

The STAR Level-3 Trigger System

C. Adler^a, J. Berger^a, M. Demello^b, T. Dietel^a, D. Flierl^a,
J. Landgraf^c, J. S. Lange^{a,1}, M. J. LeVine^c, A. Ljubicic, Jr.^c,
J. Nelson^d, D. Roehrich^e, R. Stock^a, C. Struck^a, P. Yepes^b

^a*University of Frankfurt, August-Euler-Straße 6, D-60486 Frankfurt, Germany*

^b*Rice University, Houston, Texas 77251, USA*

^c*Brookhaven National Laboratory, Upton, New York 11973, USA*

^d*University of Birmingham, Birmingham B15 2TT, United Kingdom*

^e*University of Bergen, Allegaten 55, 5007 Bergen, Norway*

Abstract

The STAR Level-3 trigger issues a trigger decision upon a complete online reconstruction of Au + Au collisions at RHIC energies. Central interactions are processed up to a rate of 50 s^{-1} including a simple analysis of physics observables. The setup of the processor farm and the event reconstruction are described as well as experiences and the proposed trigger algorithms.

Key words: Trigger; Data Acquisition; Nucleus-Nucleus Collisions; Proton-Proton Collisions

1 Introduction

The Solenoidal Tracker at RHIC (STAR) is a large acceptance detector for the study of hadronic collisions at the Relativistic Heavy Ion Collider (RHIC). The baseline detector operating in the first year of data taking (2000) consisted of a large Time Projection Chamber (TPC)(1), a Central Trigger Barrel (CTB)(2) and two Zero Degree Calorimeters (ZDC)(3). A large conventional solenoidal magnet, with a magnetic field of up to 0.5 Tesla, surrounds the experiment. The TPC is a cylindrical detector with a 4 m long axis and an inner(outer) radius of 0.5(2) m. Its drift volume is divided into two halves along the beam

¹ Corresponding Author, Email soeren@bnl.gov

direction, each half with a drift direction away from the center. Its pseudorapidity coverage is $|\eta| < 1.5$. A track traversing its full radial distance ($|\eta| < 1$) generates 45 hits in the detector. The CTB is a cylinder of 240 scintillator tiles placed around the TPC, covering the pseudorapidity region $|\eta| < 1$. It provides a fast measurement of particle multiplicity, used in the trigger. The ZDCs, placed at 18 m from the interaction point, provide information of energy flow in the direction of the incoming beams.

In addition to the baseline detectors, STAR is commissioning the following detectors for the year 2001 run: a Silicon Vertex Tracker (SVT) (4), two Forward Time Projection Chambers (FTPC) (5), an Electromagnetic Calorimeter (EMC) (6) and the endcaps of the main TPC, used as Multi-Wire Proportional Chamber (MWC) (2) for triggering, and a small Time of Flight detector (pTOF). The SVT is a cylindrical silicon drift detector placed inside the TPC, providing a good vertex resolution. The FTPCs provide tracking at $2.5 < |\eta| < 4$. The data volume produced by the drift detectors is large and therefore limits the amount of events, which can be written to tape.

RHIC can accelerate a wide range of ions, up to Au, at 100 GeV/A, and polarized protons beams at 250 GeV. So far the accelerator has only collided Au ions. For Au + Au collisions the interaction rate at RHIC is expected to reach 2.4 kHz at design luminosity. Given the data size of central events (≈ 10 MBytes), STAR can store only ≈ 2 events/s. Therefore the trigger system should select 0.1% of the interactions for further offline analysis. For minimum bias, ultra-peripheral and p + p collisions the events are smaller and can be written to tape at a higher rate of up to 30 s^{-1} . There are two main stages in the trigger system. The first stage uses fast detectors and is constituted by trigger levels 0,1 and 2 (L0, L1, L2). These fast detectors are read out on every RHIC beam crossing (107 ns). They provide information on multiplicity or global energy produced in the collision. This information is used to select events with a specific collision geometry. For example, central events (with a small impact parameter) are selected by requiring a large multiplicity or global energy. The second stage is constituted by the third level trigger (Level-3). This level collects data from tracking detectors, reconstructs tracks and makes a decision based on that information. Tracking allows a more powerful event selection based on information of single particles, rather than global characteristics of the event. Examples of Level-3 triggers are selection of J/ψ and antinuclei, and search of back-to-back tracks in ultra-peripheral events. The third level trigger has been implemented as a multiprocessor farm interconnected by a fast network (cf. Sections 2,5).

In the 2000 run, only the TPC data was accessible by the Level-3 trigger system. In 2001 the SVT and FTPC contributed cluster data. (cf. Section 10).

	τ
L0	1.5 μs
L1	100 μs
L2	5 ms
L3	200 ms

Table 1

Trigger processing times τ for the different trigger levels. After the given time a trigger level has to issue a decision. For Level-3, a processing time of 200 ms correspond to a rate of $\approx 50 \text{ s}^{-1}$ in the current setup.

2 Architecture

2.1 The STAR Trigger Architecture

The STAR trigger system is subdivided into 4 hierarchic levels. L0, L1 and L2 use fast detectors, e.g. the CTB or the ZDC. Level-3 is a software trigger which uses data of the slower detectors TPC, SVT and FTTPC. The trigger processing times (2) for the different levels are given in Table 1.

L0 processes the information available every bunch crossing and is responsible for delivering a signal to the slow detectors (TPC, SVT, FTTPC) to start the readout. L1 can abort an event during the readout (40 μs for the TPC). After the readout, the data is digitized on the detector front end electronics, which takes ≈ 10 ms for the TPC. During this digitizing time L2 can abort an event before the data is transmitted to the Data Acquisition system (DAQ).

The 10 ms digitizing time limits the TPC readout rate to 100 s^{-1} . The final DAQ rate for central collisions is $\approx 2 \text{ s}^{-1}$ corresponding to the maximum bandwidth from STAR to the tape robot at the RHIC computing facility of 20 MB/s. Instead of storing a random sample of the available collisions, Level-3 can analyze all the data (up to a rate of 50 s^{-1}) and use this additional information to decide which events should be stored.

