29 August 1995

STAR Trigger Functional Requirements
Table of Contents

3.1.4.4 Event order in selection
3.1.4.5 Latency

3.1.5 Scalers

3.1.5.1 Luminosity and polarization input
3.1.5.2 Trigger diagnostics
3.1.5.3 General

3.1.6 Upgrades

3.1.6.1 Inclusion of new detectors in the trigger.
3.1.6.2 Modification and addition of new triggering criteria.

3.1.7 Software configurable

3.1.7.1 Each run
3.1.7.1.1 Prescale values
3.1.7.1.2 Selection code
3.1.7.1.3 Lookup Tables (LUT)
3.1.7.1.4 Timing
3.1.7.2 Infrequent
3.1.7.2.1 Input data

3.1.8 QA program

3.1.8.1 Documentation
3.1.8.2 Drawings
3.1.8.3 Fault tolerance
3.1.8.3.1 Diagnostics
3.1.8.3.1.1 Demand acceptance
3.1.8.3.1.2 Simulated input for testing
3.1.8.3.1.3 Filling Slow Control Buffers
3.1.8.3.2 Electronics
3.1.8.3.3 Default position
3.1.8.3.4 Reset
3.1.8.3.5 Software
3.1.8.3.6 Radiation damage

3.2 Subsystem Specific

3.2.1 DAQ

3.2.1.1 Timing
3.2.1.2 Interface to trigger
3.2.1.3 Event ID
3.2.1.3.1 Event types
3.2.1.3.2 Uniqueness

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 azimuthal 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 external TPC (XTPC), an electromagnetic calorimeter (EMC), a time-of-flight (TOF) system, 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 micro-sec of the occurrence of the interaction

2. Locate events to an accuracy of +/-3cm in interaction diamond
3. Select events based on particles produced in $-2<\eta<2$ interval
4. 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, leading to 4 levels of time structure set by the conditions indicated.

Level 0 : We need to open the TPC amplification grid within approximately 1μs 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 1 : 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 which requires 10ms 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 10ms 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 different 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 other faster 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 Vertex Position Detector (VPD), the Central Trigger Barrel (CTB), the Multi Wire Counter (MWC) and the Veto Calorimeters (VTC). 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.

VPD: The vertex position detector consists of two units, each containing 24 VPD modules. One unit sits on each side of the interaction diamond, at a distance which is variable between 3 and 8 meters from the diamond center. Location is determined to optimize position information while minimizing interference with other outer detectors. The time difference between the first hit in the left side and the first hit in the right side gives the location of the interaction vertex to an accuracy that depends on the number of hits in each detector. For AuAu interactions simulations suggest that we achieve <1cm position resolution with 100ps resolution per module.

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 $1<\eta<2$.

VTC: at the cave exit on each side, in the crotch of the first beamline dipoles, there are two simple Pb-scint calorimeters whose purpose is to record passage of fragmentation region neutrons and protons. This provides a 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. It consists of 1200 towers of calorimeter. Each tower sends light to 1 PMT.

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 then gathers slow detector information and begins track analysis for Level 3 trigger decisions based on separate detector segments. Finally, at Level 4, 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
EC	Experiment Control subsystem
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 - 10 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
VPD	Vertex Position Detector
VTC	Veto Calorimeter
Z	Position align beamline

This document will be arranged in such a manner that for each requirement there will be three sections. The three sections are; 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	scintillator paddles 20cmx100cmx1cm
Light Detectors	Hamamatsu 2490-05 PMTs
ADC (PMT amplitude)	8 bits
High voltage	local Cockroft-Walton generator
Low voltage	+30V, +/-5V
Power supplies	20
Camac crates	1
Data transfer	PMT base - digital pipeline via 266 MHz fiber
Optical transceivers	120
Calibration system	laser on trolley distributed over fiber
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

Vertex Position Detector

parameter	**VPD**
purpose	locate vertex position
cell coverage	$\pi/8$ or $\pi/12$ in ϕ; 0.05 - 0.1 in η
occupancy	0.8 in central AuAu collisions
Channels	48 total - 24 on each side
Position resolution	<3 cm.
Method	extrapolated leading edge (dual threshold)
TDC	25 ps, 10bits
Lookup table	2 x 10 bits to 8 bits
Duty cycle	every RHIC bunch crossing
Total TDC processing time	185 ns
Photon converters	$1\lambda_{rad}$ Pb
Cherenkov radiators	1 cm thick 5 cm diameter quartz discs
Detectors	48 Hamamatsu R2490-05 PMTs
High voltage	remote from LeCroy 1440
Data transfer	anode output over BNC to digitizer on trolley
Front end digitizer boards	8 @ 6 channels each
DSM boards	4
Timing generation control board	1

