

The STAR Level-3 Trigger System *

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

^aUniversity of Frankfurt, August-Euler-Straße 6, D-60486 Frankfurt, Germany

^bRice University, Houston, Texas 77251, USA

^cBrookhaven National Laboratory, Upton, New York 11973, USA

^dUniversity of Birmingham, Birmingham B15 2TT, United Kingdom

^eUniversity of Bergen, Allegaten 55, 5007 Bergen, Norway

^fUniversity of Texas, Austin, Texas 78712, USA

^gUniversity of Heidelberg, Philosophenweg 12, 69120 Heidelberg, Germany

Abstract

The STAR level-3 trigger is a MYRINET interconnected ALPHA processor farm, performing online tracking of $N_{track} \geq 8000$ particles ($N_{point} \leq 45$ per track) with a design input rate of $R=100$ Hz. A large scale prototype system was tested in 12/99 with laser and cosmic particle events.

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

1 Introduction

The RHIC accelerator at Brookhaven National Laboratory, USA, will investigate $Au+Au$ collisions with $\sqrt{s} \leq 200$ A·GeV and $p+p$ collisions with $\sqrt{s} \leq 500$ GeV. The STAR experiment (1) is a large scale, cylindrical, symmetric 4π -detector. Physics data taking will start in 2000 with a full size TPC (Time Projection Chamber, $R_{in}=0.6$ m, $R_{out}=2$ m) with 24 sectors, 6912 pads each. TPCs are specifically suitable for detecting high density charged particle fluxes in high multiplicity nucleus-nucleus events.

¹ Corresponding Author, Email soeren@bnl.gov

2 Architecture

2.1 The STAR Trigger Architecture

The STAR trigger system is subdivided into 4 hierarchic levels. The level-0 output rate is 10^5 Hz; level-1 and level-2 as well as coincidence with the TPC gating reduce the rate by one order of magnitude each. The level-3 trigger is supposed to reduce an input rate of 10^2 Hz to the final DAQ rate of $R_{tape}=1$ Hz at an expected TPC event size of $\simeq 15$ MB. Task examples for the level-0/-1/-2 trigger stages are (a) selection of central and peripheral $Au+Au$ events based upon multiplicity (function of impact parameter) and (b) rejection of beam-gas events with a vertex far from the interaction point. The tasks of the level-3 trigger are event selections based upon the online reconstructed track parameters of each particle. Several applications have been proposed for $Au+Au$ collisions, being on the one hand high p_T trigger applications as enrichment of heavy (anti-)fragments (e.g. $\overline{He^4}$), STAR EMC (Electromagnetic Calorimeter) calibration using tagged high p_T π^- , QCD hard parton scattering (leading high p_T hadrons in jets), and c and b quark decays (e.g. high p_T leptons). On the other hand, online invariant mass reconstruction of $J/\psi, \Upsilon \rightarrow e^+e^-$ is proposed, as suppression of $c\bar{c}$ and $b\bar{b}$ production is commonly regarded as one of the most promising signatures of the quark-gluon plasma. Additionally, for $p+p$ collisions, a level-3 trigger algorithm for filtering of ≤ 700 pile-up events in the TPC (at highest Luminosity $\mathcal{L}=2 \cdot 10^{31}$ $\text{cm}^{-2} \text{s}^{-1}$) per one level-0 trigger is being developed.

2.2 The STAR DAQ Architecture

The level-3 trigger architecture is closely related to the STAR DAQ architecture, in which one VME crate is mapped onto each physical TPC sector (two TPC sectors in the first stage), containing a *Sector Broker*, i.e. Motorola MVME-2306 VME board, carrying a PowerPC 604 (300 MHz, VxWorks), as the TPC sector master controller. The Sector Broker carries a MYRINET interface (cf. Section 5) for (a) raw data transfer to the main STAR event builder and (b) connection to the level-3 track finder CPUs. Moreover, each DAQ crate also contains six VME receiver boards, each carrying three mezzanine cards with

- one Intel i960 CPUs (33 MHz, VxWorks) for (a) data formatting and (b) running the level-3 cluster finder,
- 4 MB of dual-ported VRAM for buffering and pipelining of raw data of 12 events.

sector, 432 for the whole TPC. Input to the cluster finder are zero-suppressed TPC raw data, stored in the VRAM buffer. The output cluster data, i.e. (a) cluster center-of-gravity and (b) cluster total charge (ADC sum), are sent via VME to the Sector Broker, which itself ships the data via MYRINET to the level-3 track finder CPU (expected data transfer rate of $\simeq 3$ MB/s per TPC sector).

