Proposal
Level-3 Trigger System

J. Berger, D. Röhrich, E. Schäfer[*],
D. Schmischke, R. Stock, C. Struck
University of Frankfurt

V. Lindenstruth
University of Heidelberg

M. Demello, P. Yepes
Rice University, Houston

August 18, 1998

[*]Max-Planck-Institut für Physik, Munich.

1 Data Compression and Trigger

From a trigger viewpoint the detectors in STAR can be divided into two categories: fast and slow. Fast detectors provide information for the trigger system at every RHIC crossing. Decisions at trigger levels 0, 1 and 2 are made using information from those detectors. Fast detectors are the Central Trigger Barrel (CTB), the end-caps of the TPC read as a Multiwire Proportional Chamber (MWC), the Vertex Position Detector (VPD), the Zero Degree Calorimeter (ZDC) and the Electromagnetic Calorimeter (EMC). These detectors feed the trigger with information about the multiplicity and energy produced in any given crossing.

The slow detectors are drift tracking devices and therefore need a longer time span after the collision to deliver their data. The Time Projection Chamber (TPC), the Silicon Vertex Tracker (SVT) and the Forward TPC (FTPC) are slow detectors. Their slowness is compensated by the detailed information they provide. The STAR Level-3 system is intended to take advantage of that information (up to 20-30 Mbyte/event at rates of up to 100 Hz) in order to reduce the data rate to a level that can be handled by the DAQ system, i.e. about 16 MBytes/sec. The data is then recorded onto an archival-quality medium for subsequent off-line analysis.

A key component of the proposed system is the ability to process the raw data performing track pattern recognition in real-time. It is designed to utilize the information from the TPC and fast detectors. However the system should be flexible enough to be expanded in a natural manner to include the other tracking devices.

Data reduction can be achieved in different ways:

1. Generation and application of a software trigger capable of reducing the input data stream by a factor of 100.

2. Reduction in the size of the event data by selecting sub-events. By analyzing the tracking information, regions-of-interest (ROI) can be defined. The data volume is significantly reduced by recording only summary information and raw data of the ROIs (e.g. electron tracks or a selected pp-event from pile-up).

Reduction in the size of the event data can also be achieved by compression techniques. This method may be useful in the later phase of the experiment when the detector and reconstruction performance is well understood. By online tracking and compression techniques an event size reduction by a factor of 10-20 at up to 100 Hz can be achieved [5, 6]. Nevertheless, these techniques are not to be pursued any further.

The system will be implemented as a distributed computing farm built from commercial low-cost systems and a SCI network (already used in DAQ [2] and TRIGGER [3]).

2 Physics and Detector Requirements

The Level-3 Trigger system is needed for most of the official STAR physics programs. The requirements from the various Physics Working Groups are presented in the following sections. Summarizing these requirements the Level-3 system should have full Φ coverage so that all TPC sectors can be processed, it should combine fast detector information (especially EMC) with at least TPC data, it should operate in two modes - L3-Trigger and ROI (i.e.

selection of sub-events) - and it should be able to handle event rates of up to 70 Hz (in the final configuration).

2.1 pp running [8]

For pp running Level-3 should combine tracking information with EMC data, including the shower max detector, to perform the following tasks:

- process events at a rate of 70 Hz from Levels 0-2
- remove pile-up by rejecting TPC hits not associated with the trigger event
- if there are areas in the TPC where it is not possible to remove pile-up (for example, if it is not possible to remove pile-up without the EMC), then reduce the event size by writing out track segments instead of hits
- reconstruct all jets in the event
- identify direct photons by forming energy clusters, making isolation cuts, and rejecting π^0 decays
- identify electrons by forming clusters, making isolation cuts, and requiring a matching track
- select events, subject to prescales, based on thresholds for jets, photons, and electrons, and based on the kinematics of the event.

2.2 High p_T physics [9]

Presently, Level-3 is the only way to select J/Ψ events in Au-Au. In addition, it is the only way to obtain a large enough sample of photons, electrons and jets in pA collisions. Some uses for Level-3 in High p_T physics in AA collisions are:

- select events with J/Ψ candidates to search for color screening effects. This requires the Level-3 system to be capable of performing the following tasks:
  - find electron candidates
    - find EMC and SMD clusters
    - find global tracks
    - match the above based on energy, momentum, shower shape and track quality
  - loop over electron pair candidates
    - calculate the mass with a vertex constraint
    - select events in a mass window (e.g. 2.5-4 GeV)
- select events with $c\bar{c}$ candidates to study energy loss of heavy quarks. The following calculations are needed to identify these candidates:
  - find electron candidates (as above)
  - loop over electron pair candidates (as above, with a higher mass window)

- select events with direct photon candidates to study jet quenching by photon-jet balancing as proposed by X.N. Wang et al. [10]. In order to trigger on those events the system should be able to
  - find photon candidates
    - find EMC and SMD clusters
    - match the above based on energy, momentum, and shower shape
- select events with high p_T tracks or π^0's in them to study jet quenching by the effect on inclusive spectra. The system must be able to
  - find photon candidates (as above) over a certain transverse momentum threshold.
  - loop over global tracks and select those events with a track passing some quality cuts above some transverse momentum threshold.

In pA collisions:

All of the physics topics studied in AA should also be analyzed in pA. In addition, there are the following topics:

- select events containing jets to measure initial gluon densities. This can be achieved by running the following algorithm:
  - create a grid of EMC energies in $\eta - \phi$ space
  - form a similar grid of energies of global tracks in the same $\eta - \phi$ space
  - use a standard jet algorithm to associate energies in these grids into jets

- select photon-jet and jet-jet pairs based on kinematic quantities. This is needed to insure that the whole trigger bandwidth is not occupied by events of common topologies. If that were the case, the experiment would not be live when an event with unusual topology occurs.
- select high p_T single electrons to measure initial gluon densities via heavy flavor production. In this case electrons need to be identified as previously explained and those above a p_T certain threshold need to be selected.

2.3 Peripheral Collisions [11]

There are two selection criteria, for 2 and 4 track events. For any other number of tracks, the event should be rejected, at least in this iteration. The selection criteria are as follows:

- 2 track events:
  - sum of charge should be zero
  - both tracks have pseudorapidity $|y| < 1.5$
  - both tracks pass through the interaction diamond, in both xy and z (to the limit of the tracking accuracy)
  - | vector sum of p_T for the tracks | $< 100 MeV/c$ (cut subject to change)
  - possibly (further study is required): Tracks are accolinear (in order to eliminate the few remaining cosmic ray muons)

- after a few years, a cut requiring the invariant mass of the two-tracks, assuming that they are pions, is greater than some threshold will be added.

- 4 track events:
  - sum of charge should be zero
  - all tracks have pseudorapidity $|y| < 1.5$
  - At least one of the following conditions should be fulfilled:
    * all tracks pass through the interaction diamond, in both xy and z
    * form a K_S^0 mass with another track
    * form a vertex with less than the K_S^0 mass with another track
  - $|$ vector sum of p_T for the tracks $| < 100MeV/c$.

2.4 Spectra [12]

The selection criteria are as follows:

- Trigger on track multiplicities in
  - different detectors TPC, SVT, FTPC, ...
  - specific pseudorapidity and transverse momentum bins
- PID'd tracks triggers
  - similar to previous but for candidates of specific particles
- Pair triggers

  - Kaon candidate pairs for Φ's
  - electron candidate pairs for Φ's and Ψ's

- Fragment triggers
  - d,$\overline{d}$ from momentum-dependent dE/dx thresholds.

2.5 EMC Calibration

In addition to the physics demands there is a detector specific request (EMC) which applies to AA running but is important to the pp program [8, 13]: Level-3 should be able to write out sparse events containing certain tracks and related EMC information for calibrating the EMC.

3 Online Pattern Recognition

The ultimate goal of pattern recognition is the reconstruction of an event, which is formed by different tracks measured in the various detectors and even the same tracks measured in different detectors. Therefore pattern recognition is a global, detector independent task. Nevertheless, in this proposal we will in the following only focus on the TPC detector. The problem of merging tracks from different detectors with different intrinsic resolutions and of mapping different local coordinate systems into one global system will not be discussed. Simulating and implementing a complete chain of a Level-3 Trigger which includes all STAR detector subsystems and is based on a global model goes far beyond the scope of this proposal (regarding manpower, financial resources, etc).

Pattern recognition of tracking detectors like the TPC is usually done sequentially. After finding clusters, the position and charge of the clusters are corrected for gain differences, distortions, time offsets, etc. These space points are then passed on to a track finder which builds track pieces out of clusters. Track pieces are finally merged to global vertex and non-vertex tracks.

