards	1

DSM (Data Storage and Manipulation) Boards

parameter	**DSM**
Input bits	128
Output bits	32
Gain Correction and pedestal sub LUT	16 - 8 bit or 8 - 16 bit each
Cyclic buffers for storing input data	16 - 64k x 8 bit
Level 0 Algorithms	1 FPGA
VME interface and Local Bd Control, VME CBLT and multicast support	1 FPGA
Output FIFO (tree alignment)	64k x 32 bits

Trigger Control Unit (TCU)

parameter	**TCU**
Input bits from DSM trees	16
Busy bits (from TCD)	8
Input counter	2 - 1M x 20 bits
Cyclic input buffers	1 - 65k x 16 bits
Temporary data storage	1 - 16 bits
Polarization input bits	4
Luminosity input bits	16
Luminosity scalers	2 - 4k x 20 bits
Trigger word LUT	1M x 16 bits
Trigger word counter	2 - 65k x 20 bits
Prescale table / scalers	2 - 65k x 20 bits
Token FIFO	4k x 12 bits
Response FIFO (to TCD)	8k x 28 bits
Accepted Trigger info FIFO	4k x 60 bits
History FIFO	1k x 1 bits

VME interface bd and Local Bd control	2 FPGA

Trigger Processing

parameter	**STAR Trigger System Logic**
Level 0 TCU	2 plus 1 spare
Level 0 DSM boards	36 + 4 spares
Level 0 TCD	1
Trigger clock distribution system	8 boards
Level 1 processors	10 (deferred)
Level 2 processors	10 (deferred)
Trigger-DAQ Interface	256 MBy
Level 3 processors	??
L1 control (L1CTL)	1 ≡ 2306 CPU
L1 I/O	1 ≡ 2306 CPU
L2 control (L2CTL)	1 ≡ 2306 CPU w/ 64 MBy

3.0 Requirements

3.1 - GENERIC

3.1.1 Event Selection

3.1.1.1 Centrality

Requirement: Select interactions based on centrality.

Justification: Physics suggests that the most likely event type for finding the QGP has maximum participation or maximum overlap.

Status: Implemented in the trigger via the multiplicity detectors CTB and MWC and through Level 0 multiplicity maximum selection. Correlations linking impact parameter to multiplicity and the detector response functions are all documented on the WWW. These suggest a maximum selectivity of approximately 3%. Simulations have shown that this must be variable to include lower multiplicity events that have other manifestations of plasma formation.

3.1.1.2 Structure in $M(\eta,\phi)$, $E(\eta,\phi)$

Requirement: Select events based on structure in both η and ϕ in multiplicity and energy.

Justification: Jets and asymmetric energy distributions are expected to signal possible transitions to QGP.

Status: Implemented through pixel analysis of EMC, CTB and MWC data. The theoretical basis for this analysis is documented on WWW. Event structures investigated so far include jets, QGP bubbles, volcanoes. The minimum momentum structure we can resolve with these trigger detectors is approximately 150 MeV/c as documented in a study of structure vs. momentum on the WWW.

3.1.1.3 Strangeness

Requirement: Select events having peculiar strangeness content.

Justification: Strangeness enhancement is expected to be sensitive to the final hadronization processes and may give indication of the interaction having proceeded through a QGP formation.

Status: Requires L3 tracking analysis with information from the TPC and, for much greater efficiency, from the SVT and TOF. We are presently investigating different strangeness selection algorithms for L3. These include analysis of track pointing to secondary vertices, dE/dx on track segments, and V finding in SVT.

3.1.1.4 Minimum Bias

Requirement: Allow selection of events with minimum bias in the selection procedure.

Justification: Need to normalize rate for more exotic processes.

Status: Implemented via RHIC clock and selection on signal in any trigger detector. A separate prescaler implemented in a DSM board allows us to issue a min bias event at frequencies as low a 1/month or anything higher. The FPGA tree can be programmed to allow any multiplicity threshold from 1 to 10^4 in CTB and MWC, or for any signal in ZDC, or for any other arbitrary input.

3.1.1.5 Special Criteria

Requirement: Provide channel for special trigger requests from other subsystems.

Justification: Many types of events may be needed for diagnostics.

Status: The Trigger Response FIFO on the TCU can be loaded from VME. Special trigger requests will come via software to the L1IO CPU, pass to the L1CTL CPU and onto that FIFO. The contents of this FIFO are sent to all subsystems when there is not a new L0 trigger for this crossing.

3.1.1.5.1 Subsystem Calibration events

Requirement: Provide for the handling of trigger requests from each detector system for the purpose of calibrating the detectors.

Justification: Need to calibrate each detector system independently, as well as to test the system as a whole.

Status: Implemented through special trigger request and RHIC strobe distribution scheme at Level 0. Calibrations can take the form of laser or pulser requests to trigger. The basic idea is that a subsystem will request a trigger and TRIG will wait for an opportune time and then issue the "special trigger" via the normal trigger channel. see "Calibration Event" appendix.

3.1.1.5.2 Trigger system calibrations

Requirement: Need to establish periodically the multiplicity accuracy of CTB and MWC, and response of ZDC.

Justification: Expect variations in timing and amplitude as components age or temperatures vary, etc.