The cluster finder time constraint is $\tau_{cluster} \leq 10$ ms (input rate $R=100$ Hz). Benchmarks on the i960 were performed for 600 clusters (realistic $Au+Au$ scenario) on the TPC's most inner padrow. The position resolution of $\Delta(r\phi) \simeq 37 \mu\text{m}$ and $\Delta z \simeq 13 \mu\text{m}$ could be obtained with an algorithm within $\tau_{cluster} = 7.5$ ms (absolute cluster finding efficiency $\varepsilon = 93\%$). The clusters and reconstructed centers-of-gravity are shown in Fig. 10 (*left*). If two clusters are merged, an additional deconvolution subroutine must be started, consuming 6.0 % more CPU time than in case of two separated clusters. Fig. 10 (*right*) shows the reconstructed clusters for a STAR beam-gas event, recorded in 07/99. Different TPC clock timing lead to a maximum drift timebin $t_{max} = 348$ in that case (x -axis).

4 Track Finder

Monte-Carlo simulations² predict for a central $Au+Au$ collision (impact parameter $b \leq 2.0$ fm) that the track finder algorithm must be able to fit at least $N \simeq 400$ tracks per event per TPC sector, each consisting of $N_{point} \leq 45$ points (given by the number of padrows). This shall be referred as "Au+Au benchmark event" hereafter.

The fast track finder algorithm has specifically been developed for the level-3 trigger project (7). It employs *conformal mapping*, i.e. a transformation of a circle³ into a straight line, followed by a fit with a *follow-your-nose* method. A given space point (x, y) is transformed into a conformal space point (x', y') according to the equations $x' = (x - x_0)/r^2$ and $y' = (y - y_0)/r^2$, using $r^2 = (x - x_0)^2 + (y - y_0)^2$. The transformation requires the knowledge of one point (x_0, y_0) on the track trajectory, either (a) the interaction point (vertex constraint for primary tracks) or (b) the first point associated with the track (no vertex constraint for secondary tracks). Fig. 3 shows an example of 1000 simulated level-3 trigger π^+/π^- particle tracks ($p_T \leq 1.0$ GeV/c).

² Monte-Carlo simulations were performed using the event generator HIJING 1.31, which is based on a QCD-inspired model for jet production (6).

³ In the STAR solenoid magnetic field of $B=0.5$ T charged particle tracks can be parametrised as helices, being visible as circles in an xy -projection.

point, full-duplex networks SCI and MYRINET were tested for application in the level-3 trigger. Due to missing DMA capabilities⁵, Gigabit Ethernet was not considered. SCI (Scalable Coherent Interface) (9) utilizes transfer of small data payload of 16 or 64 bytes at a high clock frequency of 200 MHz. Signals are transmitted via 18 pin differential ECL. Prototype D310 PCI-SCI interface are available from DOLPHIN (10), the C2D PMC-SCI interface was available from VMETRO (11) (product line stopped in 1999). MYRINET (12) supports variable data payloads ≤ 4 kB (LANai interface chip operating at 33 MHz PCI clock). Signals are transmitted via two different cable types, i.e. SAN cable (PMC adapter to switch, 20 pin single ended signal) and LAN cable (PCI adapter to switch, 37 pin Low Voltage Differential Signal 1.2 ± 0.4 V). Both PMC and PCI host adapters (32 bit, 2 MB memory in the level-3 trigger configuration) are available from MYRICOM (13).