3.1 Cluster Finder

The following results were obtained by simulating the L3 reconstruction chain in the STAF off-line framework, STAF. Au+Au central events generated with VENUS and HIJING were used in this analysis. In addition, an event with 1000 muons with flat distributions in p_T (0.05-10 GeV/c) and rapidity ($|\eta| < 2$) was analyzed. These events were run through the detector simulation program, GSTAR, with all physics processes switched on. Then they were processed through the TPC slow simulation package (tss). The output was then fed to the L3 and off-line cluster finder algorithms, l3cl and tcl respectively. The reconstructed hits were matched to Monte Carlo hits using geometric criteria. The number of hits reconstructed by the off-line cluster finder for the VENUS, HIJING and muon events are 303K, 145K and 57K respectively.

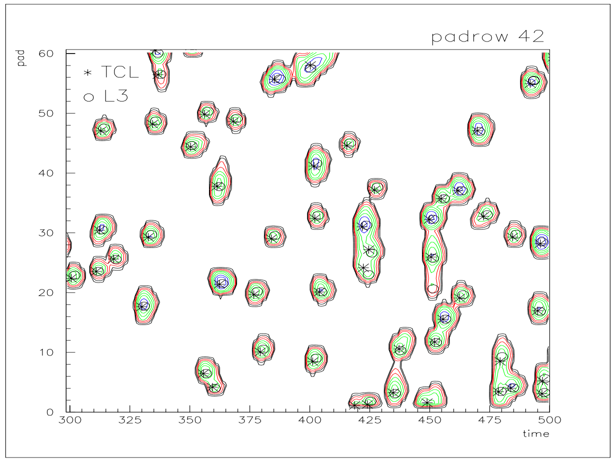

Figure 1: Hit distribution (log-scale) on Row 42 (VENUS (central Au+Au) + GSTAR (physics on) + tss). The results of the off-line (tcl) and the L3 cluster finders are also shown.

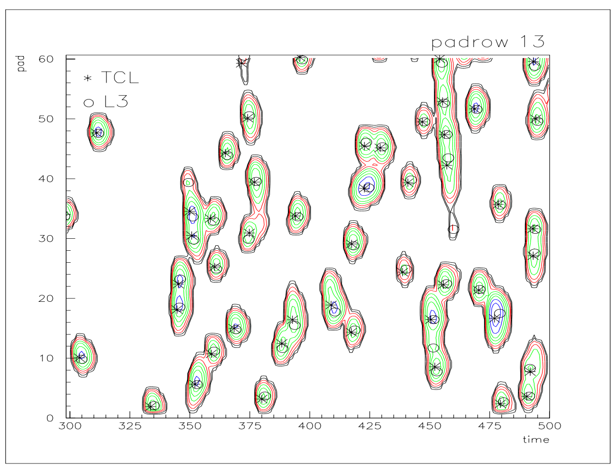

Figure 2: Hit distribution (log-scale) on Row 13 (VENUS (central Au+Au) + GSTAR (physics on) + tss). The results of the off-line (tcl) and the L3 cluster finders are also shown.

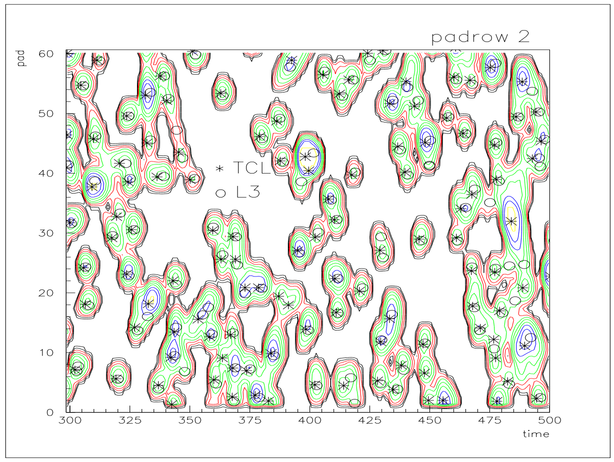

Figure 3: Hit distribution (log-scale) on Row 2 (VENUS (central Au+Au) + GSTAR (physics on) + tss). The results of the off-line (tcl) and the L3 cluster finders are also shown.

Simulations of a central Au+Au event (VENUS + GSTAR + SlowSimulator) are shown in Figs. 1, 2, 3; the results of the off-line cluster finder (tcl) and of the fast L3 cluster finder are also shown. Both algorithms deconvolute overlapping clusters in the high density region.

Figures 4, 5 and 6 show the ratio of the number of hits reconstructed by L3 and the off-line code as a function of the row number in the TPC for muon, HIJING and VENUS events respectively. L3 finds more hits in the inner sector (rows 1-13). In the outer sector, however, around 5% of the hits found by the off-line code are not reconstructed by the L3 algorithm.

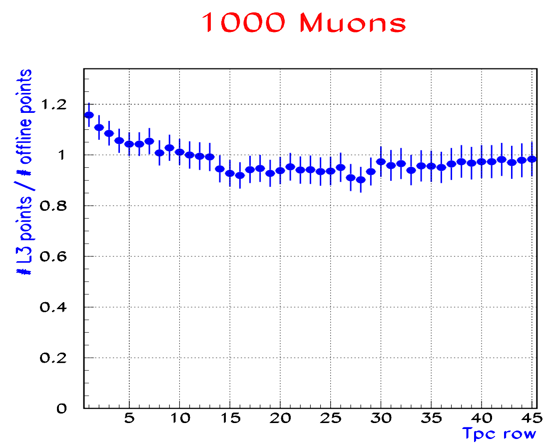

Figure 4: Ratio of hits reconstructed by L3 and the off-line code as a function of row number in the TPC for the muon event.

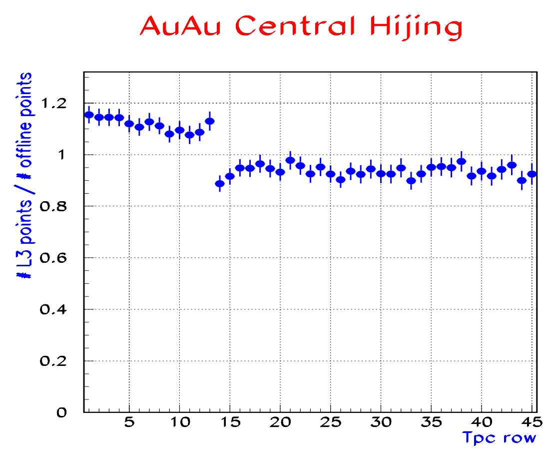

Figure 5: Ratio of hits reconstructed by L3 and the off-line code as a function of row number in the TPC for HIJING events.

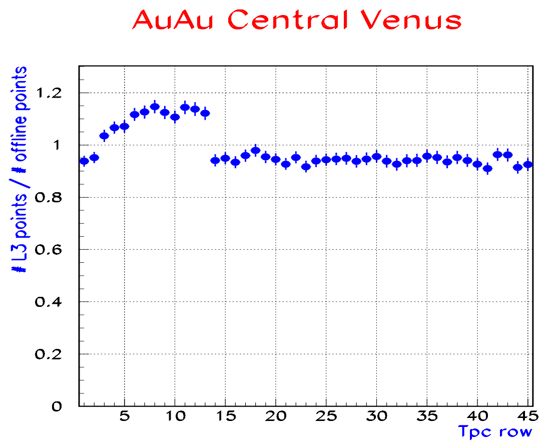

Figure 6: Ratio of hits reconstructed by L3 and the off-line code as a function of row number in the TPC for VENUS events.

The resolution of the cluster finder was analyzed by looking at the difference between the Monte Carlo and the reconstructed hits in rϕ and z. The difference rϕ, $\Delta r\phi$, (r is distance to beam axis, z=0, and ϕ is the azimuthal angle) provides information on the resolution in the plane perpendicular to the beam.

Figures 7, 8, 9 show $\Delta r\phi$ for muon, HIJING and VENUS events for different TPC row groups. As can be observed, the resolution worsens with decreasing row number and track density. For the inner sector, and especially for the VENUS event, the L3 cluster finder produces an excess of hits for negative $\Delta r\phi$. This problem is not understood at this time and needs further investigation.

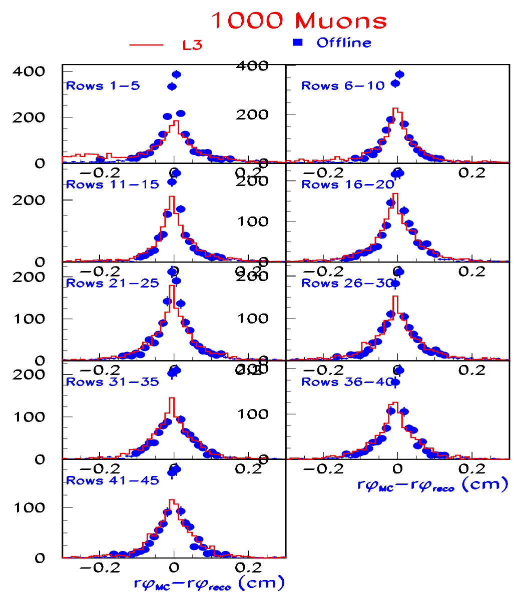

Figure 7: Resolution $\Delta r\phi$ of hits reconstructed by L3 and offline for different TPC row groups for the muon events.