Veto Calorimeter

parameter	**VTC**
purpose	verify centrality in AuAu collisions
cell coverage	tiny, near 0 degrees
occupancy	25 n for central AuAu collision
Channels	4
ADC resolution	8 bits, 20 MHz
Duty cycle	every RHIC bunch crossing
Radiator	scintillator
Converter/absorber	Pb plates
Detector	Hamamatsu 2490-05 PMT
Data transfer	analog via BNC to trolley
Optical transceivers	2
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)	64 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	40 + 4 spares
Level 0 TCD	1
Trigger clock distribution system	8 boards
Level 1 processors	10
Level 2 processors	10
Trigger-DAQ Interface	128 MBy
Level 3 processors	24
L1 control (L1CTL)	2 @ 16 MBy
L1 I/O	1
L2 control (L2CTL)	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.

Present status: Requires third level 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 Level 3 and perhaps Level 4. 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.

Present Status: Implemented via RHIC clock and selection on signal in any trigger detector. A separate prescaler in the TCU (24 bits or 10^{14} counts) 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 VTC or VPD, 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 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.

Present 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 "LASER EVENT" appendix.

3.1.1.5.2 Trigger system calibrations

Requirement: Need to establish periodically the timing characteristics of VPD, multiplicity accuracy of CTB and MWC, and response of VTC.

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

Status: The VPD is calibrated by taking a series of events and finding the minimum hit time for each PMT. We use the VPD to select events, then use the TPC to select events with a vertex at the mid point of the diamond, and then find the minimum time for each PMT. Calibration of the CTB and MWC is accomplished similarly, we gather a series of events and then use 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.

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.

Present 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.

Present status: Implemented through the VPD which can select interaction location.

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 : Implemented with VPD location measurement and possibility to select given range in Z within the DSM tree.

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.

Present 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 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. The VPD is 70% efficient for selecting the vertex for pp events and approximately 95% efficient for AuAu events (excluding EM decays).

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. Discussions concerning HV and rates on PMT in progress. Pipelines allow 10MHz operation with $<1\mu s$ latency.

3.1.3.3 $-2<\eta<2$

Requirement: Since the TPC covers the region from $-2<\eta<2$ and it is the primary source of information in the experiment, 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.

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

3.1.3.4 Multiplicity

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

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

Status: The CTB, MWC, and VPD are all capable of responding to single mip hits. The CTB uses a charge integrating ADC so has a dynamic range up to about 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). The VPD gets better as M increases, but responds well to single hits.

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.

Present 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 Event selection 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

Present 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 : We assume that RHIC will provide luminosity information. In aadditionwe will use our CTB +MWC multiplicity as scaler input. We plan to have 8k separate scalar channels integrating luminosity for each BUSY condition and for each polarization condition. Integrated luminosity for a given trigger type is found by taking the full $\mathcal{L}$ sum and subtracting the channels in which one of the relevant subsystems for that trigger type is dead. Scalar event readout is controlled by TRIG sending DAQ an event whose content is the scaler buffers. The event will consist of an event type, RDX header, and data.

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

3.1.5.2 Trigger diagnostics

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: See section 3.1.8.3.1

3.1.5.3 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.

3.1.6 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 units.

3.1.6.2 Modification and addition of new triggering 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.

Present 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: Under discussion.

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 VTC.

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.

3.1.8.3.2 Electronics

Requirement: No failure of inaccessible electronics shall disable the trigger system. (We consider electronics on the trolley accessible since we can gain trolley 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 1 VPD or VTC 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.

3.2 Subsystem Specific

3.2.1 DAQ

3.2.1.1 Timing

Requirement: Notify DAQ within 10ms of occurrence of an event including the TPC or SVT.

Justification: The slow detector digitization chains take 10ms to complete, regardless of occupancy.

Status: TRIG informs DAQ at the end of L2 processing, w/i a variable time whose maximum latency is 10ms, 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

Present Status: Documented and signed off. This interface uses a shared memory accessible via VME P3 from the trigger side and through VME P2 from the DAQ side. It is possible in this architecture for TRIG to sent TRIGGER ONLY events to DAQ, with the further requirement that these be routed to permanent storage.