2.2 The STAR DAQ Architecture

The Level-3 trigger architecture is integrated into the STAR DAQ architecture (7), in which one VME DAQ readout crate is mapped onto two physical TPC sectors or one FTTPC side or one SVT half-shell, respectively. Each DAQ crate

contains a *Detector Broker* (DET), a Motorola MVME-2306 VME board, carrying a PowerPC 604 (300 MHz, VxWorks), as the master controller. The DET carries a Myrinet(14) (cf. Section 5) interface to transfer raw data to the event builder (EVB) and connect to the Level-3 track finder CPUs. Each DAQ crate also contains 6-12 VME receiver boards (depending on detector) carrying three mezzanine cards with:

- ASICs performing zero-suppression on the raw data (cf. Section 3). This reduces the data size by a factor of 10.
- 4 MB of dual-ported VRAM for buffering and pipelining the zero-suppressed data of 12 events.
- One Intel i960 CPU (33 MHz, VxWorks) for data formatting and running the Level-3 cluster finder.

The DAQ message communication is steered by the *Global Broker* (GB), a Myrinet connected MVME-2306 board. The GB announces incoming data to the Level-3 system.

If Level-3 accepts an event, control is given to the Eventbuilder (EVB) which requests and assembles the data contributions of all available subdetectors. The EVB is a process running on a Myrinet connected SUN Enterprise 450 Server which also acts as a buffer for the data before it is transferred to the RHIC Computing Facility for storage.

2.3 The STAR Level-3 Trigger Architecture

The Level-3 trigger scheme consists of two main components:

- The *Sector Level-3* part (SL3) is mapped onto one DAQ VME readout crate. It contains the Level-3 cluster finder (cf. Section 3) and the Level-3 track finder (cf. Section 4). Cluster and track data are transferred via Myrinet. Typical data sizes for a central event are 10 MB (all STAR subdetectors, zero suppressed), 2 MB after cluster finding and 0.3 MB after track finding.
- The *Global Level-3* part (GL3) consists of 3 master CPUs, collecting all track data and making the Level-3 decision.

In the 2000 run, only TPC data were accessible to the Level-3 trigger, during the 2001 run, SVT and FTPC contribute cluster data and the trigger detector data (CTB, ZDC, MWC) is available. Furthermore, Level-3 will receive EMC data. External vertex information from different subdetectors will also be utilized to improve tracking performance (cf. section 9).

Figure 1 shows the schematic Level-3 trigger architecture, as of June 2001. In the following the main emphasis will be on the description of Level-3 processing

of STAR TPC data. Additional STAR subdetectors contributing to Level-3 will be covered in Section 10.

3 Cluster Finder

The TPC consists of 24 sectors. Each sector has 45 padrows in radial direction with each between 88 and 182 pads for a total of 5960 pads. When the TPC is read out, the accumulated charge seen by each pad is sampled every 107 ns (=1 time bin) and converted to an ADC value. There is a maximum of 512 time bins. Only ≈ 380 time bins contain valid data. While the padrow and pad define the position of a pixel in the plane perpendicular to the beam, the time bin provides information about the position along the beam.

The TPC raw data transmitted to the DAQ system consists of the ADC values of all time bins of all pads. The zero-suppression scans the ADC values for each pad for contiguous sequences of values above a certain threshold. Only sequences with a certain length are stored.

Clusters are contiguous regions in the pad-time plane for a given padrow, with *ADC* values above threshold. For a typical central Au + Au collision at $\sqrt{s_{NN}}=200$ GeV $\approx 140,000$ clusters are reconstructed.

Algorithm. The weighted mean of the ADC sequences in drift direction is calculated for a pad. The position of the center of gravity is then compared to the stored values of the previous pad. If a matching entry is found, the sequences get connected and the center of gravity in the pad direction is calculated. Only clusters with a certain minimum charge (ADC sum) and extending over more than one pad are accepted. Figure 2 shows ADC raw data and reconstructed centers of gravity for 2 different padrows in one sector.

Deconvolution. If two tracks are close enough, the generated clusters may overlap. The cluster finding algorithm provides a deconvolution step, which attempts to separate clusters produced by close tracks. This deconvolution is applied in both time and pad directions, by splitting a cluster if a local minimum in the charge distribution is found. Reconstruction of two overlapping clusters with deconvolution consumes about 6 % more CPU time than the case of two separated clusters.

Resolution. The position resolution is determined by embedding Monte-Carlo (MC) generated clusters into recorded Au + Au events. This MC simulation of the TPC includes a microscopic simulation of the electron cloud drift. The achieved values of $\sigma_{r\phi} \approx 830 \mu\text{m}$ (pad direction) and $\sigma_z \approx 1130 \mu\text{m}$ (time bin direction) are 3% worse than the STAR offline cluster finder resolution.

Efficiency. The performance of the cluster finder depends on event multiplicity and noise. The cluster finding efficiency, determined by embedding MC simulated particle tracks into real events, was found to be $\varepsilon \geq 80\%$. For a typical central Au + Au event at $\sqrt{s_{NN}}=130$ GeV, 15-20% of reconstructed clusters are not assigned to a track.

Hardware. The cluster finder algorithm runs on Intel i960 CPUs on the DAQ receiver boards (cf. Section 2.2). This parallelizes the cluster finder task onto 432 CPUs for the whole TPC (18 CPUs per sector). Input to the cluster finder are zero-suppressed TPC raw data, stored in the VRAM buffer. The output cluster data, i.e. cluster center of gravity and cluster total charge (ADC sum), are sent via VME to the *Detector Broker*, which ships the data via Myrinet to a Level-3 track finder CPU.

Timing. Figure 3 (*left*) shows the cluster finder CPU timing results as a function of number of clusters per i960 for the 15% most central Au + Au events. The average cluster finding time is $\tau_{cluster} \simeq 15$ ms. Figure 3 (*right*) shows the CPU time vs. i960 index in one DAQ crate for Au + Au events. The solid line gives the number of clusters per i960, assuming that each padrow has the same number of clusters (arbitrarily scaled). Due to the mapping of the i960 CPUs onto a sector, the number of clusters per i960 is not constant and the CPU time per i960 can vary by as much as a factor of 2.

