SVT as a Level 2 Processor

Bill Ashmanskas, L2 Trigger Review, 2001-12-07

- Synopsis
  - Is SVT or the SVT architecture suitable for making L2 decisions?
  - Are there L2 failure scenarios (plausible or implausible) for which the SVT architecture could provide a backup?
  - How realistic / how much effort?
- Outline
  - Background information
  - Past performance
  - Future returns?

- Why would you even consider SVT as a Level 2?
  - Answer: the data flow is similar
- Level 2 does this:
  - each L1A, receive data asynchronously from L1, XTRP, SVT, L2cal, RECES
  - preprocess data in parallel; fan in on 128 bit × 10MHz bus
  - reduce data to local L2 decision bit
  - deliver local L2 decision bit to TS
  - on global L2A, provide TL2D bank
  - provide scalers for accounting
- SVT does this:
  - each L1A, receive data asynchronously from L1, XTRP, SVX
  - process data in parallel; fan in on 23bit × 33MHz cables
  - deliver data to L2/Track board

- Some building blocks:
  - SVT cable: 23 bit LVDS @ 33 MHz, plus data strobe and flow control
  - Merger: 4×SVT in → 2×SVT out; respect event and packet boundaries; check BC match at end-event
  - GhostBuster: 3 × (SVT in → big Altera → SVT out)
    - has VME interface
    - sees CDF clock, L1A, etc. from J2
    - has connector to drive test data into L1/L2 interface
  - September kludge: SVT in → L2/TS handshake out
  - PCI/SVT: PCI bus ⇔ SVT in/out (+ L2/TS)
    - being assembled/tested by Franco Spinella
    - PCI interface via big, flexible Altera chip; e.g. can preprocess data, initiate transfer to/from PC memory
    - would allow a PC to act as an SVT board
  - Magical Mystery Board: Magic Bus ⇔ SVT in/out
    - principal use is as test stand data sink
    - has option to work with straight-through P3
    - proposed in October; now being prototyped
    - expect 2×PCB 12/10, 2 assembled boards 12/28

- In September, we extended SVT as follows, so that SVT commissioning and L2 commissioning could proceed in parallel
  - implemented GhostBuster board firmware to reduce final SVT data stream to a decision bit
  - built a kludge card to deliver that decision bit to TS
  - wrote a Python script to do begin-run download:
    - tell XTFA board which L1 bits to select
    - tell GB board what L2 cuts to apply (chi-square, impact parameter, curvature, number of tracks)
  - recorded per-track and per-event “accept” bits in SVTD
- If we had needed to develop this into something that could be turned over to the shift crew, we would have added this:
  - download XTFA/GB automatically via Run_Control
  - accumulate scalers in GB firmware
  - modify XTFA firmware to propagate all 64 L1 bits into SVT data stream (now maps 64 bits $\rightarrow$ 2 bits)
  - format decision data into TL2D in SVT crate CPU
  - L2 “algorithms” would be firmware based, with only L1 bits, XTRP tracks, and SVT tracks as input
    - $\Rightarrow$ better than L1 alone, but quite limited

- Timeline
  - Sept 8: proposed at L2 workshop that we allow L2 and SVT commissioning to proceed in parallel
  - Sept 11: Franco and Bill start work on SVT/TS kludge card and GhostBuster firmware
  - Sept 18: first SVT "cutting" run
  - Sept 25 - Oct 6: collected SVT data each store