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.

Present 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 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 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.

Present 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 microsec of the interaction.

Justification: We lose information at the rate of ~2% per microsec 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 microsecond 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: Special trigger requests are handled through the Trigger Response FIFO which is loaded via software from experiment control. See LASER appendix.

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.

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.

Present 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 microsecs 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 can be as short as 1 microsec.

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: Position and T0 information will be made available in TRIG header passed to DAQ.

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 from TRIG and TRIG will then queue this request at an appropriate time. Trigger must allow for software and hardware triggerable electronics calibration of the SVT.

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.

Present Status: Implemented in current trigger architecture through VPD.

3.2.4 EMC

3.2.4.1 Timing

Requirement: Maximum Latency of EMC data into EMC Level 1 shall be 50 micro sec. in the case that Level 0 comes after 10 crossings.
Maximum Latency of Shower Max Data into EMC Level 1 shall be 200 micro sec.

Justification: A) The maximum must be as short as possible to allow level 1 processing time. The specified latencies would be a significant part of the processing time in the case that EMC would contribute to aborting a TPC readout.
B) The specified latencies are the best we believe we can achieve with available or affordable technology.

Status: under discussion with EMC.

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 the following Level 0 trigger information:
1) single tower hit above each of 3 (programmable) thresholds.
2) 4 tower hits above each of 3 (programmable) thresholds. (These may be non-overlapping 2x2 combinations)
3) 16 tower hits above each of 3 (programmable) thresholds.
4) Global Et above each of 3 (programmable) thresholds.
5) Digitized Global Et of at least 8 bits.
6) Digitized energy in 12 "Isospin Patches" in the Barrel which are 1.0 in eta and 1.05 radians in phi.
7) Digitized energy in a few patches in the end cap.

These basic triggers are necessary for the following list of physics triggers, and can be implemented in less than 1 u sec. See EMC appendix.

3.2.4.2.2 L1,L2

Requirement: EMC requires calculation of the correlation between EMC phototubes and Shower Max information within the level1-level2 time frame. These correlations cross EMC phi boundaries and TPC sector boundaries by one EMC nearest neighbor. EMC shall put the pedestal-subtracted phototube data for events accepted by Level 1 into a dual-port memory for use by Level 2.

Justification: The correlation between EMC and Shower max should be done as soon as possible because the rate of events into Level 3 will be too large to handle if it is not done.

Status: We specifically reject
a) direct gamma candidate events which have an obvious pi-zero in the tower,
b) events which do not pass a chi-square cut as being direct gamma or electron,
c) events in which the direct gamma candidate is too close to boundaries (outside fiducial cuts)
d) events which have a "splash" of hits in multiple areas of a single shower max box and are not reconstructable off-line.

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.

3.2.4.2.3.2 Sector based DAQ

Requirement: Upon a Level2 accept, EMC shall supply the EMC data to 24 TPC Sector Crates for correlation with TPC tracks and/or passing to DAQ .

Justification: The amount of processing to do the correlation with tracks is large and can best be handled by multiple processors operating on independent parts of the TPC data.

Status: : The data path for EMC to DAQ is through the 24 sector crates.

3.2.4.2.3.3 Electron identification

Requirement: reconstruct electrons in the trigger with sufficient resolution to make E(emc) vs P(tpc) cuts.

Justification: to get the trigger rates down to acceptable levels.

Status: Under discussion and simulation.

3.2.4.2.1 Calibrations - standard path

3.2.4.2.2 Busy - standard path

3.2.4.3 Rates

Requirement:
Maximum average rates into Level 1 from EMC shall be less than 10 Khz.
Maximum average rates into Level 2 from EMC shall be less than 2 KHz.
Maximum average rates into Level 3 from EMC associated with TPC shall be less than 100 Hz.
Maximum average rates into Level 3 from EMC not associated with TPC shall

be less than 1 KHz.

Justification: Monte Carlo Studies of high Luminosity pp have indicated that the above rates will allow the physics to be done. The above assumes that the jet rates are prescaled with multiple thresholds, except no prescale for the highest energy jets.

3.2.5 TOF

3.2.5.1 Timing - no requirements

3.2.5.2 Interface to trigger

Requirement : Get multiplicity data from TOF into Level 0

Justification : TOF displaces our primary detector, the CTB, in trays where we place TOF detectors.