4 Track Finder

Algorithm. The track finder algorithm (8) employs *conformal mapping*, i.e. a transformation of a circle² into a straight line, followed by next-hit-on-track finding using a *follow-your-nose* method. A given space point (x, y) is transformed into a conformal space point (x', y') according to the equations

$$x' = (x - x_0)/r^2 \tag{1}$$

$$y' = (y - y_0)/r^2, \tag{2}$$

using $r^2 = (x - x_0)^2 + (y - y_0)^2$. This transformation requires the knowledge of a point (x_0, y_0) on the track trajectory, either the interaction point (vertex constraint for primary tracks) or the first point associated with the track (no vertex constraint for secondary tracks). Once a track is found, it can be refit

² In the STAR solenoid magnetic field, charged particle tracks can be parametrized as helices, being circles in an xy -projection onto the TPC bending plane perpendicular to the z -direction of the magnetic field.

with a helix model in real space. In the current version, the Level-3 track finder operates with a vertex constraint in conformal space.

Track finder parameters (e.g. assumed hit errors) were optimized with the 2000 data in order to improve the track finding efficiency and p_T resolution.

Transverse Momentum Resolution

The transverse momentum resolution reflects the quality of the event reconstruction. It is also crucial for most of the Level-3 trigger algorithms.

The resolution was determined by embedding MC simulated particle tracks into real Au+Au events. A MC track was considered found, if a track was reconstructed having at least 23 hits (out of a maximum of 45 hits) from a MC track assigned to it.

The transverse momentum is determined by fitting a circle to the projection of the hits onto the bending plane. Including the the x-y position of the event vertex into the fit improves the resolution significantly. However if the assumed x-y vertex position is not accurate, the resolution will degrade considerably. The beam position provides sufficient precision to be used as x-y-vertex of the events. The uncertainty in the vertex is given by the transverse size of the beam which is in the order of $\sigma_{xy}=0.5$ mm.

Figure 4 shows the transverse momentum resolution, $\Delta p_T/p_T$ for Level-3 with and without vertex constraint, and for the STAR offline for the same set of tracks. For $p_T=5$ GeV/c particles $\Delta p_T/p_T \approx 10\%$ for Level-3 with vertex constraint, whereas the offline analysis chain reaches a resolution of 9%.

Efficiency

The reconstruction efficiency ϵ , is an essential component in determining the absolute abundance of a certain event signature. Using the embedding method, ϵ is defined as the ratio of the number of matched tracks over the number of embedded tracks. Again a track is considered matched if it shares at least 23 hits with an embedded track.

As explained in section 2.3, Level-3 performs local tracking in each of the TPC sectors. Therefore tracks crossing sector boundaries will be split. The central membrane of the TPC, dividing the TPC in two halves in the longitudinal direction is especially important for this matter. The reason is that the width of the distribution of the longitudinal vertex position σ_z is large. In the year 2000 with $\sigma_z \approx 90$ cm it was far above the design value of ≈ 25 cm. The further away the event vertex is from the middle of the TPC, the more tracks are split between the two halves of the TPC.

Figure 5 shows the efficiency versus the transverse momentum for the Level-3 and offline tracking. Results for Level-3 are presented for all events, and for events with a vertex position less than 10 cm away from the center of the TPC. For the latter case the efficiency is above 80% and only slightly worse than the offline reconstruction.

Hardware. The current system consists of 36 Compaq DS-10 (21264 ALPHA CPU(9)) 466 MHz and 12 DS-10 (21264 ALPHA CPU) 600 MHz machines running Linux (2.2.x kernel). At the time of the first Level-3 CPU purchase in April 1999, the ALPHA processor was the only processor fulfilling the timing requirements for track finding of ≤ 100 ms (track finding on an ALPHA DS-10 466 MHz was 1.6 times faster than on a Pentium III 450 MHz).

Timing. Figure 6 shows the track finding CPU time vs. the number of tracks. The corresponding Au + Au event were triggered by L0 as 15% most central. For these events the Level-3 track finding requires on average 70 ms (per 2 TPC sectors), using the 466 MHz ALPHA processors. Approximately the same time is used by a 800 MHz Pentium processor, while a 600 MHz Pentium takes $\approx 30\%$ longer. With 48 installed track finder CPUs (SL3) in year 2001, the maximum processing rate $R=12-14$ s $^{-1}$ can be quadrupled to $\approx 50-60$ s $^{-1}$.

To identify a particle species, the specific energy loss dE/dx in the detector gas is calculated. This consumes an additional overall $\Delta\tau \leq 1.5$ ms.

For data taking in 2001, coordinate corrections will be applied online to all hits. They correct $E \times B$ field distortions as well as geometrical misalignments. The corrections are determined offline (1) and coded as a look-up table in the transformation from TPC native coordinates (padrow, pad, and drift time bin) to Cartesian coordinates. This improvement does not consume additional CPU time.

For future applications of matching TPC tracks to other subdetectors (cf. Section 10) routines for helix extrapolation are prepared ($\tau \approx 4$ ms for $N_{track}=1000$).

Software Environment. The track finder runs as part of a process which is also responsible for the message handling and data transfer. It is implemented as a multi-threaded program with 4 threads. One thread interacts with the Myrinet device driver putting received messages into a receive queue and buffering outgoing messages. Another thread does the actual message handling (interpretation, invocation of the appropriate action, and responding). The third thread consists mostly of the track finder itself, and is invoked as soon as there is cluster data to process. A fourth thread controls the other threads with timeout timers and communicates with the run control. The main emphasis here is on stability (since a single point failure in the DAQ system can stall the whole system) and error tolerance (occasional communication errors or hardware problems, should be handled gracefully so that data taking

can continue).

The multiple thread design was chosen to enable the different tasks to run quasi-parallel (e.g. to enable the message handling thread to receive a message or data while the track finder thread is busy with an event). This could also be implemented as multiple processes, using shared memory for data. However Direct Memory Access (DMA) to/from shared memory was not possible with earlier versions of the Myrinet driver software.

The same basic software structure was also used for the other parts of the Level-3 system, such as GL3.

Dynamic data distribution scheme. In the 2000 run, the mapping between an SL3 node and a Detector Broker was fixed one-to-one. With 48 SL3 CPUs serving 14-16 Detector Brokers in year 2001, a set of SL3 CPUs will be determined dynamically by the Global Broker for every event. Initially the track finder CPUs will be selected in a sequential fashion. In the future, information about processing times and CPU loads will be used for load balancing.