Status: Calibration of the CTB and MWC is accomplished by gathering a series of events and then using the TPC to track onto each slat of the CTB or into each pixel of the MWC and compare our counts to the crossing track count. Each CTB slat has its own LED to check for aliveness.

3.1.1.5.3 Cosmic Rays

Requirement: Provide for a cosmic ray (CR) trigger for STAR

Justification: Need to be able to trigger on cosmic rays to check tracking efficiency in the TPC.

Status: This is presently implemented in the design of the Level 0 trigger algorithms. Free-running the pseudo-RHIC clock pulse generator leads to a 10% efficiency for direct CR detection using the CTB slats. Other trigger entry points are also possible for this type of trigger, such as loading the trigger response FIFO from an external coincidence.

3.1.2 Event rejection

3.1.2.1 Beam-gas interactions

Requirement: The trigger must recognize beam-gas interactions as distinct from beam-beam interactions and reject such events.

Justification: These constitute an interaction background and must be recognized as such.

Status: These events may have a signature in the MWC or ZDCs (asymmetric response for high multiplicity events) or be rejected after tracking in L3.

3.1.2.2 Interaction location

Requirement : Select events occurring within +/-3 cm of any particular location in the interaction diamond.

Justification : Need to test selected portions of the TPC.

Status : Deferred ≡ no longer a requirement.

3.1.2.3 After initial selection

Requirement : Once an event is selected at Level 0 we need a mechanism to reject the event if further analysis warrants this.

Justification : We want to be liberal in issuing Level 0 triggers so that we do not overly constrain the selection process thereby missing otherwise interesting events.

Status : Rejections are handled within the context of the command word structure in the trigger-subsystem communications. The abort command will be followed by the ID (token) of the specific event being rejected. Rejection can be performed by the Level 1 or 2 controllers without notification of DAQ. Rejection at Level 3 is handled by the DAQ subsystem. We are presently investigating entropy content, statistical selectivity, and neuron functions to establish requirements for an event to pass L1 and

L2. We are using isolation cuts with EMC data to select events having W and Z intermediaries. We are quantifying TPC cluster and track parameter information for use in L3 rejection and in coupling information from fast and slow detectors.

3.1.3 Dynamic Range

3.1.3.1 Projectile

Requirement: The trigger must operate for pp, pA, and AA collisions where A spans the range from 1 to 238.

Justification: While the largest A systems are expected to be the most likely systems to form the plasma, we will need to study the full range of projectile-target combinations to understand the approach to plasma formation. In addition, it is envisioned that there will be a strong polarization program for the pp interactions and STAR will make spin studies a major goal.

Status: Implemented in the present level structure using combinations of standard trigger detectors with the addition of an EMC.

3.1.3.2 Rate

Requirement: Trigger must operate at rates to 10^7 interactions per second

Justification: This is the RHIC crossing rate

Status: All trigger detectors are digitized at 10 MHz.

3.1.3.3 $-2<\eta<2$

Requirement: The trigger must select events based on the particles produced in this interval.

Justification: We need to select events where we have sufficiently powerful detector coverage to have a chance to understand the interactions. The TPC covers the region from $-2<\eta<2$ and is the primary source of information in the experiment.

Status: Implemented through the CTB ($-1<\eta<1$) and MWC ($-2<\eta<-1$; $1<\eta<2$).

3.1.3.4 Multiplicity

Requirement: Operate at multiplicity (M) of a few to many thousand.

Justification: the pp and the peripheral collisions programs produce typically M~2-5 while the AuAu collisions produce M~2000 in our trigger detectors.

Status: The CTB and MWC are all capable of responding to single mip hits. The CTB uses a charge integrating ADC so has a dynamic range >100. The MWC uses hit counting so begins to saturate when multiple particles pass through a single wire region or at rates of $dN/d\eta>2k$ or so (4k wires in single η unit).

3.1.4 Interfaces

Individual subsystem interfaces are discussed in section 3.2

3.1.4.1 RHIC Strobe Distribution

Requirement: The trigger shall provide the subsystems with a signal with a defined phase with respect to the RHIC bunch crossing stable to +/- 1ns.

Justification: This is required to synchronize the ADC gates and event time bases for the experiment.

Status: This is presently designed as a functional block of the trigger but this question also has ties to event synchronization between the detectors and experiment control and is currently under discussion. Presently this distribution is accomplished through the trigger-clock-distribution (TCD) system. (See TCD requirements document)

3.1.4.2 Subsystem BUSY status

Requirement: Trigger must provide means for subsystems to input a busy condition that will prohibit further triggers to that subsystem.

Justification: Hardware and software resources may periodically become unavailable for a variety of reasons and this information must be correlated throughout STAR to measure cross sections for various processes. In addition, Trigger must avoid sending trigger signals to any subsystem that is not ready to accept them.

Status: Busy lines from each subsystem hardware and software components enter trigger through the DSM into Level 0 decisions and into TCU for scaler bookkeeping.

3.1.4.3 Subsystem Notification

Requirement: Trigger must inform all detector subsystems of each event as it is selected at Level 0, Level 1, and Level 2.