Status : We take TOF paddles in groups of 25 (i.e., half tray) and put each PMT signal through a discriminator to determine whether it was hit or not, giving each PMT 1 bit. Then we sum the 25 bits to return one 8 bit word that is sent to Level 0 over the same electronics used by the CTB that would have resided in that tray.

3.2.5.2.1 L3

Requirement: Combine TOF data with TPC data to make trigger based on PID.

Justification: PID requires TOF

Status: Implemented in sector based DAQ.

3.2.5.2.2 Calibrations - standard path

3.2.5.2.3 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.

Present Status: Under discussion with slow controls group. Provisions made in the trigger architecture either through shared memory access or special event type inserted in data stream.

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: Under discussion with EMC, TOF, CTB and VPD people.

3.2.7 Experiment Control

3.2.7.1 Commands

Requirement: Trigger shall establish a protocol with EC for handling commands such as CONFIGURE, START, STOP, ABORT, PAUSE, RESUME.

Justification: EC must be the overall experiment coordinator.

Status: Under discussion with EC. We expect the configure 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: EC 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 EC of its internal status upon request.

Justification: This allows EC 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.

Present 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 nsec ~ 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 microsecs, which is small compared to the normal 10ms 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 microsecs 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.

EMC Appendix:

June 21, 1995.

Summary Document of EMC trigger data flow

Bill Christie, BNL

This document is in response to a request from the STAR trigger group to supply a data flow document for the STAR Electromagnetic calorimeter. I've attempted to pull together information for this document from the STAR EMC Requirements Document (dated Aug. 10, 1994, and from various other sources.

Level 0

The EMC will provide the following data to the trigger hardware for the level 0 trigger decision:

* Single Tower Energy - There will be 3 bits provided on single tower energies. These 3 bits correspond to whether or not 3 different, programmable, single tower energy thresholds have been exceeded.

* 2x2 (4) Tower Sums - There will be 3 bits provided on the energy in 4 tower sums. These 4 tower sums will either be all possible contiguous 4 tower sums or else all non-overlapping sets of 4 tower sums. These 3 bits correspond to whether or not 3 different, programmable, 4 tower sum energy thresholds have been exceeded.

* 4 x 4 (16) - There will be 3 bits provided on the energy sum for all possible configurations of 2x2 sums making 4x4 tower sums. These 3 bits correspond to whether or not 3 different, programmable, 16 tower energy thresholds have been exceeded.

* Global "Et" - There will be 3 bits provided on the energy sum from all towers in the barrel EMC. The quantity Et has been put in quotes to denote the fact that this global sum is the sampled (up to ~ 6% max) energy deposited in the EMC due to both electromagnetic and hadronic incident energy. These 3 bits correspond to whether or not 3 different, programmable, Global energy thresholds have been exceeded.

* Absolute Global "Et" - A digitized Global "Et" value of 12 to 16 bits.

* Isospin patches - The digitized Et in patches of towers will be passed to the trigger at level 0. Whereas at L1 and above the size of these patches will be variable, this L0 information will most likely be for fixed tower sums. At this point our presumption is that "square bins" with lengths of something like 1 in eta by 1 in phi (patches of 100 towers in the barrel, ~ 90 towers in the endcap) are probably appropriate.

* Luminosity monitor - There will be ~ 3 bits passed to the trigger for use in luminosity monitoring. The idea is that these bits correspond to something like a fairly low (e.g. 3 GeV) high tower threshold. These bits would generate a 0 action word for the trigger and would be used to increment scalars.

Level 1

On a level 0 accept the EMC will ship the digitized pulse heights (ADCs) for all of the barrel and endcap phototubes to a memory location on the electronics carts. This will be done using one optical fiber link for every 20 phototubes (pmt) in the initial implementation and may go up to 60 phototubes/fiber if depth segmentation is eventually used (N.B. with the possible upgrade to 3 depth segments for each tower, one gets 60 ADC values for each 6 degree phi wedge in the EMC barrel). With the Barrel (1200 towers x 3 tubes/tower) and Endcap (540 towers x 3 tubes/tower) one will have a maximum of 87 fibers. At 25 ns per 16 bit ADC value this transfer will take a maximum of ~2 us (60 towers/fiber x 25ns = 1.5 us, 20 towers/fiber x 25ns = 0.5 us).

