MODSIM for CDFII Trigger/DAQ

- Code written and ran by Dr. Henry Kasha
- Previous results presented:
  - CDF notes 2306, 3495, 4213
    - # of L2 & scan buffers, pipelined L2, abort gaps
  - DAQ meeting: October 12, 2000
    - silicon digitization time (7 or 8 bits)
  - Previous studies aimed at `full' performance
- `Real' time simulation
  - C++ based
  - State machines, queues, etc.

Goals for today:

- Explain briefly what we have in place now
  - Clarify the meaning of the available inputs
  - Clarify the present rules
- Discussion
  - How well does the code correspond to reality?
  - Understand if and define how to change
    - the rules
    - the inputs
- … then we can turn to comparisons with benchmark runs and other studies, etc.
- And hopefully help provide some insights

Input parameters

- L1 accept rate:
  - raw FRED rate (i.e. unaffected by dead-time)
  - example: 10 kHz (appropriately randomized)
- Number of L2 buffers: 4
- L2 accept rate:
  - actually sets the L2 acceptance factor
  - example: 300 Hz => 3% are accepted
- Level 2 load time:
  - [4 - 15, 35 - 41, 45 - 65, 88 - 93]
    - [ID's in PJW L2 spreadsheet from the L2 review]
    - example: 10 μsec
- Level 2 processing time:
  - [96 - 104]
  - example: 20 μsec + exponential tail of 2 μsec

Silicon and SVT

- Silicon digitization time
  - [12 - 16, 20 - 21]
  - example: 10 μsec
- Silicon readout time: includes both r-ϕ & r-z
  - [22 - 24]
  - example: 5.5 μsec + exponential tail of 0.5 μsec
- SVT processing time
  - additional time after r-ϕ readout
  - to finish processing & get the silicon information to L2
  - [69 - 85]
  - example: 13 μsec

For L2 Accepts

- Scanning time
  - = readout time = DONE handshake time
  - example: 900 μsec + exponential tail of 100 μsec
- Number of scan buffers
  - where the event sits in VRBs awaiting HEVB
  - default: 4
- Switching time
  - latency for moving an event out of VRB into L3
  - example : 950 μsec + exponential tail of 50 μsec

Silicon Rules

- The silicon process (handles one L2 buffer at a time)
  - a) Check if there is a full L2 buffer
  - b) Digitize
  - c) When digitization is complete, begin the silicon readout
    - silicon readout can be aborted (by L2R; after digitization)
  - d) When readout is completed (or aborted) go back to a)
- The capacitor is returned to the pipeline after digitization
  - as soon as a (mini-)abort gap is seen
  - and send L1Done signal to TS
  - this is handled quasi-asynchronously wrt digitization/readout
- Simulation treats all of the silicon as one entity
  - knows nothing about differences between L00, SVXII, ISL
  - nominally assumes r-ϕ and r-z are each ½ of the readout

Level 2 rules

- Loading:
  - a) Check if there is a FULL L2 buffer
  - b) Wait the time for loading
  - c) Update the queue for L2 processing
  - d) Go back to a)

- This aspect of the loading is meant to cover everything but the silicon
  - Simulation assumes SVT is not held up waiting for information from XFT
- Strict FIFO ordering based on L1A assumed

Level 2 rules

- Processing:
  - a) Check if there is an event in the queue to process
  - b) Check to see if silicon/SVT information is available
    - wait for digitization + r-ϕ silicon readout + SVT processing
  - c) Wait the `L2 processing time'
  - d) Check if the L1Done condition is satisfied
  - e) `Process' the event
    - For rejects
      - send L2 Reject message (relevant to SRC)
      - free the L2 buffer (FULL --> EMPTY)
    - For accepts
      - promote the L2 buffer (FULL --> PENDING)
  - f) Go back to a)

L1DONE Rules

- 1. #L2 Decisions - #L1Dones <= 1
  - TS requires <=0
    - i.e. must get L1Done back for current buffer
  - In principle this rule could be just for L2A
    - If SRC could look ahead at queued L2 Rejects and return the capacitor prior to digitization
    - TS would need to treat L2A & L2R separately
    - In principle a help if L2R doesn't depend on SVT
- 2. #L1 Accepts - #L1Dones < 4
  - No FREE L2 buffer unless this condition holds
- Can't evade these rules: SVX3 chip is FIFO

Minimum Time for L2 Rejects

- Some systems require a minimum time between L1A and L2R
  - TDCs (2 μsec)
  - SMXR (6 μsec)
  - The TS has a fixed minimum time in place for testing with auto (TS) L2R or alternating L2A/L2R
    - Check with Andrew if this time is `added' in when L2 is making the decisions

Rules after Level 2 Accept

- Scanning/Readout
  - a) Check if there is a pending L2 buffer
  - b) Check if there is an empty scan buffer
  - c) Move the event from the L2 buffer to the scan buffer
    - PENDING → BUSY
    - this interval set by the DONE handshake time
  - d) Update the input queue for the HEVB
  - e) Free the L2 buffer
    - BUSY --> EMPTY
  - f) Go back to a)

- Switching (this is a (very) poor man's version of HEVB)
  - a) Check if there is an event in the switching queue
  - b) Move the event (this interval is the switching time)
  - c) Go back to a)

Output information

- Number of crossings
- Number of FRED Level 1 Accepts
- Number of events
  - lost (no buffer)
  - processed by L2
  - accepted by L2
  - scanned
  - switched
- Dead-time:
  - L2, Readout, L2-or-Readout , BUSY/WAIT, TS bookkeeping
- L2 and VRB buffer state occupancies
  - How many buffers are EMPTY, FULL, PENDING, & BUSY

Known to be missing in simulation:

- TL2D formation time
  - Longer for L2A than for L2R
- Mixed trigger table (L2R without SVT)
- Full L1-Done handling (2nd rule)
- `Readout Lists' (including RECES)
- Your favorite
  - Tunable parameter
  - Bit of output