

Online System and Level 3 Trigger

Technical Review of BaBar Improvement Plans

SLAC

October 11, 2000

Rainer Bartoldus

Outline

- Online System Overview
- OEP/Level 3 Farm
- Performance Projections
- Conclusions




TM and © Laurent de Brunhoff

Online System Overview

Will not talk about components that don't scale with luminosity, such as Run Control, Detector Controls, Online Databases.

The Event Path



Dataflow (ODF)

- Transports and assembles event data from Front End Elements (FEEs) to the processing farm
- Provides framework for Feature Extraction (FEX) code (detector specific)

Online Event Processing (OEP)

- Receives data from DAQ crates and passes them on
- Provides framework for Level 3 Trigger (L3) and Fast Monitoring

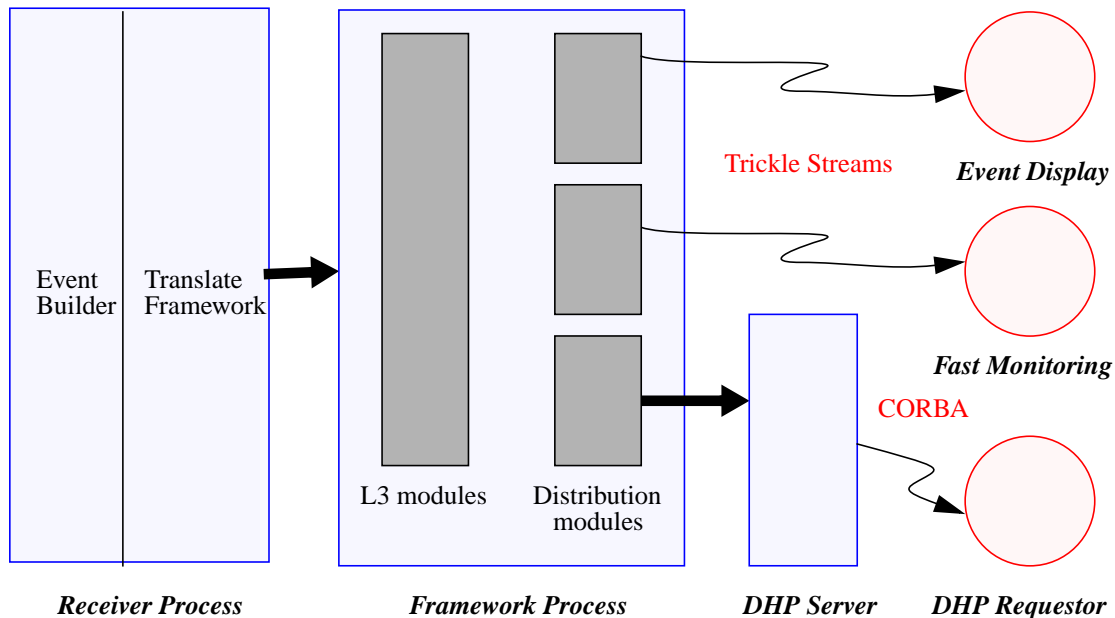
Intermediate Store (IS)

- Buffers events in the eventuality of down-stream failure
- Decouples Prompt Reconstruction from dead-time path

Logging Manager (LM)

- Collects events from multiple nodes and *writes* files to IS
- *Reads* from IS and fans out to Prompt Reconstruction farm

Online Event Processing



Responsibilities

- Supports the Level 3 Trigger (L3)
- Provides for Fast Monitoring and Event Displays
- Provides for Distributed Histogram service (DHP)

Hardware

- 333 MHz *Sun Ultra-5*, 512 MB memory, running *Solaris*
- 78 nodes installed, 32 used for L3, another 32 for Fast Monitoring (soon), 5 for calibration, 9 for subsystem tasks

Level 3 Trigger

Characteristics

- First part of the DAQ system to see complete events
- Runs on the output rate of the L1 (hardware) trigger
- Processes data from drift chamber (DCH) and calorimeter (EMC): fast tracking and clustering
- Logging decision: Decides what goes to “tape”

Design Requirements for $3 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$

- Reduce L1 output rate from 2 kHz to 120 Hz (including calibration samples)
- Provide high efficiency for B physics ($> 99\%$) and other physics processes, *e.g.*, charm, tau and two-photon physics
- Execute fast algorithms: < 10 ms per event (given 50 % CPU on 32 nodes)
- Support online calibration and monitoring

Strategy

- Define open trigger based on track or cluster topologies (want to minimize bias and stay open for New Physics)
- Orthogonality provides robustness and understandability
- Rejection of Bhabha events (>50 nb) with veto algorithm