5 Network

The Level-3 trigger system requires a high bandwidth network connection between PMC adapters (on the sector broker VME boards) and PCI adapters (in the SL3 and GL3 nodes). In an early stage Myrinet(10) and SCI(11) were considered (12)(13). Myrinet was finally chosen due to stability issues. Gigabit Ethernet was not considered due to its lack of DMA capabilities³. Myrinet is a high-performance, low-latency commercial network build by Myricom(14). It uses a switch based topology and 1.28 Gbit/s full duplex links.

Data transfers. Currently 32 bit and 64 bit Myrinet PCI and PMC cards are used. Measured values for point-to-point bandwidth are $\approx 70(120)$ MB/s between two 32(64) bit adapters. For a system of 12 SL3 CPUs with 32 bit adapters sending to one GL3 CPU equipped with a 32(64) bit PCI adapter, values of $\approx 74(90)$ Mb/s were measured. Myrinet latency⁴ is in the order of $L \leq 50(15)$ μ s for 32(64) bit cards. There are two main steps in the data transfer in Level-3 which are potentially time consuming:

- Transfer of cluster data from the DET to the SL3 CPU's. The typical size

³ Background DMA data transfers do not consume any CPU time (except for bus arbitration and interrupt handling). Thus, the processor is free for tasks like track finding or data formatting.

⁴ The Latency L was measured here as half the message round-trip time for the limit of zero message size.

of the TPC cluster data per SL3 CPU for a central Au + Au collision is 100-150 kB, leading to a transfer time of $\tau \leq 2$ ms.

- Transfer of track data from the SL3 CPUs to the GL3 CPU. The typical size of the TPC track data for a central Au+Au collision is 0.3 MB, leading to a transfer time $\tau \leq 4$ ms.

Thus, overall data transfers only require $\leq 5\%$ of the total Level-3 processing time.

Message transfers were determined to require a negligible amount of time, i.e. $t_{message} \leq 2$ ms for event processing with zero data size.

6 Global Level-3 Trigger

The GL3 CPUs perform track data collection from all SL3 machines, followed by a Level-3 trigger algorithm based on event characteristics and finally issue the Level-3 yes/no decision to the event builder. Up to 32 decision algorithms can be run simultaneously, the yes/no decision being issued as logical OR.

For later cross section determination, GL3 keeps scalers with the number of processed and triggered events and the number of decisions for each algorithm.

Currently, three different machines/platforms are used as GL3 CPUs: a Pentium III 600 MHz, a dual Pentium III 800 MHz and a Compaq XP1000 (21264 ALPHA) 500 MHz. Multiple events can be processed at the same time by different GL3 CPUs. The contributions of a single event are analyzed by one CPU only.

Algorithms run in a C++ framework, which provides tools for accessing track information easily. The identical code runs in the STAR offline framework. This allows the analysis of already recorded or simulated data for testing and performance studies.

Algorithms in year 2000. Two Level-3 trigger algorithms were run successfully in the year 2000 RHIC run.

- *z-vertex trigger for central Au + Au events.* At the STAR interaction point, the collision z vertex in the 2000 run had a wide spread of $\sigma_z \approx 90$ cm. Most of the non-centered events were not useful for physics analysis. Thus, the Level-3 trigger selected events with $|z_{vertex}| \leq 75$ cm, corresponding to a rejection factor of 2-3. The Level-3 trigger was switched on only during the beginning of the accelerator stores, when L0 trigger rates for central events were too high to save each event. The Level-3 trigger was operating for about 3 days with an input (output) rate of 8-12(3-4) s^{-1} . During that time about

100k central Au + Au events were recorded. The achieved vertex resolution for central Au + Au events was $\sigma_z=1.3$ cm and $\sigma_{x,y}=0.2$ cm. Figure 8 shows vertex distributions (*top left*).

- *Ultra-peripheral Au + Au collisions* (15). $\gamma + \gamma$ and $\gamma +$ pomeron interactions, as induced in very peripheral Au + Au collisions, produce only a few tracks in the detector ($N_{track}=2-6$). However, abundant low multiplicity event background is also present, to be rejected by a Level-3 trigger algorithm. Cosmic ray background was rejected by requiring the transverse position of the found vertex to be within 15 cm of the beam axis. Beam-gas and beam-beampipe interactions were reduced by selecting events with a longitudinal position of the vertex within the limits of the TPC.

Proposed algorithms for year 2001. For the 2001 run the following Level-3 algorithms have been proposed:

- heavy anti-fragments (16), e.g. ${}^4\overline{He}$. Trigger on negatively charged particles with track parameters and dE/dx information indicating a particle charge of $Z < -1$.
- high $p_T (\geq 3 \text{ GeV}/c)$ particles traversing the STAR RICH detector. For enriching these events, the GL3 extrapolates TPC track helices onto the RICH $60 \times 60 \text{ cm}^2$ plane at a radius of $R=2.40$ m.
- Invariant mass for $J/\psi \rightarrow e^+e^-$ using partially installed EMC and $\Upsilon \rightarrow e^+e^-$ (using a high $p > 3 \text{ GeV}/c$ candidate pre-cut). The suppression of quarkonia is commonly regarded as a promising signature of the quark-gluon plasma. Figure 7 shows the GL3 CPU time as a function of the number of tracks for Au + Au events when only tracks with $p_T > 1.5 \text{ GeV}/c$ and $N_{hit} \geq 25$ hits are used. Due to track-track pair looping the dependence is quadratic. The typical number of candidate tracks for $\Upsilon \rightarrow e^+e^-$ is $N_{track} \leq 20$, thus the required CPU time $\tau_{global} \leq 2$ ms.

Additionally, it has been proposed to use Level-3 for EMC calibration (17).

For $p+p$ collisions Level-3 will be used for data reduction (18). At the highest Luminosity ($\mathcal{L}=2 \cdot 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$) up to 700 p+p collisions occur within the readout time of the TPC. Collisions happening before a valid trigger leave tracks in the TPC that drift to the readout planes before the triggered collision happens. In the same manner collisions happening after the valid trigger leave tracks in the TPC that are also read out. This results in an estimated number of 7000 tracks in the TPC for one event that only produced ≈ 10 tracks. The plan is to identify the vertex of the triggered event, select those tracks that point to this vertex and only write out the clusters belonging to those tracks.