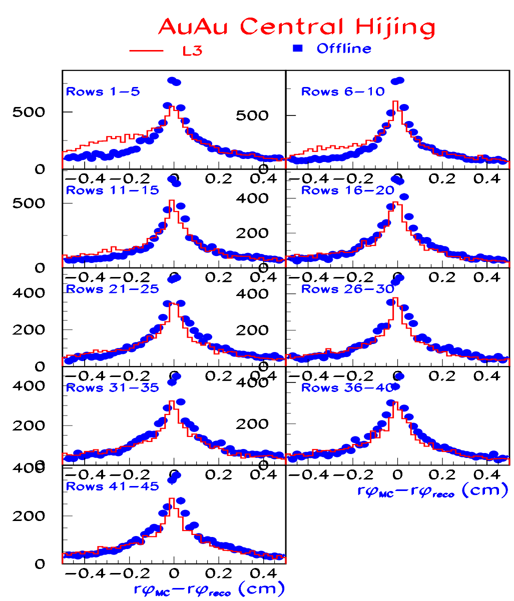

Figure 8: Resolution $\Delta r\phi$ of hits reconstructed by L3 and offline for different TPC row groups for HIJING events.

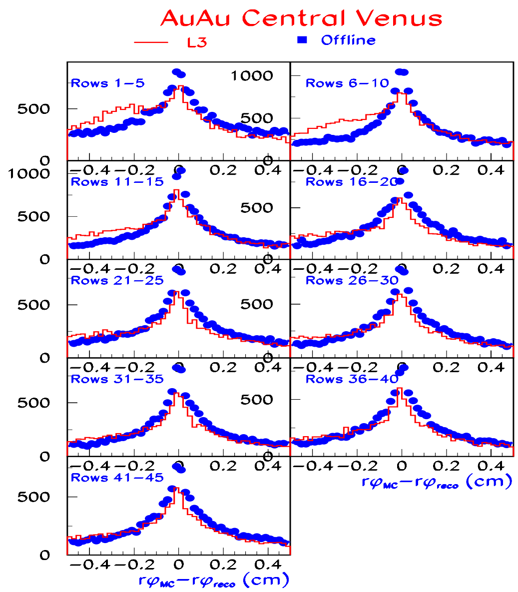

Figure 9: Resolution $\Delta r\phi$ of hits reconstructed by L3 and offline for different TPC row groups for VENUS events.

Figures 10, 11, 12 show the RMS of $\Delta r\phi$ for muon, HIJING and VENUS events as a function of row number. It is important to note that the quantity plotted is RMS and not the result of a gaussian fit. This choice was made because the distributions have large tails for the VENUS and and HIJING events. For the muon event the distributions do not have large tails and the RMS and the width of a gaussian are similar. In that case the RMS is close to 500 microns.

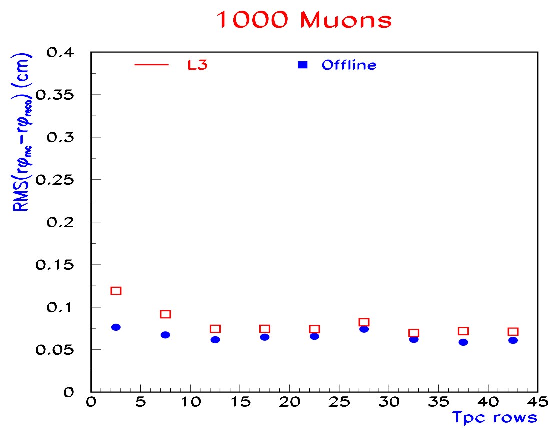

Figure 10: RMS of the $\Delta r\phi$-distribution of hits reconstructed by L3 (open symbols) and offline (filled symbols) as a function of the TPC row number for the muon events.

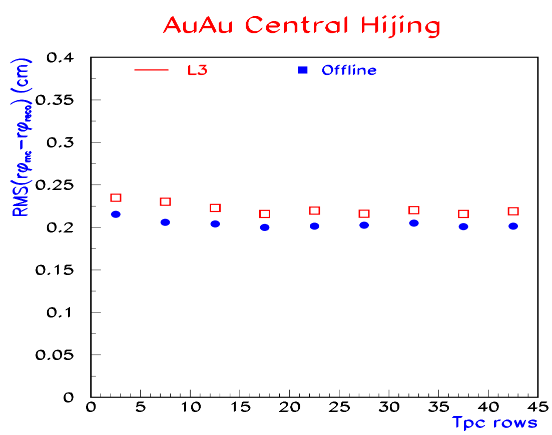

Figure 11: RMS of the $\Delta r\phi$-distribution of hits reconstructed by L3 (open symbols) and offline (filled symbols) as a function of the TPC row number for HIJING events.

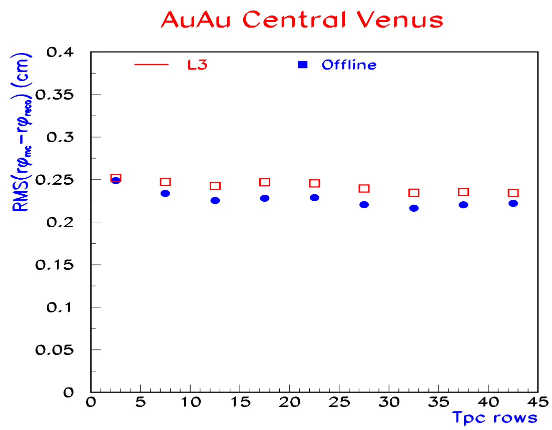

Figure 12: RMS of the $\Delta r\phi$-distribution of hits reconstructed by L3 (open symbols) and offline (filled symbols) as a function of the TPC row number for VENUS events.

Figures 13, 14, 15 show Δz for the muon, HIJING and VENUS events for different TPC row groups. Δz is the difference in z between Monte Carlo and reconstructed hits.

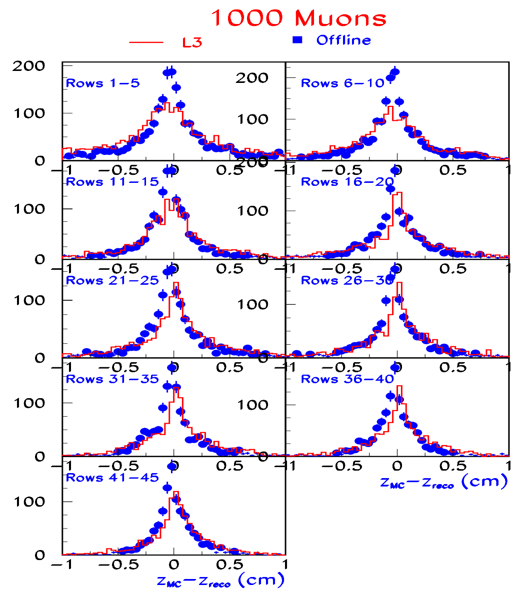

Figure 13: Resolution Δz of hits reconstructed by L3 (open symbols) and offline (filled symbols) for different TPC row groups for the muon events.

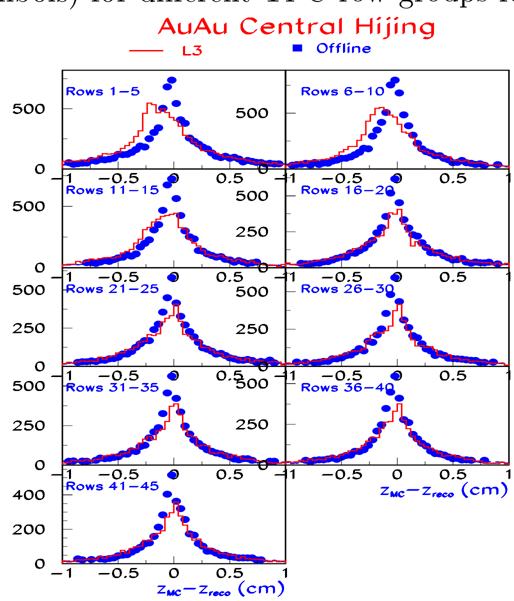

Figure 14: Resolution Δz of hits reconstructed by L3 (open symbols) and offline (filled symbols) for different TPC row groups for HIJING events.

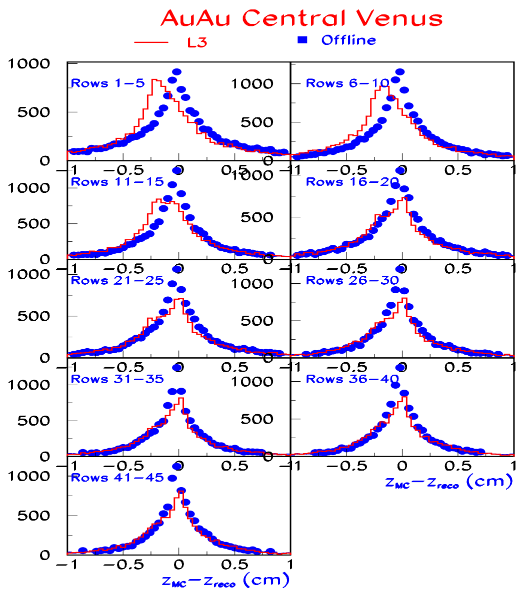

Figure 15: Resolution Δz of hits reconstructed by L3 (open symbols) and offline (filled symbols) for different TPC row groups for VENUS events.