In the initial implementation of the electronics (without the longitudinal segmentation) and using 1 us (could be longer) for the level 0 decision, ~0.5 us for the data transfer, and ~1.5 us for various latencies, this data will be available for level 1 trigger processing in ~ 3 us. With the upgraded longitudinal segmentation the data would be available in ~ 4.0 us (1 us L0 + 1.5 us data xfer + 1.5 us for various latencies).

The SMAX data will also be passed off of the detector via Glink fibers on L0 accept. The data passed per fiber for the SMAX is ~ 4 times that passed for the EMC phototubes. As such this data will be available for processing off of the detector ~ 16 us after the event.

The total volume of EMC data passed off of the detector for each L0 accept will be ~ 10.44 kB for the towers (60 (pmt/fiber) x 87 (fibers) x 2 (B) = 10.44 kB) and ~ 40 kB for the SMAX.

It is during the L1 trigger processing that the finer bins can/will be used to look for significant isospin fluctuations. As all of the CTB/TOF data and all of the EMC data is available for calculations, one can look at the ratio of energy in the EMC to the charged particle multiplicity in whatever detail necessary.

With the EMC and CTB data one can also look for likely jet event candidates at L1. The characteristic that one would be looking for is a localized region of energy with an area (in eta-phi) of something like 1.5 units**2. The discriminating power of this calculation is yet to be determined.

One also has significant discriminating power at L1, using the EMC and CTB/TOF data to reject events which passed an electron or direct gamma trigger at L0. This is done by applying an isolation cut around the high tower that triggered the L0 trigger, where one uses either just the energy in the surrounding EMC towers or perhaps some combination of the energy in the surrounding towers and the multiplicity of charged tracks in the region of interest.

Between the time when these rather fast calculations are done at L1 and the time when tracks become available from the L3 tracking in the TPC, one could perform chi-squared cuts on the shape of the showers in the SMAX to select or reject events consistent with electrons or direct gammas.

SOME INFORMATION ON STAR EMC TRIGGER RATES
(mainly for high Lum. pp running)
and some assumptions.
D. Underwood
This is to go with the note about EMC trigger information
Written by Bill Christie June 21, 1995

process	level 0 trig.	good physics
W root(s)=500 pt > 20 GeV	4 Hz	.03 Hz
Gamma root(s)=200 pt > 5 GeV	~ 100 Hz	~ 1 Hz
Jet root(s)=500 pt(EM) > 10	> ~ 100 Hz	> ~ 100 Hz
Jet pt(EM) > 5	~ 2000 Hz	~ 2000 Hz (but keep 100 Hz high pt ones)

Note that the EMC only triggers on a fraction of the Jet energy.

HARDWARE LIMITS

input		output
10 ^ 7 Hz	LVL 0	10^4 Hz
	LVL 1	limit by grid reset, FPGA processing
	LVL 2	100 Hz gives 100% TPC readout deadtime
	LVL 3	computing
	Tape	1 to 100 Hz, depends on TPC data

ASSUMPTIONS

1) L0 has no deadtime
2) A L0 accept triggers readout of EMC data through Glink fibers into token addressable memory
3) there are 60 fibers reading the barrel EMC data and 60 more for barrel shower max and 24 reading the end cap EMC and 24 more for end cap shower max.
4) the L0 physics accept rate is 2000 Hz for EMC and 200 Hz for SM plus perhaps 10000K just EMC events, some of which have nul action word and are used to scale luminosity, and some of which go to tape.

(We assume that the 2000/sec includes the jet triggers along with the gamma and electron triggers.. We can save a lot of processing and data transfer by only reading in sh. max for gamma or electron triggers.)

5) For jet-jet we need the whole TPC, all sectors,
For gamma-jet we need the half opposite the gamma plus an area around the gamma to make isolation energy cuts.
For W we need isolation cuts and might try a crude Et. These are a small fraction of all triggers anyway.
For isospin we make better corrections for charged particles by using tracking for counting tracks per tower.
6) We don't yet know if we need a lot of the TPC raw data or just cluster information for the tracks found in level 3.
This depends on what momentum resolution is achievable without raw TPC information.

Requirement:
EMC requires that the VPD counters be in Level 0.
Justification:
The Luminosity monitor and Physics events must have the same vertex cuts in order to do normalization. A larger part of the diamond than can be accommodated by SVT is needed.