7 3D Event Display

In order to visualize and browse a large number of events quickly, a fast 3D event display was developed. It allows the display of 130,000 clusters and 5000 tracks (central Au+Au collision) as well as mouse controlled operations like rotating and zooming in quasi-realtime without significant delay. Figure 9 shows a central Au+Au collision in the event display with TPC, FTPCs and SVT.

Data is received from the Level-3 sub-event builder CPU (l3evp, cf. Figure 1, *top*) via Ethernet. The l3evp CPU itself gets the cluster and track data from the GL3 CPUs via Myrinet. The C++ event visualization program is based on the *OpenGL* graphics library. The graphical user interface has been designed using the *Qt* widget⁵ library (19). The events are displayed in the STAR main control room on a SONY 42" flat panel plasma screen. A sample of all events passing through DAQ is displayed a few seconds after being triggered.

8 Control and Monitoring

Run Control. The Level-3 trigger system is fully integrated into the STAR run control system. This enables the STAR shift crew to control Level-3 by entering parameters in a Graphical User Interface (GUI). Those parameters can be technical ones, e.g. the prescale for a specific algorithm, or parameters defining an algorithm (e.g. p_T cut).

Performance monitor. The current L3 input and processing rates for each SL3 and GL3 CPU are displayed on a terminal window, along with other parameters like CPU load or memory usage.

Online monitor. Level-3 reconstructed events are used by an online histogramming tool, in order to display event characteristics (e.g. event vertex), physics observables (e.g. p_T distribution) and track finding quality observables (e.g. N_{hit} on track) on a event-by-event basis. Figure 8 shows the online display for a typical Au + Au event ($\sqrt{s_{NN}}=130$ GeV). During data taking, the system is used in the STAR control room for data quality assurance.

⁵ A widget library provides graphical elements for the design of a Graphical User Interface.

9 Vertex Finding

Two different types of vertex finding can be distinguished:

- Vertex finding as Level-3 trigger algorithm (e.g. enriching TPC centered events) is described in Section 6.
- Vertex finding as input for the Level-3 track finder, using a given vertex as origin for generating the η/φ -slices (search slices for next hit-on-track).

An input vertex for the Level-3 track finder can be gained by three different methods (ordered according to improving z-vertex resolution):

- (1) For the 2001 run, STAR uses timing information from the ZDCs to determine the longitudinal (z) position of the interaction vertex. This information is sent to Level-3 together with the trigger detector data. The vertex resolution is $\sigma_z \approx 10$ cm without slewing⁶ correction and ≈ 5 cm with slewing correction.
- (2) The Level-3 trigger will receive data from the VPD (vertex position detector). The expected timing resolution is $\Delta t \approx 50$ ps, providing a vertex resolution $\sigma_z \approx 2$ cm.
- (3) SVT clusters can be used for z vertex determination without track finding. A proposed topological algorithm (20) takes an a priori assumed vertex (e.g. $z=0$) and iteratively updates the z vertex by optimizing the number of “triplets” (a triplet being defined as 3 hits in the 3 SVT hits, pointing to the new vertex with the same straight line). The estimated position resolution for simulated central Au + Au events (no noise) is $\sigma_z \leq 300$ μm . A GL3 algorithm is under development, the required CPU time is estimated to be $t_{SVT} \leq 10$ ms.

10 Integration of additional subdetectors

In the 2000 run, only TPC data were used in the Level-3 trigger. In 2001 three more subdetectors are integrated.

Silicon Vertex Tracker (SVT)

The SVT is a drift detector with 103,680 anode readouts, each of them sampled in 128 time bins. Cluster finding is performed on 72 i960 CPUs. The data is sent by two sector brokers. The system was tested in April 2001 for cosmic events. Track finding was not performed, but a SVT geometry coordinate

⁶ Slewing is a time error due to the pulse height dependence of a discriminated time signal.

transformation was applied and the SVT hits were displayed on the Level-3 event display (cf. Section 7). For the test, cluster finding was performed with the same algorithm as for the TPC. However, SVT clusters have a bipolar shape in time bin direction (negative after-pulse, $\sim 5\%$ of the total amplitude), to be accounted for in a future stage.

The long term goal is the combined track finding of TPC hits (resolution $\sigma_z \simeq 1100 \mu\text{m}$) and SVT hits (resolution $\sigma_z \simeq 20 \mu\text{m}$). Thus, assuming an accurate TPC/SVT alignment, the achievable vertex resolution is estimated to be $\sigma \simeq 100 \mu\text{m}$. This improved z-vertex resolution can be used to increase the efficiency of the p+p pile-up filter (21) (cf. Section 6).

Forward TPC (FTPC).

The two STAR FTPCs each have $2.5 \cdot 10^6$ pixels ($10 \text{ padrows} \times 960 \text{ pads} \times 256 \text{ time bins}$). Simulations predict that for a central Au + Au collision, $\sim 30\%$ of all generated particles are in the FTPC acceptance. The typical occupancy is $\leq 30\%$ (inner padrows) and $\leq 4\%$ (outer padrows). As the TPC ($0 \leq |\eta| \leq 2.0$) and FTPCs ($2.5 \leq |\eta| \leq 4.0$) do not overlap in pseudorapidity no combined track finding is possible.

EMC (Electromagnetic Calorimeter).

For the year 2001 data taking, a partial STAR EMC with a coverage of $\Delta\varphi = 144^\circ$, $0 \leq \eta \leq 1$ is being installed. For the proposed Level-3 applications, the EMC is important for

- lepton identification for triggering on $J/\Psi \rightarrow e^+e^-$, and
- filtering of p+p pile-up events (18) (cf. Section 6).

In both cases a GL3 algorithm will match the TPC tracks to EMC hits.

Moreover, an EMC calibration on Level-3 is proposed (17). The algorithm selects high $p_T (\geq 1 \text{ GeV}/c)$ tracks and extrapolates them to the EMC. Most of the time those tracks will leave the signal of a minimum ionizing particle (MIP) in the corresponding tower. The position of the MIP peak in the energy spectrum of the tower is then used to calibrate the detector.