Justification: Level 0 notification is what initiates data collection for each detector subsystem. Higher level notification is required to keep clear buffer allocations for each subsystem

Status: Accomplished with trigger distribution queue in TCU.

3.1.4.4 Event order in selection

Requirement : There is no requirement on event order notification.

3.1.4.5 Maximum latency

Requirement: There is no requirement on maximum latency at this time.

3.1.5 Scalers

3.1.5.1 Luminosity and polarization input

Requirement : Record luminosity during livetime of each trigger type and each beam polarization.

Justification : Need to compute cross sections for physics interpretation of each trigger type.

Status : deferred ≡ no longer in trigger scope.

**OPEN ISSUE - what gating do we need on the scaler information? Do we need to record only for hits "in the diamond"?

3.1.6 Upgrades

3.1.6.1 Inclusion of new detectors in the trigger

Requirement: We must allow for new detector systems to be triggered and to participate in the triggering decision.

Justification: New detector systems will be added to the STAR detector complement to enhance the ability of STAR to find new physics.

Status: This is implemented in the present architecture. The detector subsystem interface allows for the inclusion of new detector subsystems of both fast and slow variety. It allows entry points for the request of triggers and for participation in the triggering decision throughout the trigger sequence from Level 0 to 4. Entry at L0 is through the DSM boards. At L1 and L2 it is through shared memory. At L3 and L4 it is through global control CPUs.

3.1.6.2 Modification and addition of new trigger criteria.

Requirement: The trigger shall provide mechanisms that allow the triggering criteria to be modified and the triggering capability upgraded.

Justification: As our knowledge of the RHIC physics environment increases, we will need to modify the criteria used to define which events are accepted to maximize the acceptance of "interesting" events.

Status: This has been implemented in the trigger through extensive use of an open and scalable architecture employing Field Programmable Gate Array decision engines at Level 0 and generic CPU processing at Level 1, 2, 3.

3.1.7 Software configurable

3.1.7.1 Each run

3.1.7.1.1 Prescale values

Requirement: Need to be able to load prescale values for each trigger type.

Justification: We want to study a wide variety of event types and want to be able to control the rate at which types are collected.

Status: Implemented for Level 0. We need to implement a similar scheme for events surviving each trigger level.

3.1.7.1.2 Selection code

Requirement: Need to be able to change selection code at all trigger levels.

Justification: Will have different types of runs as well as evolution in our understanding of event selection.

Status: Implemented through VME connection to FPGAs at Level0 and to generic CPUs at L1, L2, and L3.

3.1.7.1.3 Lookup tables

Requirement: Need to be able to correct input data for gain, pedestals, etc.

Justification: Detectors cover different regions in η,ϕ and electronics have different quiescent operating points.

Status: Implemented through LUT at input of DSM for each TRIG detector.

3.1.7.1.4 Timing

Requirement: Trigger must configure fully in less than 1 minute.

Justification: Begin run sequence must complete in arbitrarily short time.

Status: All code and tables can be loaded in a few seconds now.

3.1.7.2 Infrequent

3.1.7.2.1 Input data

Requirement: Need to tell system to read stored data or accept new data.

Justification: Needed to test operation of trigger.

Status: Implemented through MUX on DSM boards.

3.1.8 QA Program

3.1.8.1 Documentation

Requirement: All components, hardware and software, must be accompanied by descriptive and operating documents.

Justification: Need to provide coherence over personnel changes and failing memories.

Status: Implemented.

3.1.8.2 Drawings

Requirement: Drawings must adhere to STAR standards for drawings.

Justification: Need experiment wide compatibility.

Status: Implemented.

3.1.8.3 Fault Tolerance

3.1.8.3.1 Diagnostics

Requirement: Provide information to determine whether trigger selection is as instructed.

Justification: Need to tell if trigger is operating correctly.

Status: Scalers at each decision point. Registers holding FIFO status for all TCU FIFOs. L1 and L2 controllers write status of event analysis for each event into data stream.

3.1.8.3.1.1 Demand acceptance

Requirement: We need to specify that a selected event be accepted regardless of the outcome of Level 1,2,3, and 4 analysis.

Justification: Need to be able to test various aspects of the selection process.

Status: Implemented through special trigger input lines for Level 0. A bit set at L0 will become part of a control word, which is interrogated at all subsequent levels before issuing any reject. Under discussion with DAQ using control word for level 3.

3.1.8.3.1.2 Simulated input for testing

Requirement : All trigger electronics must have provision for test data presented at the input on the same path that real data follows.

Justification : Need to be able to debug hardware and software in a controlled fashion.

Status : Implemented through MUX selectable data paths at input to all DSM boards as well as on TCU. We also can pass through any event or set of events with a "reject override" option in the Level 1 and Level 2 queue controllers. Simulation timing data can be used to stop VPD TDCs to test the analog section of that device, excluding discriminators and PMT. LED pulsers are used to test the full analog chain of VPD, CTB, and ZDC.

3.1.8.3.1.3 Filling test buffers for access via slow controls

Requirement: Provide mechanism whereby individual event raw data is moved to buffers for access via slow control path.

Justification: Will be necessary to test individual front end boards.

Status: A software settable flag will be passed to the trigger-clock-distribution boards to indicate that the next event will be routed to intermediate storage, regardless of its subsequent action.