Figures 16, 17, 18 show the RMS of Δz for muon, HIJING and VENUS events as a function of row number.

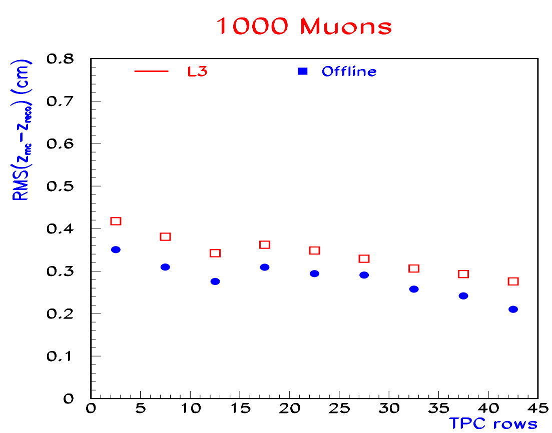

Figure 16: RMS of the Δz-distribution of hits reconstructed by L3 and offline as a function of the TPC row number for the muon events.

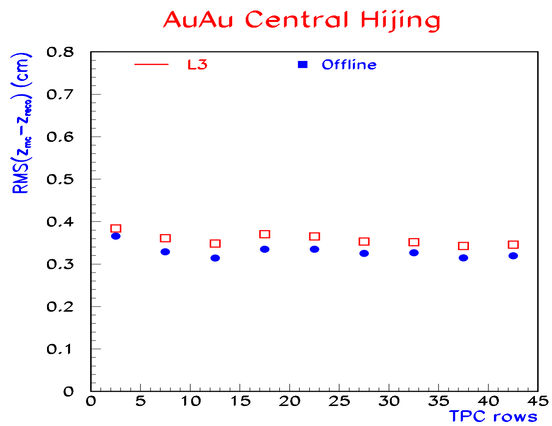

Figure 17: RMS of the Δz-distribution of hits reconstructed by L3 and offline as a function of the TPC row number for HIJING events.

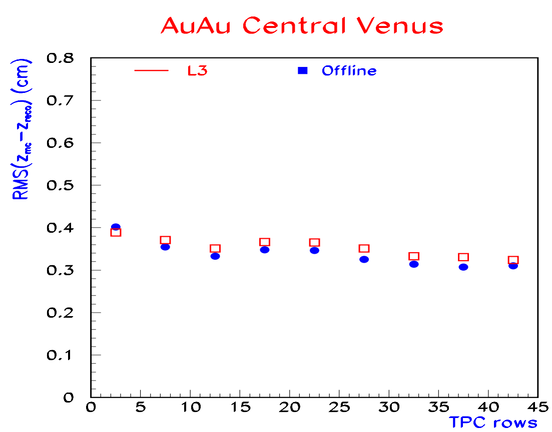

Figure 18: RMS of the Δz-distribution of hits reconstructed by L3 and offline as a function of the TPC row number for VENUS events.

The quality of the charge determination by the L3 cluster finder and the measurement of the specific ionisation has to be studied.

The fast cluster finder running on the i960 processors of the receiver boards has been optimized for speed [15]. Preliminary timing tests showed that cluster finding without deconvolution can be done in 10-20 msec on a 33 MHz version of the i960. An evaluation of the timing of the online cluster finder including deconvolution on a 66 MHz i960 (as used in the DAQ implementation) is shown in Fig. 19. The figure shows the time an i960 processor needs to process its corresponding padrows (typically 2-3) for all 18 i960 CPUs. The processors with lower numbers are responsible for the inner sector. The number of clusters per processor ranges from 350 for the outermost padrows to 500 in the inner sector. Cluster finding without deconvolution can be done in less than 10 msec, with deconvolution of overlapping clusters in 20-35 msec. Cluster finding in less than 20 msec should be possible by further optimization of the code and by omitting timebins in η-regions where no tracking is possible due to the shortness of the tracks.

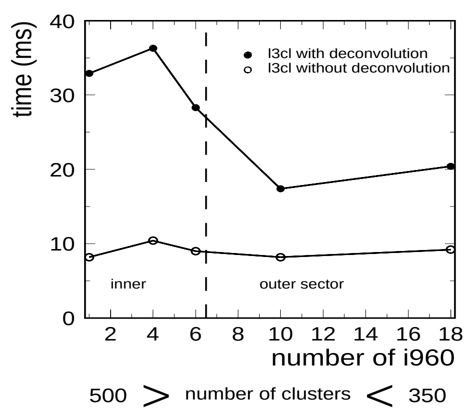

Figure 19: Timing results of the online cluster finder without (open symbols, lower curve) and with (filled symbols, upper curve) deconvolution running on i960s at 66 MHz.

3.2 Track Finder

In the following we will focus on the TPC detector, because the TPC produces most of the data volume and its information suffices to reconstruct tracks, calculate momenta and vertices and identify particles by dE/dx. The second step in the sequential pattern recognition scheme is a track finder. The tracker combines a number of space points to form track segments. Track segments are then merged to form vertex and non-vertex tracks. The track finder currently used is based on an algorithm described in [16, 17]. Its main features are an optimized data organization and a conformal mapping to speed up fitting procedures.

The track finder performance was studied using the same events as described in the previous section. This time they were processed through the fast TPC simulation chain (HIJING/VENUS + GSTAR(physics on) + Fast-Simulator). In the following results the tracker considers all tracks as primary.

The track finding efficiency was defined as the number of reconstructed tracks divided by the number of tracks contained in the TPC. Tracks contained in the TPC are those tracks that produce at least 10 hits in the detector. Reconstructed tracks are those with at least 10 reconstructed hits. In addition, any other reconstructed segment of the same Monte Carlo track should have less than 10 hits.

Figures 20 and 21 show the track reconstruction efficiency, ϵ_{tr}, for the L3 and offline trackings. The first panel shows ϵ_{tr} as a function of transverse momentum in the central region ($|\eta < 1|$). ϵ_{tr} as a function of rapidity for $p_T > 0.4$ GeV is depicted on the second panel. As can be observed the efficiency drops at forward rapidities, especially for the VENUS event. This effect is not surprising since the L3 tracking was tuned with tracks at central rapidities. Further work is needed to understand and solve this drop in efficiency.

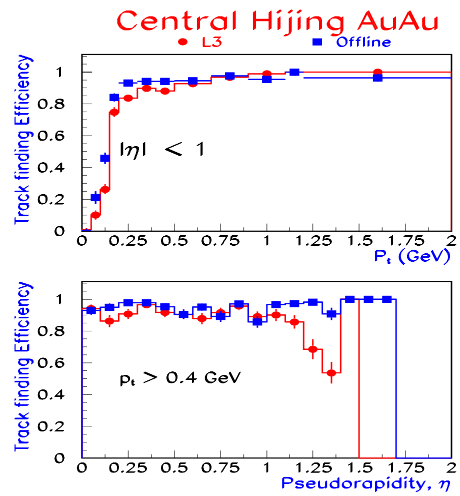

Figure 20: Track reconstruction efficiency for the L3 and the offline tracking for HIJING events.

Figures 22 and 23 depict the momentum resolution as a function of transverse momentum for the HIJING and VENUS events respectively. A vertex constraint was imposed. As can be observed the momentum resolution is worse in the lowest bin. This is believed to be due to reconstruction effects, since low momentum tracks are more challenging to reconstruct. All the other points are well below 2% resolution.

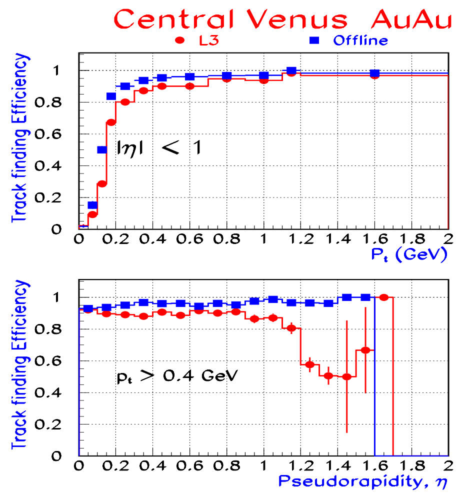

Figure 21: Track reconstruction efficiency for the L3 and the offline tracking for VENUS events.