11 Summary

During the year 2000 run the STAR Level-3 trigger system has shown that real time event reconstruction of high multiplicity Au+Au collisions is possible. Its filtering capabilities were used for topological applications (z-vertex trigger) as well as a physics selection (enrichment of ρ candidates in ultra peripheral

Au+Au collisions). It also proved to be a useful and intuitive tool for online QA monitoring, by displaying reconstructed data taken with the TPC in realtime.

New detectors are incorporated into the system in the year 2001 run. The EMC information combined with tracking is expected to allow the identification of electromagnetic particles at Level-3. The system will provide visual monitoring of their performance during data taking.

The additional processors enable the system to run at a rate of up to 50 s^{-1} for central Au-Au collisions in 2001. This takes into account the increased beam energy which results in an increased particle multiplicity. Compared to the taping rate of $2\text{-}4 \text{ s}^{-1}$, this enables Level-3 to scan a factor of ≈ 20 more events than can be written to tape, thus improving the ability of STAR to measure rare probes like Υ or ${}^4\overline{He}$.

12 Acknowledgments

This work was supported by the *Bundesministerium für Bildung und Forschung*, Germany (contract #06OF 840 I) and the U.S. Department of Energy (contract #DOE DE-FG-93ER40772).

References

- [1] The STAR Time Projection Chamber.
- [2] The STAR Trigger System.
- [3] C. Adler *et al.*, “The RHIC Zero Degree Calorimeter”, Nucl. Instr. Meth. **A461** (2001) 337.
- [4] The STAR Silicon Vertex Tracker.
- [5] The STAR Forward Time Projection Chamber.
- [6] The STAR Electro-Magnetic Calorimeter.
- [7] A. Ljubicic, Jr. *et al.*, “The STAR Experiment’s Data Acquisition System”, IEEE Trans. Nucl. Sci. **47** (2000) 99.
- [8] P. Yepes, “A Fast Track Pattern Recognition”, Nucl. Instrum. Meth. **A380** (1996) 582.
- [9] <http://www.compaq.com>.
- [10] J. Landgraf *et al.*, “The Implementation of the STAR Data Acquisition System using a Myrinet Network”, IEEE Trans. Nucl. Sci. **48** 3 (2001).
- [11] SCI (Scalable Coherent Interface) Standard Compliant, ANSI/IEEE 1596-1992.
- [12] J. S. Lange *et al.*, “The STAR Level-3 Trigger System”, Nucl. Instr. Meth. **A453** (2000) 397.
- [13] J. Berger *et al.*, “The Proposed Level-3 Trigger System for STAR”, IEEE Trans. Nucl. Sci. **47** (2000) 358.
- [14] <http://www.myri.com>.
- [15] S. Klein, for the STAR Collaboration, “Observation of $Au + Au \rightarrow Au + Au + \rho^0$ and $Au + Au \rightarrow Au^* + Au^* + \rho^0$ ”, nucl-ex/0104016. To appear in Proc. 17th Winter Workshop on Nuclear Dynamics, Park City, UT, Mar 10-17, 2001.
- [16] C. Adler *et al.*, “ \bar{d} and ${}^3\overline{He}$ production in $\sqrt{s_{NN}}=130$ GeV Au + Au collisions”, nucl-ex/0108022. To appear in Phys.Rev. Lett.
- [17] T. M. Cormier *et al.*, “STAR Barrel Electromagnetic Calorimeter Absolute Calibration using *Minimum Ionizing Particles* from Collisions at RHIC”, hep-ex/0107081. To appear in Nucl. Instrum. Meth. A.
- [18] W. B. Christie, “Associating the correct tracks to the trigger Event in STAR for high luminosity proton on proton running”, STAR Note 169.
- [19] <http://www.trolltech.com>.
- [20] D. J. Prindle, “SVT Event Reconstruction Algorithms”, STAR Note 153 (1994).
- [21] V. L. Rykov, “STAR SVT in High Luminosity pp Collisions at RHIC (How Useful Is the SVT for Doing Spin Physics in STAR)”, STAR Note 237 (1996).

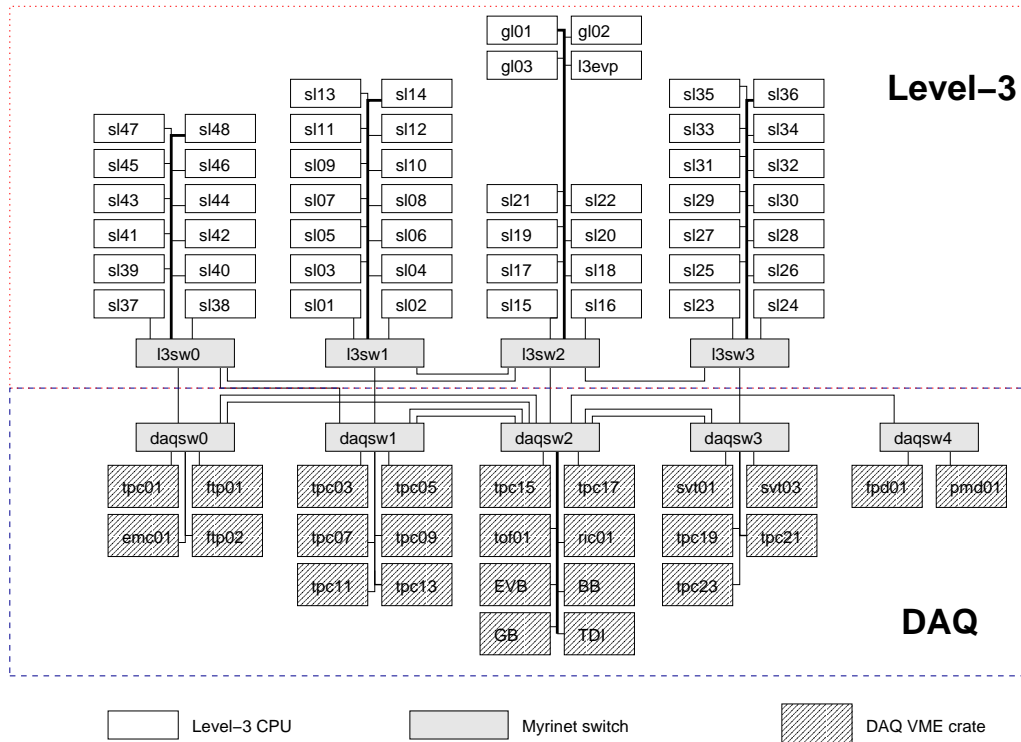


Fig. 1. STAR Level-3 trigger system architecture for year 2001 data taking. From top to bottom: Global level-3 (GL3, 3 CPUs, ALPHA/Pentium) and track finder (SL3, 48 ALPHA CPU), network (Myrinet switches), Level-3 cluster finder (VME readout crates, each 36 i960 CPUs).