3.1.8.3.1.4 Trigger counters

Requirement: Provide scalers to keep track of all triggers formed, issued, and rejected.

Justification: Need to compute cross sections for various physics processes and to monitor health of the trigger.

Status: The arrays holding counts of each trigger type accepted and rejected at each level can be read by a L1 CPU and made available through the L1IO CPU.

3.1.8.3.1.5 General

Requirement: Provide path for forming one and two dimensional histograms of fast detector data

Justification: Need to look at experiment performance as well as to find areas of interest on the “wings” of distributions.

Status: Implemented in Level 2.

3.1.8.3.2 Electronics

Requirement: No failure of inaccessible electronics shall disable the trigger system. (We consider electronics on the platform accessible since we can gain platform entrance at end of fill cycle for limited time.)

Justification: RHIC will limit access to variety of areas where electronics might be located.

Status: No cave resident electronics can affect more than 80 MWC wires, 1 CTB tray or ZDC PMT. We require RHIC resident spares for each remaining component (DSM, TCU, L1CPU, L2CPU).

3.1.8.3.3 Defaults.

Requirement: Default condition for any given crossing will be no event. Default condition will require positive selection at each level of trigger for event to proceed.

Justification: Choose that the default condition not swamp other STAR resources.

Status: Level 0 requires positive selection signal. Level 1 and 2 must issue explicit acceptances before next stage is initiated. Note that the default action can be changed through the configuration file.

3.1.8.3.4 Resets

Requirement: Trigger must power up in reproducible specified state.

Justification: System may hang in any state.

Status: All trigger hardware and software settable parameters are loaded at begin run. DSM boards power up in a LOAD state, and will not operate until explicitly loaded.

3.1.8.3.5 Software:

Requirement: Trigger software must perform constant self consistency checks for error detection. In case of an error, backup procedures have to be in place to allow graceful recovery.

Justification: If trigger stops, the experiment stops.

Status: Level 1 and Level 2 analysis processes are allowed to hang without interrupting more than one event in the data stream. Other steps are under discussion.

3.1.8.3.6 Radiation damage

Requirement: All components survive the RHIC environment for at least 10 years.

Justification: Expected lifetime of STAR system.

Status: Need estimate of rad bkgd at each det.

3.1.8.4 Pre Installation Testing

Requirement: All trigger components must be tested with documented test results prior to installation at STAR.

Justification: standard QA practice.

Status: Each CTB slat and all electronics modules are fully tested and documented on the WWW.

3.2 Subsystem Specific

3.2.1 DAQ

3.2.1.1 Timing

Requirement: Notify DAQ within 5 ms. of occurrence of an event including the TPC or SVT.

Justification: The slow detector digitization chains take 5ms. to complete, regardless of occupancy.

Status: TRIG informs DAQ at the end of L2 processing, w/i a variable time whose maximum latency is 5 ms. of the occurrence of an event.

3.2.1.2 Interface to trigger

Requirement: The trigger and DAQ shall mutually agree on one interface standard for the efficient passing of trigger event and detector data.

Justification: This is necessary for the functioning of the experiment

Status: Documented and signed off.

3.2.1.3 Event ID

3.2.1.3.1 Event types

Requirement: The trigger shall provide a mechanism whereby both the multiple detector events and single detector events are synchronized.

Justification: In order to reconstruct full events we need to be sure that the full events built by DAQ are the same event as seen by each detector.

Status: The plan is to distribute a series of bits that describe the status of the present event through the TCD. These bits will include a command word, followed where appropriate by a trigger word, followed by event ID (token) and header information. These bits will be distributed to the individual data components within each subsystem by the subsystem distribution net.

3.2.1.3.2 Uniqueness

Requirement: Trigger must provide an event ID that is unique over the life of the experiment.

Justification: We do not want any ambiguity in discussing a particular event.

Status: There is a 64 bit word whose content may be a UT clock or the contents of a RHIC strobe counter included in each event header.

3.2.1.4 Subsystem coupling in events

Requirement: Trigger must be able to initiate triggers involving any combination of STAR detectors.

Justification: Physics can be gleaned from single detectors like the EMC as well as from fully coupled events that include all detector components.

Status: This is implemented with the inclusion of the rule that detectors that wish to have the option of running decoupled from the TPC must provide a data buffer scheme for internal storage.

3.2.2 TPC

3.2.2.1 Timing

Requirement: Trigger must provide event notification to the TPC within ~1.5 μs of the interaction.

Justification: We lose information at the rate of ~2% per μs delay due to the drifting of electrons in the TPC out of the active volume.

Status: Level 0 informs all subsystems within a configuration set parameter that can be as short as 1 μs after the interaction.

3.2.2.2 Interface to trigger

Requirement: Trigger must provide a path for using TPC tracking information in making a decision whether to keep a given event.

Justification: We can acquire many more TPC events than we can record on tape, so we must be able to use the TPC data itself in making a final trigger decision.

Status: Implemented through both Level 3 and Level 4 trigger processors.

3.2.2.2.1 Calibrations

Requirement: Provide for laser and other types of calibrations of the TPC.

Justification: Need straight tracks of known properties to calibrate the TPC response.