The basic architectures differ significantly, i.e. SCI utilizes a ring topology, MYRINET a switch topology. Thus, e.g. SCI requires the Sector Broker to carry two PMC adapters (different rings), as while MYRINET requires the existence of (16-port) switches between the PMC and the PCI side. Fig. 10 shows the bandwidth and latency⁶ as a function of buffer size for both MYRINET and SCI (PCI-PCI point-to-point DMA benchmark). Typical buffer sizes for the level-3 trigger are $b=128$ byte for messages and $b\geq 20$ kB for data transfers. In both cases, the bandwidth is limited by PCI bus to $R\approx 60-70$ MB/s. In case of SCI, the maximum bandwidth is already achieved for small buffer sizes $b\geq 64$ bytes (corresponding to SCI payload), but saturates at a lower level ($R=62$ MB/s). However, former D310 versions also achieved $R=72$ MB/s (cf. Fig. 6 in (2)), only a recent hardware revision lead to a bandwidth reduction of $\Delta R/R\approx 14\%$. For both cases, the CPU usage is as low as $\leq 12\%$ due to DMA. In case of SCI the (one-way) latency of $L=2-3$ μ s is smaller than in case of MYRINET ($L\approx 30-40$ μ s) due to an extra MYRINET software layer. The latency limit for the STAR DAQ design is $L\leq 100$ μ s. The final decision for the usage of MYRINET was driven by (a) long-term test stability issues⁷, (b) hardware revision status and availability and (c) free availability of MYRINET driver software as "open source" for numerous platforms (e.g. Linux/Intel, Linux/ALPHA, VxWorks).

⁵ Background DMA (Direct Memory Access) data transfers do not consume any CPU time (except for bus arbitration and interrupt handling). Thus, the processor is free for tasks as track finding or data formatting.

⁶ Latency $L=t_1-t_2$ is defined by the time t_1 for issuing an interrupt (e.g. end-of-package) on the receiver and the sender acknowledge time t_2 .

⁷ In a ring topology, any failure (e.g. a single faulty cable) affects the complete system, in a switch topology only a single point-to-point connection.

- (2) While the STAR TPC and the STAR magnet ($B=0.5$ T) were switched on, a test of the complete⁸ level-3 trigger prototype system (432 Intel i960, 12 Sector-L3, one Global-L3) was performed with laser⁹ and cosmic particle events. Event displays for one event of each type (using the 3D event display, cf. Section 7) are shown in Fig. 10.

For single cosmic particle tracks, a very high track finding efficiency $\epsilon \geq 95\%$ was achieved, although loose χ^2 cuts (in order to be able to find tracks with z_{vertex} even outside the TPC) lead to a number of reconstructed tracks $N_{track} \geq 1$ in $\sim 30\%$ of all cases. In total, 15,000 level-3 specific events (cluster and track data) were recorded successfully. Further cosmic data taking is planned for 01-03/2000. The start of the RHIC physics program is envisaged for 04/2000.

9 Summary

The STAR level-3 trigger system will perform online track finding for high multiplicity $Au+Au$ collider events ($N_{track} \geq 8000$, $N_{point} \leq 45$ per track), utilizing high performance ALPHA 21264 CPUs and high bandwidth MYRINET network interfaces (expected data transfer rate $\simeq 52$ MB/s). A global level-3 CPU will perform tasks like e.g. online invariant mass reconstruction and issue an accept/reject decision with a design input rate $R=100$ Hz ($R \leq 25$ Hz for a prototype system in 12/99).

10 Acknowledgements

This work has been supported by the *Bundesministerium für Bildung und Forschung*, Germany (contract #06OF 840 I). We would like to thank MYRI-COM, Inc. for excellent technical and documental support. J. S. Lange would like to thank the Alexander-von-Humboldt organization for support of INSTR99 conference participation.

⁸ The step from one up to 12 Sector-L3 CPUs also marks at the same time the step from synchronous to asynchronous mode (arriving data sequence by random).

⁹ For laser events the track finder was not operational, because straight tracks result in unreasonable $p_T \rightarrow \infty$ values.