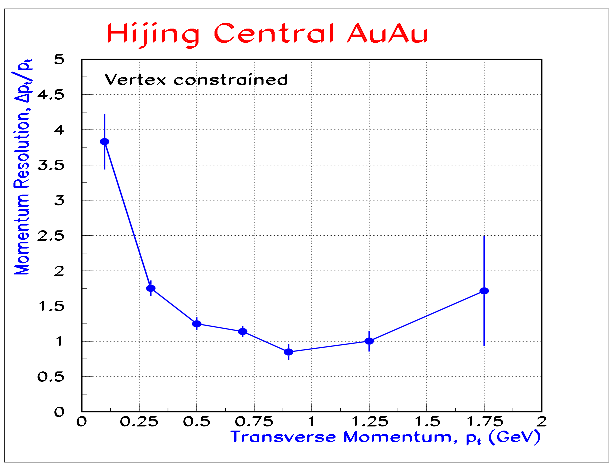

Figure 22: L3 tracking momentum resolution as a function of transverse momentum for HIJING events.

The track finder has been ported to C++ and runs on a PentiumPro processor under WindowsNT. Preliminary timing measurements were performed

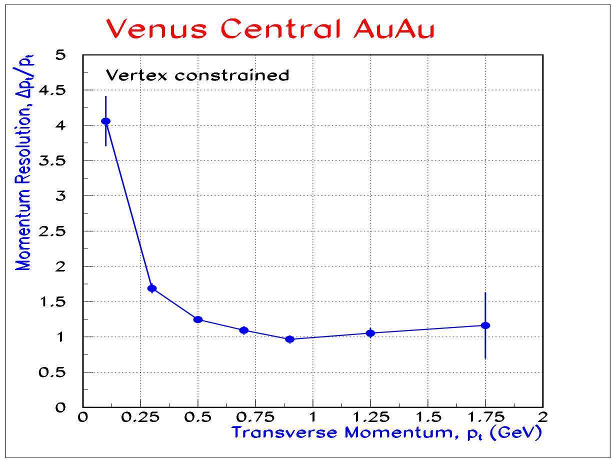

Figure 23: L3 tracking momentum resolution as a function of transverse momentum for VENUS events.

using (smeared) GEANT space points as an input. Fig. 24 shows the timing results for different track multiplicities.

The performance of the fast tracker can be estimated for different CPUs currently available (or in the near future). 350 tracks per sector were assumed, the processing time for track reconstruction was measured on the PentiumPro system. Based on the published SPEC benchmarks the performance for track reconstruction was estimated for the other platforms. More benchmarks on various platforms are planned.

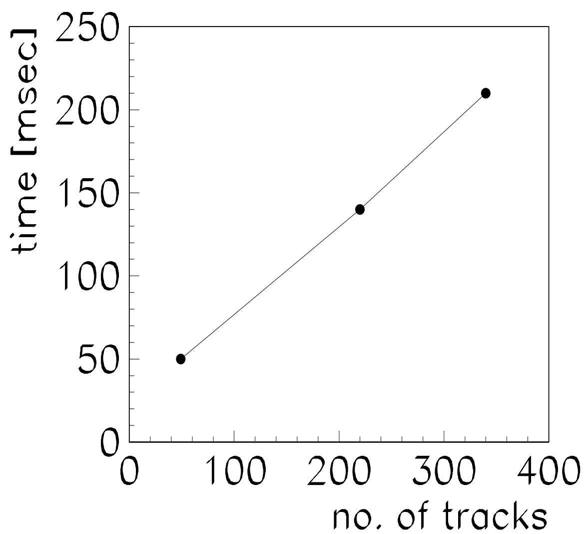

Figure 24: Timing results of the online track finder running on a 200MHz PentiumPro as a function of the input track multiplicity in one sector.

Table 1: Measured (PentiumPro) and estimated (all other platforms) performance of the Level-3 fast tracking code running on various processors.

task	Pentium Pro 200 MHz	Pentium II 400 MHz	PowerPC 750 266 MHz	Alpha 21164 600 MHz	Alpha 21264 667 MHz
SPEC-95 int	8.7	15.8	12.4	18	44
SPEC-95 fp	6.7	12.4	8.4	27	66
fast tracker	200 msec	110 msec	150 msec	80 msec	30 msec

3.3 Combined Chain of L3 cluster finder and tracker

A study of the perfomance of the complete chain (L3 cluster finder and tracker) is under way. Preliminary results are shown in Figures 25, 26 and 27. Efficienices are around 80% and stable.

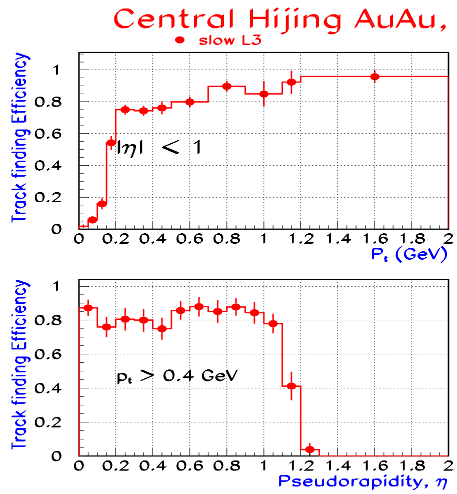

Figure 25: Efficiency of the combined online pattern recognition chain for HIJING events.

3.4 Calibrations and Distortion Corrections

Calibrations and corrections for various distortions (geometrical, ExB etc.) can be done on the raw data level (ADC-values), space point and track level. Since applying corrections and calibrations is usually not complicated but understanding and modeling of distortions are problematic, a discussion of this topic has to be postponed until the first data have been recorded. We assume that corrections and calibrations are easy to apply to the L3 chain.

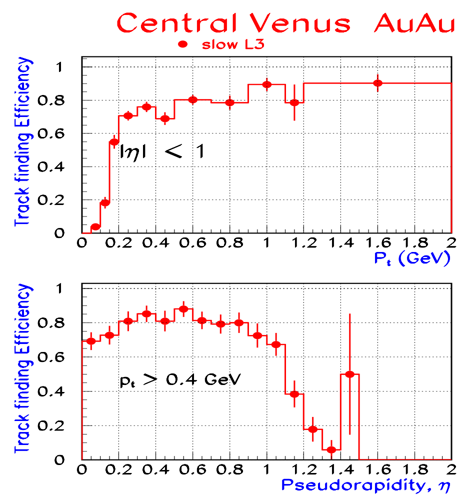

Figure 26: Efficiency of the combined online pattern recognition chain for VENUS events.

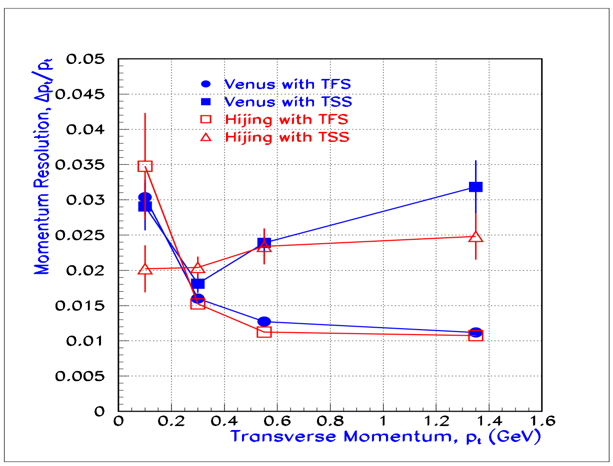

Figure 27: Resolution of the combined online pattern recognition chain.

3.5 Track Merging

Track segments found in different TPC sectors (and later in the SVT) are merged to form global tracks. Hits in the EMC are associated to tracks. The merging of tracks from different TPC sectors was included in the simulation results presented above. Merging tracks is probably one of the main tasks of global L3. The performance and timing of the merging has to be studied. However it is not expected to be a significant fraction of the tracking time.

3.6 Alternative Approaches

If track densities especially in the inner pad-rows get too high, clusters start to overlap so that a simple cluster finder cannot recognize or resolve merged clusters. Pattern recognition methods for track finding on raw data like template matching or adaptive generalized Hough transforms [18, 19] have to be employed. These methods have been used in NA35 and tested in NA49, but have not yet been adapted to the STAR TPC detector.

4 Modes of Operation

The Level-3 system can operate in two modes: Trigger mode and data reduction mode. The latter offers two different strategies: defining ROIs and selecting sub-events or using compression techniques. The general idea behind these compression techniques is presented in this document. However further development will be postponed until the detector performance is well understood.

4.1 Level-3 Trigger

The Level-3 Trigger (L3) is responsible for deriving a trigger decision based on the information of all detectors, especially the tracking detectors TPC, SVT and FTPC. The L3 Global system collects the output data streams of the L3 Sector systems, merges the tracking information of the different tracking detectors and combines the results with the remaining detectors e.g. EMC. L3 Global then accepts or rejects events with a sensitivity of 1 out of 100. The crucial inputs to L3 Global are the reconstructed tracks (track summaries) which have been reconstructed online by the L3 Local systems. The task of L3 Local is pattern recognition, L3 Global derives a physics trigger decision from the results.

4.2 Reduction of the Size of the Event Data

The raw data volume per event can be reduced by compression techniques or by selecting sub-events. These sub-events contain only the raw data information of a few tracks.