Requirement:
EMC shall make correlations with CTB/TOF at two stages of the trigger. a) in Level 0 using hit counts from CTB slats and EMC isospin patches (approximately 1 unit in eta by 1 unit in phi) b) in late Level 1 or in level 2 using analog multiplicity from CTB with EMC PMT data in programmable patch sizes.
Justification:
The best physics calculations possible at each trigger stage are made in order to reduce the rate into the next stage of the trigger.

Requirement:
It shall be possible to process built events to the following extent before writing to tape:
a) Corrections to EMC to get neutral energy when there are

charged tracks hitting the calorimeter must be made.
This involves both tracks within a sector and crossing sectors.
b) Energy based isolation cuts around direct gamma candidates
must be made using both neutral energy and charged tracks.
This involves both tracks within a sector and crossing sectors.
c) Jet energies must be calculated using neutral energy plus TPC tracks.
This involves both tracks within a sector and crossing sectors.
Justification:
The bandwidth to tape may be exceeded if we are not sufficiently selective in choosing quality events. We cannot do the physics if a high percentage of good events are lost due to the bandwidth to tape being saturated by non-useful events.
This is most critical in high Luminosity pp running, but may apply in other cases as well.

Requirement:
It shall be possible to histogram the results of calculations in level 1 and level 2 for event rates much higher than the Level 3 event rate.
Justification:
The EMC physics goal here is to look at the correlation between the number of charged particles and the associated Et. (I discuss first integrating over the full acceptance of the CTB/EMC - $|eta|<1$. Note that this discussion is NOT directed at 'jets', rather at global, 'soft/thermal' physics.) There are some simple arguments that the mean Pt/particle should change as one crosses phase-transition boundaries (eg hadrons -> mixed -> plasma). E(EMC)/N(CTB) gives a reasonable approximation to this. What we would like to do is to plot this ratio on an event-by-event basis to see if there are slope changes, or changes in the width of the distribution as E(EMC) changes.

Because of the low energy of most of the charged particles that enter the calorimeter (low for usual uses of a sampling calorimeter), the STAR EMC response for these kinds of events has significant contributions from both charged particles and photons/electrons. As Bill Christie has shown in simulations (SN???) this behavior and the large statistical samples available in a single event can be used with the CTB multiplicity to make a statistical separation of the observed 'energy/pulse height' into its charged and electromagnetic components. This has led us to suggest that the ratio E(EMC)/N(CTB) is also sensitive to possible variations in the neutral/charged pion production. This has also been simulated by Christie (SN???).

We note that both the EMC and the CTB are 'fast' devices (response times in the 10 nsec range), and that in total the channel count is less than 2000 (1200 EMC; ~400 CTB). A single event consists of < 4KB. The total interaction rate for AuAu events is expected to be about 5K/sec, and the 'central collision' rate is usually taken to be about 10% of that, or 500/sec. We would like to record all of the central collisions in our histogram and a roughly equal number of the non-central events for a total rate of about 1K/sec.

The bandwidth requirement is thus about 4MB/sec. Note that this bandwidth is in some sense 'internal' in that it does not go to tape. We need to be able to examine the histogram as it is accumulating (from some workstation), and to write it out to some tape at an interval of an hour or so. (It could be written to the data tape with other end-of-run data, for example.)

How to justify this in terms of statistics? Note that the resolutions in E(EMC) and in N(CTB) for central collisions are approximately equal at a sigma of 3%. Using a rule of thumb that one should have 5 bins or more per full width of a peak, we find a bin width in the vicinity of the CC data of about 0.7%. Thus we require about 200 bins in each of two dimensions - 4 10**4 in total. Not all of these bins are occupied; let's assume ~10% are. Then we have

(10**3/sec)/(4 10**3) = 0.25/sec/bin = 900/bin/hr = 20000/bin/day

or a result with statistical error of just under 1%/bin/day.

This only provides an estimate; it does not account for the gaussian shape of the distribution, nor of the need to probe the tails of the distribution. That must be left to more simulations.

It should be strongly emphasized that this kind of study will permit us to rapidly (days?) determine interesting thresholds on the EMC/CTB ratio which can then be used to select events with full detector readout. Without this mechanism, we would have to wait for acquisition of a similar sample from the TPC data stream, 1000 times as long.

This does not pretend to be a complete document. The connections with trigger and DAQ need to be worked out. (Action words, busies, prescales at level-0, etc)