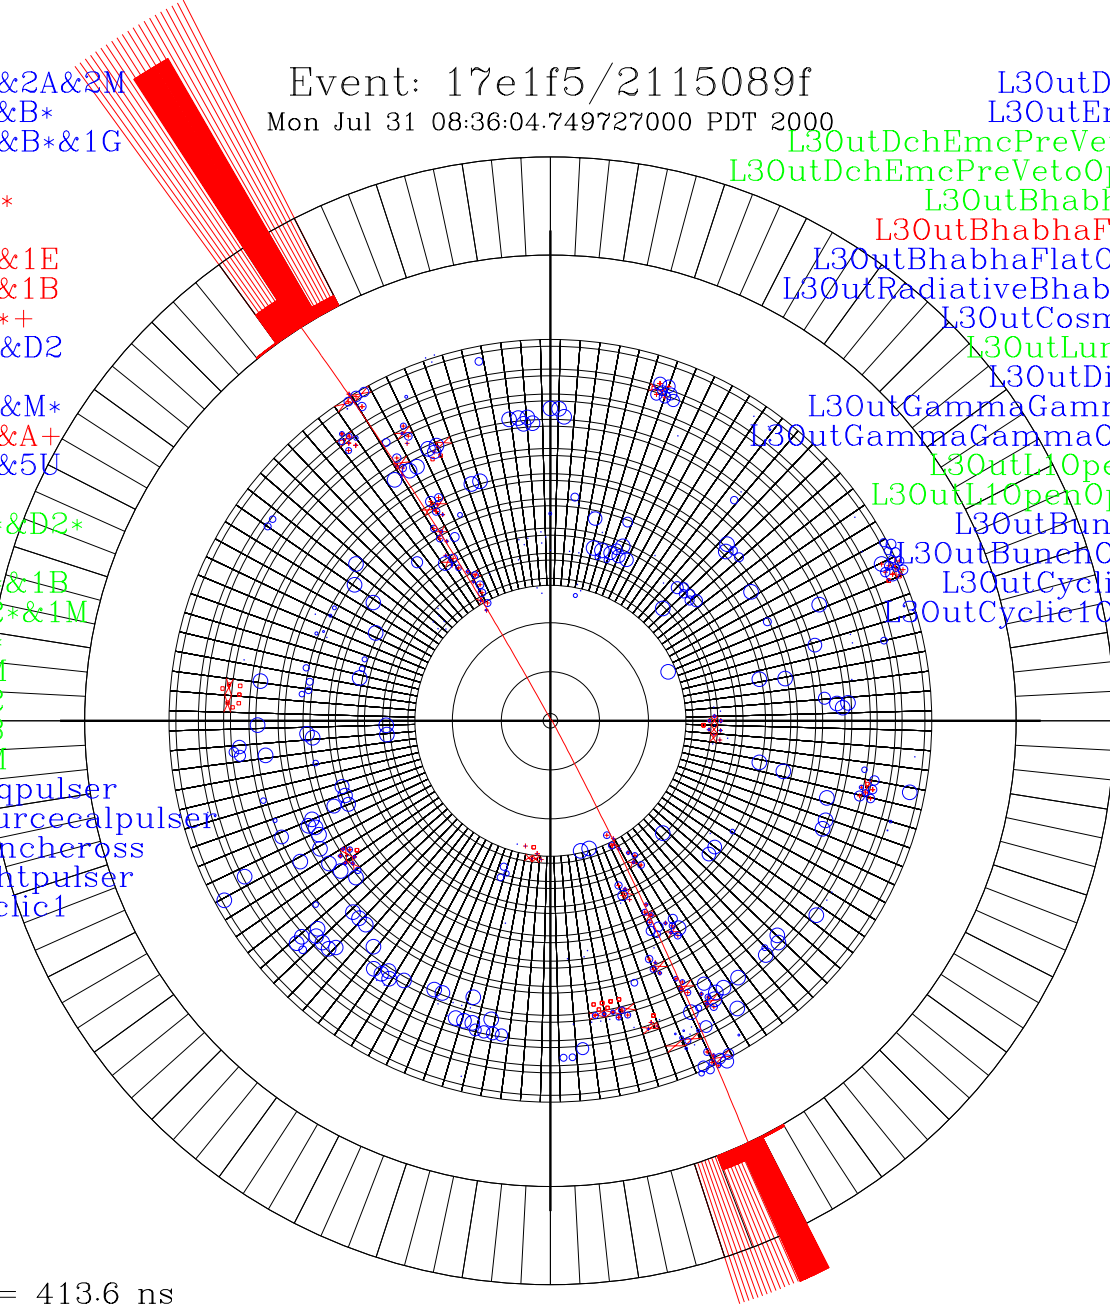
Bhabha Event

Event: 17e1f5/2115089f

Mon Jul 31 08:36:04.749727000 PDT 2000

- 0 3B&2A&2M
- 0 3A&B*
- 0 3B&B*&1G
- 1 2E
- 1 EM*
- 1 G*
- 1 D2&1E
- 1 1Y&1B
- 1 D2*+
- 0 3M&D2
- 0 4M
- 0 3M&M*
- 1 M*&A+
- 0 M*&5M
- 1 1Y
- 1 M*&D2*
- 0 3M
- 1 M*&1B
- 1 D2*&1M
- 1 M*
- 1 2M
- 1 D2
- 1 1B
- 1 1M
- 0 daqpulser
- 0 sourcealpulser
- 0 buncheross
- 0 lightpulser
- 0 cyclic1