4.2.1 Selection of Sub-events - Regions-of-Interest

Sub-events or regions-of-interest can be defined on the basis of tracking information including a rough PID and the knowledge of the interaction vertex. All raw data inside these regions go on tape, all other data are dropped. In the case of pp interactions the extraneous TPC hits from pile-up are dropped. For the lepton measurements the data volume can be reduced to candidate e^+e^- tracks, which would yield a few tracks per event (e.g. for the Φ) [20]. Based on the track information Global L3 selects tracks segments as seed for a ROI. Local L3 systems define the ROIs by deriving lists of selected pad numbers and time bins and the corresponding i960. The following table illustrates the different applications of L3 as a trigger and the selection of a ROI:

Table 2: Estimated rates of $J/\Psi \rightarrow e^+e^-$ and $\Phi \rightarrow e^+e^-$. 100 Au+Au collisions per second and a L3-Trigger selectivity of 1:100 are assumed [5, 21].

	$J/\Psi \rightarrow e^+e^-$ $p_T(e) > 1.5$ GeV/c	$\Phi \rightarrow e^+e^-$
Signal/event	10^{-5}	0.02
S/B	1:58	1:50
Background/event	$6 \cdot 10^{-3}$	1
Method	L3-Trigger	ROI
Signal/year	10^4	$2 \cdot 10^7$
S_{eff}/year with L3	10^2	$2 \cdot 10^5$
S_{eff}/year without L3	1	$2 \cdot 10^3$

4.2.2 Data Compression

General data compression techniques have been applied to TPC data, i.e. both the NA49 raw data [6] and the STAR simulation data [14]. Lossless transformations like variable length codes (e.g. Huffman coding) or even lossy compression methods like vector quantization can only compress the ADC data by factors of 2 to 5. A further reduction in data volume is possible using data modeling techniques: only quantized differences to a data model – cluster and local track model – are stored. This results in reduction factors of about 20 [5]. The compression feature of the Level-3 System can be important in the later phase of the experiment when the detector performance is well understood. Therefore any further effort into this capability will be postponed. Nevertheless the basic ideas will be presented in the following.

The relevant information given by a tracking detector are the local track parameters and the clusters belonging to this track segment. The local track model is a helix; the knowledge of the track parameters helps to describe the shape of the clusters in a simple model [7].

The pattern recognition reconstructs clusters and associates them with local track segments. Note that pattern recognition at this stage can be redundant, i.e. clusters can belong to more than one track and that track segments can overlap. Once the pattern recognition is completed, the track can be represented by helix parameters. These are: curvature R, starting point (X, Y, Z), dip angle, azimuthal angle, track length, average charge, χ^2 of the helix fit and the number of clusters belonging to this track segment (see Table 3).

In a second step, the deviation of the cluster centroid position from the track model (residuals), the deviation from the average charge and deviations from the expected shape (based on the track parameters) are calculated for each cluster. These numbers are then quantized by a non-linear transfer function adapted to the detector noise and detector resolution (see Table 4). Remaining clusters can be optionally kept as raw data arrays. Depending on the TPC occupancy this will reduce the achieved compression. A study of this effect is under way.

The compression method discussed above allows for a later second pass of calibration and distortion corrections, track and vertex finding and fitting and dE/dx analysis; no relevant data is lost.

Table 3: Track parameter

parameter	size
curvature R	4 Byte (float)
begin X	4 Byte (float)
begin Y	4 Byte (float)
begin Z	4 Byte (float)
dip angle	4 Byte (float)
azimuthal angle	4 Byte (float)
track length	2 Byte (integer)
(average) cluster charge	2 Byte (fixed point)
χ^2	2 Byte (fixed point)
number of clusters	1 Byte (integer)
sum	31 Byte

Table 4: Cluster parameter

parameter	size
Flag empty cluster	1 Bit
Δtime	6 Bit
Δpad	6 Bit
Δcluster charge	7 Bit
Δshape	4 Bit
sum	24 Bit (3 Byte)

The following example illustrates that the relevant information of a central Au+Au event in the TPC can be stored in 1.3 MByte:

- typical event: 8000 tracks with 45 clusters each
- track data: $8000 \cdot 32$ Byte = 256,000 Byte
- cluster data: $8000 \cdot 45 \cdot 3$ Byte = 1,080,000 Byte

In total 1.3 MByte per event are needed. This corresponds to a data reduction factor of about 15.

5 L3-Requirements

Level-3 will allow event selection before event building, based on summary information extracted from each detector. Summary information consists of, for example, track parameters and space points found in the TPC. Complementary to the information derived from the detector itself, each Level-3 processor will have summary information derived from the trigger detectors available.

All Level-3 summary information for all events, both accepted and rejected events, is stored. Storage is either the main data stream or an auxiliary local L3-stream. L3 will provide a local data storage allowing access to all L3 information for monitoring and diagnostic purposes.

In addition to the summary information (e.g. tracks), Level-3 will provide a reduced event containing local track segments and their corresponding cluster information in a compressed representation. Summary information (e.g. tracks, vertices) and/or compressed event data is available for online monitoring, trigger decision and can go on tape.

The concept of the Level-3 system is an hierarchical structure of scalable multiprocessors. On the local level, computing farms are connected to each physical sector crate of DAQ. One global system collects all information from the local nodes and makes trigger decisions.

5.1 Data input

The input data to the Level-3 system are detector information from the tracking devices TPC, SVT and FTPC and from the trigger detectors. Raw data means ADC-data after zero suppression by the receiver board ASIC/i960, cluster data means results of the cluster finder running on the i960s.

Raw data from the fast detectors, CTB, MWC, VPD, ZDC and EMC is sent to the local processors. This is needed in order to select tracks from the triggered crossing.

- TPC, SVT, FTPC:
  raw data - zero suppressed
  cluster data
  mixture of raw data (inner pad-rows) and cluster data (outer pad-rows).
- Raw data from fast detectors.

5.2 Tasks

Level-3 Local and Level-3 Global have three tasks each:

Level-3 Local

- I/O:
  Collect cluster data and/or raw data from the i960s
  Move track summaries and/or compressed data to the Level-3 Global
  Transmit ROI to the i960s

- Track reconstruction:
  Use cluster data
  Use raw data
  Combine EMC data and track parameters

- ROI building:
  Build list of selected pads, time bins and corresponding i960s from track parameters

- Monitoring (slow control for system integrity):
  Watchdog task
  External control via ethernet.

Level-3 Global

- I/O:
  Collect track summaries and/or compressed data from the Level-3 Local
  Communicate Level-3 trigger decision
  Communicate Level-3 ROI definition to the Level-3 Local
  Move track summaries and/or compressed event to the event builder or local L3 data stream for monitoring purposes

- Event reconstruction, trigger mode:
  Merge local track segments
  Match tracks from different detectors
  Find vertices
  Bookkeeping and histogramming
  Make trigger decision

- Event reconstruction, ROI mode:
  Select sub-event and identify contributing track segments
- Event reconstruction, compression mode:
  Build compressed event
- Monitoring (slow control for system integrity):
  Watchdog task
  External control via ethernet.

5.3 Data Volume

Based on the current version of the slow simulator module TSS and the event generators VENUS and VNI, the typical event size can be estimated to be 20-30 MByte. TPC occupancy ranges between 18% and 43% (Fig. 28).

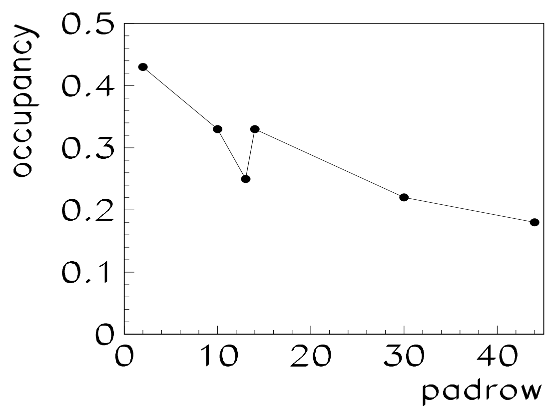

Figure 28: TPC occupancy as a function of the pad-row for a VNI event. The discontinuity is due to the transition from inner sector to outer sector.

Table 5: Number of tracks and clusters and size of raw data of a central VENUS Au+Au event (GEANT physics on/off) and a central VNI Au+Au event (GEANT physics on).