Status: We anticipate 2 types of lasers (one for each end), as well as SAS pulser and the ground plane pulser, plus reading out an event stored in memory for a total of 5 separate calibration type events. Special trigger requests are handled through the Trigger Response FIFO that is loaded via software from experiment control. See LASER appendix and Calibration Event Appendix. We expect two inputs to DSM from TPC one for each laser. These are OR′ed to decide whether to issue a laser pulser for the TPC. Which laser actually fired may be encoded as one of the bits of the DAQ command, perhaps. The SAS pulser requests are expected to come through software only. Ground-plane pulser requests also come only through software. A hardware path exists through the TCD which send up to 4 bits to the last DSM board to request various types of special triggers for each subsystem.

3.2.2.2.2 Busy

Requirement: Block sending triggers to the TPC when it is in a busy state.

Justification: The TPC has limited buffer space and has a maximum rate at which it can be triggered.

Status: TPC will inform TRIG when it is not in a condition to accept further triggers. This is done through the TCD, which takes signals from DAQ as well as from its internal registers to set a busy bit for the TPC.

3.2.2.2 Pre and post history

Requirement: Provide a crossing history both preceding and following a Level 0 trigger that includes the TPC.

Justification: TPC acting alone needs to have all crossings within +/-1 cm upstream and downstream clearly identified for track content so that it can select the real event vertex.

Status: implemented with pileup counter in Level 0 and 1. see "HISTORY.DOC" appendix

3.2.3 SVT

3.2.3.1 Timing

3.2.3.1.1 Notification

Requirement: Trigger must inform SVT of an event within ~3 microseconds of the interaction.

Justification: The SVT scrolls data onto an SCA which holds ~4.5 µs of data. It takes some time to switch the array from acquire to digitize. This time needs to be fixed and deterministic within a run, so that the TCD can implement a fixed delay sequence to notify the RDO boards.

Status: Level 0 informs all subsystems within a configuration set time interval that is presently set to 1.5 µs.

3.2.3.1.2 T0 correction

Requirement: Timing information from VPD must be made available to level 3 processors.

Justification: to allow for t0 correction of the space-point finding.

Status: deferred ≡ no longer in trigger scope.

3.2.3.2 Interface to trigger

Requirement: Trigger must provide a path for using SVT tracking information in making a decision whether to keep a given event.

Justification: We can acquire many more SVT events than we can record on tape, so we must be able to use the SVT data itself in making a final trigger decision.

Status: Implemented through both Level 3 and Level 4 trigger processors.

3.2.3.2.1 Calibrations

Requirement: Trigger must allow for electronics calibration of the SVT.

Justification: Electronics drift and pedestals/gains may change.

Status: SVT will request a special trigger through the TCD to the DSM0 and TRIG will then queue this request at an appropriate time. Software trigger requests are sent to the L1IO CPU for queuing on the response FIFO.

3.2.3.2.2 Busy

Requirement: Block sending triggers to the SVT when it is in a busy state.

Justification: The SVT has limited buffer space and has a maximum rate at which it can be triggered.

Status: SVT will inform TRIG when it is not in a condition to accept further triggers.

3.2.3.2.3 Diagnostics

Status: Currently there is a trigger3 command ("executed as trigger0 but copy event also into local buffer"), but nothing equivalent for pulser events. We were thinking that it would be good to have this control implemented via a slow controls interface on the readout cards, so that individual RDO boards could be instructed to store the next event in local memory. This would also allow each detector subsystem to individually determine, if they want to store the next event or not. We would implement this as a logical "OR" with the next trigger command that would generate an event. The trigger command would also reset this (slow control) switch, so that only one event would be stored. Replay of this event can then proceed with the "pulser3" trigger command.

3.2.3.3 Select Interaction Vertex Position

Requirement: Select events occurring within +/- 3 cm of the center of the interaction diamond.

Justification: The Silicon Vertex Tracker spans only 8 cm around the diamond center.

Status: deferred ≡ no longer in trigger scope.

3.2.4 EMC

3.2.4.1 Timing

Requirement: none

3.2.4.2 Interface to trigger

3.2.4.2.1 L0

Requirement : Make Level 0 trigger decisions based on EMC data

Justification : EMC will be the primary trigger detector for pp and will be useful in isospin studies in AA as well.

status: Data is input to trigger Level 0 through the data storage and manipulation boards. EMC shall supply 12 bits per tower, 300 towers total, with 6 bits for the tower sum and 6 bits representing the maximum hit in the tower.

3.2.4.2.2 L1,L2

Requirement: none

3.2.4.2.3 L3

3.2.4.2.3.1 Vertex position

Requirement: EMC requires knowledge of the beam centroid position at the time of each event to within 0.3 mm.

Justification: A momentum resolution for TPC tracks of better than 12% at 40 GeV (rms) is needed in pp running. This is for electrons from W and Z. There is a similar requirement for leading particles in jets in pp and pA running. Also, full acceptance from 80% of the interaction diamond is needed at eta from -1 to +2.

Status: Implemented through tracking TPC data at L3. The SVT may not be useful for every event in this case.