- L3OutDch off
- L3OutEmc off
- L3OutDchEmcPreVeto pre
- L3OutDchEmcPreVetoOpr pre
- L3OutBhabha pre
- L3OutBhabhaFlat on
- L3OutBhabhaFlatOpr off
- L3OutRadiativeBhabha off
- L3OutCosmic off
- L3OutLumi pre
- L3OutDiag off
- L3OutGammaGamma off
- L3OutGammaGammaOpr off
- L3OutL1Open pre
- L3OutL1OpenOpr pre
- L3OutBunch off
- L3OutBunchOpr off
- L3OutCyclic1 off
- L3OutCyclic1Opr off



L3 T0 = 413.6 ns
 L1AT0 = 0 ns
 MC T0 = 0 ns

2 tracks, 2 clusters

Level 3 Trigger Performance

Processing Time

- Current algorithms take about 8.5 ms per event on average
- Dominated by tracking, processing of DCH data, clustering, processing of EMC data
- Time has grown several times with improvements to the reconstruction tools, mainly tracking resolution and efficiency
- Have so far managed to compensate with extensive code optimization (but this will saturate)

Rate Limit

- Have developed a simple model to estimate L3 processing rates from test runs with pulsers and relaxed L1 prescales
- At the limit, the L3 process is typically running at ~72 % CPU, ~11 % used by OEP receiver, rest spent in kernel.
- At 8.5 ms this translates into ~2.7 kHz (or 85 Hz per node)
- L1 fluctuations mandate a cushion of at least 20 % to avoid hitting the deadtime regime
- In factory mode operation, rather want to run comfortably at half of the maximum (or less)
- More detailed studies to characterize the system performance are under way (performed beam-time rate limit tests by running on fewer nodes last night)

Upgrade Options

Beyond 32 nodes

- Would allow a gradual increase of CPU power
- Needs work on the Dataflow side to remove event level limitation
- Brings risk of considerable penalty in operation and maintainance
- Have experienced bad scaling behavior when going to 32 (online databases)

Solaris vs Linux

- Will explore PC-based Linux nodes as an alternative in the coming weeks
- Have to resolve system incompatibilities:
Byte-order conversion (little->big endian),
Solaris system calls, semaphores, shared memory, IPC
- Would open an almost continuous range for choosing a cost/performance optimum (rather than the few discrete steps available with *Sun*)

Online Farm in 2001

Plan

- Having exceeded the design specs allows us to keep running beyond $3 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ in 2001
- Have decided not to upgrade the farm for next year
- Negotiated a fallback solution for a replacement of the *Sun Ultra-5's* with 32 *T1's* (may gain ~30 % CPU)
- Will use fallback, if beam conditions in the beginning of next year turn out much worse (vacuum scrubbing)

Consequences

- Increased luminosity will eat up headroom for L1
- Not much potential for doing a better job in L3
- Will need to focus on refinements and wait for an upgrade to be able to perform a significant step

Online Farm in 2002

Input Rates

- Rate budget will be driven by the success of the DCZ trigger upgrade
- If L1 upgrade won't be ready, will need at least a factor of 2 to cope with $1 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ (without L3 improvements)
- Otherwise may be able to defer upgrade until end of 2002, but see below
- Decision point is spring 2001

Output Rates

- At $1.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ leakage of the Bhabha veto and other QED events will make up 50 % of the logging rate
- Will impose a considerable burden for Prompt Reconstruction (35 % of *its* current CPU budget)
- Will even cause non-negligible tape costs for “garbage” events
- Would need another factor of 2 for significant improvements in the L3 physics performance
- Early upgrade seems reasonably balanced against downstream costs

Conclusions

At present

- In its first year the BaBar Online System (i.e. Trigger and DAQ) has met the challenge of taking data at record luminosities up to $2.5 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$
- At high operational efficiencies (>95 %) we manage to take 150 pb^{-1} of data per day
- Current L3 performance exceeds design specs of 2 kHz L1 and will allow us to run beyond design luminosity next year

In the future

- Need to re-generate headroom for L1 rate
- DCZ trigger upgrade can help for 2002 run
- Will definitely need to upgrade the online farm for 2003
- However, logging rates can warrant an earlier performance upgrade to be able to reduce downstream costs
- Still have to define rate requirements for $>3 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$