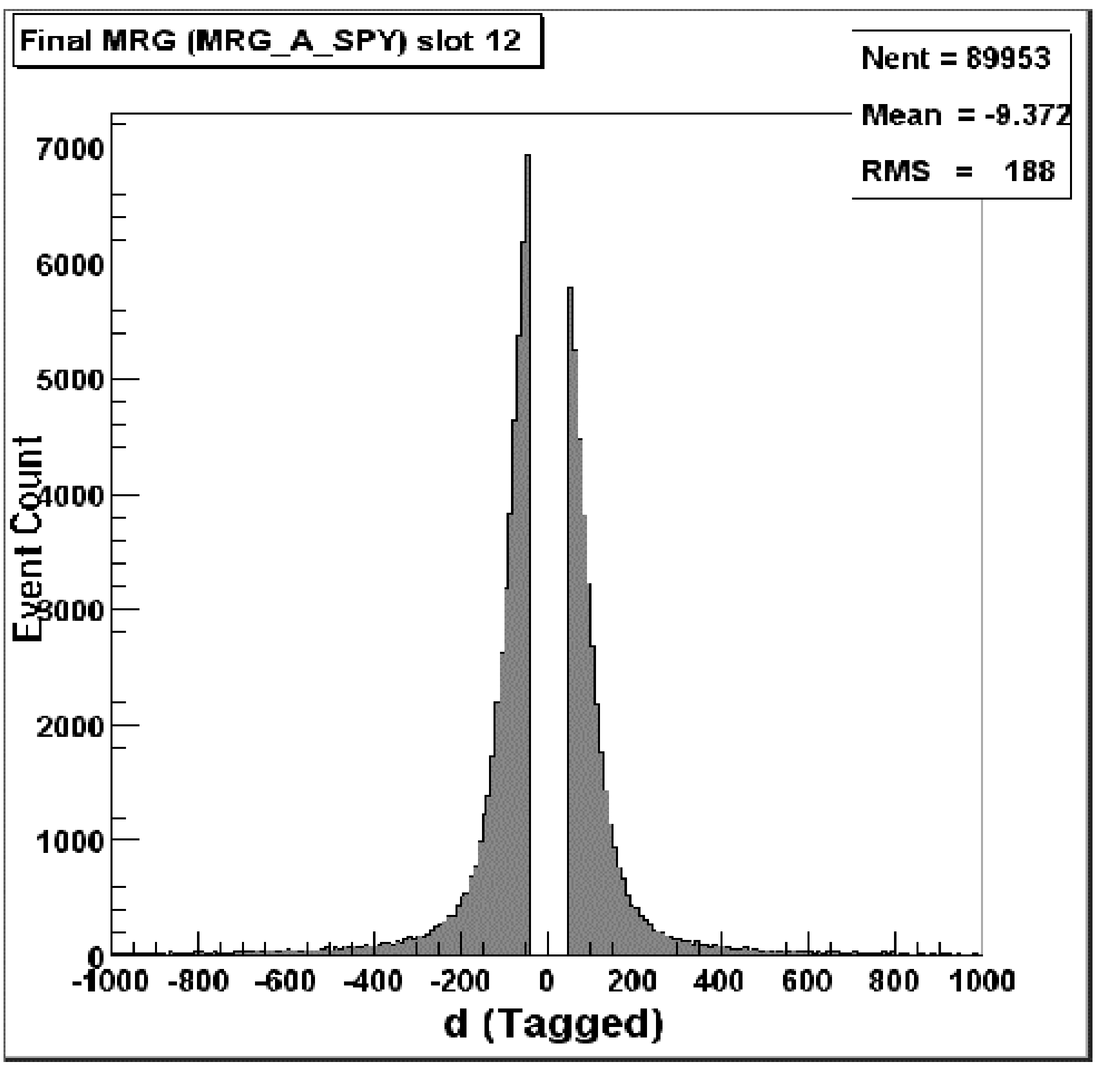

- Benefits
  - work on L2 decision crate was able to proceed on its own schedule, without distraction from SVT group
  - problem with SVX handling of L2 rejects was identified sooner
  - SVT-specific issues (e.g. beamspot correction, timing) were addressed sooner, because solution was needed sooner
  - collected useful data for SVT trigger studies