3.2.4.2.1 Calibrations - standard path

3.2.4.2.2 Busy - standard path

3.2.5 TOF

3.2.5.1 Timing - none

3.2.5.2 Interface to trigger - none

3.2.5.2.1 Calibrations - standard path

3.2.5.2.2 Busy - standard path

3.2.6 Slow Controls (SC)

3.2.6.1 Data archiving

Requirement: Archival data from the trigger detectors and trigger processor shall be made available to the slow controls interface for monitoring and archiving

Justification: Long term monitoring of detector performance and detection of fault conditions are a necessary part of the trigger/experiment design.

Status: Under discussion with slow controls group. Provisions made in the trigger architecture through message system.

3.2.6.2 Detector control and monitoring

Requirement : Need to monitor and set HV for each PMT.

Justification: Need to control gains and operation.

Status: SC has provided an excellent interface to the LeCroy HV system.

3.2.7 Experiment Control or Online

3.2.7.1 Commands

Requirement: Trigger shall establish a protocol with EC for handling commands such as DisabledToEnabled, EnabledToReady, ReadyToRunning, RunningToReady, ReadyToEnabled.

Justification: ONL must be the overall experiment coordinator.

Status: Under discussion with ONL. We expect the EnabledToReady command to lead to complete configuration of the trigger system, which will be done at the start of each run. This command sets up a list of files that are then executed under trigger control to load prescales, FPGA code, L1 and L2 algorithms.

3.2.7.1.1 Configuration

Requirement: ONL must be able to configure trigger for specific runs.

Justification: Physics and calibration demands require coordination by EC.

Status: Under discussion. Issues like rate control and calibration runs will be dealt with through configuration files.

3.2.7.2 Handshakes

Requirement: Trigger must be able to inform ONL of its internal status upon request.

Justification: This allows ONL to coordinate the experiment.

Status: Under discussion.

HISTORY Appendix:

last update 1 Feb 1995

Requirement: Trigger must provide a summary of the multiplicity in each crossing for a full drift time preceding and following each triggered event that involves the TPC. Trigger must be able to reject any event based on the total multiplicity in any preceding or following crossing up to the full drift time of the TPC.

Justification: Individual tracks in the TPC point back to a primary vertex on the beamline with an accuracy (sigma,rms) of no better than 1 cm. Thus, even when the system is running smoothly we will need to reject events with vertices closer than a few cm (250ns in drift time) to the interaction of interest. At turn on, it is necessary to guarantee clean single events in the TPC to allow software development to proceed with relative ease.

Status: The trigger provides a multiplicity for each crossing as part of the Level 0 decision process. This multiplicity word is stored just like all other trigger words in a RHIC synchronous pipeline having a length of 100ms. (1 M long). If we record the 2 byte word for each preceding and each following crossing for an interval of 100 micro-sec, this adds 4kB to the event.
It is also possible to read out the whole trigger data set (2kB max) for each of these crossings, or for any subset of this number. This is the crucial issue to be decided. How much history is desired?

from S.Klein:
I think that we should divide the drift time into two periods. The first, ~ 250 ns. ~ 1 cm, is the period during which it is possible to, based on tracking, not tell that there are two vertices, or even assign some fraction (1%?) of the tracks to the wrong primary vertex. I would argue fairly strongly that these events should be rejected unless they otherwise look extremely interesting.

The second time period is the full TPC drift time before and after the event. Here, the main effect is to increase the TPC occupancy, and hence degrade the tracking. Here, there needs to be some subdivision, because a 10 microsecond time difference gives a fairly serious 87% overlap, while 70 microseconds is only 12%, a hopefully negligible amount. But, in this period, overlaps are generally less serious, and also it doesn't seem necessary to record multiplicity data from each individual crossing - saving total multiplicity in 1 microsecond increments should be OK. I'd really like to see some estimates by the tracking people as to how much extra, out of time tracks hurt the tracking.

from M.Botlo:
maybe you should ask EMC on this: I could imagine that they would like to have more history info (i.e. more crossings before/after). there could have been a TPC triggered event but since the TPC integrates over many crossings, EMC could have seen another event of interest to them. if this event is not too far away (in time) from the TPC event, they have gained an event for them (almost) for free - if they have trigger info for this particular crossing (or this + 1 before + 1 after).

LASER EVENT Appendix:

last update: 31 May 1995

from H.Wieman:

The laser has a large latency time but it can be made to fire precisely timed with respect to the RHIC clock by externally triggering. The operation would be as follows. A pretrigger is issued to the laser to fire the flash lamp. It then takes about 200 microseconds of flash lamp time for the population inversion to build up to a maximum. A trigger is then issued to the Q switch and the laser pulse occurs probably within a few to tens of ns of when the Q switch trigger was issued. The laser pulse time jitter relative to the Q switch trigger is about 1 ns. This should make any trigger system happy and corresponds to only 60 microns of z jitter in the track location. The 200 microsecond delay between flash initiation and Q switch should be held constant to the per cent level just to maintain uniform pulse to pulse performance of the laser.

from H.Crawford:

Trigger plan: TPC will issue a special trigger request by raising a level on the line for special triggers from the TPC. TRIG will note this request and issue a "special trigger" through the normal trigger distribution scheme when the system is ready to accept such a trigger. TPC will take this trigger, decode the command and trigger words to realize that this is the trigger it requested, and then initiate the actions that lead to an appropriate event. TRIG will inform DAQ after a delay appropriate for this trigger type. Since the laser latency is only 200 µs, which is small compared to the normal 5ms. it takes to process a TPC event, we could use the same delay that is normal for informing DAQ of a trigger. We assume that the normal BUSY circuits will be employed, again since the 200 µs does not represent any large time perturbation on the system.