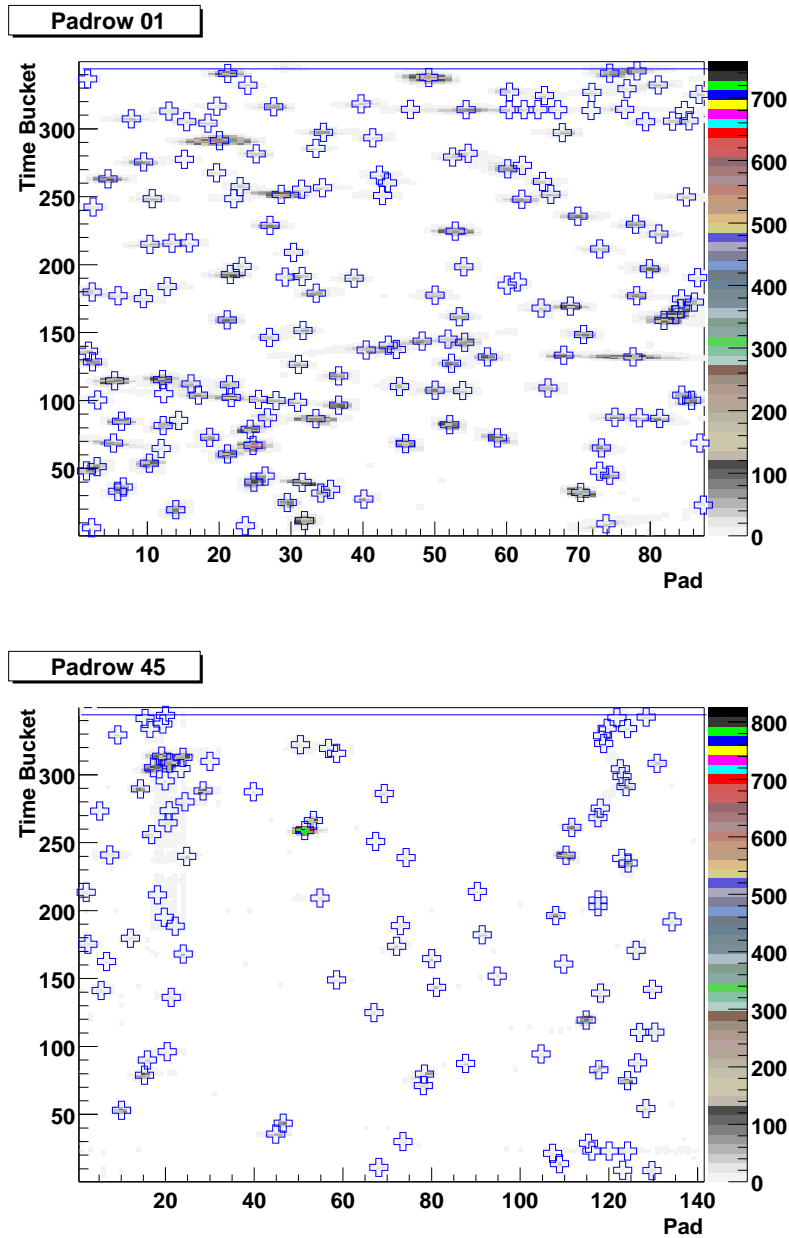


Fig. 2. Example of raw data and clusters for a typical central Au + Au event, $\sqrt{s_{NN}}=130$ GeV (~ 130000 TPC clusters in total). Two padrows (pad-drift plane) for one TPC sector are shown. *Top*: Innermost padrow (#1). *Bottom*: Outermost padrow (#45). The shades indicate the ADC value of the zero-suppressed raw data. The symbols mark the clusters found by the Level-3 cluster finder.




Fig. 3. Timing results of the cluster finder on the Intel i960 CPUs. *Left*: CPU time vs. number of clusters per i960 (cf. Section 2.2). *Right*: CPU time vs. i960 within one TPC sector for Au + Au events at $\sqrt{s_{NN}}=130$ GeV. The solid line shows an expected number of clusters per i960 assuming that each padrow has the same number of clusters (arbitrarily scaled, only for qualitative comparison).




Fig. 4. Momentum resolution $\Delta p_T/p_T$ for Level-3 (*squares*) and STAR offline (*circles*) reconstruction, determined by embedding MC simulated particle tracks into real Au + Au events at $\sqrt{s_{NN}}=130$ GeV. See text for details.




Fig. 5. Track reconstruction efficiency for Level-3 (*squares*) determined by embedding MC simulated particle tracks into real Au + Au events at $\sqrt{s_{NN}}=130$ GeV. Offline reconstruction efficiency (*crosses*) is shown for comparison. For details see text.




Fig. 6. CPU time used by the track finder algorithm (2 TPC sectors in parallel) vs. the number of tracks for Au + Au events ($\sqrt{s_{NN}}=130$ GeV).