TPC	central VENUS Au+Au physics off	central VENUS Au+Au physics on	central VNI Au+Au physics on
tracks (total)	6200	11800	9900
tracks ($TPC - hits > 10$)	4700	9000	8500
tracks ($TPC - hits > 40$)	3200	6000	4700
tracks per sector ($TPC - hits > 10$)	200	375	350
clusters	191000	480000	380000
clusters per sector	8000	20000	16000
clusters per i960	440	1100	900
pixels	13000000	22000000	18000000
raw data (total)	13 MByte	22 MByte	18 MByte
raw data (total,incl. overhead)	17 MByte	29 MByte	24 MByte
raw data per sector	550 kByte	900 kByte	750 kByte
cluster data per sector	200 kByte	330 kByte	220 kByte
track summary per sector	6 kByte	12 kByte	11 kByte
track summary (total)	150 kByte	300 kByte	270 kByte
compressed data (total)	850 kByte	1.6 MByte	1.3 MByte

5.4 Time Budget

The TPC receiver boards will provide sufficient buffering for up to twelve uncompressed raw events. Therefore, given the maximum TPC event rate of 100 Hz, any event processing that exceeds 120 msec will result in increased dead time, unless zero-suppressed events are stored in the i960 memory. Stated differently, independent of the number of Level-3 processors operating in parallel on separate events, the total elapsed time for dealing with a single event at all stages of the Level-3 processing should not exceed 120 msec. The choice of the local L3 CPUs and the complexity of the pattern recognition algorithm should match this requirement. However, the time limit can be relaxed if hit sequences are copied from input memory into i960 working memory. This would increase the number of events which can be buffered on the receiver boards. The number of local L3 CPUs working in parallel on separate events is determined by the requirement to reconstruct 100 events per second, i.e. the average time spent on an event should not be longer than 10 msec.

5.5 Level-3 Track Reconstruction Performance

Level-3 trigger decisions are based on real-time processing of the data produced separately by each detector. For the TPC, SVT and FTPC this corresponds to local track recognition (Local L3). Track segments are combined with each other and with EMC and other trigger data by the Global L3 system.

- Trigger mode:
  use reliable, long tracks (e.g. $>$ 20 hits on track)

- Compression mode:
  use all tracks (>5 hits on track),
  use tracking information as a model for hit-compression.

5.6 Data Output

Level-3 Local

- Trigger mode: track summary
- Compression mode: track summary + compressed space points.

Level-3 Global

- Trigger mode: trigger decision + event summary
- Compression mode: compressed event.

6 Architecture

The trigger system is a distributed, symmetric, scalable multiprocessor system. It consists in the STAR baseline configuration of 12 (later of 24 TPC-sectors and 4 SVT-'sectors'and 6 FTPC-sectors) local processing clusters, a fast network (SCI) connecting the local sector units to the global network and a global processing cluster. Processing power for a sector crate is supplied by a farm of e.g. ALPHA-CPUs.

Table 6: Level-3 Hardware

Computing Nodes	Network
PC architecture	SCI
PCI-Bus	ring topology
1-4 CPUs	PCI-SCI adapter
32 MByte/CPU	- (computing node)
Ethernet adapter	PMC-SCI adapter
- (monitoring/booting)	- (sector broker)
optional hard disk	
- (booting)	

6.1 Sector L3 System

The Level-3 local systems consist of a SCI-ring connecting 1...n computing nodes to the sector broker.

6.2 Global L3 System

The Level-3 global system consists of a SCI-ring connecting 1...n computing nodes to the global Level-3 broker. The Level-3 global nodes will be equipped with at least 128 MByte of memory.

6.3 L3 Local-Global Network

If the traffic on the DAQ SCI-ring is so high that it prevents the Level-3 system from delivering track summeries in time, a second SCI-ring which connects the local and the global L3-rings is a backup solution.

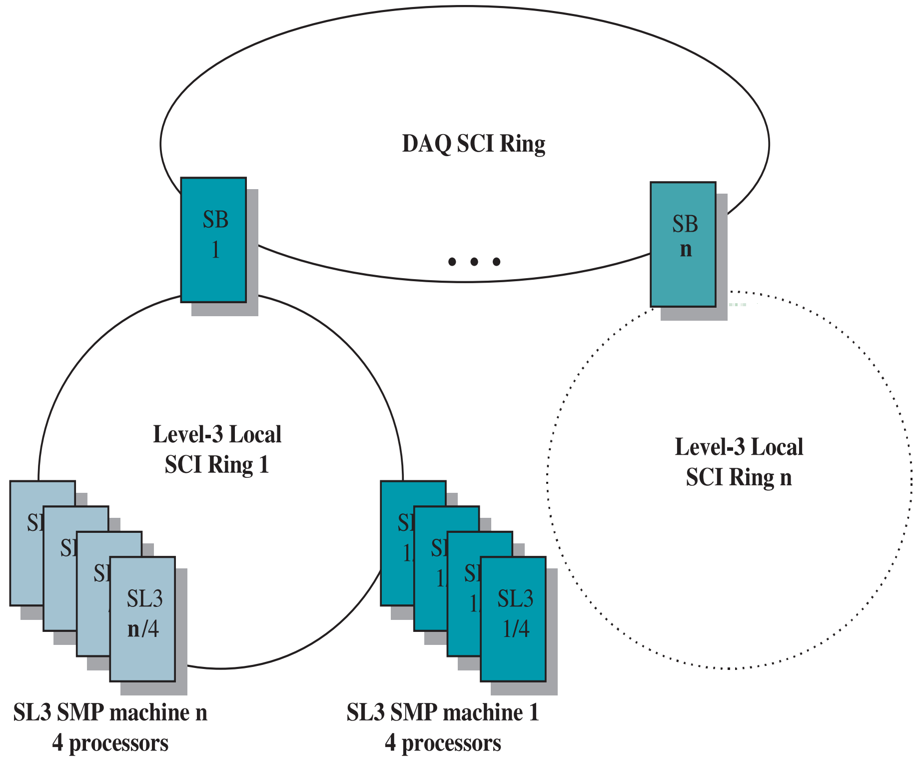

Figure 29: Architecture of a local L3 system. Each sector broker (SB) is connected to the local L3 system, which consists of one or more machines with up to 4 CPUs each.

7 OS and Software

The operating system for the computing nodes is WindowsNT; Solaris, Linux and VxWorks are alternatives. In computationally intensive applications it is essential to use the best available compilers which produce efficient code. At the moment Microsoft cc-compilers offer the best performance for PC platforms. Furthermore an adaption to the most recent processor flavors is required. SMP machines must be supported by the OS. This limits the available OS' to Linux SMP, Solaris, WindowsNT and VxWorks + SMP ex-

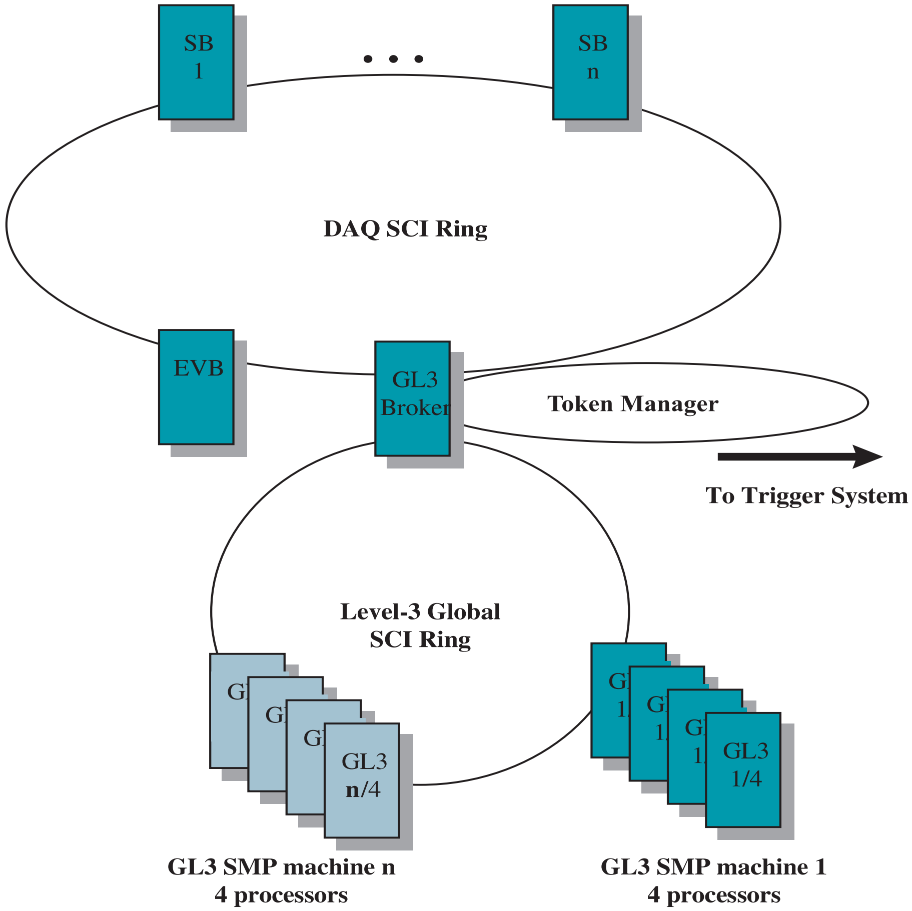

Figure 30: Architecture of the global L3 system.

tension. The OS must be inexpensive, RT licenses which cost more than 500 DM per processor cannot be fitted into the limited budget. This leaves the choice between Linux SMP and WindowsNT. Another factor in system development is the cost of implementing software. At this point it is highly recommended to use an advanced commercial development environment, such as WindowsNT and MS Visual C++. Despite the fact that NT is not a real-time system, its SCI-interrupt latency is sufficiently short (Fig. 31).