Howard has said that he may want to take a series of such triggers. Presumably this can be accomplished by simply leaving the special trigger request line high so that as soon as the system is not busy, another special trigger would be issued.

Comments?

From Howard Wieman - 1 Feb 95:

Me again, we will probably issue separate commands for each laser flash because we have to run the laser at roughly 10 Hz when it is operating. If we miss a few flashes because of a busy it is not a problem. We will flash the lamp but not open the Q switch. Howard W

From S.Klein 31 May 1995

The "laser pulse time jitter relative to the Q switch trigger is about 1 ns." is not true for a real application. The actual output pulse timing changes by 20-30 ns as the laser amplitude changes. This necessitates inclusion of a TDC to tell the relative time of the laser firing with respect to the RHIC strobe. There is a potential complication for the laser firing crossing RHIC strobe boundaries.

W/ S.Klein -- need writeup of recent changes in laser

The STAR Laser Controller generates two basic signals: the laser flashlamp and the laser Q switch.

The main input is the 100 ns. period RHIC clock. This is divided down to provide a 10.0 +/- 0.5 Hz flashlamp signal. This signal is always on. The flashlamp signal is TTL, 200 ns. wide.

The flashlamp signals also goes to a programmable delay, which is adjustable from 120 to 200 microseconds in 1 microsecond steps. This delay is VME programmable, and must also be
VME readable. Because the delay must remain constant to a few ns., it must be entirely digital. It must also be possible to disable the Q switch via VME.

The output of the delay is then synchronized with the RHIC strobe with a simple register. The output of this register goes into a programmable delay line, the Data Delay Devices PDU-13256F-.5 or equivalent. This delay line is programmable from 0 to 127.5 ns. in 0.5 sec steps. The output of the delay line provides a Q switch signal, a 200 nsec wide TTL pulse. It also provides two auxiliary outputs, a differential ECL TDC start (50 ns. wide minimum) and a trigger request, a 110 ns. wide differential LVDS signal.

This delay line is part of a feedback loop that keeps the laser timing constant; it is controlled by the signal from the laser diode that monitors when the laser actually fired.

The TTL signal from the laser monitoring diode is buffered, and goes two places. First, it goes to a 50 nsec minimum width differential ECL output to a TDC. Second, it goes to a time to amplitude converter (TAC). The TAC could be a simple set-reset flip flop driving a RC network, or something slightly more complicated.

The TAC start signal is the synchronization register; the diode provides the stop signal. This time difference is what must be held constant. If the measured time rises, a shorter delay time must be selected; if the measured time decreases, a longer delay is needed. The TAC output is monitored by two comparators, one that detects if the amplitude (and hence time) gets too high; the other detects if the amplitude is too low. Both of these comparators compare the TAC with voltages generated by 8 bit DAC′s, with a range such that they are sensitive to TAC times in the 30 to 150 ns. range.

The output of the comparator is sensed an appropriate time after the diode signal is received (100 ns.?), giving time for the comparators to settle. This sampling can be done with simple AND gates; if the high comparator fired, or the low one didn′t, then the 8 bit up/down counter that controls the delay line must be incremented or decremented. The state of this counter should be readable (but not writable) via VME.

If no diode signal is received, then the counter setting must not change.

When the module is first powered on, it must come up with the Q switch outputs (including Q switch, TDC and trigger request) disabled.

The module should be in the 6U form factor. The base address must be selectable by switches.

Calibration Events Appendix:
8/29/98 HjC, LG, JE, EJ, ZM

Calibration Strawman

This note describes the steps taken to produce a calibration trigger for the CTB and the TPC separately. We recognize two types of special trigger requests, one made through a cron job for a main data run, and another initiated by a collaborator (the Requester) requesting a series of special triggers for a separate data stream. We describe here the Requester path. The cron path can be viewed as a simple replacement of the Requester with the CronJob. We assume that the CronJob runs in the Trigger Clock Distribution CPU and always takes priority, and in fact locks out the Requester until a run is finished.

Items explicitly recognized as discussion points are flagged by **.

1. The Requester requests permission for CTB LED trigger events from ONLINE.
   We believe that ONL should have the role of data-stream-content coordinator.

2. ONL checks compatibility with existing runs. If OK, ONL requests appropriate LED firing mask files from Requester.

3. Requester sends mask to ONL.

4. ONL requests TRIG to load appropriate registers in DSM0[1] FPGA or TCU LUT[2] to allow CTB triggers. This could be through the ReadyToEnable, EnableToReady, command sequence.

At this point, ONL can step out of the loop. ONL does not need to be involved until it receives a request from Requester to terminate the calibration session, or ONL decides to remove permission for special triggers for some other reason.