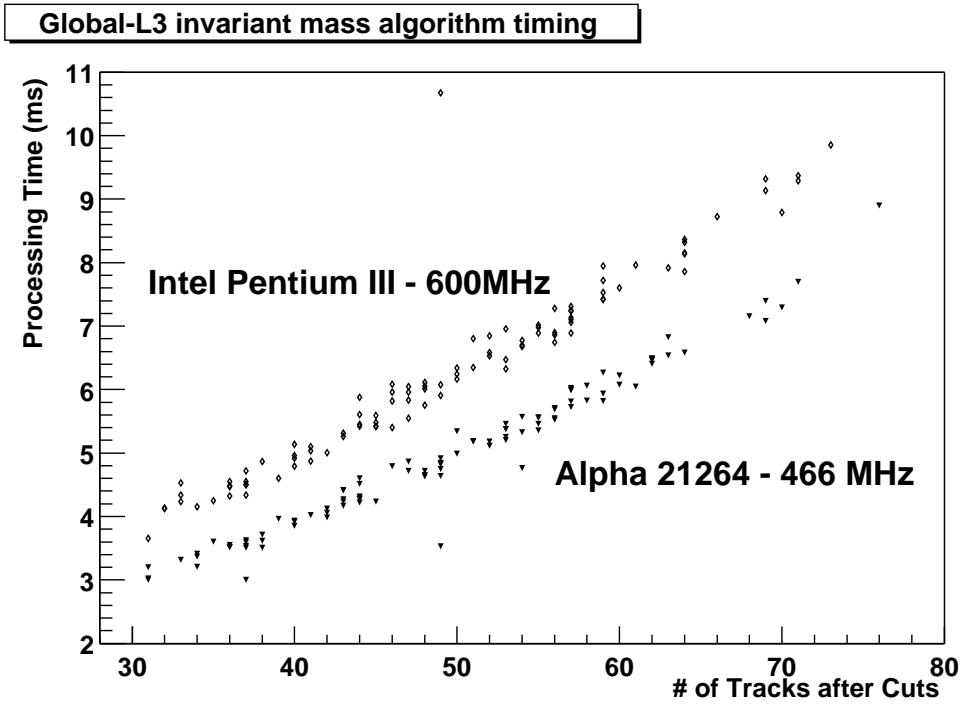


Fig. 7. Example of a GL3 algorithm timing. CPU time used by an invariant mass algorithm vs. number of candidate tracks. For details see Section 6.

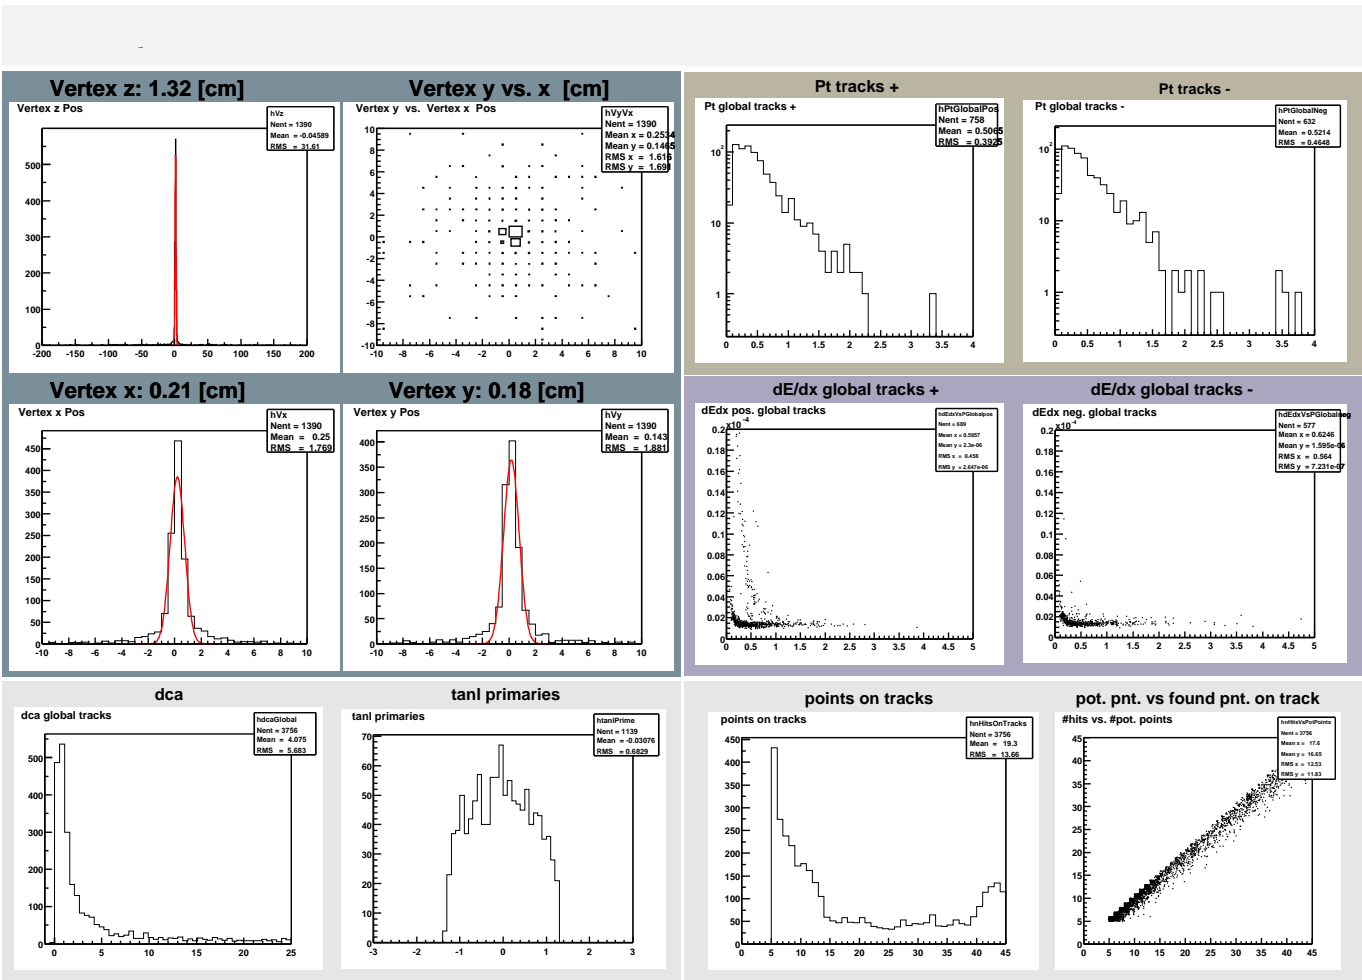


Fig. 8. Level-3 Online monitor. Example of a suite of plots that can be displayed online. For details see Section 8.




Fig. 9. The Level-3 event display showing a Au + Au collision at $\sqrt{s_{NN}}=200$ GeV. In the TPC only tracks are displayed. The smaller cylinders on both sides are the FTPCs with clusters. The SVT is visible in the middle of the TPC.