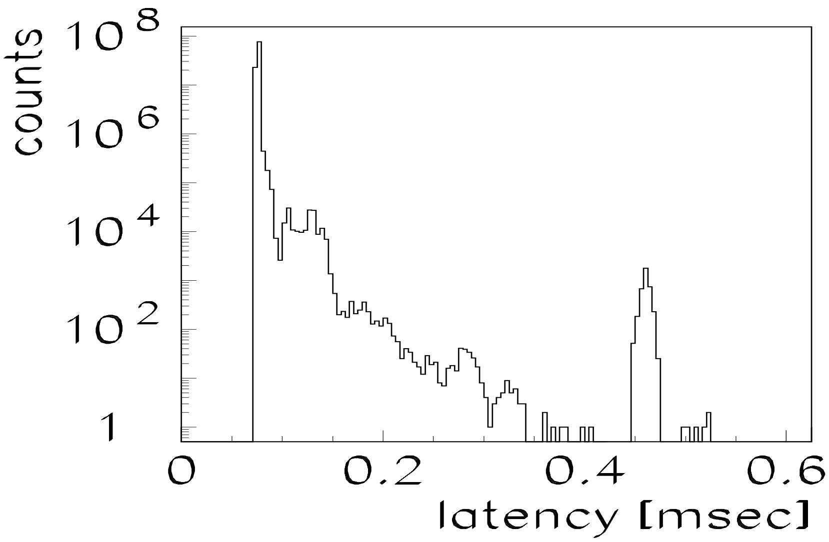

Figure 31: SCI-interrupt latency of a 150 MHz Pentium running WindowsNT.

8 Interfaces

The Level-3 system interfaces with the DAQ system: at the local level to the sector broker and on the global level to the Level-3 Global broker. User programs, i.e. global trigger selection algorithms, are interfaced to the Level-3 global code.

8.1 Interface to DAQ and TRIGGER

As outlined in section 6 and Figure 12, the L3 system interfaces with DAQ. However interfaces with other STAR subsystems such as run control and possibly slow controls (for system integrity) are anticipated. Another interface type will be used for on-line monitoring of the L3 performance at full speed, which will be presented in a seperate document.

A DAQ/Local-L3 prototype interface was developed and implemented [22]. However no interface document according to the "STAR Interface requirement document standard" has been defined yet. The L3 working group plans to propose such interface documents prior to the implementation of the system. The experience gained by designing the existing L3 prototype will be useful.

8.2 User Interface

There are two ways to implement the Level-3 global system:

- Implementation as one executable. A framework C++ class exists that contains the track list, the vertex list, some other lists of global interest, that need to be defined, and an analysis module list. The main() function must be modified for each new analysis module (supplied by the user) in a way that an instance of each analysis module (analysis class) is created.
- Implementation in separate executables. The framework exists as described above but the connection between the framework and the analysis modules is achieved via global events and shared memory.

9 Simulation Tools

The two fundamental pieces of the L3 reconstruction chain have been implemented as modules of the official STAR off-line framework, STAF (Standard Analysis Framework).

The module l3cl interfaces the TPC raw data tables with the L3 cluster finder to be run on the i960 processors. There are two versions of L3 track finder. They are interfaced to STAF through the modules l3t and ftf. ftf runs the original L3 tracking, originally written in C and later moved to C++. L3t interfaces to a newer version of the algorithm implemented in a more modern C++. Even though the algorithm used by both pieces of code is very similar, ftf has some features, like track refit, not implemented in l3t. The track reconstruction is performed in each sector independently. When all sectors are processed, track pieces from different sectors are matched together.

10 Implementation

The baseline DAQ implementation consists of 12 TPC sector crates, each crate housing 2 TPC sectors (no SVT and no FTPC). In this scenario the event rate is limited to 50 Hz by VME bandwidth. Since the estimated event size has grown by a factor of three due to the use of different event generators and a detailed simulation of the TPC response, the VME bandwidth further limits the event rate for central Au+Au collisions to about 20 Hz. The time budget which can be spent on a single event therefore increases to about 50 msec. One local Level-3 computing node with 1-2 ALPHA-21264 CPUs seems to be sufficient to handle this rate. Upgrading the DAQ system to 24 TPC sector crates will allow an event rate of 50-100 Hz. The Level-3 system then has to be doubled. Adding SVT and FTPC would require roughly duplicating the system again.

Table 7: Implementation. For the computing nodes ALPHA 21264 CPUs are assumed. The price is estimated to be $12.5k ($5.5k for a 600MHZ 21164 system + $7k upgrade to a 21264 CPU). One node per sector crate in the baseline scenario would limit the trigger rate to about 15Hz. It is further assumed that by the year 2000 the price for a 21264 system will be the same as it is now for a 21164 PC.

	Baseline (1999)	Upgrade 1	Upgrade 2
DAQ	12 TPC sector crates - -	24 TPC sector crates + $204k DAQ money	4 SVT + 6 FTPC sector crates - -
Level-3 local PMC/SCI adapter Computing node PCI/SCI adapter	 12·$3k=$36k 12·$12,5k=$150k 12·$1k=$12k	 12·$3k=$36k 36·$5,5k=$198k 36·$1k=$36k	 10·$3k=$30k 20·$5,5k=$110k 20·$1k=$20k
Level-3 global PMC/SCI adapter Computing node PCI/SCI adapter	 $3k $20k $1k	 - $20k $1k	 - $ 20k $1k
sum	$222k	$291k	$181k

11 Resources and Funding

The resources and funding for the project until March 31, 2000 are given in the following tables. They suffice to equip the baseline configuration (12 TPC sector crates) with one local computing node each, allowing a maximum trigger rate of 15 Hz. Funding for upgrading the system to run at higher rates and for including other detectors must come from STAR.

Table 8: Funding from IKF

Personell	1 Postdoc at BNL 1 PhD student at IKF 1 PhD student at BNL/Yale
Travel	40 kDM per year
Investment	430 kDM + 80 kDM (already spent)

Table 9: Funding plan

	1997	1998	1999
Personell	1 PhD student	2 PhD students 1 Postdoc	1 PhD student 1 Postdoc
Travel	40 kDM	40 kDM	40 kDM
Investment:			
Prototype	80 kDM		
Level-3 Local		30 kDM	360 kDM
Level-3 Global			40 kDM

References

[1] STAR Conceptual Design Report, PUB-5347, LBL (1992).

[2] 'Design and Implementation of the STAR Experiment's DAQ', A. Ljubicic et al., Proceedings RT'97, Beaune (1997).

[3] 'Trigger Level 1 and 2 architecture', V. Lindenstruth et al.,
http://rsgi01.rhic.bnl.gov/star/starlib/doc/www/html/
trg_l/doc/trig_level1_2.html,
'Trigger Level 0 Architecture', K. Marks et al.,
http://rsgi01.rhic.bnl.gov/star/starlib/doc/www/html/
trg_l/doc/Trigl0.ps.

[4] ALICE Technical Proposal, CERN/LHCC 95-71, LHCC/P3 (1995).

[5] J. Berger, diploma thesis, University of Frankfurt, 1998.

[6] 'The Data Compression Book', M. Nelson and J.L. Gailly, M&T Books, New York, 1996; J. Berger et al., IKF-HENPG/1-98, to be published.

[7] H. Appelshäuser, PhD thesis, University of Frankfurt, 1996.

[8] G. Eppley, private communication, 1998.

[9] T. LeCompte, private communication, 1998.

[10] X.N. Wang et al., *Phys. Rev. Lett* **77** (1996) 231.

[11] S. Klein, private communication, 1998.

[12] W. Llope, private communication, 1998.

[13] V. Rykov, private communication, 1998.

[14] STAR DAQ Review, BNL, 1997.

[15] 'New (faster) version of a cluster finder', D. Schmischke, http://www.rhic.bnl.gov/STAR/html/daq_l/public/software/cluster.html.

[16] 'Status of the Fast Finder of Tracks', P. Yepes, STAR Note 246, 1995.

[17] 'A Fast Track Pattern Recognition', P. Yepes, STAR Note 248, 1996; *Nucl. Instr. Meth.* **A380** (1996) 582.

[18] F. Pühlhofer, D. Röhrich, R. Keidel, *Nucl. Instr. Meth.* **A263** (1988) 360.

[19] D. Brinkmann et al., *Nucl. Instr. Meth.* **A354** (1995) 419.

[20] STAR Computing Requirements Task Force Report, Draft 3.1, 1997.

[21] 'Dielectron Backgrounds and Their Impact On J/Ψ Physics', T.J. LeCompte, STAR Note SN0287 (1997).

[22] 'Sector Broker Software', D. Schmischke, http://daq.star.bnl.gov/~ daq/Milestone1/documents/ds/sb-doku.pdf;
'Sector Level-3 Software Draft', D. Schmischke, http://daq.star.bnl.gov/~ daq/Milestone1/documents/ds/sl3-doku.pdf.