5. Requester establishes communications with TRG through the TLC[3] Sun station

6. Requester asks that a file-specified CTB Digitizer Board (CDB) mask be invoked to arm individual LED arming switches.

7. TRG informs Requester that LEDs are armed.

8. Requester sends request for CTB LED trigger to TLC.

9. TLC sets request level for RHIC Clock Control (RCC) board.

10. RCC sends trigger pulse to CDBs.

---

[1] DSM0 is the last DSM board in the tree.
[2] Trigger Control Unit LookUp Table
[3] Trigger Local Controller - this is the firewall between trigger processors and the world. This model now assumes that TRG controls the TCD (Trigger Clock Distribution) CPU.

11. CDBs issue LED pulser signals to armed LEDs. We expect it to take approximately 10 ticks to establish stability in the LEDs. If the CDB simply continues to send gates, then we will get 10 crossings in a row which have varying amounts of charge integrated in the PMTs. Thus, a trigger LED event can actually be 10 crossings in a row. We assume that these LED events will be taken in special separate runs when STAR is not otherwise on, since these events are only to determine slat aliveness. Charge calibration is performed using isolated single tracks as found in the TPC.

**We can make this clean for the CTB in at least three ways: 1. have a dedicated run for aliveness; 2. use a BUSY bit to lock out other action during the CTB calibration; 3. run this calibration for the CTB at the beginning of each run. The problem only exists because of the 10 crossing stabilizing time for the LEDs and is not expected to affect other pulser type triggers.

12. A CDB sends a delayed signal to DSM so that it enters DSM0 with CDB ADCs from Calibration Event.

13. DSM0 checks compatibility (config data) and, if compatible with other run parameters, sets special trigger request (STR) flag and sends special trigger value to TCU.

14. TCU computes Trigger Word (TW) based on TCU input bits (8 of the 20 inputs arrive via the P3 backplane and indicate STAR BUSY status). TRG checks prescale for this TW, and, if 1, places TW and TAW on its information FIFO for distribution on P3.

15. Event propagates through DAQ like normal event.

16. Repeat steps 7-15 to fire same set of LEDs or 6-15 to change mask.

17. Requester informs ONL that task is finished.

18. ONL requests TRG to alter configuration accordingly.

19. TRG informs ONL that CTB calibration events are no longer active.

TPC Calibration events:

We assume there are two types of TPC calibration events, one which is slow and indeterminate in time, such as a "laser event" and one which is fast, such as a simple injection of charge onto the FEE input (called "pulser event"). There are actually 4 such calibration event types defined, but it is sufficient here to distinguish fast and slow types. We assume there are two paths for the calibration event to be requested, one through the cron job and one through the Requester, as above.

1. The Requester sends request to ONL to establish TPC calibration events in this run. This may be accomplished by entering appropriate values in the TCU LUT.

2. ONL checks compatibility with other types of triggers employed in the existing run.

3. If compatible, ONL sends request to TRG to configure for TPC calibration events in the data stream through the Ready-to-Enable and Enable-to-Ready sequence.

4. ONL sends TPC permission to generate calibration trigger requests.

Note that steps 1-4 are essentially identical to the steps 1-4 above.

5. TPC establishes communication path to TLC.

6. TPC sets up its own internal states appropriately (e.g. turns on the laser)

7. TPC sends request to TLC to fire a TPC calibration event.

8. TLC sends request on to TCD CPU.

9. TCD CPU sends request on to TPC TCD module. Note that the TCD module can determine the time of the request to the DSM0 if this is critical. The TCD CPU is assumed to be under the control of the SC group. We assume that all cron jobs for calibration events will be run in the CPU.

We assume that each STAR subsystem will have one TCD module whose mezzanine board controls communication of trigger information to the subsystem hardware. Each TCD mezzanine will have a set of up to 4 lines that go directly to the DSM0 inputs. These are set to high or low levels to inform trigger what type of special trigger is being requested. We assume that "collisions" involving simultaneous requests from different subsystems can be handled by appropriate priority configuration in the TCU LUT.

10. TPC TCD sets appropriate levels on its 4 lines to DSM0.

11. DSM0 recognizes the STR-TPC request and, if the crossing it is processing is otherwise empty (perhaps literally - we can demand that it correspond to a particular bunch number) and if the TPC Calibration Trigger bit is set in its register configuration, then it places the Special Trigger Request flag as input to the Trigger Control Unit (TCU). The remaining 12 bits of the TCU input can then indicate which type of trigger is being called special.

11.a. We need a clear understanding of how TPC is going to determine whether TRG is doing as it expects.

12. The TCU takes the input word, combines it with the BUSY bits describing STAR (not necessarily one per subsystem - we have 8 BUSY bits in our thinking), and decides whether to issue the corresponding special trigger by the value in the trigger word lookup table addressing the appropriate prescale value. TRG either issues the trigger request with the appropriate Trigger Action Word (TAW) or it doesn't.

13. When TPC TCD sees the expected TAW it lowers the STR lines to the DSM0 and proceeds to take control of the event, issuing either laser pulses or charge injections as it so desires.

14 When TPC has collected enough calibration data, it sends a request to ONL to remove permission for TPC calibration events in this run.

11. ONL then requests TRG to alter its configuration file to remove TPC calibration event permissions.