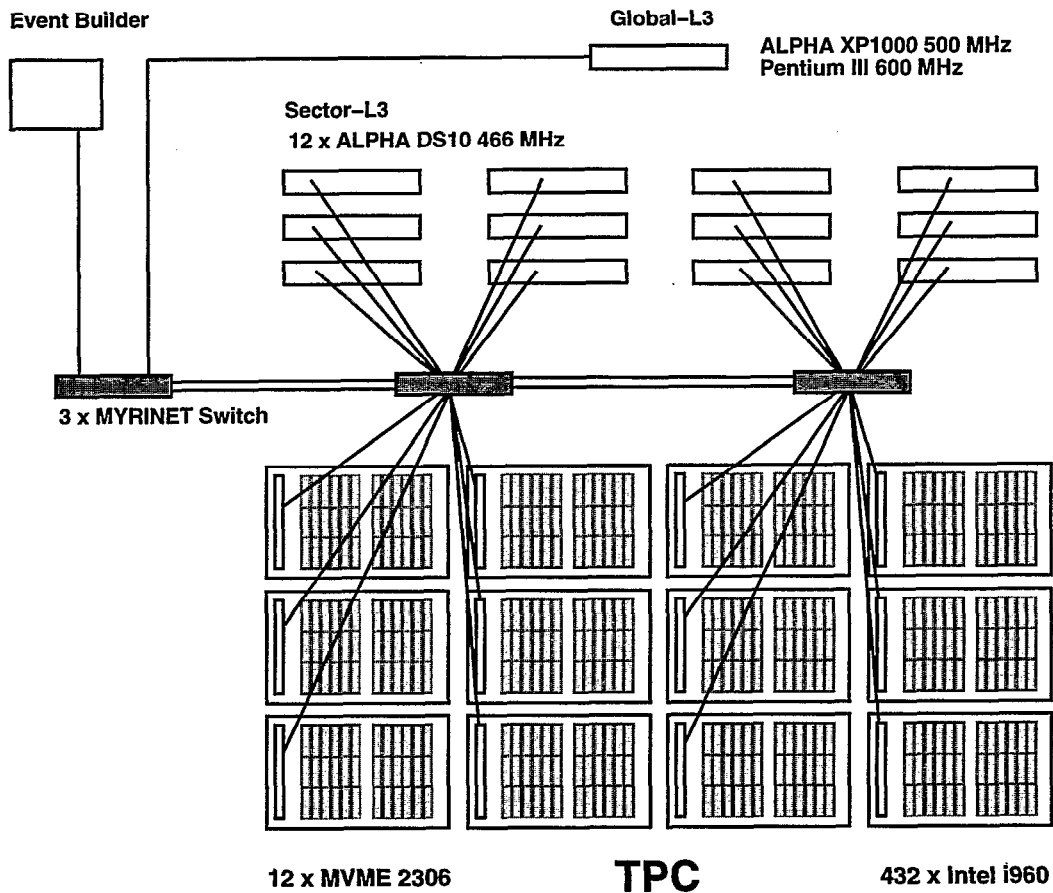


Fig. 1. STAR level-3 trigger system architecture as used in the system test 12/99 (cf. Section 8). Event building of level-3 specific events has been performed locally on the Global-L3 CPU, integration into the STAR event building is foreseen.




Fig. 3. Level-3 track finder example. 1000 simulated π^+/π^- particle tracks in the STAR TPC are shown. Tracking was performed on an ALPHA XP1000 within $\tau_{track}=9.6$ ms per TPC sector.

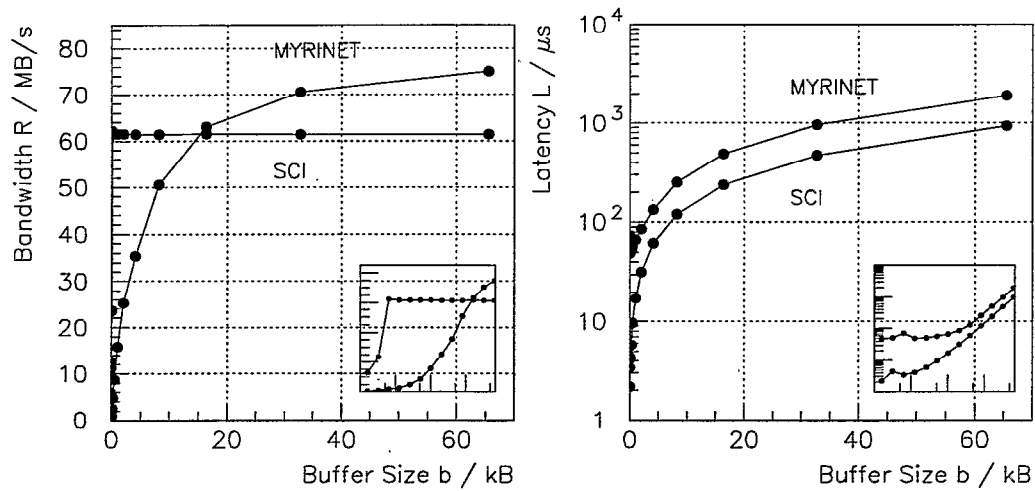


Fig. 5. Bandwidth (left) and Latency (right) as a function of buffer size for both MYRINET and SCI (Level-3 PCI-PCI point-to-point DMA, logarithmic buffer size axis in small pictures).