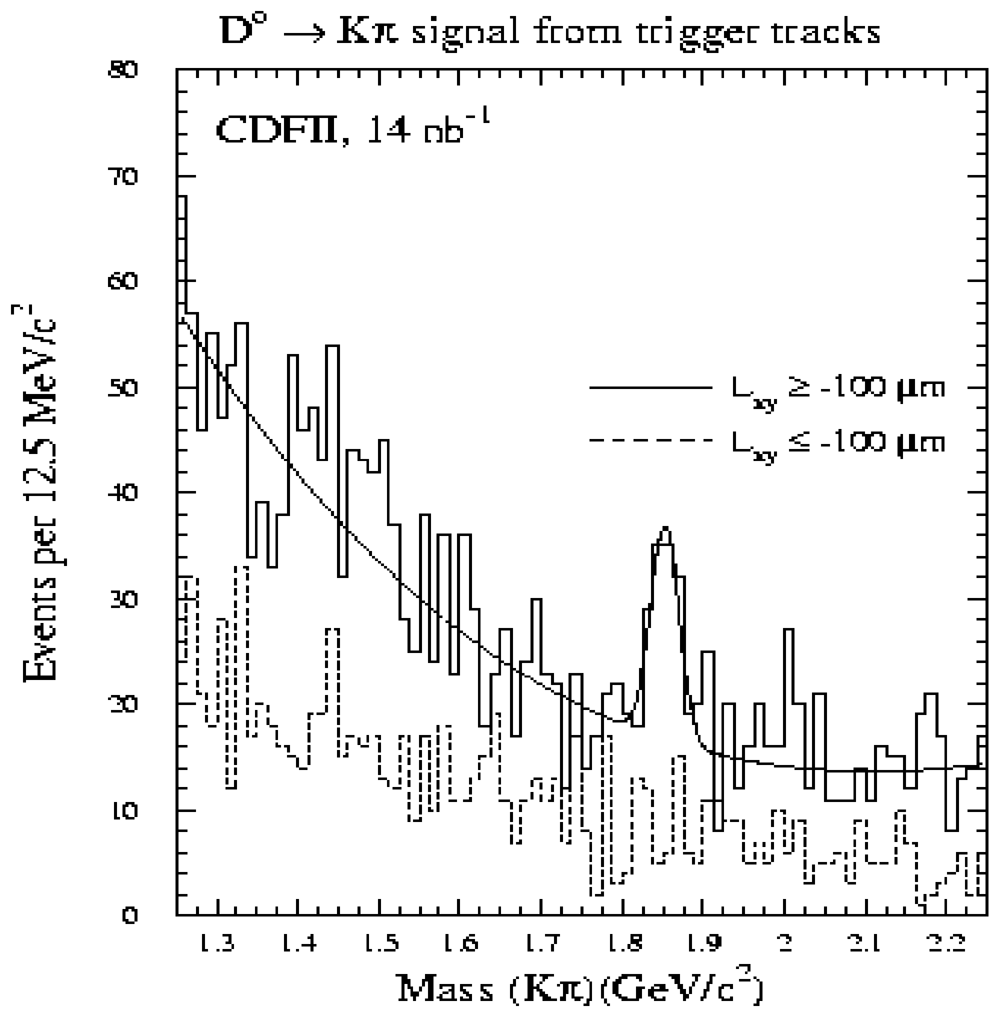

- Remarks
  - ability to commission subsystems in parallel is important for rapid progress
  - SVT hardware has tremendous flexibility
  - good building blocks can make both commissioning and improvisation much easier
  - it is not too hard to improvise ways to move trigger data around
  - moving the bits around is only part of the job; with workarounds, there is usually a price to pay in monitoring software, offline code, trigger tables, etc.
- With that in mind, I have been asked to consider ways in which SVT hardware could provide L2 fallback options
  - None of them come for free: effort needed to develop to state usable for physics data
  - I think the most realistic L2 that can be operating smoothly on January 11 is the baseline system, thanks to heroic L2 commissioning effort
  - To develop a backup that can be switched on with short notice, one needs to know in advance precisely what problem one is trying to insure against

- Scenario 1: SVT-based triggers go directly to TS
  - Task would be to make the pre-shutdown "SVT special run" capability usable for normal running
  - It is hard to imagine needing to run this in parallel with the Alpha (though it can be done):
    - TrackList board is the most aggressively tested data path into L2 (e.g. used for Magic Bus tests)
    - Ted's test stand should be ready for TrackList board in early Jan (TrackList will be used to test the test stand!)
    - Plenty of spare TrackList boards
  - It may still be useful as a simple way to do SVX/SVT rate tests
    - Can be used as-is for this
  - Could keep CDF running if L2 crate has a hardware failure that requires many hours to debug
    - Would need special trigger table
    - Most L1 triggers would be L2-auto; only simple track-based triggers at L2 (SVT, XTRP), firmware-coded
    - Few days' work to implement R_C auto-download
    - If you want L3 to use TL2D instead of L1 bits, then
      - 1 week effort by Simone Donati on XTFA firmware
      - 1 week effort by me to implement scalers, fake TL2D
    - Additional firmware effort if algorithms do more than count SVT tracks passing simple cuts

- Scenario 2: Interface boards work, Alpha works, Magic Bus only works at low rate, no arbitration
  - Using MMB, can divide MB traffic onto 1 slow bus + 3 single-slot "buses"
  - One Magic Bus (maybe a D0-style short bus?) contains L1 board (single word per event), 3 RECES boards (slow, programmed I/O), and one MMB
  - A straight-through slot (which exists e.g. on TDC backplane or an SVT backplane) connects CLIST to a second MMB
  - A straight-through slot connects ISOList to a third MMB
  - These three SVT-cable data streams, plus XTRP and SVT, are merged into one SVT cable (95MB/s bandwidth) either with a Merger board or by daisy-chaining MMBs
  - A fourth MMB sends merged data stream into a single Alpha on a straight-through slot
  - I estimate that this would take the full month of January to implement, assuming MMB works on arrival (we would also need to stuff the other two MMB PCBs, and we'd lose our test stand)
  - Do recent Magic Bus tests rule out the need for this?

- Scenario 3: Everything works, but (for whatever reason) there are too few spare Alphas
  - On a ~ few month time scale, a system based on commodity PCs could be prototyped
  - A MMB could control a Magic Bus in a test stand, or spy on the Magic Bus in a running L2 crate
  - A copy of L2 data flow is transmitted to SVT cable (on-board chip can buffer 13KB, well over the typical ~500 bytes/event; average data flow is ~20MB/sec at 40kHz)
  - SVT/PCI board would get data into a PC for evaluation; data to be read into a bank would be sent on SVT cable into a TrackList or GB board
  - Significant development effort would be needed:
    - Implementing MMB data capture
    - Testing SVT/PCI board; writing device driver and firmware
    - Handling operating system and real-time issues to keep latency low
    - Porting L2 algorithm code, R_C download code